Proceedings of the 2004 Winter Simulation Conference
R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, eds.

SIMULATION OPTIMIZATION USING BALANCED EXPLORATIVE AND EXPLOITATIVE SEARCH

Andrei A. Prudius
Sigrún Andradóttir

School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205, U.S.A.

ABSTRACT

We present a new random search method for solving simulation optimization problems. Our approach emphasizes the need for maintaining the right balance between exploration and exploitation during various stages of the search. Exploration represents global search for promising solutions within the entire feasible region, while exploitation involves local search of promising subregions. Preliminary numerical results are provided that show the performance of the method applied to solve deterministic and stochastic optimization problems.

1 INTRODUCTION

Consider an optimization problem of the form

$$\max_{\theta \in \Theta} f(\theta) = \mathrm{E}[X_\theta], \tag{1}$$

where $f : \Theta \to \mathbb{R}$ is the objective function, Θ is the discrete feasible region, and X_θ is a random variable whose distribution depends on the decision parameter θ, for all $\theta \in \Theta$. We further assume that for any given solution $\theta \in \Theta$, the expected value in (1) cannot be computed exactly, but instead can only be estimated via a "black-box" simulation procedure. This feature, together with the necessity to optimize, make the problem in (1) especially difficult (see Fu 2002).

In past decades, there has been a growing interest in solving discrete simulation optimization problems. The recent work includes several newly developed random search methods, such as the stochastic ruler methods of Yan and Mukai (1992) and Alrefaei and Andradóttir (2001, 2004), the stochastic comparison methods of Gong, Ho, and Zhai (1999) and Andradóttir (1999), the simulated annealing algorithms of Gelfand and Mitter (1989), Gutjahr and Pflug (1996), and Alrefaei and Andradóttir (1999), and the nested partitions methods of Shi and Ólafsson (2000) and Pichitlamken and Nelson (2003). For recent overviews on the topic, including discussion of simulation optimization techniques other than random search, the reader is referred to Andradóttir (1998), Fu (2002), and references therein.

This paper focuses on developing a new random search approach called Balanced Explorative and Exploitative Search (BEES) for solving simulation optimization problems. We also discuss the tradeoffs between searching the entire feasible region (exploration) and locally searching promising subregions (exploitation). The proposed approach suggests adaptive use of exploration and exploitation during various stages of the search. For more comprehensive development of the approach, convergence proofs, and additional numerical results, the interested reader is referred to Prudius and Andradóttir (2004).

The remainder of the paper is organized as follows. In Section 2, we present a BEES algorithm for optimizing deterministic functions and provide preliminary numerical results. In Section 3, we modify the BEES method of Section 2 to handle noisy responses and illustrate its performance via a numerical example. Concluding remarks are given in Section 4.

2 DETERMINISTIC OPTIMIZATION

In this section, we discuss the BEES approach for optimizing deterministic objective functions. As its name indicates, this approach attempts to maintain the appropriate balance between exploration and exploitation. Note that if a good subregion within Θ has been identified, then it is sensible to spend more effort searching locally for better solutions. By contrast, if a good subregion has not been found, then more priority should be given to exploring the search space Θ. This observation suggests that it is appropriate to switch from exploration to exploitation at some point during the search. Unfortunately, it can be difficult to determine when this switch should take place. Observe that if this switch is performed either too early or too late in the search, then the convergence to the optimal solution(s) might be slow. Our approach takes a different perspective. We propose maintaining the appropriate balance between exploitation and exploration during the various stages of the search. We now present a BEES algorithm in Section 2.1 and provide numerical results for this algorithm in Section 2.2.

2.1 BEES Algorithm for Deterministic Problems

In this section, we describe a specific version of the BEES method. This approach adaptively alternates between sampling from different distributions. One sampling distribution aims at exploring the entire feasible region (exploration or global search), while another class of distributions (one for each feasible point) aims at searching promising subregions (exploitation or local search). The search nature (global or local) is reviewed each time k points have been sampled. Let f^* be the function value of the best solution θ^* found so far (in this case maximum, see equation (1)) and let f_l^* be the function value of the best point found the last time local search was performed. Define Δ to be the improvement in the function value between the current and preceding reviews and D to be the Euclidian distance between the points where the corresponding function values were achieved. Then the pseudo-code for how the sampling distribution is updated is given in Algorithm 1 below. If the flag LocalSearch is true, then local search is performed. Otherwise, global search is done. Note that the sampling distribution update procedure requires two thresholds, namely the distance threshold d and the improvement threshold δ.

