

THE OPTQUEST APPROACH TO CRYSTAL BALL SIMULATION OPTIMIZATION

Fred Glover
James P. Kelly
Manuel Laguna

Graduate School of Business
University of Colorado
Boulder, Colorado 80309-0419, U.S.A.

ABSTRACT

The area of integrating simulation and optimization has recently undergone remarkable changes. New advances are making available applications of simulation that previously had been considered infeasible or beyond the scope of current technology to handle.

This paper describes recently developed computer software that effectively integrates Crystal Ball simulation and optimization. We demonstrate the ability of the system to find optimal and near optimal solutions in minutes for applications where an exhaustive examination of relevant alternatives requires days or months. The scaling of reduced time applies equally to settings where simulation runs require greater time. As a result, for a selected practical time limit, the new approach provides the opportunity to obtain decisions and scenarios whose quality greatly exceeds the quality available in the past.

1 INTRODUCTION

The ability to guide a series of simulations in the most effective way, instead of blindly itemizing scenarios (with the hope that at least one of those itemized will be one that is most desirable to implement) has been a long standing goal. The integration of simulation and optimization is putting this goal within practical reach.

Now, practical software exists that is capable of interfacing simulation and special search processes, to guide a series of simulations to uncover optimal or near optimal scenarios. Applications include the goals of finding:

- the best configuration of machines for production scheduling;
- 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 practical objectives.

Current advances not only open new doors for simulation, but also extend the areas to which optimization can be applied. The great advantage of simulation, which has been lacking in traditional optimization, is the ability to handle uncertainties and complex interactions (at a level that can scarcely be formulated in standard optimization models). The marriage of simulation and optimization offers a way to overcome this limitation.

Earlier attempts to create methods for optimizing simulations have largely been based on ad hoc approaches, or have relied on the user to run through a cumbersome "seat of the pants" analysis. Alternatively, they have been based on stochastic approximation designs whose main focus involves the analysis of convergence behavior in an infinite time frame. Not surprisingly, the results of such efforts have left a great deal to be desired from a practical standpoint. As a step in the direction of greater rigor within a finite time horizon, a systematic catalog of all possible alternatives may be examined by complete enumeration algorithms. Although this approach

guarantees optimal solutions, it has very limited application. As an example of an exceedingly simple setting where enumeration may be practicable, suppose that a simulation model depends on only two input factors, as for example, a portfolio consisting of two possible investments. If a feasible investment strategy would allow from \$1,000 to \$10,000 to be invested in each investment and we restrict ourselves to multiples of \$1000, then 100 simulation runs are needed to enumerate all possibilities. If each simulation is relatively short (e.g., 3 seconds), then the entire process could be done in 5 minutes of computer time. However, if instead of 2 investments we allow a very modest increase to consider 5 different investments, then enumerating all alternatives to find an optimal one would require $10^5 = 100,000$ simulations, or approximately 3.5 days of computer time. Of course, most simulation settings are not really so simple as the one described. It is easily possible for complete enumeration to take weeks or even months to carry out.

Recent developments in the area of optimization have led to the creation of intelligent search procedures capable of finding optimal or near optimal solutions to complex problems with large solution spaces, by exploring only a small fraction of the possible alternatives. In particular, the system we describe in this paper is the result of implementing a search technology known as *scatter search*, by customizing its operation to the context of optimizing simulations. The system searches for the best possible solution to an optimization problem, defined on a set of input factors to a simulation model.

2 CLASSICAL METHODS AND META-HEURISTICS

To give a background for understanding the rationale underlying our approach, we provide a brief review of classical optimization perspectives and the emerging role of meta-heuristics.

Formally, optimization deals with finding the best (optimal) solution to problems that in general can be expressed in the form of an objective function (to be optimized) and a set of constraints (which restrict the values of the decision variables). The best known optimization tool is linear programming, which model assumes that the objective function and constraints can be expressed using linear functions. Linear programming techniques are able to find optimal solutions to problems without the need to evaluate all possible alternatives. Models with thousands and even millions of variables can be solved with reasonable amounts of computer time.

