

INTEGRATING OPTIMIZATION AND SIMULATION: RESEARCH AND PRACTICE

Michael C. Fu

Van Munching Hall
University of Maryland
College Park, MD 20742 U.S.A.

John S. Carson

AutoSimulations
1355 Terrell Mill Rd, Bldg 1470, Ste 200
Marietta, GA 30067 U.S.A. U.S.A.

Charles R. Harrell

Manufacturing Engineering
Brigham Young University
Provo, UT 84062 U.S.A.

James P. Kelly

College of Business
University of Colorado
Boulder, CO 80309 U.S.A.

Sigrún Andradóttir

Dept. of Ind. & Sys. Engineering
Georgia Institute of Technology
Atlanta, GA 30332 U.S.A.

Fred Glover

College of Business
University of Colorado
Boulder, CO 80309 U.S.A.

Yu-Chi Ho

Division of Applied Sciences
Harvard University
Cambridge, MA 02138 U.S.A.

Stephen M. Robinson

Dept. of Industrial Engineering
University of Wisconsin-Madison
Madison, WI 53706 U.S.A.

ABSTRACT

The integration of optimization and simulation has become nearly ubiquitous in practice, as most discrete-event simulation packages now include some type of optimization routine. This panel session's objective was to explore the present state of the art in simulation optimization, prevailing issues for researchers, and future prospects for the field. The composition of the panel included views from both simulation software developers and academic researchers. This Proceedings paper begins with a brief overview of some issues, introduced by the chairman and organizer of the session, followed by the position statements of the panel members, which served as a starting point for the panel discussion.

1 INTRODUCTION AND OVERVIEW

Integrating "optimization" routines (the reason for the quotes will be explained shortly) into simulation packages has become almost a necessity for commercial providers of discrete-event simulation software. This is, however, a fairly recent development. One way to see this is to compare the newest (third) editions of two of the most popular textbooks for simulation courses – Law and Kelton (2000)

and Banks, Carson, Nelson and Nicol (2000) – with their previous editions. For example, all of the currently available software routines for performing optimization listed in Law and Kelton (2000, p.664) — AutoStat, OptQuest, OPTIMIZ, SimRunner2, and WITNESS Optimizer (Three of these commercial packages are represented on the panel.) — were not in existence at the time of the earlier printings. The goal of these routines is to seek improved settings of user-selected system parameters with respect to the performance measure(s) of interest. However, unlike in mathematical programming software packages, the user has no way of knowing if an optimum has actually been reached (hence the quotations around optimization at the beginning of this paragraph).

The term "simulation optimization" has itself become more widespread; for example it is one of the new entries in the updated 2nd edition of the *Encyclopedia of Operations Research and Management Science* (Gass and Harris 2000), published in November of this year. We distinguish between the simulation optimization focused on by the panel and that of choosing from a given large set of alternatives, where statistical ranking & selection (multiple comparison) methods can be applied (e.g., Goldsman et al. 1999). The primary difference is that in the setting we consider, a constraint set (possibly unbounded and uncountable) is provided, over

which the algorithm seeks improved solutions, whereas in the ranking & selection setting, a fixed set of alternatives is provided a priori. In the former case, the focus is on the searching mechanism, whereas in the latter, statistical considerations are paramount. Clearly statistics must also come into play if any convergence results are to be rigorously established for the search algorithms.

The five packages listed previously all use metaheuristics from combinatorial optimization based on evolution strategies such as genetic algorithms, tabu search, scatter search (see Glover, Kelly, and Laguna 1999), with some adaptation of other techniques taken from the deterministic optimization literature, e.g., neural networks and simulated annealing (even though the latter is probabilistic in nature, it has been primarily applied to deterministic problems). On the other hand, the research literature in simulation optimization (refer to Andradóttir 1998 or Fu 1994) is dominated by continuous-parameter stochastic approximation methods (with more recent work on random search methods for discrete parameters problems), which concentrate on local search strategies based on a single point, versus the group or family of points adopted by many of the strategies above. Yet another approach is the so-called sample path optimization approach, which like the first set of strategies adopts deterministic algorithms, but instead of combinatorial approaches, uses nonlinear programming algorithms. These implementations actually exploit the wide availability of code for these algorithms (e.g., Gürkan, Özge, and Robinson 1999).