Algorithm 1: Sampling Distribution Update Procedure

if LocalSearch=true **then**
  if $\Delta \leq \delta$ **then**
    $f_l^* \leftarrow f^*$
    LocalSearch $\leftarrow$ false
  end if
else
  if $\Delta \leq \delta$ **then**
    if $f^* - f_l^* \geq \delta$ **then**
      LocalSearch $\leftarrow$ true
    end if
  else
    if $D \leq d$ **then**
      LocalSearch $\leftarrow$ true
    end if
  end if
end if

We now briefly discuss the motivation for the procedure. Observe that the switch from local to global search occurs only if the improvement in the objective function value between successive reviews is small (less than δ). Hence, if local search is not yielding much improvement in the objective function value, then there is little merit in continuing searching locally. On the other hand, if local search is making good progress, then the search will stay local.

On the other hand, there are two ways in which the BEES algorithm switches from global to local search. The first way occurs when the perceived improvement Δ is small, but substantial improvement has been achieved (larger than δ) since the last switch from local to global search. The second way occurs when the improvement between successive reviews is large and the distance D is small (this is sensible because the improvement has been local in nature and hence, switching to local search may be beneficial). The pseudo-code for this BEES algorithm is given in Algorithm 2 (one iteration of this algorithm corresponds to one execution of the statements inside the while loop of Algorithm 2).

Algorithm 2: BEES Algorithm

counter $\leftarrow 0$
LocalSearch $\leftarrow$ false
Sample a solution θ from global distribution
Evaluate objective function at θ
Let $f_l^* \leftarrow f(\theta)$ and $\theta^* \leftarrow \theta$
while Stopping criterion is not satisfied **do**
  if LocalSearch=true **then**
    Sample a solution θ from local distribution corresponding to θ^*
  else
    Sample a solution θ from global distribution
  end if
  Evaluate objective function at θ
  Update θ^* and f^* (if needed)
  counter $\leftarrow$ counter+1
  if counter=k **then**
    Update Δ and D
    Update search nature (use Sampling Distribution Update Procedure)
    counter $\leftarrow 0$
  end if
end while
Present θ^* as the estimate of the optimal solution

2.2 Deterministic Example

Consider the optimization problem (1) with

$$f(\theta) = \max\{f_1(\theta), f_2(\theta), 0\} \tag{2}$$

and $\Theta = \{\theta = (\theta_1, \theta_2) \in \mathbb{N}^2 : 0 \le \theta_1, \theta_2 \le 49\}$, where $f_1(\theta) = -(0.4\theta_1 - 5)^2 - 2(0.4\theta_2 - 17.2)^2 + 7$ and $f_2(\theta) = -(0.4\theta_1 - 12)^2 - (0.4\theta_2 - 4)^2 + 4$. This objective function is of interest because it has two hills of different heights (4 and 6.96), located relatively far apart (the hill of height 4 is centered at (30, 10) and the hill of height 6.96 is centered at (12, 43) and (13, 43)), and separated by a flat valley (of height 0).

We now describe the implementation details of the BEES method specified in Algorithm 2 applied to solve the optimization problem (2). The global sampling distribution is the uniform distribution on the feasible space Θ and the local sampling distribution for each $\theta \in \Theta$ is the uniform distribution on $N(\theta^*)$, where $N(\theta) = \{(x_1, x_2) \in \Theta \setminus \{\theta\} : |x_i - \theta_i| \le 1 \text{ for } i = 1, 2\}$. Finally, the algorithm is terminated after a fixed number of objective function evaluations have been performed.

The performance of this BEES algorithm with parameters $k = 20$, $\delta = 0.01$, and $d = 5$ was compared to the Simulated Annealing (SA) algorithm with constant temperature (see Alrefaei and Andradóttir 1999). We used two different neighborhood structures and refer to the resulting implementations of the SA method as Global SA and Local SA. The neighborhood structure for Global SA is such that each solution is a neighbor of every other solution. By contrast, the neighborhood of a solution θ is $N(\theta)$ for Local SA. The temperature parameter is set to 0.1 and 1, respectively for the Global and Local SA algorithms. The initial state is selected randomly for all three algorithms. The performance of the algorithms was compared based on 100 replications. Figure 1 shows the average performance of the three approaches as the simulation effort increases.