Evidently, however, not all business and industrial problems can be expressed by means of a linear objective and linear equalities or inequalities. Many complex systems may not even have a convenient mathematical representation, linear or non-linear. Techniques such as linear programming and its cousins (non-linear programming and integer programming) generally require a number of simplifying assumptions about the real system to be able to properly frame the problem. One usual simplifying assumption is to disregard the so called “statistical fluctuations” of the system. For instance, a production process may be modeled by assuming that production times do not vary, permitting a single time estimate to be used for modeling purposes. The advantage, of course, is that once the problem has been formulated, well-established techniques are often able to find an optimal solution, provided the formulation is “congenial.” However, there are problems such as production scheduling where even the deterministic versions remain hard to solve. This is due to the combinatorial nature of these problems, as illustrated by the situation where the goal is to determine an optimal order in which to process a set of jobs. Even for fewer than a hundred jobs to be ordered, the number of alternative configurations is astronomical. (Seventy jobs would require longer than the age of the universe to enumerate using all of today’s computers working at once!)


Due to the realization that practical problems are not going to become any easier to solve, and that practitioners are seeking solutions to increasingly more complex problems, researchers have actively developed solution procedures that are referred to as *meta-heuristics*. Heuristics have been used for many years to provide approximate solutions to complex problems. For example, a production heuristic may be to give priority to jobs with the shortest estimated processing time. Depending on the context, this heuristic (or processing rule) may actually work fairly well. However, in some other situations the results may be disastrous, with dire consequences for equipment utilization, production lead times, and work-in-process inventory.

The alternatives seem at first to be less than encouraging: either to seek optimal solutions to simplified problems or to seek sub-optimal and possibly very poor solutions to complex systems. The area of meta-heuristics arose with the goal of providing something better, based on integrating high-level intelligent procedures and fast computer implementations. Numerous successful applications emerged, but with an accompanying limitation: typically they required highly problem-specific designs, customizing the solution procedures to each particular case. Every time a new problem surfaced, a new procedure needed to be developed. Meta-heuristic approaches

are based on general principles, but they also owe their efficiency to the knowledge of characteristics particular to each situation. In this sense, there is no separation between the model and the solution procedure. In fact, the solution procedure may be seen as a way of modeling the problem.


Of course, it is preferable to separate the solution procedure from the system we are trying to optimize if such a separation can be achieved successfully. The disadvantage of this “black box” approach (see Figure 1.1), is that the optimization procedure is generic and does not know anything about what goes on inside of the box.

Figure 1.1 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 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.

Figure 1.2 Coordination Between Optimization and Simulation.



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 1.2). 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 some termination criterion is satisfied (usually given by a limit expressing the user’s preference for the amount of time to be devoted to the search). The underlying components of our method, scatter search and tabu search, are briefly sketched in the next two sections.

3 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 early 1970s, at the University of Michigan (Holland 1975). Parallel to the development of GAs, Fred Glover, at the University of Colorado, established the principles and operational rules for tabu search (TS) and a related methodology known as *scatter search* (Glover 1977).

Scatter search has some interesting commonalities 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 (Glover 1994) is designed to operate on a set of points, called *reference points*, that constitute good solutions obtained from previous solution efforts. The approach systematically generates linear combinations of the reference points to create new points, each of which is mapped into an associated feasible point. Tabu search is then superimposed to control the composition of reference points at each stage. Tabu Search (see e.g., Glover and Laguna 1993) 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 in many cases find a globally optimal solution.

Similarities are immediately evident 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, several contrasts between these methods may be noted. The early GA approaches were 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, in the sense of being indifferent to choices among alternatives. However, the approach is designed to incorporate strategic probabilistic biases, taking 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 then may 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 combination of these factors creates a highly effective solution process. The incorporation of such designs is responsible for endowing our system with the ability to solve complex simulation-based problems with unprecedented efficiency.

