Proceedings of the 2002 Winter Simulation Conference
E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, eds.

SIMULATION OPTIMIZATION

Sigurdur Ólafsson
Jumi Kim

Department of Industrial and Manufacturing Systems Engineering
Iowa State University
2019 Black Engineering
Ames, IA 50011, U.S.A.

ABSTRACT

Simulation optimization has received considerable attention from both simulation researchers and practitioners. In this tutorial we present a broad introduction to simulation optimization and the many techniques that have been suggested to solve simulation optimization problems. Both continuous and discrete problems are discussed, but an emphasis is placed on discrete problems and practical methods for addressing such problems.

1 INTRODUCTION

Many systems in areas such as manufacturing, supply chain management, financial management, are too complex to be modeled analytically. Discrete event simulation has long been a useful tool for evaluating the performance of such systems. However, a simple evaluation of performance is often insufficient and a more exploratory process may be needed in the form of simulation optimization. Simulation optimization is the process of finding the best values of some decision variables for a system where the performance is evaluated based on the output of a simulation model of this system. There has been a great deal of work on simulation optimization in the research literature, and more recently optimization routines has been incorporated into several commercial simulation packages.

Techniques for simulation optimization vary greatly depending on the exact problem setting. In this tutorial we take the underlying structure of the decision variables, that is discrete or continuous, to be the primary distinguishing factor. There also appears to be a significant gap between those methods that have been studied extensively in the research literature and those that are commonly use in practice. In this tutorial we survey methods used for both continuous and discrete optimization and discuss which of these methods have been successfully implemented as part of commercial software packages. Previous review of simulation optimization include Andradóttir (1998), Azadivar (1999), Swicher et al. (2000), and April et al. (2001).

The remainder of the paper is organized as follows. In Section 2 we establish a common framework for simulation optimization problems and present the notation to be used. Section 3 surveys techniques for optimizing continuous decision variables, and Section 4 does the same for discrete decision variables. In Section 5 we discuss simulation optimization software, and finally Section 6 contains some concluding remarks.

2 PROBLEM SETTING

Simulation optimization is optimization where the performance is the output of a simulation model, and the problem setting thus contains the usual optimization components:

- Decision variables,
- objective function, and
- constraints.

We denote the decision variables as θ and the constraints are represented by these variables having to be contained in some feasible region Θ, that is $\theta \in \Theta$. The objective function is a real valued function defined on these variables $f : \Theta \to \mathbf{R}$, but due to the complexity and stochastic nature of the underlying system an analytical expression does not exit for $f(\cdot)$ and it must be estimated using a function of the stochastic simulation output, say $X(\theta)$, that we write as a function of the decision variables. Typically, this might be an unbiased estimate of the true objective function, that is $f(\theta) = E[X(\theta)]$, but we will not be concerned with this here nor how this performance output relates to the simulation output variables.

Various simulation optimization techniques can be classified based on the nature of the feasible region. If it is a continuous set, that is $\Theta \subset \mathbf{R}^n$, then it may be appropriate to use a gradient based search method such as stochastic approximation. If it is finite and fairly small, say

$\Theta = \{\theta_1, \theta_2, ..., \theta_m\}$, where $m < 30$ then it is possible to use ranking and selection methods, whereas if it is finite but combinatorially large a metaheuristic may be appropriate. Those and other methods are surveyed in the following sections.

3 CONTINUOUS DECISION VARIABLES

We start by considering the situation where the underlying variables are continuous, that is $\Theta \subset \mathbf{R}^n$ is uncountable and infinite. This is perhaps the most studied problem setting in the research literature but we will only briefly consider those here as our main focus is on discrete methods.

3.1 Stochastic Approximation

Stochastic approximation (SA) is the iterative process of moving from one solution to another based on moving in the direction an estimate of the gradient. This process is analogous to the steepest descent gradient search in nonlinear optimization, but as there is no analytical expression for the objective function there is of course also no such expression for the gradient.

In mathematical notation, the general step of a SA algorithm proceeds as follows. Let $\theta^{(k)}$ be the current solution for the decision variables, the basic SA algorithm is of the following form (for minimization):

$$\theta^{(k+1)} = \Pi\left(\theta^{(k)} - \alpha_k \hat{\nabla} f\left(\theta^{(k)}\right)\right).$$

Here $\hat{\nabla} f\left(\theta^{(k)}\right)$ is an estimate of the gradient, α_k is the step size, and Π is a projection onto the feasible region, $\Pi\colon \mathbf{R}^n \to \Theta$.