In putting together the panel, a conscious effort was made to cover the spectrum of disparate perspectives, including both academic researchers and commercial software vendors (note that a large number of the panelists wear both hats, which is a healthy sign for integrating research into practice by way of commercial software development), and representing each of the major approaches to simulation optimization. Roughly speaking, in light of the discussion above, these approaches can be categorized as follows:

- gradient-based and random search algorithms (mainly stochastic approximation);
- evolutionary algorithms and metaheuristics (including genetic algorithms and tabu search);
- mathematical programming-based approaches (mainly the sample path method);
- statistical search techniques, such as sequential response surface methodology.

To recap, in terms of software implementation, the overwhelming majority of the available routines are based on the second approach. On the other hand, a perusal of the *simulation* research literature would find a stark reversal of this situation, i.e., the other approaches (especially the first approach) are much better represented in archival journals

on simulation. Indeed, other than in the Winter Simulation Proceedings, one would be hard-pressed to find published examples of metaheuristics in the simulation literature.

Why is this the case? you might ask. There appears to be two major barriers: either the algorithms that are implemented are not *provably* convergent, or the use of simulation is secondary, i.e., the simulation model is merely treated as a black box, in which case it seems more appropriate that the algorithm be published in the *Journal of Heuristics* than in the *ACM Transactions on Modeling and Computer Simulation*.

The next section provides the position statements of the individual panelists. I end this section with a brief summary of my own view of the primary challenges faced in truly integrating optimization into simulation (Fu 2001):

- providing some measure of goodness (other than just improvement over the starting point, which most packages provide) for the metaheuristics that dominate the commercial field (see also Carson and Glover/Kelly statements);
- developing practical and effective implementation of algorithms with proven convergence properties that dominate the research literature (see also Andradóttir and Robinson statements).

In addition, I believe the prospects for the so-called ordinal optimization approach of Ho et al. (1992, 2000) have yet to be fully exploited in simulation optimization. The key idea behind this approach is that it is much easier to determine approximate order than precise estimation. In other words, by treating the simulation model as a black box, as most metaheuristic approaches do, there is an immense waste of simulation replications used to obtain precise estimates at parameter settings whose poor relative performance becomes apparent with just a few replications. Related to this avoiding wasted simulation philosophy is the idea of factor screening (see Harrell statement).

2 POSITION STATEMENTS

Each panelist was asked to provide a short statement summarizing his/her assessment of the state of the art in simulation optimization, both in research and in practice. In particular, the panelists were encouraged to expound upon what they felt are recent major advances and most urgent needs in the area.

2.1 Sigrún Andradóttir, Georgia Institute of Technology

The field of simulation optimization is concerned with the use of simulation to design and optimize systems (Andradóttir 1998). This is fundamentally a challenging problem because using simulation to estimate the performance of a

single system design often requires a substantial amount of computer time, and determining the optimal system design is obviously a more difficult problem because of the need to evaluate the system performance for several different designs. Different approaches to solving simulation optimization problems have been developed, ranging from heuristic approaches to rigorous methods that provide some performance guarantees (e.g., convergence with probability one to the set of global optimal solutions as the computational budget grows, assurance that the solution provided by the method is near-optimal with specified probability, etc.). The class of rigorous simulation optimization methods includes vastly different approaches, such as statistical methods (e.g., ranking, selection, and multiple comparison approaches), methods that use gradient estimates for continuous parameter optimization (e.g., stochastic approximation and sample path approaches), random search methods (e.g., simulated annealing), etc.