4 TABU SEARCH BASICS

One way of intelligently guiding a search process is to forbid (or discourage) certain solutions from being chosen based on information that suggests these solutions may duplicate, or significantly resemble, solutions encountered in the past. In tabu search, this is often done by defining suitable attributes of moves or solutions, and imposing restrictions on a set of the attributes, depending on the search history. Two prominent ways for exploiting search history in TS are through *recency* and *frequency* memories. 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.

Short and long term components based on recency and frequency memory can be used separately or together in complementary TS search strategies. Note that this approach operates by implicitly modifying the neighborhood of the current solution. Tabu search in general includes many enhancements to the scheme sketched here, and we refer the interested reader to Glover and Laguna (1993) or Glover (1996). The details of the short-term and long-term adaptive memories, and a recovery strategy for both intensifying and diversifying the search are discussed in the following section.

5 AN OVERVIEW OF THE ALGORITHM

We assume that a solution to the optimization problem can be represented by a n -dimensional vector \mathbf{x} , where x_i may be a real or an integer bounded variable (for $i = 1, \dots, n$). In addition, we assume that the objective function value $f(\mathbf{x})$ can be obtained by running a related simulation model that uses \mathbf{x} as the value of its input factors. Finally, a set of linear constraints (equality or inequality) may be imposed on \mathbf{x} .

The algorithm starts by generating an initial population of reference points. The initial population may include points suggested by the user, and it always includes the following midpoint:

$$x_i = l_i + (u_i - l_i)/2,$$

where u_i and l_i are the upper and lower bounds on x_i , respectively. Additional points are generated with the goal of creating a diverse population. A population is considered diverse if its elements are “significantly” different from one another. We use a distance measure to determine how “close” a potential new point is from the points already in the population, in order to decide whether the point is included or discarded.

Every reference point \mathbf{x} is subjected to a feasibility test before it is evaluated (i.e., before the simulation model is run to determine the value of $f(\mathbf{x})$). The feasibility test consists of checking (one by one) whether the linear constraints imposed by the user are satisfied. An infeasible point \mathbf{x} is made feasible by formulating and solving a linear programming (LP) problem. The LP (or mixed-integer program, when \mathbf{x} contains integer variables) has the goal of finding a feasible \mathbf{x}^* that minimizes the absolute deviation between \mathbf{x} and \mathbf{x}^* .

The population size is automatically adjusted by the system considering the time that is required to complete one evaluation of $f(\mathbf{x})$ and the time limit the user has allowed the system to search. Once the population is generated, the procedure iterates in search of improved outcomes. At each iteration two reference points are selected to create four offspring. Let the parent-reference points be \mathbf{x}_1 and \mathbf{x}_2 , then the offspring \mathbf{x}_3 to \mathbf{x}_6 are found as follows:

$$\begin{aligned}\mathbf{x}_3 &= \mathbf{x}_1 + d \\ \mathbf{x}_4 &= \mathbf{x}_1 - d \\ \mathbf{x}_5 &= \mathbf{x}_2 + d \\ \mathbf{x}_6 &= \mathbf{x}_2 - d\end{aligned}$$

where $d = (\mathbf{x}_1 - \mathbf{x}_2)/3$. The selection of \mathbf{x}_1 and \mathbf{x}_2 is biased by the values $f(\mathbf{x}_1)$ and $f(\mathbf{x}_2)$ as well as the tabu search memory functions. An iteration ends by replacing the worst parent with the best offspring, and giving the surviving parent a tabu-active status for given number of iterations. In subsequent iterations, the use of two tabu-active parents is forbidden.

5.1 Restarting Strategy

In the course of searching for a global optimum, the population may contain many reference points with similar characteristics. That is, in the process of generating offspring from a mixture of high-quality reference points and ordinary reference points member of the current population, the diversity of the population may tend to decrease. A strategy that remedies this situation considers the creation of new population.

Our implementation of a restarting mechanism has the goal of creating a population that is a blend of high-quality points found in earlier explorations (we call these the *elite* points) complemented with points generated in the same way as during the initialization phase. The restarting procedure, therefore, injects diversity through newly generated points and preserves quality through the inclusion of elite points.

