Proceedings of the 1995 Winter Simulation Conference
ed. C. Alexopoulos, K. Kang, W. R. Lilegdon, and D. Goldsman

A NEW SEARCH ALGORITHM FOR DISCRETE STOCHASTIC OPTIMIZATION

Mahmoud H. Alrefaei

Department of Industrial Engineering
University of Wisconsin - Madison
1513 University Avenue
Madison, WI 53706, U.S.A.

Sigrún Andradóttir

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

ABSTRACT

We present a new method for finding a global optimal solution to a discrete stochastic optimization problem. This method resembles the simulated annealing method for discrete deterministic optimization. However, in our method the annealing schedule (the cooling temperature) is kept fixed, and the mechanism for estimating the optimal solution is different from that used in the original simulated annealing method. We state a convergence result that shows that our method converges almost surely to a global optimal solution under mild conditions. We also present empirical results that illustrate the performance of the proposed approach on a simple example.

1 INTRODUCTION

Consider the problem of optimizing an objective function over a discrete feasible set of parameters in situations where the objective function does not have a closed form expression, so that its values have to be estimated or simulated. In mathematical notation, this problem can be represented as

$$\min_{x \in \mathcal{S}} f(x), \tag{1}$$

where $\mathcal{S}$ is a discrete set and $f : \mathcal{S} \to \Re$ is a deterministic function whose evaluations all include some noise. Often f is the expected performance of a complex stochastic system, so that for all $x \in \mathcal{S}$,

$$f(x) = E[h(x, Y_x)], \tag{2}$$

where Y_x is some random variable that depends on the parameter x and h is a deterministic function.

In the deterministic case, when the values of the function f can be evaluated easily, then one can use a discrete deterministic optimization technique such as the branch-and-bound method or the simulated annealing method to solve the optimization problem (1). But when the function evaluations are difficult to come by and include noise, deterministic optimization techniques can not be used to locate the optimal solution. Instead a specialized method should be used to deal with functions with noisy evaluations.

For objective functions of the form given in equation (2), when independent and identically distributed observations $Y_x(1), Y_x(2), ..., Y_x(n)$ of the random variable Y_x can be generated for all $x \in \mathcal{S}$, and when $\mathcal{S}$ is finite, one could think of replacing the original optimization problem (1) by an approximate optimization problem

$$\min_{x \in \mathcal{S}} \overline{f_n(x)}, \tag{3}$$

where $\overline{f_n(x)} = \frac{1}{n} \sum_{i=1}^{n} h(x, Y_x(i))$, for all $x \in \mathcal{S}$, and then use a deterministic optimization technique to solve this approximation of the original optimization problem. But in order to guarantee convergence, the sample size n may have to be very large, which implies that it may require too much computer time to get a good estimate of the solution. In addition, it can be difficult to determine how large n should be for (3) to be a good approximation of the original optimization problem (1) - (2).

The approach outlined above (using equation (3) to approximate the original optimization problem (1) - (2)) would be more convenient if the number of alternatives were small, say less than twenty. In such situations, methods for selecting the best system can be used to solve the discrete stochastic optimization problem (1) - (2). Goldsman, Nelson and Schmeiser (1991) present a brief overview of three methods for selecting the best system: interactive analysis, ranking and selection, and multiple comparisons. For more details on ranking and selection and multiple comparison procedures, see Goldsman and Nelson (1994). Other related work has appeared in the literatures on the multi-armed bandit problem and on learning automata, see for example Lai and Robbins (1985), Devroye (1976), and Yakowitz and

Lugosi (1990).

More recently, Yan and Mukai (1992), Gong, Ho, and Zhai (1993), and Andradóttir (1995, 1996) have proposed new methods for discrete stochastic optimization. These methods all generate a Markov chain on the feasible set $\mathcal{S}$ of the optimization problem (1). However, the specifics of this Markov chain, as well as the approach used to estimate the solution, differ between the methods. Both Yan and Mukai (1992) and also Gong, Ho, and Zhai (1993) show that under certain conditions, their methods converge in probability to a global solution of the underlying optimization problem. On the other hand, Andradóttir proves that under certain conditions, the method proposed in Andradóttir (1995) converges almost surely to a local optimizer of the objective function, whereas the method proposed in Andradóttir (1996) converges almost surely to a global optimizer.