Much of the research on simulation optimization to date is concerned with methods that require a certain amount of sophistication on the part of the user, in terms of understanding both the details of the optimization approach being used and the nature of the stochastic processes underlying the simulation. Moreover, in order to guarantee convergence, these approaches are sometimes rather conservative, which can lead to slow convergence in practice. This has led many practitioners to use heuristic approaches instead, which are designed to provide an answer relatively quickly but without assurances about the quality of that answer.

To bridge the gap between simulation researchers and practitioners when it comes to system design and optimization, it is important to develop efficient, easy to use, and rigorous approaches for solving simulation optimization problems. In particular, it is important to develop general purpose methods that are suitable for solving a wide range of simulation optimization problems and hence are appropriate for inclusion in general purpose simulation languages. It is also important to develop special purpose methods and software that can be used to solve important special classes of simulation optimization problems by even a novice simulation practitioner, and that take advantage of the structure of the underlying optimization problem to achieve faster convergence than the general purpose approaches described above.

The previous paragraph addresses the need for efficient simulation optimization methods designed for use by practitioners that are not experts on simulation optimization. The continued development of more sophisticated approaches that can be adapted by experts in the simulation optimization area to exploit special structure to solve specific optimization problems with high efficiency is also important. Such research is likely both to lead to optimization methods and software suitable for use by even novice simulation practitioners, and also to have applications to

solving stochastic optimization problems that do not lie within the traditional field of simulation optimization.

2.2 John Carson, AutoSimulations

From a simulation practitioner's and software vendor's perspective, it appears the market is crowded with far too many competing optimization methods for a layman to sort them all out. We have genetic algorithms, evolution strategies, tabu search, scatter search, neural networks, simulated annealing, not to mention combinations of these along with the old-fashioned gradient methods. What's a non-expert to do? While many commercial products now offer optimization, what method should they use? Does it matter? Which is the best? Which will give the best answer with the least computing time?

For a simulation software vendor, the challenge is to figure out the best method to implement in their output analysis package. For marketing reasons, most if not all vendors have implemented some form of optimization.

What is needed are thorough review, evaluations, guidelines and recommendations from researchers to simulation vendors. (There are very few practitioners who would implement an optimization routine; it must go from research to commercial software to practice.) What's needed most at this time, I feel, is for an unbiased researcher to test all methods under a variety of types of models, and then communicate the results to the simulation community.

2.3 Fred Glover and James P. Kelly, University of Colorado and OptTek Systems

1. Simulation optimization applies to an extremely broad area of practical applications, well beyond the range that most current researchers and practitioners are aware of.
2. These application domains afford a wealth of opportunities for treating critical issues of uncertainty and complexity – issues that have been incompletely handled by the theory and attempted practical implementations of the past.
3. Ironically, traditional statistics and optimization do not provide the major point of access to simulation optimization. Although these classical domains are strongly relevant, they operate as a “supporting player” to the area of metaheuristics, which takes a dominant role in integrating simulation and optimization (e.g. Glover, Kelly, and Laguna 1999).
4. It is no coincidence that simulation optimization is only now emerging as a prominent practical tool, as innovations in metaheuristics have reached a point where they allow real world problems to be solved that have previously been beyond reach. Even so,

the developments of “mainstream metaheuristics” are not sufficient to provide the best methods for integrating simulation and optimization. Special approaches based on considerations not previously envisioned are making new levels of performance possible in this highly challenging area – performance that exceeds the outcomes of “mainstream metaheuristics” by wide margins, in some cases two to three orders of magnitude.

5. An important recent development in simulation optimization is the emergence of applications beyond what we normally view as “simulation based”. Any setting that involves expensive or highly complex operations to evaluate a proposed solution is a candidate for being treated by the methodology of simulation optimization. Consequently, applications are springing up in realms of business and industry that traditionally are not conceived as linked to simulation.
6. Areas that bear special watching in the future include the “glamour technologies” as well as the hard core business and industrial areas – technologies such as DNA sequencing and assembly, web-based analysis, and telecommunications. Included in this list of applications that offer special future promise is the design and tuning of algorithms themselves.