5.2 Adaptive Memory and the Age Strategy

Some of the points in the initial population may have poor objective function values. Therefore, they may never be chosen to play the role of a parent and would remain in the population until restarting. To additionally diversify the search, we increase the “attractiveness” of these unused points over time. This is done by using a form of long-term memory that is different from the conventional frequency-based implementation.

In particular, we introduce the notion of “age” and define a measure of “attractiveness” based on the age and the objective function value of a particular point. The idea is to use search history to make reference points not used as parents “attractive,” by modifying their objective function values according to their age.

At the start of the search process, all the reference points \mathbf{x} in a population of size p have zero age. At the end of the first iteration, there will be $p-1$ reference points from the original population and one new offspring. The ages of the $p-1$ reference points are made one and that of the new offspring zero. The process then repeats for the subsequent iterations, and the age of every reference point increases by one in each iteration except for the age of the new population member whose age is initialized to zero. (A variant of the above procedure sets the surviving parent’s age also to 0.)

Each reference point in the population has an associated age and an objective function value. These two values are used to define a function of attractiveness that makes an old high-quality point the most attractive. Low-quality points become more attractive as their age increases.

5.3 Neural Network Accelerator

This strategy is designed to increase the power of the system's search engine. The concept behind embedding a neural network is to "screen out" values \mathbf{x} that are likely to result in a very poor value of $f(\mathbf{x})$. 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. Engaging the neural network accelerator is an option to the user. When the neural network is used, 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 once again 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 $f(\mathbf{x})$ for \mathbf{x} -values found during the search, 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.

6 THE OPTQUEST SYSTEM

A commercial implementation of the system described above has been recently released under the name of OptQuest for Crystal Ball (1998). In its current version, OptQuest has been specifically customized to help users find optimal input parameter settings to simulation models built with Crystal Ball. (Crystal Ball is registered trademark of Decisioneering, Inc.) In order to use OptQuest the user first creates a Crystal Ball spreadsheet model. Once the simulation model has been created, an "OptQuest" option can be selected within Crystal Ball to access the optimization procedure.

A Portfolio Allocation Problem

We illustrate the operation of some of the features of OptQuest with a simulation optimization for a portfolio allocation problem.

Problem description

An investor has \$100,000 to invest in four assets. Below is a list of the assets' expected annual returns and associated distributions, and the minimum and maximum amounts the investor is comfortable allocating to each investment.

Figure 6.1 Portfolio Allocation Data

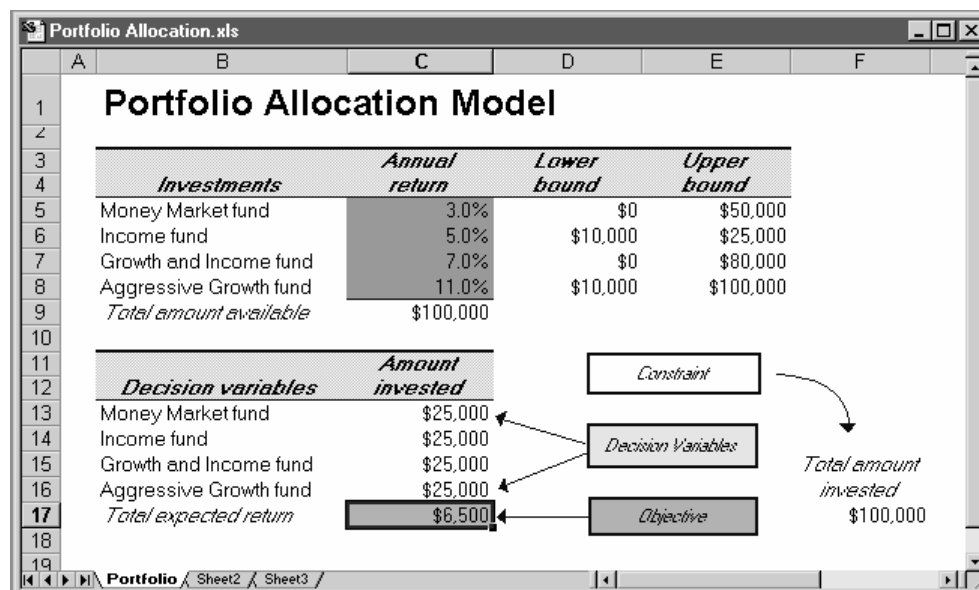
Investment	Annual return	Distribution	Distribution Parameters	Lower bound	Upper bound