Extending back to the fundamental work of Robbins and Monro (1951) and Kiefer and Wolfowitz (1952), stochastic approximation has received a great deal of attention. The asymptotic convergence of SA can be guaranteed under certain conditions, which typically involve letting the step size go to zero but at a sufficiently slow rate, e.g. $\lim_{k\to\infty} \alpha_k = 0$ but $\sum_k \alpha_k = \infty$.

The simplest way to estimate the gradient $\hat{\nabla} f(\theta) = \left[\hat{\nabla} f(\theta_1), ..., \hat{\nabla} f(\theta_n)\right]$ is by looking at some small change $\Delta\theta_i \in \mathbf{R}$ in each of the decision variables and using the finite differences, either one-sided

$$\hat{\nabla} f(\theta_i) = \frac{X(\theta_i + \Delta\theta_i) - X(\theta_i)}{\Delta\theta_i}$$

or two-sided

$$\hat{\nabla} f(\theta_i) = \frac{X(\theta_i + \Delta\theta_i) - X(\theta_i - \Delta\theta_i)}{2\Delta\theta_i},$$

$i = 1, 2, ..., n$. The one-sided estimate requires $n + 1$ simultaneous simulations of the performance measure, and the two-sided estimate requires $2n$ such simulations, a considerable computational effort when optimizing complex systems.

The computational efficiency and convergence properties of SA can be dramatically improved with a direct estimate of the gradient instead of finite differences and a great deal of research has been devoted to developing such methods. The most common types of approaches for such direct estimation are perturbation analysis (Glasserman, 1991, Ho and Cao, 1991) and likelihood ratio (Glynn, 1989; Rubinstein, 1991, Rubinstein and Shapiro, 1993).

3.2 Other Methods

Although gradient search, and in particular those based on stochastic approximation, appear to have received the most attention in continuous simulation optimization literature, several alternatives have been suggested. Here we will only briefly mention a couple of these alternatives.

One such approach is the sample path method presented by Gurkan et al. (1994). The basic idea here is to fix one particular sample path, at which point the problem becomes deterministic and the powerful machinery of mathematical optimization can be applied directly. The method then iteratively moves towards an optimal solution by considering one sample path at time. Under certain conditions, this approach can be shown to converge almost surely (Robinson, 1996).

Response surface methodology (RSM) is a well studied statistical approach that attempts to find a functional relationship between the input variables and the output function. This can be applied to simulation optimization with the output function being the stochastic output of simulation model (Kleijnen, 1998). More recently low cost response surface methods have been proposed for simulation optimization (Allen and Yu, 2000).

4 DISCRETE DECISION VARIABLES

When the feasible region is countable or countable finite then the methods of the previous sections do not usually apply, although a few efforts have been made to apply for example SA to discrete problems. In this case, we need to distinguish between problems where the feasible region is small and a complete enumeration is possible and problems where it is impossible to evaluate every alternative and some search method must be included to determine which solutions should be evaluated.

4.1 Statistical Selection

When the optimization involves selecting the best of a few alternatives, that is $\Theta = \{\theta_1, \theta_2, ..., \theta_m\}$, where m is relatively small, then it may be possible to evaluate every solution and compare the performance. In the deterministic context this would be straightforward, but since the performance must be estimated based on the stochastic simulation output some further analysis is needed to compare alternative solutions. Numerous different approaches have been developed to address this problem, including subset selection, indifference-zone ranking and selection (R&S), multiple comparisons procedures (MCP), and decision theoretic methods.

Subset selection or screening does not attempt to find the optimal solution but simply to reduce the feasible region to a (small) subset of solutions. Early work developed techniques that apply when the simulation output is normal with common variance and the same number of simulation observations are used for each solution (Gupta, 1956; 1965). These assumptions are rarely satisfied for simulation outputs and although many methods are quite robust with respect to the normality assumption, the assumption of common variance is quite restrictive and new techniques have been developed that do not require this assumption (Nelson et al., 2001).

The subset selection procedures do not find a single best (optimal) solution, but this can be accomplished using R&S methods. The most common approach is to define an indifference zone δ for the performance and develop a procedure that selects a solution with performance that is within δ units of the optimal performance with a given probability, that is, if θ^* is the optimal solution and θ is the selection solution then

$$\text{Prob}\left[\left|f(\theta) - f(\theta^*)\right| < \delta\right] \geq 1 - \alpha,$$

where $1 - \alpha$ is the desired probability. To achieve this guarantee, a two-stage procedure that prescribes how many simulation estimates are needed for each alternative is commonly applied (Dudewicz and Dalal, 1975; Rinott, 1978). A discussion of alternative indifference-zone procedure can be found in standard simulation texts such as Law and Kelton (2000). As they have complimentary functions, a natural approach is to combine subset selection for screening with R&S for selection of a specific solution (Nelson et al., 2001).