2.4 Charles Harrell, Brigham Young University and ProModel Corporation

I believe that simulation optimization has generally been well received and, for the most part, intelligently applied. The seamless integration of optimization with commercial simulation software has certainly contributed to its success (e.g., Price and Harrell 1999). Optimization relieves much of the trial-and-error approach to experimentation and can even reveal superior solutions that may not be intuitively obvious. Simulation optimization is not a replacement for traditional experimentation and output analysis, however, and should continue to be viewed as only part of an overall output analysis methodology.

Much of the future improvement needed in simulation optimization pertains to speed and validity of the results. Much of the speed improvement can be obtained by taking better advantage of multi-processing technology which would enable multiple experiments to be run concurrently. Another way for improving the efficiency of the optimization is to integrate it with statistical factor screening techniques and ranking and selection techniques. For example, after a user defines a goal (objective function) and decision variables (input factors or controllable variables), a factor screening utility could be run first to “weed out” input factors that do not significantly influence the value of the objective function. Then the optimization could be run on

the reduced search space. This would be useful in cases when the user believes that a certain input factor influences the objective function but does not know for sure.

Research needs to be conducted in the use of techniques for improving search algorithms so they are able to quickly discern whether the mean performance of different solutions is truly (significantly) different as solutions are evaluated during the search. Likewise, more work is needed to develop methods for automatically determining the end of the warm-up period for non-terminating simulations. This is necessary because the length of the warm-up period for a model can change as new scenarios are automatically generated and simulated by the optimizer. Optimizers usually assume the same warmup period for each scenario, even though it may actually be different. Determining the end of the warm-up period on the fly is difficult and there is no foolproof technique that works in all cases. However, it stands to reason that some reasonably accurate automated technique would be better than doing nothing at all.

New innovations in simulation optimization are continually being developed (e.g., Bowden and Hall 1998). The thoughts presented here represent just a few of the areas where significant improvements are still waiting to be made. As the technology progresses and practitioners become more aware of its benefits and educated in its use, it will become an increasingly integral part of the simulation experimentation process.

2.5 Yu-Chi Ho, Harvard University

There is no question that simulation is the only general purpose and generally applicable modeling tool for truly complex systems, natural or human made. If simulation models are used for design and optimization (as opposed to validation) purposes, then the users are faced with some fundamental limitations on computation. Among these are:

1. the $1/(\text{simulation length})^{1/2}$ limit — confidence interval cannot decrease faster than this;
2. combinatorial explosion of the search space or the curse of dimensionality;
3. lack of structure for many of the search spaces, which implies that one cannot do better than blind search on the average.

Any one or combination of the above can render a direct attack on the performance optimization via simulation models infeasible. The thesis of my panel remarks is that in the face of such basic difficulties some strategic re-direction of effort is required. Two major thrusts are proposed (Ho et al. 2000). First, we need to lower or soften our goals. Instead of asking “the best for sure” settle for “the good enough with high probability”. This is implicit in many tools of computational intelligence such as Genetic Algorithm,

Fuzzy Logic, etc., and explicitly in Ordinal Optimization (Ho et al. 1992). Second, we need in our process of optimization to constantly strike a balance between "breadth vs. depth" or "search for better designs vs. making sure the current designs are indeed good enough". Simulation resources should be constantly applied with the above trade off in mind. As an example consider the tradeoff between using the "Nested Partition" method for "breadth" search, and the "Optimal Computing Budget Allocation" method for deciding which design performance to pursue in "depth" (i.e., narrow its confidence interval via increasing its simulation length).

We submit that these two considerations define a general framework within which we can address the integration of simulation and optimization for computationally complex systems.