It is obvious from Figure 1 that the BEES method performs considerably better than both versions of the simulated annealing algorithm. This numerical example supports the idea that the search should be well balanced in the use of exploration and exploitation. Observe that Global SA essentially performs explorative search only, while Local SA is exploitative in nature. The substantially worse behavior of Local SA in this example can be explained by the fact that an initial solution might be far away from the subregions containing good designs and the SA algorithm might take a while identifying a good subregion using the local neighborhood structure.

3 STOCHASTIC OPTIMIZATION

This section extends the BEES algorithm described in Section 2 to optimizing stochastic objective functions. The modifications of this BEES algorithm are described in Section 3.1 and a numerical example is provided in Section 3.2.

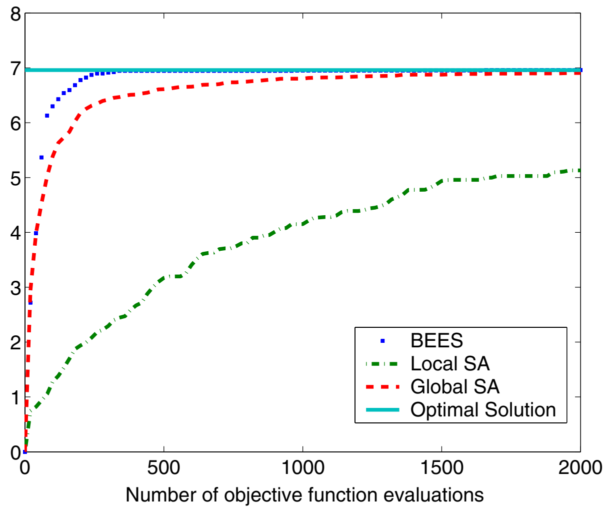

Figure 1: Average Objective Function Value at the Estimated Optimal Solution for the Deterministic Problem

3.1 BEES Algorithm for Stochastic Problems

We first discuss how we modify the BEES algorithm described in Section 2. Define the best point θ^* at any time as the solution with the best estimated objective function value and let Δ be the improvement in the estimated objective function value between the two most recent reviews. By estimated objective function at any solution, we mean the cumulative average of all observations taken at this solution. Algorithm 2 is generalized as follows. First, the best point θ^* is sampled with probability α, and with probability $1 - \alpha$, the sampling is done as explained in Section 2.1 (depending on the nature of the search). Second, the local (global) search is conducted for k_l (k_g) iterations before attempting to switch to global (local) search (by invoking Algorithm 1). Typically, the parameters k_l and k_g satisfy $k_l \ge k_g$. Third, we modify Algorithm 1 to switch to local search if global search has been conducted for g consecutive reviews. Finally, whenever a solution is sampled, m simulation replications are conducted at it.

As the estimator of the optimal solution in this BEES approach, we propose to use the estimator of Andradóttir (2004). More specifically, $\theta \in \Theta$ is chosen to be the estimate of the optimal solution if it has the best estimated objective function value among solutions which have been replicated at least n^γ times, where n is the iteration number and $0 \le \gamma < 1$ (and hence the estimated optimal solution may be different from θ^* if $\gamma > 0$). If the set of systems which have been sampled at least n^γ times is empty, then the estimate of the optimal solution is the solution θ^*.

3.2 Stochastic Example

The test problem used here is the same as in (2) but with white noise added. The noise is a normally distributed random variable with mean 0 and variance 50. Observe that the standard deviation of the noise is roughly the same as the range of the objective function values. This makes the response surface highly noisy and hence this problem is relatively difficult to solve.

For the numerical studies, the additional parameters of the BEES algorithm under consideration were set as follows: $k_l = 25$, $k_g = 20$, $g = 5$, $\alpha = 0.3$, $m = 10$, and $\gamma = 0.5$. As before, the SA method is implemented as described by Alrefaei and Andradóttir (1999) with the parameter values given in Section 2.2. The number of replications at the current and candidate solutions conducted in each iteration of the Local and Global SA approaches was set equal to 10. Again the performance of the algorithms was compared based on 100 replications. Figure 2 shows the average performance of the estimated optimal solution at each point in time as the simulation effort increases. The considerably worse performance of the Global SA approach in this example can be explained by the fact that the estimator of the optimal solution is very noisy in this case. The remaining conclusions are similar to those in Section 2, except that convergence is slower for all three algorithms due to the noise in the estimated objective function values.