Another approach to selecting the best solution are MCPs that calculate simultaneous confidence intervals for $f(\theta_i) - f(\theta^*)$, $i = 1, ..., m$, where θ^* is as before the optimal solution (Hochberg and Tamhane, 1987). These procedures are actually closely related to the R&S procedures as indifference-zone procedures can automatically provide such confidence intervals with the width of the interval corresponding to the selected indifference zone (Matejcik and Nelson, 1995; Nelson and Matejcik, 1995).

Most of the statistical selection procedures mentioned above involve a two stage process where in the first stage the mean and the variance of each solution is estimated and those estimates used to determine how many more simulations are needed to make the desired selection. In implementing such methods a key issue is how much effort to put into the first stage. If it is too little an inaccurate estimate may prescribe much more simulation in the second stage than is really needed and vice versa too much effort in the first stage may spend more simulation time on each solution than was needed. More recently there has been considerable effort devoted to developing sequential procedures that solve this problem and these methods have been found to perform very favorably to the sequential procedures (Kim and Nelson, 2001; Chen et al., 1997).

A completely different approach to selecting the best system is the decision theoretic framework of Chick and Inoue (1999, 2000) that uses Baysian analysis to develop both two-stage and sequential statistical selection procedures.

4.2 Random Search

When it is not possible to evaluate every solution using a statistical selection procedure, some procedure must be used to determine which solutions are to be considered and simulated. This is most often some type of a random search approach.

Random search typically involves an iterative process where in each iteration the search progresses to a new (possibly better) solution in the neighborhood of the current solution.

0. Select and initial solution $\theta^{(0)}$ and simulate its performance $X\left(\theta^{(0)}\right)$. Set $k = 0$.
1. Select a candidate solution θ^c from the neighborhood $N(\theta^{(k)})$ of the current solution and simulate its performance $X(\theta^c)$.
2. If the candidate θ^c satisfies the acceptance criterion based on the simulated performance, then let $\theta^{(k+1)} = \theta^{(k)}$; otherwise let $\theta^{(k+1)} = \theta^{(k)}$.
3. If stopping criterion is satisfied terminate the search; otherwise let $k = k + 1$ and go back to Step 1.

The various random search methods that have been proposed in the literature can be though of as specifying the neighborhood structure, how to select a candidate, the acceptance criterion, and the stopping criterion.

4.3 Metaheuristics

Various metaheuristics have been suggested for simulation optimization. Such methods include genetic algorithms,

simulated annealing, tabu search, and neural networks. Although these methods are generally designed for combinatorial optimization in the deterministic context and may not have guaranteed convergence, they have been quite successful when applied to simulation optimization.

Simulated annealing (SA) can be though of within the framework of the random search describe above and can be adapted for simulation optimization (Haddock and Mittenhall, 1992). Starting with an initial solution, SA moves from one solution to the next, hopefully converging on the global optimum. All such random search methods may, however, get stuck at a locally optimal solution, and SA attempts to rectify this by accepting inferior solutions with certain probability and thus allowing the search to escape local optima. Thus, the main innovation of the SA approach is in Step two of the generic algorithm, where in the k-th iteration a randomly selected candidate θ^c is accepted with the probability (assuming a minimization problem):

$$Prob\left[\text{Accept } \theta^c\right] = \begin{cases} 1, & L\left(\theta^c\right) < L\left(\theta^{(k)}\right) \\ e^{\frac{L(\theta^c) < L\left(\theta^{(k)}\right)}{T_k}}, & \text{otherwise.} \end{cases}$$

In other words, the candidate solution is always accepted if it is better but it is also sometimes accepted even if it is inferior. The probability of accepting the inferior candidate is higher if the difference in performance is small, and it is higher if the constant T_k, called the temperature, is high. Usually, this temperature is allowed to decrease as the search progresses, the idea being that after a while no big moves up hill should be allowed and eventually no moves should be made to an inferior solution. However, in the context of simulation optimization there are indications that a constant temperature search may work as well or better (Alrefaei and Andradóttir, 1998).