Money market fund	3%	uniform	minimum: 2% maximum: 4%	\$0	\$50,000
Income fund	5%	normal	mean: 5% standard deviation: 5%	\$10,000	\$25,000
Growth and income fund	7%	normal	mean: 7% standard deviation: 12%	\$0	\$80,000
Aggressive growth fund	11%	normal	mean: 11% standard deviation: 18%	\$10,000	\$100,000

The source of uncertainty in this problem is the annual return of each asset. The more conservative assets, the Income and Money Market funds, have relatively stable annual returns, while the Aggressive Growth fund has higher volatility. The decision problem, then, is to determine how much to invest in each asset to maximize the total expected annual return while maintaining the risk at an acceptable level, and keeping within the minimum and maximum limits for each investment.

The spreadsheet for this problem is shown below.

Figure 6.2 Portfolio Allocation Spreadsheet Model

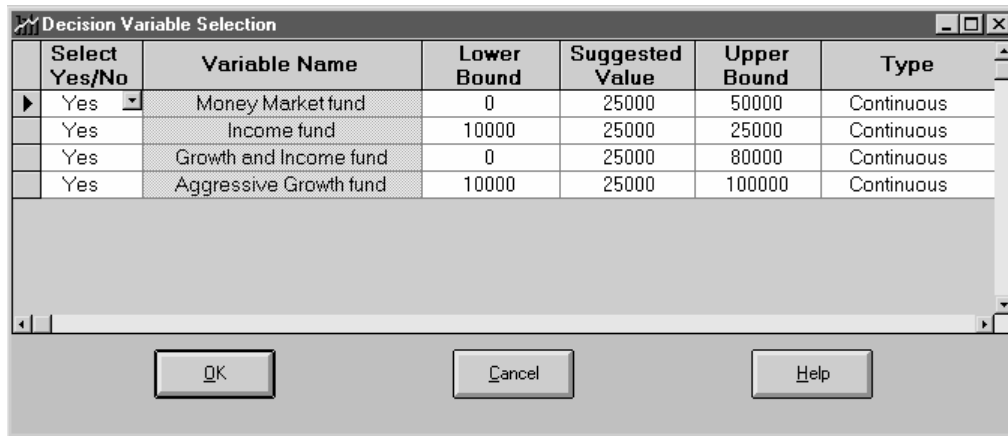


In this example, problem data are specified in rows 5 through 9. Model inputs (the values of the decision variables), the model output (the forecast objective), and the constraint (the total amount invested) are on the bottom half of the spreadsheet.

This model already has the assumptions and forecast cells defined in Crystal Ball.

Within OptQuest, the decision variables are displayed as shown in Figure 6.3.