In this paper we introduce a new method for discrete stochastic optimization that resembles the simulated annealing method for discrete deterministic optimization. In particular, the proposed approach generates a Markov chain on the state space $\mathcal{S}$ of the underlying optimization problem (1) that strongly resembles the Markov chain generated by the simulated annealing algorithm. However, the proposed algorithm differs from the simulated annealing algorithm in that it uses a constant, rather than decreasing, cooling temperature (see the discussion in Sections 2 and 3). Also, our new method employs the mechanism proposed by Andradóttir (1995, 1996) to estimate the optimal solution, which is different from the mechanism used by the original simulated annealing approach.

This paper is organized as follows. In Section 2 we give a brief description of the simulated annealing algorithm for discrete deterministic optimization. In Section 3 we present the proposed discrete stochastic optimization algorithm and state conditions under which this algorithm is guaranteed to converge almost surely to a global optimal solution of the optimization problem (1) – (2). Section 4 contains empirical results for the proposed method, and some concluding remarks are given in Section 5.

2 THE SIMULATED ANNEALING ALGORITHM

Simulated annealing is a random search algorithm that is designed to find a global optimizer of a given deterministic function f over a finite set $\mathcal{S}$ in situations where the objective function f may have many local minima. Before stating the details of the algorithm, we need the following definitions and assumptions:

Definition 1 *For each $x \in \mathcal{S}$, there exists a subset $N(x)$ of $\mathcal{S} \setminus \{x\}$, which is called the set of neighbors of x.*

Assumption 1 *For any $x, x' \in \mathcal{S}$, x' is reachable from x; i.e., there exists a finite sequence, $\{n_i\}_{i=0}^{l}$ for some l, such that*

$$x_{n_0} = x, x_{n_l} = x', x_{n_{i+1}} \in N(x_{n_i}), i = 0, 1, 2, ..., l-1.$$

Definition 2 *Let $R : \mathcal{S} \times \mathcal{S} \to [0, 1]$ be a nonnegative function that satisfies*

1. *$R(x, x') > 0 \Leftrightarrow x' \in N(x)$, and*
2. *$\sum_{x' \in \mathcal{S}} R(x, x') = 1$.*

Then $R(x, x')$ is called the probability of generating x' from x.

Assumption 2 *The neighbor system $\{N(x) : x \in \mathcal{S}\}$ and the transition probability function R are symmetric, i.e.,*

1. *$x' \in N(x) \Leftrightarrow x \in N(x')$, and*
2. *$R(x, x') = R(x', x)$.*

Now we state the original simulated annealing algorithm. Note that $\{T_m\}$ is a sequence of positive scalars, and that for all $x \in \Re$, $[x]^+ = x$ when $x \geq 0$ and $[x]^+ = 0$ otherwise.

Algorithm 1

Step 0: *Select a starting point $X_0 \in \mathcal{S}$.*

Step 1: *Given $X_m = x$ choose a candidate $Z_m \in N(x)$ with probability distribution*

$$P[Z_m = z | X_m = x] = R(x, z),$$

where $N(x)$ and $R(x, z)$ are defined in Definitions 1 and 2.

Step 2: *Given $Z_m = z$, generate $U_m \sim U[0, 1]$, and set*

$$X_{m+1} = \begin{cases} z & \text{if } U_m \leq p_m, \\ x & \text{otherwise,} \end{cases}$$

where

$$p_m = \exp\left[\frac{-[f(z) - f(x)]^+}{T_m}\right].$$

Step 3: *Let $m = m + 1$. Go to Step 1.*

Note that the simulated annealing algorithm allows "hill-climbing" moves (to go from x to x' with $f(x) < f(x')$) in order to avoid local minima. Also, the higher the "cooling temperature" T_m, the more likely it is that a "hill-climbing" move will be made. The initial temperature and the rate of decrease of the temperature are important parameters that affect the speed of convergence of the simulated annealing algorithm and the quality of the final configuration.

Let

$$\mathcal{S}^* = \{x \in \mathcal{S} : f(x) \leq f(y) \text{ for all } y \in \mathcal{S}\}$$

be the set of global solutions to the optimization problem (1). To show that Algorithm 1 converges in probability to an element of $\mathcal{S}^*$, consider the undirected graph with the states in $\mathcal{S}$ forming the nodes and the neighborhood structure $\{N(x) : x \in \mathcal{S}\}$ forming the edges; i.e., if $x' \in N(x)$, then the edge (x, x') belongs to this undirected graph. The distance $d(x, x')$ between two nodes x and x' is defined to be the length (number of edges) of the minimum path from x to x'. Let

$$\mathcal{S}_m = \{x \in \mathcal{S} : f(y) \leq f(x) \text{ for all } y \in N(x)\}$$

be the set of all the points that are local maxima for the cost function. Define the radius of the graph

$$r = \min_{x \in \mathcal{S} \setminus \mathcal{S}_m} \max_{x' \in \mathcal{S}} d(x, x'),$$

and let

$$L = \max_{x \in \mathcal{S}} \max_{y \in N(x)} |f(y) - f(x)|.$$

The following result follows directly from Theorems 4.2 and 5.1 and Propositions 3.2 and 5.1 in Mitra, Romeo, and Sangiovanni-Vincentelli (1986).