Tabu search can also be placed within the framework of general random search (Glover and Laguna, 1997). One of the unique features of this approach is a restriction of the neighborhood $N(\theta^{((k)})$ of the current solution $\theta^{(k)}$ as certain solutions are made tabu. Specifically, solutions are tabu if they require the reverse move of a recently made move, which forces the search to continue when it might otherwise get stuck at a local optimum. Although maintaining a list of tabu moves may be considered the main feature of the method, it has numerous other properties. This includes for example long term memory that allows the search to restart at a previously found good solution with a new list of tabu moves that forces a different search direction from this good starting point. More details on using tabu search for simulation optimization can be found in Glover, Kelly, and Laguna (1999) and April et al. (2001).

Genetic algorithms (GA) and other evolutionary methods are again similar to the generic random search but work with a population of solution rather than a single solution. Thus, in the k-th iteration, a point $\theta^{(k)}$ is actually a set of solutions and the neighbors of $\theta^{(k)}$ are constructed by operating on these solutions jointly. The most common such operators are cross-over and mutation. The cross-over operation typically takes two solutions from the set $\theta^{(k)}$ that have relatively good performance and combine them to make two new solutions. This is meant to resemble an evolutionary process where two fit individuals are allowed to reproduce to generate offspring that resemble the parents. The mutation operator, on the other hand, takes a single high performing solution and alters it slightly. From this it should be clear that the main innovative contribution of GA when placed in the context of general random search is a novel construction of a neighborhood based on natural selection principals.

Another random search metaheuristic is the nested partitions (NP) method of Shi and Ólafsson (1997). This method takes a global approach to simulation optimization and generates iterative partitions of the entire feasible region. That is, in the k-th iteration there is some subset $\sigma(k) \subseteq \Theta$ that is considered the most promising ($\sigma(0) = \Theta$), and the method attempts to narrow the search be looking at subsets $\sigma_i(k) \subset \sigma(k)$, $i = 1, 2, ..., M$ of this region while simultaneously also looking at the surrounding region $\Theta \setminus \sigma(k)$. Thus, it focuses the computational effort while simultaneously maintaining a global perspective. If one of the subsets, say $\sigma_l(k)$, is found to be best this becomes the most promising region in the next iteration ($\sigma(k+1) = \sigma_l(k)$), but if the surrounding region is found to be best the method backtracks ($\sigma(k+1) = \sigma(k-1)$). In terms of the general random search framework, each point $\theta^{(k)}$ thus corresponds to a subset $\sigma(k)$ and the neighbors are either subsets constructed according to a specified partitioning method or the subset that was partitioned to created $\sigma(k)$:

$$N\left(\theta^{(k)} = \sigma(k)\right) = \{\sigma_1(k), ..., \sigma_M(k), \Theta \setminus \sigma(k)\}.$$

Eventually, the subset become singletons and by assuring that the correct move is made with a given probability in each iteration, it is possible to guarantee that when this happens a sufficiently good solution has been found with a satisfactory high probability (Ólafsson, 1999; Ólafsson and Kim, 2001). Statistical selection methods, such as those reviewed in Section 4.1, can be used to determine the amount of sampling needed from each region to assure a proper selection in each iteration. Specifically, given an indifference zone δ and a probability of eventually selecting a singleton that has performance within δ units of the optimal performance is given, the probability by which the correct selection in each iteration must be made is calculated. Given this probability the statistical selection procedure is used to prescribe how many solutions are sampled from each of the subregions and the surrounding region.

The methods described in this section have all been successfully applied to industrial problems, and many have been incorporated into standard simulation software packages. This issue of simulation optimization in practice will be discussed next.

5 SIMULATION OPTIMIZATION IN PRACTICE

Recently there has been considerable research focused on how to combine simulation and optimization in practice (Fu et al., 2000; April et al., 2001; Ólafsson and Kim, 2001; Fu, 2002). Although simulation optimization has been an active area of research for considerable length of time, then except for statistical selection methods that simply compare all alternatives, optimization packages have only been incorporated into commercial simulation software in the last decade. Examples of such optimization packages include ProModel's SimRunner (Harrel and Price, 2000) and AutoMod's AutoStat (Bitron, 2000) that use evolutionary and genetic algorithms, SIMUL8's OPTIMIZ that uses neural networks, and OptQuest package, which works with simulation software that includes Arena and Crystal Ball, and uses scatter search, tabu search, and neural networks (Glover, Kelly, and Laguna, 1999).

