Proceedings of the 1999 Winter Simulation Conference
P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, eds.

NEW ADVANCES FOR WEDDING OPTIMIZATION AND SIMULATION

Fred Glover
James P. Kelly
Manuel Laguna

School of Business, CB 419
University of Colorado
Boulder, CO 80309, U.S.A.

ABSTRACT

Many real world problems in optimization are too complex to be given tractable mathematical formulations. Multiple nonlinearities, combinatorial relationships and uncertainties often render challenging practical problems inaccessible to modeling except by resorting to simulation – an outcome that poses grave difficulties for classical optimization methods. In such situations, recourse is commonly made to itemizing a series of scenarios in the hope that at least one will give an acceptable solution. Consequently, a long standing goal in both the optimization and simulation communities has been to create a way to guide a series of simulations to produce high quality solutions, in the absence of tractable mathematical structures.

Applications include the goals of finding:

- the best configuration of machines for production scheduling
- the best integration of manufacturing, inventory and distribution
- the best layouts, links and capacities for network design
- the best investment portfolio for financial planning
- the best utilization of employees for workforce planning
- the best location of facilities for commercial distribution
- the best operating schedule for electrical power planning
- the best assignment of medical personnel in hospital administration
- the best setting of tolerances in manufacturing design
- the best set of treatment policies in waste management

and many other objectives.

Recent innovations for integrating metaheuristics, classical optimization and artificial intelligence have produced a practical software system, called OptQuest, that is capable of guiding a series of simulations to uncover optimal or near optimal solution scenarios.

The advances represented by this new system open new doors for simulation as well as extending the areas to which optimization can be applied. Within the brief span since it has come into existence, OptQuest has been used in several thousand real world applications that combine simulation and optimization, and has also been used to determine optimal parameters for other computer based decision tools, to increase their effectiveness.

1 INTRODUCTION

It is widely acknowledged that simulation is a powerful computer-based tool that enables decision-makers in business and industry to improve operating and organizational efficiency. The ability of simulation to model a physical process on the computer, incorporating the uncertainties that are inherent in all real systems, provides an enormous advantage for analysis in situations that are too complex to represent by “textbook” mathematical formulations.

In spite of its acknowledged benefits, however, simulation has suffered a limitation that has prevented it from uncovering the best decisions in critical practical settings. This limitation arises out of an inability to evaluate more than a fraction of the immense range of options available. Practical problems in areas such as manufacturing, marketing, logistics and finance typically pose vast numbers of interconnected alternatives to consider. As a consequence, the decision making goal of identifying and evaluating the best (or near best) options has been impossible to achieve in many applications.

2 NATURE OF THE CHALLENGE

Theoretically, the issue of identifying best options falls within the real of optimization. Until quite recently, however, the methods available for finding optimal decisions have been unable to cope with the complexities and uncertainties posed by many real world problems of the form treated by simulation. In fact, these complexities and uncertainties are the primary reason that simulation is chosen as a basis for handling such problems. Consequently, decision makers have been faced with the "Catch 22" that many important types of real world optimization problems can only be treated by the use of simulation models, but once these problems are submitted to simulation there are no optimization methods that can adequately cope with them. In short, there has not existed any type of search process capable of effectively integrating simulation and optimization. The same shortcoming is also encountered outside the domain of simulation, as situations increasingly arise where complex (realistic) models cannot be analyzed using traditional "closed form" optimization tools.

Recent developments are changing this picture. Advances in the field of metaheuristics – the domain of optimization that augments traditional mathematics with artificial intelligence and methods based on analogs to physical, biological or evolutionary processes – have led to the creation of a new approach that successfully integrates simulation and optimization, embedded in a computer software system called OptQuest. The availability of this new system opens the door to handling decision-making problems in business and industry that could not be adequately approached in the past.

We first describe the organization of this system and the special algorithms that underlie its design, in Sections 4-7. In Section 8 we provide an illustrative computational comparison of the system's performance. Finally, in sections 9 and 10 we identify additional features of the system and benefits that have emerged in its applications.