Theorem 1 *Under Assumptions 1 and 2, if the sequence $\{T_m\}$ satisfies*

$$T_m = \frac{\gamma}{\log(m + m_0 + 1)}, \quad m = 0, 1, 2, \ldots,$$

where m_0 is any parameter satisfying $1 \leq m_0 < \infty$ and $\gamma \geq rL$, then

$$\lim_{m \to \infty} P\{X_m \in \mathcal{S}^*\} = 1.$$

3 THE PROPOSED ALGORITHM FOR NOISY FUNCTIONS

The original simulated annealing algorithm (Algorithm 1) is not designed to solve optimization problems that are stochastic in the sense that the evaluations of the objective function f (see equations (1) and (2)) involve noise. Gong, Ho and Zhai (1993) mention that in order for the original simulated annealing algorithm to be applied to solve such discrete stochastic optimization problems, one needs to obtain accurate estimates of the function values, and that this will cost too much computer time. Haddock and Mittenthal (1992) implemented this basic idea. However, in order to reduce the required computer time, they "employed a more rapid temperature decreasing (i.e., heuristic) annealing schedule" (so that standard convergence results such as Theorem 1 do not apply).

We propose another algorithm which does not require accurate estimates of the function values at every iteration. This algorithm strongly resembles the original simulated annealing approach, and in particular, it has the "hill-climbing" feature which allows it to escape from local minima. However, our method uses the criterion of Andradóttir (1995, 1996) in determining the estimate of the optimal solution. In particular, the state that the algorithm has visited most often at any given time will be the estimate of the optimal solution. We will need the following assumption:

Assumption 3 *The temperature T is a positive (constant) real number. In addition, $\{K_m\}$ is a sequence of positive integers satisfying $\lim_{m \to \infty} K_m = \infty$.*

Now we state the proposed simulated annealing algorithm for noisy functions. Note that after m iterations, X_m is the current state of the Markov chain generated by the algorithm, for all $x \in \mathcal{S}$, $V_m(x)$ is the number of times the Markov chain $\{X_n\}$ has visited state x in the first m iterations, and X_m^* is the state that the Markov chain $\{X_n\}$ has visited most often in the first m iterations.

Algorithm 2

Step 0: *Select a starting point $X_0 \in \mathcal{S}$. Let $V_0(X_0) = 1$, and $V_0(x) = 0$, for all $x \in \mathcal{S}$, $x \neq X_0$. Let $m = 0$ and $X_m^* = X_0$.*

Step 1: *Given $X_m = x$, choose a candidate $Z_m \in N(x)$ with probability distribution*

$$P[Z_m = z | X_m = x] = R(x, z),$$

where $N(x)$ and $R(x, z)$ are defined in Definitions 1 and 2.

Step 2: *Given $Z_m = z$, generate independent observations $Y_z(1), Y_z(2), \ldots, Y_z(K_m)$ of Y_z and $Y_x(1), Y_x(2), \ldots, Y_x(K_m)$, of Y_x (see equation (2)). Evaluate $\overline{f_m(x)}$ and $\overline{f_m(z)}$, where $\overline{f_m(s)} = \frac{1}{K_m} \sum_{i=1}^{K_m} h(s, Y_s(i))$ for $s = x, z$.*

Step 3: *Given* $Z_m = z$, *generate* $U_m \sim U[0,1]$, *and set*

$$X_{m+1} = \begin{cases} z & \text{if } U_m \leq p_m, \\ x & \text{otherwise,} \end{cases}$$

where

$$p_m = \exp\left[\frac{-\left[\overline{f_m(z)} - \overline{f_m(x)}\right]^+}{T}\right].$$

Step 4: *Let* $m = m+1$, $V_m(X_m) = V_{m-1}(X_m) + 1$, *and* $V_m(x) = V_{m-1}(x)$, *for all* $x \in \mathcal{S}$, $x \neq X_m$. *If* $V_m(X_m) > V_m(X^*_{m-1})$, *then let* $X^*_m = X_m$; *otherwise let* $X^*_m = X^*_{m-1}$. *Go to Step 1.*

We now state a convergence result for the proposed algorithm. This theorem is proved in Alrefaei and Andradóttir (1995).