2.6 Stephen M. Robinson, University of Wisconsin-Madison

As my experience with simulation optimization may have been somewhat different from that of others in this panel, I might start with a few words about that. My early optimization research experience was in the deterministic optimization community. I have been involved for the last 10-12 years with particular kinds of stochastic optimization problems, including some originating in decision-making problems faced by the military. Originally these were mainly of types that could be formulated satisfactorily by using scenarios, with subsequent reduction to large structured deterministic problems. However, more recently some of these problems, and others coming from manufacturing, have required the use of simulation because they could not be satisfactorily handled with scenarios.

My research group has worked mostly with a particular variety of simulation optimization method, which we usually call sample-path optimization (e.g., Gürkan, Özge, and Robinson 1999), but that has also been given other names, such as the stochastic counterpart method. We have had some success with these methods for particular problem classes by combining them with efficient deterministic optimization methodology, for both smooth and nonsmooth (convex) problems. In addition, progress had been made on the theoretical underpinnings of this methodology, and here I might particularly mention the work of Alexander Shapiro and his group. However, there are serious computational issues that are not currently resolved.

One of most critical of these issues is that of the efficient computation of gradients of performance functions with respect to the decision variables. In a number of applied problems this computation can be done efficiently by infinitesimal perturbation analysis (IPA) or some of its extensions (e.g., Fu and Hu 1997). However, the information needed for application of IPA is not usually accumulated

by production simulation software packages. This means that to test our methods we have had to hand-code the simulations, rather than being able to use efficient modern software. We have tried to generate IPA gradient information by adapting automatic differentiation methods to work on simulation code, but to date we have not had much success with this because of the one-sidedness of the derivatives involved (they are basically directional derivatives). I think this is a real deficiency, because no matter how good an optimization procedure is, people in industry are unlikely to use it if it requires a lot of additional coding work.

I would like to see the producers of production software make a serious effort to provide IPA add-ins to their code that could be used for computation of derivatives. Having these conveniently available would mean that a user could employ an efficient optimization package in combination with the simulation software, without the need for recoding. They cannot do so effectively now because of the coding barrier.

REFERENCES

- Andradóttir, S. 1998. Simulation optimization. Chapter 9 in *Handbook of Simulation: Principles, Methodology, Advances, Applications, and Practice*. ed. J. Banks. New York: John Wiley & Sons.
- Banks, J., J.S. Carson, B.L. Nelson, and D.M. Nicol. 2000. *Discrete event systems simulation*. 3rd edition. Prentice Hall.
- Bowden, R.O., and J.D. Hall. 1998. Simulation optimization research and development. In *Proceedings of the 1998 Winter Simulation Conference*, ed. D.J. Medeiros, E.F. Watson, J.S. Carson, M.S. Manivannan, 1693-1698. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.
- Fu, M.C. 1994. Optimization via simulation: a review. *Annals of Operations Research* 53: 199-248.
- Fu, M.C. 2001. Optimization for simulation: theory vs. practice. manuscript under preparation as Feature Article for *INFORMS Journal on Computing*.
- Fu, M.C., and J.Q. Hu. 1997. *Conditional Monte Carlo: Gradient estimation and optimization applications*. Kluwer Academic Publishers.
- Gass, S.I., and C.M. Harris, eds. 2000. *Encyclopedia of operations research and management science*. 2nd edition. Kluwer Academic Publishers.
- Glover, F., J.P. Kelly, and M. Laguna. 1999. New advances for wedding optimization and simulation. In *Proceedings of the 1999 Winter Simulation Conference*, ed. P.A. Farrington, H.B. Nembhard, D.T. Stoffock, and G.W. Evans, 255-260. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.
- Goldsman, D., B.L. Nelson, T. Opicka, and A.A.B. Pritsker. 1999. A ranking and selection project: experiences

from a university-industry collaboration. In *Proceedings of the 1999 Winter Simulation Conference*, ed. P.A. Farrington, H.B. Nembhard, D.T. Stuflock, and G.W. Evans, 83-92. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Gürkan, G., A.Y. Özge, and S.M. Robinson. 1999. Sample-path solution of stochastic variational inequalities. *Mathematical Programming* 84: 313-333.