As one can see from the examples above, commercial simulation optimization packages that do search in addition to comparison of solutions are currently dominated by metaheuristic approaches. Thus, in simulation optimization practice, such methods appear to take precedence over other methods that have received more attention by the academic research community and may have more appealing convergence properties. The reasons for this are undoubtedly multiple, but some explanation include that convergence properties such as asymptotic convergence have limited relevance in practice, and the metaheuristics are generally fast, robust, and generate multiple alternative solutions while focused on finding the optimal solution.

However, we believe that combining the robustness of metaheuristics with the established methods for guaranteeing performance is valuable and will be integral part of future simulation practice. There are several ways in which one can imaging achieving this goal. For example, Boesel, Nelson, and Ishii (1999) and Boesel, Nelson, and Kim (2001) combine genetic algorithm for searching for solutions with rigorous statistical selection for comparing the solutions that are generated. The solutions generated by the GA search are first screened using subset selection and then the best solution is selected using a two-stage indifference-zone procedure. A different approach incorporating the same tools is taken by Ólafsson (1999) that uses a two-stage indifference-zone procedure to guarantee that the nested partitions random search method converges with a fixed probability to within an indifference zone of the optimal solution. Note that in addition to providing a convergence guarantee this relaxes the goal of the optimization from finding the optimal solution to finding a 'good enough' solution with a given probability, a goal that may be more reasonable in practice. Similar goal softening is also part of the ordinal optimization paradigm (Ho et al., 1992; Ho et al., 2000). Finally, using the nested partitions framework, metaheuristics such as GA and tabu search can be incorporated to speed the search, while simultaneously retaining the performance guarantees (Ólafsson and Kim, 2001). For much more discussion on simulation optimization theory and practice see Fu (2002).

6 CONCLUSION

Simulation optimization is an active field of research and is also increasingly being used in practical simulation applications and being incorporated into simulation software tools. In this tutorial we have given a broad overview of simulation optimization but with an emphasis on problems with discrete decision variables.

As should be apparent, there is still somewhat of a gap between the academic work on simulation optimization, which historically has focused on gradient-based approaches and convergence proofs, and practical implementations of simulation optimization, which primarily implement metaheuristics. This difference has, however, been widely recognized and we have indicated some new effort in bridging the gap.

REFERENCES

Allen, T. and L. Yu. 2000. "Low Cost Response Surface Methods for and from Simulation Optimization," in *Proceedings of the Winter Simulation Conference*, 704-714.

Alrefaei, M.H. and S. Andradóttir. 1998. "A Simulated Annealing Algorithm with Constant Temperature for Discrete Stochastic Optimization," technical report, Department of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA.

Andradóttir, S. 1998. "A Review of Simulation Optimization Techniques," in D.J. Medeiros, E.F. Watson, J.S. Carson, and M.S. Manivannan (eds.), *Proceedings of the 1998 Winter Simulation Conference*, 151-158.

April, J., F. Glover, J. Kelly, and M. Laguna. 2001. "Simulation/Optimization using "Real-World" Applications," in *Proceedings of the 2001 Winter Simulation Conference*, 134-138.

Azadivar, F. 1999. "Simulation Optimization Methodologies," in *Proceedings of the 1999 Winter Simulation Conference*, 93-100.

Bitron, J. 2000. "Optimizing AutoMod Models with AutoStat," *AutoFlash*, February 2000.

Chick, S.E. and K. Inoue. 1999. "A Decision-Theoretic Approach to Screening and Selection with Common

Random Numbers," in *Proceedings of the Winter Simulation Conference*, 603-610.

Chick, S.E. and K. Inoue. 2000. "New Results on Procedures that Select the Best System using CRN," in *Proceedings of the Winter Simulation Conference*, 603-610.

Dudewicz, E.J. and S.R. Dalal. 1975. "Allocation of Observations in Ranking and Selection with Unequal Variances," *The Indian Journal of Statistics*, 37B, 28-78.

Fu, M.C. 2002. "Optimization for Simulation: Theory vs. Practice," *INFORMS J. on Computing*, to appear.

Fu, M.C., S. Andradóttir, J.S. Carson, F. Glover, C.R. Harrell, Y.C. Ho, J.P. Kelly, and S.M. Robinson. 2000. "Integrating Optimization and Simulation: Research and Practice," in *Proceedings of the 2000 Winter Simulation Conference*, 610-660.

Glasserman, P. 1991. *Gradient Estimation via Perturbation Analysis*, Kluwer Academic Publishers, Boston, MA.