Theorem 2 *Under Assumptions 1, 2, and 3, the sequence* $\{X^*_m\}$ *generated by Algorithm 2 converges almost surely to an element of* $\mathcal{S}^*$ *(in the sense that there exists a set* A *such that* $P(A) = 1$ *and for all* $\omega \in A$, *there exists* $M_\omega > 0$ *such that* $X^*_m(\omega) \in \mathcal{S}^*$ *for all* $m \geq M_\omega$*).*

4 APPLICATION OF THE PROPOSED ALGORITHM

Now we apply the proposed algorithm to solve a simple discrete stochastic optimization problem. In equation (1), let $\mathcal{S} = \{1, \ldots, 10\}$, and let $f(x) = E[Y_x]$ for all $x \in \mathcal{S}$, where Y_x is a uniform random variable on the interval $\mu(x) \pm 0.5$, for all $x \in \mathcal{S}$ and $\mu(1), ..., \mu(10)$ are $0.3, 0.7, 0.9, 0.5, 1.0, 1.4, 0.7, 0.8, 0.0$, and 0.6, respectively. We will apply Algorithm 2 to solve this optimization problem with a number of different choices of the parameters T, $\{K_m\}$, $\{N(x) : x \in \mathcal{S}\}$, and $\{R(x,x') : x \in \mathcal{S} \text{ and } x' \in N(x)\}$. In particular, we use two different values of the temperature T: $T = 0.1$ and $T = 1$. We also use two choices for the sequence $\{K_m\}$. In the first one we let $K_m = \lfloor 2\log(m+3) \rfloor$ for all m, where $\lfloor x \rfloor$ denotes the integer part of x for all $x \in \Re$. Hence, this sequence $\{K_m\}$ increases very slowly in m. The second choice is $K_m = 1 + \lfloor m/10 \rfloor$ for all m, so this sequence increases more rapidly in m than the first choice. Also, we use two different neighborhood structures $\{N(x) : x \in \mathcal{S}\}$. The first neighborhood structure is

$$N(x) = \{x \pm 1 \pmod{10}, x \pm 2 \pmod{10}\} \quad (4)$$

for all $x \in \mathcal{S}$. In this case, we let $R(x,x') = 1/4$ for all $x \in \mathcal{S}$ and $x' \in N(x)$. On the other hand, the second neighborhood structure is given by

$$N(x) = \{x \pm 1 \pmod{10}\} \quad (5)$$

for all $x \in \mathcal{S}$. In this case, we let $R(x,x') = 1/2$ for all $x \in \mathcal{S}$ and $x' \in N(x)$. Note that in the first neighborhood structure (4), we have one local minimum at $x = 4$ and one global minimum at $x = 9$. On the other hand, in the second neighborhood structure (5) we have three local minima at $x = 1, 4$, and 7 and one global minimum at $x = 9$. Since the second neighborhood structure has more local minima, we expect that Algorithm 2 will converge more slowly in this setting than when the first neighborhood structure is used.

Tables 1 through 4 show the results obtained by applying Algorithm 2 to solve this optimization problem with the choices of parameters described above. In particular, the tables show how many of one hundred replications have converged to the true global optimizer as a function of the number of iterations of Algorithm 2 that have been completed. The tables also give, for each choice of parameters, the average number of observations (computed from the one hundred replications) that were needed before the algorithm converged.

Table 1: The performance of the proposed method when $K_m = \lfloor 2\log(m+3) \rfloor$ for all m and the neighborhood structure is given in equation (4)

Iteration	$T = 0.1$	$T = 1$
100	93	59
200	100	72
500	100	89
1,000	100	98
2,000	100	100
3,000	100	100
Average number of observations	310	6,590

From the results given in Tables 1 through 4, we conclude that in this example it is better to have larger neighborhoods $N(x)$, where $x \in \mathcal{S}$, and smaller temperatures T (the results for the neighborhood structure (4) are better than those for the neighborhood structure (5), and the results for $T = 0.1$ are better than those for $T = 1$). In addition, we find that the number of iterations needed for the algorithm to converge to a global minimizer does not seem to depend heavily on the rate of increase of the sequence $\{K_m\}$, but when the sequence $\{K_m\}$ increases very rapidly (e.g., when $K_m = 1 + \lfloor m/10 \rfloor$ for all m), the