Ho, Y.C., C.G. Cassandras, C.H. Chen, and L.Y. Dai. 2000. Ordinal optimization and simulation. *Journal of Operations Research Society* 51: 490-500.

Ho, Y.C., R. Sreenivas, and P. Vakil. 1992. Ordinal optimization of DEDS. *Discrete Event Dynamic Systems: Theory and Applications* 2: 61-88.

Law, A.M., and W.D. Kelton. 2000. *Simulation modeling and analysis*. 3rd edition. McGraw-Hill.

Price, R.N., and C.R. Harrell. 1999. Simulation modeling and optimization using ProModel. In *Proceedings of the 1999 Winter Simulation Conference*, ed. P.A. Farrington, H.B. Nembhard, D.T. Stuflock, and G.W. Evans, 208-214. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

AUTHOR BIOGRAPHIES

MICHAEL FU <mfu@rsmith.umd.edu> is a Professor in the Robert H. Smith School of Business, with a joint appointment in the Institute for Systems Research, at the University of Maryland. He received degrees in mathematics and EE/CS from MIT, and a Ph.D. in applied math from Harvard University. His research interests include simulation and applied probability modeling, particularly with applications towards manufacturing systems, inventory control, and financial engineering. He teaches courses in applied probability, stochastic processes, simulation, computational finance, and supply chain logistics & operations management, and in 1995 was awarded the Maryland Business School's Allen J. Krowe Award for Teaching Excellence. He is a member of INFORMS, IEEE, and IIE, and received the 1999 IIE Operations Research Division Award. He is currently the Simulation Area Editor of *Operations Research*, and serves on the editorial boards of *Management Science*, *INFORMS Journal on Computing*, *IIE Transactions*, and *Production and Operations Management*. He is co-author of the book, *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*, which received the INFORMS College on Simulation Outstanding Publication Award in 1998.

SIGRÚN ANDRADÓTTIR <sa@isye.gatech.edu> is an Associate Professor in the School of Industrial and Systems Engineering at the Georgia Institute of Technology. Previously, she was an Assistant Professor

and later an Associate Professor in the Departments of Industrial Engineering, Mathematics, and Computer Sciences at the University of Wisconsin – Madison. She received a B.S. in Mathematics from the University of Iceland in 1986, an M.S. in Statistics from Stanford University in 1989, and a Ph.D. in Operations Research from Stanford University in 1990. Her research interests include stochastic optimization, simulation, and applied probability. She presently serves as Associate Editor for *Operations Research*, *IIE Transactions*, *Stochastic Models*, and the *ACM Transactions on Modeling and Computer Simulation*. She is a member of INFORMS and served as Editor of the *Proceedings of the 1997 WSC*.

JOHN S. CARSON II <john_carson@autosim.com> is the Regional Consulting Manager for the Atlanta office of AutoSimulations, a Brooks Automation company. With ASI since 1994, he has over 20 years experience in simulation in a wide range of application areas, including manufacturing, distribution, warehousing and material handling, transportation and rapid transit systems, port operations and shipping, and medical/health care systems. His current interests center on the simulation of transportation systems, train systems, bulk and liquid processing, and software for simulation output analysis. He has been an independent simulation consultant, and has taught at Georgia Tech, the University of Florida, and the University of Wisconsin-Madison. He is the co-author of two textbooks, including the widely-used *Discrete-Event System Simulation*, Third Edition, 2000. He holds a Ph.D. in Industrial Engineering and Operations Research from the University of Wisconsin-Madison, and is a member of IIE and INFORMS.