Figure 6.3 OptQuest Decision Variable Selection Window

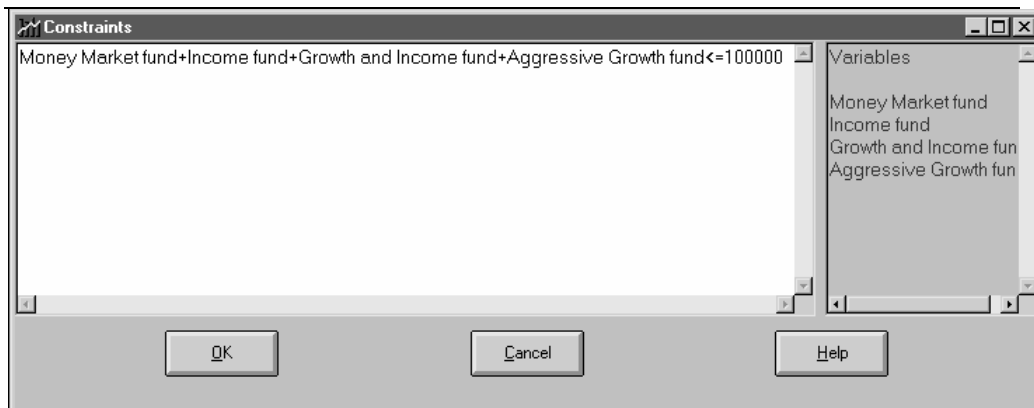


Every decision variable defined in the Crystal Ball model appears in the Decision Variable Selection window. The first column indicates whether the variable has been selected for optimization.

The other columns show the bounds, initial value, and type for each variable.

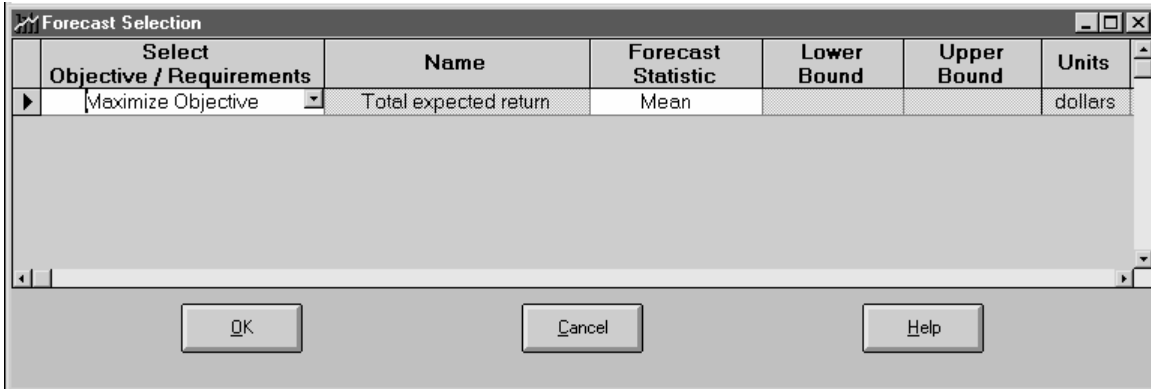
The Constraints window lets you specify any restrictions you can define with the decision variables. The constraint in this model limits the initial investment to \$100,000. The right side of the Constraints window lists the selected decision variables. The constraint window is shown in Figure 6.4.

Figure 6.4 Constraints Window



OptQuest requires that you select one forecast statistic to be the objective to minimize or maximize. In addition to defining an objective, you can define optimization requirements (bounds on simulation outputs).

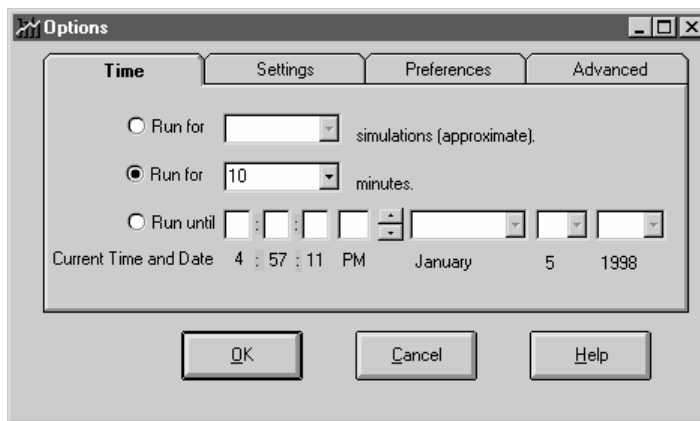
Figure 6.5 Forecast Selection Window



The goal for this example is to maximize the mean of the only forecast cell, as shown in Figure 6.5. For many problems, the mean (expected value) of the forecast is the most appropriate statistic to optimize—but it need not always be. For example, if an investor wants to maximize the upside potential of his portfolio, he might want to use the 90th or 95th percentile as the objective. The results would be solutions that have the highest likelihood of achieving the largest possible returns. Similarly, to minimize the downside potential of the portfolio, he might use the 5th or 10th percentile as the objective to minimize the possibility of large losses. You can use other statistics to realize different objectives.