4 GENERAL ORGANIZATION

In a general purpose optimizer such as OptQuest, it is preferable to separate the solution procedure from the system we are trying to optimize. A potential disadvantage of this "black box" approach (see Figure 1), derives from the fact that the optimization procedure is generic and does not know anything about what goes on inside of the box.

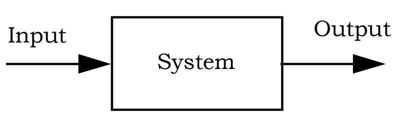

Figure 1: System as a Black Box

The clear advantage, on the other hand, is that the same optimizer can be used for many systems.

Our approach is an implementation of a generic optimizer that overcomes the deficiency of black box systems of the type illustrated in Figure 1, and successfully embodies the principle of separating the method from the model. The optimization problem is defined outside the system, which is represented in this case by a simulation model. Therefore, the simulation model can change and evolve to incorporate additional elements, while the optimization routines remain the same. Hence, there is a complete separation between the model that represents the system and the procedure that is used to solve optimization problems defined within this model.

The optimization procedure uses the outputs from the simulation model which evaluate the outcomes of the inputs that were fed into the model. On the basis of this evaluation, and on the basis of the past evaluations which are integrated and analyzed with the present simulation outputs, the optimization procedure decides upon a new set of input values (see Figure 2). The optimization procedure is designed to carry out a special "non-monotonic search," where the successively generated inputs produce varying evaluations, not all of them improving, but which over time provide a highly efficient trajectory to the best solutions. The process continues until an appropriate termination criterion is satisfied (usually given by the user's preference for the amount of time to be devoted to the search). Our method has three main components, scatter search, tabu search, and neural networks, whose roles are briefly sketched in the next three sections.

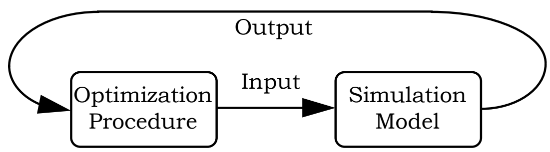

Figure 2: Coordination Between Optimization and Simulation

5 SCATTER SEARCH

Two of the best-known meta-heuristics are *genetic algorithms* and *tabu search*. Genetic Algorithm (GA) procedures were developed by John Holland in the late 1960s and early 1970s (Holland 1975). Parallel to the development of GAs, Fred Glover established the principles and operational rules for tabu search (TS) and a related methodology know as *scatter search* (Glover 1977).

Scatter search has some interesting commonalties with GA ideas, although it also has a number of quite distinct features. Several of these features have come to be incorporated into GA approaches after an intervening period of approximately a decade, while others remain largely unexplored in the GA context.

Scatter search is designed to operate on a set of points, called *reference points*, that constitute good solutions obtained from previous solution efforts. Notably, the basis for defining "good" includes special criteria such as diversity that go beyond objective function evaluations. The approach systematically generates combinations of the reference points to create new points, each of which is mapped into an associated feasible point. The combinations are generalized forms of linear combinations, accompanied by processes to adaptively assign integer values to discrete variables. Tabu search is then superimposed to control the composition of reference points at each stage. Tabu Search has its roots in the field of artificial intelligence as well as in the field of optimization. The heart of tabu search lies in its use of adaptive memory, which provides the ability to take advantage of the search history in order to guide the solution process. In its simplest manifestations, adaptive memory is exploited to prohibit the search from reinvestigating solutions that have already been evaluated. However, the use of memory in our implementation is much more complex and calls upon memory functions that encourage search diversification and intensification. These memory components allow the search to escape from locally optimal solutions and guide the quest for a globally optimal solution.