Glover, F., J.P. Kelly, and M. Laguna. 1999. "New Advances for Wedding Optimization and Simulation," in *Proceedings of the 1999 Winter Simulation Conference*, 255-260.

Glover, F. and M. Laguna. 1997. *Tabu Search*, Kluwer Academic Publishers, Norwell, MA.

Gurkan, G. A.Y. Ozge, and S.M. Robinson. 1994. "Sample-Path Optimization in Simulation," in the *Proceedings of the 1994 Winter Simulation Conference*, 247-254.

Haddock, J. and Mittenhall. 1992. "Simulation Optimization using Simulated Annealing," *Computers and Industrial Engineering*, 22, 387-395.

Harrel, C.R. and R.N. Price. 2000. "Simulation Modeling and Optimization using Promodel," in *Proceedings of the 2000 Winter Simulation Conference*, 197-202.

Ho, Y.C. and X.R. Cao. 1991. *Perturbation Analysis of Discrete Event Dynamic Systems*, Kluwer Academic Publisher, Norwell, MA.

Ho, Y.C., C.G. Cassandras, C.H. Chen, and L.Y. Dai. 2000. "Ordinal Optimization and Simulation," in *Journal of Operational Research Society* **51**, 490-500.

Ho, Y.C., R. Screenivas, and P. Vakili. 1992. "Ordinal Optimization of Discrete Event Dynamic Systems," *Discrete Event Dynamic Systems*, 2,61-88.

Kiefer, J. and J. Wolfowitz. 1952. "Stochastic Estimation of the Maximum of a Regression Function," *Annals of Mathematical Statistics*, 23, 462-466.

Matejcik, F.J. and B.L. Nelson. 1995. "Two-Stage Multiple Comparisons with the Best for Computer Simulation," *Operations Research*, 43, 633-640.

Ólafsson, S. 1999. "Iterative Ranking-and-Selection for Large-Scale Optimization," in P.A. Farrington, H.B. Nembhard, D.T. Sturrock, and G.W. Evans (eds.), *Proceedings of the 1999 Winter Simulation Conference*, 479-485.

Ólafsson, S. and J. Kim. 2001. "Towards a Framework for Black-Box Optimization," in *Proceedings of the 2001 Winter Simulation Conference*,

Rinott, Y. 1978. "On Two-Stage Selection Procedures and Related Probability-in-Equalities," *Communications in Statistics*, A7, 799-811.

Robbins, H. and S. Monro. 1951. "A Stochastic Approximation Method," *Annals of Mathematical Statistics*, 22, 400-407.

Robinson, S.M. 1996. "Analysis of Sample-Path Optimization," *Mathematics of Operations Research*, 21, 513-528.

Rubinstein, R.Y. 1991. "How to Optimize Discrete-Event Systems for a Single Sample Path by the Score Function Method," *Annals of Operations Research*, 27, 175-212.

Rubinstein, R.Y. and A. Shapiro. 1993. *Discrete Event Systems: Sensitivity Analysis and Stochastic Approximation using the Score Function Method*, John Wiley & Sons, New York.

Shi, L. and S. Ólafsson. 1997. "An Integrated Framework for Deterministic and Stochastic Optimization," in S. Andradóttir, K.J. Healy, D.H. Withers, and B.L. Nelson (eds.), *Proceedings of the 1997 Winter Simulation Conference*, 358-365.

Swisher, J.R., P.D. Hyden, S.H. Jacobson, and L.W. Schruben. 2000. "A Survey of Simulation Optimization Techniques and Procedures," in J.A. Jones, R.R. Barton, K. Kang, and P.A. Fishwick (eds.), *Proceedings of the 2000 Winter Simulation Conference*, 119-128.

AUTHOR BIOGRAPHIES

SIGURDUR ÓLAFSSON is an assistant professor in the Department of Industrial and Manufacturing Systems Engineering at Iowa State University. He received a B.S. in Mathematics from the University of Iceland in 1995, and an M.S. and a Ph.D. in Industrial Engineering from the University of Wisconsin - Madison in 1996 and 1998, respectively. His research interests include applied probability, stochastic optimization, and simulation. He is a member of IIE and INFORMS.

JUMI KIM is a Ph.D. candidate in the Department of Industrial and Manufacturing Systems Engineering at Iowa State University. She received a B.S. in Industrial Engineering from the Kangnung National University of Korea, and a M.S. in Industrial Engineering from the Seoul National University of Korea. Her research interests are in simulation optimization.