The final step in the setup of an optimization is to specify the optimization time as shown in Figure 6.6.

Figure 6.6 OptQuest Options Window



The Time tab lets you specify the total time that the system searches for the best solutions for the decision variables. You can either enter the number of minutes to run an optimization or a date and time for the process to stop. The default optimization time is 10 minutes.

At this point, the optimization process is begun. The final results are shown in Figure 6.7.


Figure 6.7 OptQuest Solution Results

OptQuest for Crystal Ball					
					Optimization File
					UnNamed.opt Crystal Ball Simulation
Optimization is Complete					
Simulation	Maximize Objective Total expected return Mean	Money Market fund	Income fund	Growth and Income fund	Aggressive Growth fund
1	6379.50	25000.0	25000.0	25000.0	25000.0
2	7044.59	8085.46	10044.3	62551.8	19318.5
4	8483.32	2974.36	17288.3	24342.6	54443.1
9	9602.11	7128.42	11162.0	567.317	81142.3
59	10076.4	2244.22	10000.00	0.00000	87755.8
Best: 138	10252.0	0.00000	10000.00	0.00000	90000.0

The last line in the Status And Solutions window shows the best solution found by OptQuest. All the money is allocated to the fund that has the highest return—the Aggressive Growth fund—with the exception of the minimal amount in the Income fund that the investor required.

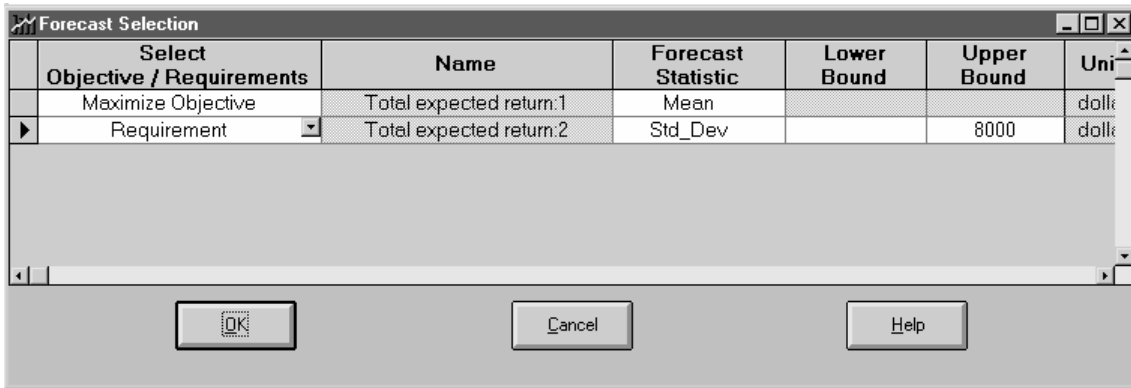
The investor's strategy maximized the return of the portfolio, but at a price: high risk due to high volatility and little diversification. Is this really what the investor wanted? To find out, the investor must interpret the results by analyzing the distribution of profits associated with the best solution. This distribution is shown in Figure 6.8.

Figure 6.8 Portfolio Allocation Forecast Chart



In portfolio management, controlling the variability of the solution to minimize risk can be just as important as achieving large expected returns. Suppose that this same investor wants to reduce the uncertainty of returns for the portfolio, while still attempting to maximize the expected return. You might want to find the best solution for which the standard deviation is much lower, say below \$8,000. This adds a requirement that the standard deviation of the expected returns must be less than \$8,000 for a solution to be considered feasible (Figure 6.9).