It is interesting to observe similarities and contrasts between scatter search and the original GA proposals. Both are instances of what are sometimes called "population based" approaches. Both incorporate the idea that a key aspect of producing new elements is to generate some form of combination of existing elements. On the other hand, GA approaches are predicated on the idea of choosing parents randomly to produce offspring, and further on introducing randomization to determine which components of the parents should be combined. By contrast, the scatter search approach does not correspondingly make recourse to randomization, particularly in the sense of being indifferent to choices among alternatives. Instead, the approach is designed to incorporate strategic responses, both deterministic and probabilistic, that take account of evaluations and history. Scatter search focuses on generating relevant outcomes without losing the ability to produce diverse solutions, due to the way the generation process is implemented. For example, the approach includes the generation of new points that are not convex combinations of the original points. The new points constitute forms of *extrapolations*, endowing them with the ability to contain information that is not contained in the original reference points.

Scatter search is an *information driven* approach, exploiting knowledge derived from the search space, high-quality solutions found within the space, and trajectories through the space over time. The incorporation of such designs is responsible for enabling our system to solve complex simulation-based problems with unprecedented efficiency.

A fuller description of scatter search and the range of applications where it has proved effective can be found in Glover (1999).

6 TABU SEARCH BASICS

Tabu search operates by identifying key attributes of moves or solutions, and imposing restrictions on subsets of these attributes, depending on the search history. Two prominent ways for exploiting search history in TS are through *recency* memory and *frequency* memory. Recency memory is typically (though not invariably) a short-term memory that is managed by structures or arrays called "tabu lists," while frequency memory more usually fulfills a long term search function. A standard form of recency memory discourages moves that lead to solutions with attributes shared by other solutions recently visited. A standard form of frequency memory discourages moves leading to solutions whose attributes have often been shared by solutions visited during the search, or alternately encourages moves leading to solutions whose attributes have rarely been seen before. Another standard form of frequency memory is defined over subsets of elite solutions to fulfill an intensification function that reinforces the inclusion of special attributes of these solutions within new solutions.

Short and long term components based on recency and frequency memory are used separately and together in complementary TS search strategies. This approach operates by implicitly modifying the neighborhood of the current solution. The introduction and exploitation of these adaptive memory strategies within tabu search distinguishes it from other metaheuristic approaches, and endows it with an ability to *learn* how to make its way effectively through solution spaces. A comprehensive treatment of tabu search and a broad survey of its applications is given in the book by Glover and Laguna (1997). The combination of this framework with the complementary population-based approach of scatter search yields a method of remarkable power for problems that unite the concerns of simulation and optimization.

7 NEURAL NETWORK ACCELERATOR

The OptQuest system also embodies a neural network component to enhance its function. The concept behind embedding a neural network is to "screen out" solutions that are likely to be poor, without the necessity to submit these solutions to a full evaluation by the simulation routine to determine their quality. In other words, the neural network is a prediction model that helps the system accelerate the search by avoiding simulation runs whose results can be predicted as inferior. To carry out this process, information is collected about the objective function values obtained by different optimization variable settings. This information is then used to train the neural network during the search. The system automatically determines how much data is needed and how much training should be done, based on both the time to perform a simulation and the optimization time limit provided by the user.

The neural network is trained on the historical data collected during the search and an *error* value is calculated during each training round. This error refers to the accuracy of the network as a prediction model. That is, if the network is used to predict an objective function $f(\mathbf{x})$ for $\mathbf{x}$-vectors found during the search (where the function may not be written in explicit mathematical form, but is implicitly determined by the simulation), then the error indicates how good those predictions are. The error term can be calculated by computing the differences between the known $f(\mathbf{x})$ and the predicted $\hat{f}(\mathbf{x})$ objective function values. The training continues until the error reaches a minimum prespecified value.

The neural network accelerator can be used at several risk levels. The risk is associated with the probability of discarding $\mathbf{x}$ when $f(\mathbf{x})$ is better than $f(\mathbf{x}_{best})$, where $\mathbf{x}_{best}$ is the best solution found so far. The risk level is defined by the number of standard deviations used to determine how close a predicted value $\hat{f}(\mathbf{x})$ is of the best value $f(\mathbf{x}_{best})$. A risk-averse user would, for instance, would only discard $\mathbf{x}$ if $\hat{f}(\mathbf{x})$ is at least three standard deviations larger than $f(\mathbf{x}_{best})$, in a minimization problem. Such options are designed for the convenience of users of the OptQuest system, and can also be handled by allowing the system to make its own standard default choices.