Table 2: The performance of the proposed method when $K_m = 1 + \lfloor m/10 \rfloor$ for all m and the neighborhood structure is given in equation (4)

Iteration	$T = 0.1$	$T = 1$
100	98	56
200	100	63
500	100	87
1,000	100	97
2,000	100	100
3,000	100	100
Average number of observations	163	23,572

Table 3: The performance of the proposed method when $K_m = \lfloor 2\log(m+3) \rfloor$ for all m and the neighborhood structure is given in equation (5)

Iteration	$T = 0.1$	$T = 1$
100	72	40
200	76	56
500	98	70
1,000	99	84
2,000	100	94
3,000	100	98
4,000	100	100
Average number of observations	2,580	15,402

algorithm usually (but not always) spends more computer time in generating the observations.

5 CONCLUSION

We have presented a new method for solving discrete stochastic optimization problems that resembles the simulated annealing method for discrete deterministic optimization. This method converges *almost surely* to a *global* solution of the underlying optimization problem. Its performance depends on the choice of a number of parameters, including the temperature, the neighborhood structure, and the number of observations obtained in the different iterations. Our preliminary numerical experience indicates that small temperatures and large neighborhood structures seem to result in better performance. On the other hand, the performance of the algorithm seems less sensitive to how many observations are drawn in each iteration (although it appears to be better to draw a relatively small number of observations per iteration). However, more research on how the parameters of the proposed method should be selected in different situations is needed.

Table 4: The performance of the proposed method when $K_m = 1 + \lfloor m/10 \rfloor$ for all m and the neighborhood structure is given in equation (5)

Iteration	$T = 0.1$	$T = 1$
100	79	53
200	85	59
500	97	72
1,000	99	88
2,000	100	98
3,000	100	98
4,000	100	100
Average number of observations	5,822	96,314

ACKNOWLEDGMENTS

The research of the first author was supported by the Jordan University of Science and Technology. The research of the second author was supported by the National Science Foundation under Grant No. DDM-9210679.

REFERENCES

Alrefaei, M. H., and S. Andradóttir. 1995. Simulated annealing for discrete stochastic optimization. Working paper.

Andradóttir, S. 1995. A method for discrete stochastic optimization. To appear in *Management Science*, 1995.

Andradóttir, S. 1996. A global search method for discrete stochastic optimization. To appear in the *SIAM Journal on Optimization*, 1996.

Devroye, L. 1976. On the convergence of statistical search. *IEEE Transactions on Systems, Man and Cybernetics* 6:46–56.

Goldsman, D., B. L. Nelson, and B. Schmeiser. 1991. Methods for selecting the best system. In *Proceedings of the 1991 Winter Simulation Conference*, eds. B. L. Nelson, W. D. Kelton, and C. M. Clark, 177–186. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Goldsman, D., and B. L. Nelson. 1994. Ranking, selection and multiple comparisons in computer simulation. In *Proceedings of the 1994 Winter Simula-*

tion Conference, eds. J. D. Tew, S. Manivannan, D. A. Sadowski, and A. F. Seila, 192–199. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Gong, W. B., Y. C. Ho, and W. Zhai. 1993. A stochastic comparison algorithm for discrete optimization with estimation. Preprint.

Haddock, J., and J. Mittenthal. 1992. Simulation optimization using simulated annealing. *Computers and Industrial Engineering* 22:387–395.

Lai, T. L., and H. Robbins. 1985. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics* 6:4–22.

Mitra, D., F. Romeo, and A. Sangiovanni-Vincentelli. 1986. Convergence and finite-time behavior of simulated annealing. *Advances in Applied Probability*, 18:747–771.

Yakowitz, S. J., and E. Lugosi. 1990. Random search in the presence of noise, with applications to machine learning. *SIAM Journal on Scientific and Statistical Computing* 11:702–712.

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

AUTHOR BIOGRAPHIES

MAHMOUD HASAN ALREFAEI is pursuing a joint Ph.D. degree in Mathematics and Industrial Engineering at the University of Wisconsin – Madison. He received his B.S. and M.S. degrees in Mathematics from Yarmouk University in Jordan. His interests include discrete stochastic optimization and simulated annealing. He is presently a member of INFORMS and AMS.

SIGRÚN ANDRADÓTTIR is an Associate Professor of Industrial and Systems Engineering at the Georgia Institute of Technology, on leave from the University of Wisconsin – Madison where she has held a faculty position since 1990. 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. She presently serves as Associate Editor for *IIE Transactions* and *Stochastic Models*. She is a member of INFORMS.