Figure 6.9 Objective and Requirement Specification



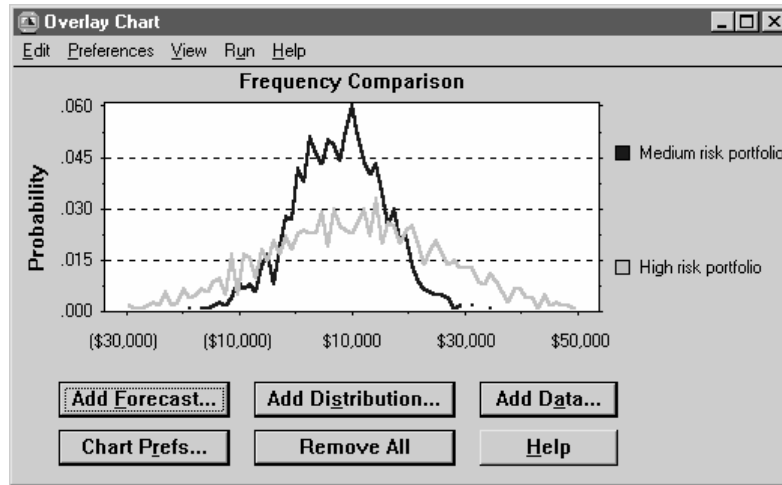
The results of the new search are shown in Figure 6.10

Figure 6.10 Results with Risk Control

Simulation	Maximize Objective Total expected return:1 Mean	Requirement Total expected return:2 Std_Dev	Money Market fund	Income fund	Growth and Income fund	Aggressive Growth fund
1	10361.0	16286.9	0.00000	10000.00	0.00000	90000.0
2	5072.91	3218.73	46142.0	24894.9	18963.0	10000.00
3	6484.38	7329.75	17157.8	14187.5	58654.7	10000.00
7	6881.64	7666.32	36495.3	15907.3	5833.30	41764.1
29	6934.59	7190.01	30651.2	17558.7	13614.4	38053.4
30	6948.35	7311.53	30667.6	17426.1	12541.3	38930.4
31	6951.05	7299.57	30665.8	17434.2	12725.2	38832.6
36	6953.59	7296.53	30665.0	17436.6	12806.7	38802.1
40	6954.73	7295.11	30664.7	17437.3	12844.4	38787.8
Best: 50	6961.88	7258.70	30662.7	17445.9	13404.5	38486.8

This solution has significantly reduced the variability of the total expected return, even though it now has a lower mean return. The portfolio achieved this by finding the best diversification of conservative and aggressive investments. Thus, the investor must make the trade-off between higher returns and higher risk or lower returns and lower risk. Figure 6.11 compares the two solutions.

Figure 6.11 Comparison of Solutions



7 CONCLUSIONS

In this paper we have described recent advances in optimization and simulation technologies that have made possible the development of a system that can effectively perform the task of optimizing simulations. We showed the benefits of adapting the approach known as scatter search in this context, while complementing this methodology with appropriate tabu search elements.

We have also demonstrated the capabilities of a commercial software package that successfully implements the notion of optimizing a complex system (in this case represented by a simulation model). The software package includes a scatter search / tabu search module, a mixed integer programming solver, and a procedure to configure and train neural networks. All these tools are integrated by a user-friendly interface.

We are currently undertaking additional research in the area of optimizing simulations with the goal of producing systems that are still more effective for dealing with the growing complexity of practical problems. The importance of integrating the complementary realms of optimization and simulation assures that future advances will have a high impact on real world applications

REFERENCES

- Glover, F. 1977. Heuristics for integer programming using surrogate constraints. *Decision Sciences* 8:156-166.
- Glover, F. 1994. Genetic algorithms and scatter search: unsuspected potentials. *Statistics and Computing* 4:131-140.
- Glover, F. 1996. Tabu search and adaptive memory programming—advances, applications and challenges. *Interfaces in Computer Science and Operations Research*, eds. Barr, Helgason and Kennington, Kluwer Academic Publishers, 1997.
- Glover, F. and M. Laguna. 1993. Tabu search. In *Modern Heuristic Techniques for Combinatorial Optimization*, ed. C. Reeves, 60-150. Blackwell Publishers: Oxford.
- Holland, J. H. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor.
- OptQuest for Crystal Ball User's Guide. Decisioneering Inc. Denver, CO and Optimization Technologies, Inc. Boulder, CO, 1998.