8 ILLUSTRATIVE COMPARISONS

To provide an indication of the ability of OptQuest's search procedure to obtain high quality solutions, we compare OptQuest to a version of a well known and highly reputed genetic algorithm solver, called Genocop, which has been adapted to be able to handle linear constraints and a nonlinear objective defined over real-valued and integer-valued variables. We denote this version of Genocop by Genocop/C (for "Genocop with Constraints"). Genocop/C is incapable of handling problems with uncertainties, however, and its special focus precludes it from solving problems generated by the marriage of optimization and simulation.

Thus, our interest is to determine the effectiveness of OptQuest in competition with a method designed to excel in a particular area which is more specialized than the area OptQuest is designed to handle. We have performed our comparisons on problems that have been selected in the Genocop literature to demonstrate the performance of this system.

Table 1 shows the results. The first eight nonlinear problems involve 6 to 13 variables and 1 to 9 constraints. The last ten nonlinear problems involve 20 variables and 10 constraints. In addition, the first seven problems alternate between containing real variables and integer variables, while problem 8 is a mixed problem containing both types of variables. The last ten problems alternate between mixed and real variable problems. The objective is a maximization objective in each case.

OptQuest vs. Genocop/C

Problem	Genocop/C Solution	OptQuest Solution	Genocop/C Evaluations	OptQuest Evaluations	OptQuest Improvement
1	213.0	213.0	2,520	13	0.0
2	213.0	213.0	19,600	13	0.0
3	47.8	47.8	35,000	5,047	0.0
4	15.0	15.0	3,080	13	0.0
5	15.0	15.0	34,370	13	0.0
6	39.0	39.0	2,800	13	0.0
7	39.0	39.0	7,350	386	0.0
8	0.4	0.4	30,310	1,145	0.0
9	384.4	394.8	317,030	13	10.4
10	378.6	394.8	338,030	521	162.2
11	830.9	884.8	340,690	167	53.9
12	869.8	884.8	335,650	562	15.0
13	7089.7	8343.7	218,540	1,064	1254.0
14	8287.7	8695.0	333,480	370	407.3
15	335.9	394.8	100,170	174	58.9
16	366.7	394.8	101,500	429	28.1
17	3955.3	4109.6	128,660	1,104	154.3
18	2031.9	4109.6	254,870	646	2077.7

Table 1

For all problems, OptQuest obtained solutions at least as good as those found by Genocop/C, and uniformly obtained superior solutions on all ten of the larger problems. (The amounts by which OptQuest improved over the Genocop/C solutions are shown in the "OptQuest Improvement" column.)

The columns labeled "Evaluations" indicate the number of evaluations (iterations) required by the methods to obtain the solutions reported in the table. These columns disclose that OptQuest required substantially fewer evaluations to obtain its reported solution – in many cases, two orders of magnitude fewer evaluations. This type of difference becomes especially important for problems that exhibit the level of complexity and generality found in the simulation setting, where a single evaluation can consume vastly more computer time than in these problem examples. A solution system that not only finds high quality solutions, but also requires dramatically fewer evaluations than methods that represent the "state-of-the-art" in optimization, is essential for dealing with the challenge of effectively integrating optimization and simulation.

9 ADDITIONAL SYSTEM FEATURES

Complementing the features previously mentioned, OptQuest gives the decision maker a valuable basis for additional control, by making it possible to specify a variety of important relationships to govern the determination of optimal decisions, such as:

- Ranges of key parameters
- Budget limitations
- Machine capacities
- Minimum and maximum lot sizes
- Limits on hours worked

In particular, it is possible to include any set of conditions that can be represented by a mixed integer programming formulation. This gives the system the useful advantage of being able to create *a model on top of a model* – i.e., effectively allowing mixed integer programming and simulation to operate in concert. Building on this, OptQuest then determines the strategic options that are investigated under its guidance. The resulting search isolates scenarios that yield the highest quality outcomes for profits, costs and risks, according to the criteria selected by the decision maker.