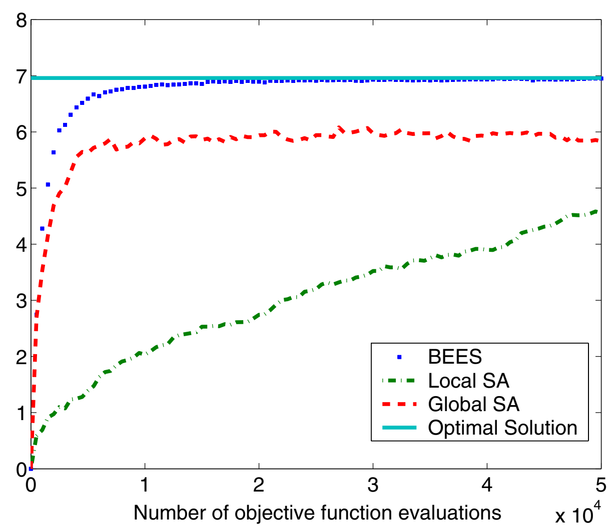

Figure 2: Average Objective Function Value at the Estimated Optimal Solution for the Stochastic Problem

4 CONCLUSIONS

This paper has presented a new random search method for simulation optimization. The approach emphasizes the need for maintaining the right balance between exploration and exploitation during various stages of the search. Preliminary numerical show promising performance of the method. More numerical studies are of course required to understand better the behavior of the approach. We are also interested in exploring whether the performance of existing methods for simulation optimization can be improved by incorporating ideas from this paper.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 0000135, Grant No. 0217860, and Grant No. 0400260.

REFERENCES

Alrefaei, M. H., and S. Andradóttir. 1999. A simulated annealing algorithm with constant temperature for discrete stochastic optimization. *Management Science* 45 (5): 748-764.

Alrefaei, M. H., and S. Andradóttir. 2001. A modification of the stochastic ruler method for discrete stochastic optimization. *European Journal of Operational Research* 133 (1): 160-182.

Alrefaei, M. H., and S. Andradóttir. 2004. Discrete stochastic optimization using variants of the stochastic ruler method. Working paper.

Andradóttir, S. 1998. Simulation optimization. In *Handbook of simulation: Principles, methodology, advances, applications, and practice*, ed. J. Banks, 307-333. New York: Wiley.

Andradóttir, S. 1999. Accelerating the convergence of random search methods for discrete stochastic optimization. *ACM Transactions on Modeling and Computer Simulation* 9 (4): 349-380.

Andradóttir, S. 2004. Simulation optimization with countably infinite feasible regions: Efficiency and convergence. Working paper.

Fu, M. C. 2002. Optimization for simulation: Theory vs. practice. *INFORMS Journal on Computing* 14 (3): 192-215.

Gelfand, S. B., and S. K. Mitter. 1989. Simulated annealing with noisy or imprecise energy measurements. *Journal of Optimization Theory and Applications* 62 (1): 49-62.

Gong, W.-B., Y.-C. Ho, and W. Zhai. 1999. Stochastic comparison algorithm for discrete optimization with estimation. *SIAM Journal on Optimization* 10 (2): 384-404.

Gutjahr, W. J., and G. Ch. Pflug. 1996. Simulated annealing for noisy cost functions. *Journal of Global Optimization* 8 (1): 1-13.

Pichitlamken, J., and B. L. Nelson. 2003. A combined procedure for optimization via simulation. *ACM Transactions on Modeling and Computer Simulation* 13 (2): 155-179.

Prudius, A. A., and S. Andradóttir. 2004. Balanced explorative and exploitative search for simulation optimization. Working paper.

Shi, L., and S. Ólafsson. 2000. Nested partitions method for stochastic optimization. *Methodology and Computing in Applied Probability* 2 (3): 271-291.

Yan, D., and H. Mukai. 1992. Stochastic discrete optimization. *SIAM Journal on Control and Optimization* 30 (3): 594-612.

AUTHOR BIOGRAPHIES

ANDREI A. PRUDIUS is a Ph.D. candidate in the School of Industrial and Systems Engineering at the Georgia Institute of Technology. He received a B.S. in Industrial Engineering from Boğaziçi University in Istanbul, Turkey, in 2001. His research interests are in simulation, and especially in simulation optimization. His e-mail address is <aprudius@isye.gatech.edu>.

SIGRÚN ANDRADÓTTIR is a Professor in the School of Industrial and Systems Engineering at the Georgia Institute of Technology. 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 stochastic processes. Her e-mail address is <sa@isye.gatech.edu>.