FRED GLOVER <fred.glover@colorado.edu> is the MediaOne Chaired Professor in Systems Science at the University of Colorado, Boulder. He has authored or co-authored more than 280 published articles and four books in the fields of mathematical optimization, computer science, and artificial intelligence, with particular emphasis on practical applications in industry and government. In addition to holding editorial and advisory posts for many international journals, he is founder and Editor-in-Chief of the *Journal of Heuristics*. He is the recipient of the distinguished von Neumann Theory Prize, as well as of numerous other awards and honorary fellowships, including those from the American Association for the Advancement of Science, the NATO Division of Scientific Affairs, TIMS, ORSA, the Decision Sciences Institute, the U.S. Defense Communications Agency, the Energy Research Institute, the American Assembly of Collegiate Schools of Business, Alpha Iota Delta, and the Miller Institute for Basic Research in Science. He serves on the advisory boards of several organizations and is co-founder of Optimization Technologies, Inc.

CHARLES R. HARRELL <harrellc@byu.edu, charles@promodel.com> is an Associate Professor of Manufacturing Engineering at Brigham Young University, and founder and chairman of ProModel Corporation in Orem, Utah. He received his B.S. in Manufacturing Engineering Technology from Brigham Young University, M.S. in Industrial Engineering from the University of Utah, and Ph.D. in Manufacturing Engineering from the Technical University of Denmark. Prior to forming ProModel, he worked in simulation and systems design for Ford Motor Company and Eaton Kenway Corporation. He is a senior member of IIE and SME. He is co-author of the book, *Simulation Using ProModel*, 1999.

YU-CHI (LARRY) HO <ho@hrl.harvard.edu> is the T. Jefferson Coolidge Chair in Applied Mathematics and the Gordon McKay Professor of Systems Engineering in the Division of Applied Sciences at Harvard University. He received his S.B. and S.M. degrees in Electrical Engineering from MIT and his Ph.D. in Applied Mathematics from Harvard University. His current research interests lie at the intersection of Control System Theory, Operations Research, and Computational Intelligence. He has published over 140 articles and three books. He is on the editorial boards of several international journals and is the founder and Editor-in-Chief of the international journal on *Discrete Event Dynamic Systems*. He is the recipient of various fellowships and awards including the Guggenheim (1970), the IEEE Field Award for Control Engineering and Science (1989), the Chiang Technology Achievement Prize (1993), the Bellman Control Heritage Award (1999) of the American Automatic Control Council, and the ASME Rufus Oldenburger Award (1999). He is a Life Fellow of IEEE, and was elected a member of the U.S. National Academy of Engineering in 1987 and a foreign member of the Chinese Academy of Sciences and the Chinese Academy of Engineering in 2000.

JAMES P. KELLY <James.Kelly@colorado.edu, kelly@OptTek.com> is an Associate Professor of Operations Management in the College of Business at the University of Colorado, Boulder. He received his bachelors and masters degrees in Chemical Engineering from Bucknell University in 1981, and his doctoral degree in Applied Mathematics and Operations Research from the University of Maryland in 1990. His current research interests are in the area of heuristic combinatorial optimization and applied artificial intelligence. The area of simulation optimization has been his prime focus in recent years. He is a member of INFORMS and principal editor from the *Journal of Heuristics*. He is currently CEO of OptTek Systems, Inc.

Madison <smrobins@facstaff.wisc.edu>. He received his B.A. in Mathematics from the University of Wisconsin, M.S. in Mathematics from New York University, and Ph.D. in Computer Sciences from the University of Wisconsin. His research interests are in quantitative methods in managerial economics, especially deterministic and stochastic optimization and methods to support decision under uncertainty. He is a member of INFORMS, the Mathematical Programming Society, SIAM, and IIE. He received the 1997 George B. Dantzig Prize, sponsored by the Mathematical Programming Society and SIAM.

STEPHEN ROBINSON is Professor of Industrial Engineering and Computer Sciences at the University of Wisconsin-