10 DECISION MAKING CONSEQUENCES– IN SUMMARY

The OptQuest system brings the intelligence of latest metaheuristic innovations to software for corporate decision-making, and opens up new possibilities for integrating optimization and simulation models in business and industry. The solution technology in this system represents the outcome of over two decades of research, and its application gives decision makers the ability to go significantly beyond conventional decision-making approaches.

As demonstrated by thousands of practical applications covering a diverse array of areas, OptQuest provides decisions whose quality and utility cannot be matched by standard simulation or optimization packages. Such an ability is essential for effective planning in competitive and uncertain environments. Additional information on this system can be obtained by containing the authors, or by accessing the web page www.optquest.com/~optinfo.

REFERENCES

Glover, F. 1977. Heuristics for integer programming using surrogate constraints. *Decision Sciences* 8:156-166.

Glover, F., J.P. Kelly and M. Laguna. 1996. New advances and applications of combining simulation and optimization. *Proceedings of the 1996 Winter Simulation Conference*, J.M. Charnes, D.J. Morrice, D.T. Brunner, and J.J. Swain (Eds.), 144-152.

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

Glover, F. 1999. Scatter search and path relinking. *New Methods in Optimization*. D. Corne, M. Dorigo and F. Glover, Eds. McGraw Hill.

Holland, J.H. 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor.

Kelly, J., B. Rangaswamy and J. Xu. 1996. A scatter search-based learning algorithm for neural network training. *Journal of Heuristics* 2:129-146.

Laguna, M. 1997. Optimizaing complex systems with OptQuest. *Research Report*, University of Colorado.

AUTHOR BIOGRAPHIES

FRED GLOVER is the MediaOne Chaired Professor in Systems Science at the University of Colorado, Boulder. He has authored or co-authored more than two hundred eighty 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 journals in the U.S. and abroad, Dr. Glover has been featured as a National Visiting Lecturer by the Institute of Management Science and the Operations Research Society of America and has served as a host and lecturer in the U.S. National Academy of Sciences Program of Scientific Exchange. He is also a founder and Editor-in-Chief of the recently launched Journal of Heuristics. Professor Glover 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, the Institute of Management Science, the Operation Research Society, 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.

MANUEL LAGUNA is Associate Professor of Operations Management in the College of Business and Administration and Graduate School of Business Administration of the University of Colorado at Boulder. He received master's and doctoral degrees in Operations Research and Industrial Engineering from the University of Texas at Austin. He has done extensive research in the interface between computer science, artificial intelligence and operations research to develop solution methods for problems in areas such as production planning and scheduling, routing and network design in telecommunications, combinatorial optimization on graphs, and optimization of simulations. Dr. Laguna has more than forty publications, including articles in scientific journals such as Operations Research, Management Science, European Journal of Operational Research, Computers and Operations Research, IIE Transactions, and the International Journal of Production Research. He is the co-author of Tabu Search; the first book devoted to this innovative optimization technology. He is principal editor of the Journal of Heuristics, is in the editorial board of Combinatorial Optimization: Theory and Practice and has been guest editor for the Annals of Operations Research. Dr. Laguna is a member of the Institute for Operations Research and the Management Science, the American Society for Quality, the Institute of Industrial Engineering, and the International Honor Society Omega Rho.

JAMES P. KELLY is an Associate Professor of Operations Management in the College of Business and Administration and Graduate School of Business Administration at the University of Colorado in Boulder. He received his bachelors and masters degrees in Chemical Engineering from Bucknell University in 1981. He worked four years as a systems analyst and project manager at Texas Instruments in Dallas, Texas. He returned to graduate school and received his doctoral degree in Applied Mathematics and Operations Research from the University of Maryland in 1990. While at the University of Maryland, his research into methods for providing statistical confidentiality was funded by a series of grants from the United States Bureau of the Census. 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. Dr. Kelly has published several papers on topics such as tabu search and simulated annealing in various journals such as Operations Research, European Journal of Operational Research, and the INFORMS Journal on Computing. He is a member of the INFORMS. He is principal editor for the Journal of Heuristics.