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

ITERATIVE RANKING-AND-SELECTION FOR LARGE-SCALE OPTIMIZATION

Sigurdur Ólafsson

Department of Industrial and Manufacturing Systems Engineering
Iowa State University
205 Engineering Annex
Ames, IA 50010, U.S.A.

ABSTRACT

We develop a new algorithm for simulation-based optimization where the number of alternatives is finite but very large. Our approach draws on recent work in adaptive random search and from ranking-and-selection. In particular, it combines the nested partitions method for global optimization and Rinott's two-stage ranking-and-selection procedure. We prove asymptotic convergence of the new algorithm under fairly mild conditions.

1 INTRODUCTION

Optimization over a large but finite feasible region is often a very difficult task. This is true even in the deterministic context, and for stochastic systems the difficulty is exacerbated by the added randomness. Oftentimes discrete event simulation is the only tool available for optimizing such systems. This area has received considerable attention and comprehensive reviews of simulation-based optimization may be found in Jacobson and Schruben (1989), Fu (1994), and Andradóttir (1998). Here we will only mention directly related research. When the number of alternatives is finite and relatively small then *ranking-and-selection* and *multiple-comparison* methods (Goldsman and Nelson 1998) are typically applied. These methods evaluate the performance of each alternative and use statistical methods to guarantee that the objective, that is the selection of the best alternative, is accomplished with a given probability. Classical methods includes for example Rinott's two-stage procedure (Rinott 1978), and more recent work includes that of Matejcik and Nelson (1995), Chick (1997), and Chen et al. (1998). When the number of alternatives becomes somewhat larger, however, then these methods become too computationally intensive and other random search methods, that only consider a fraction of all the alternatives, must be applied.

In a recent paper, Shi and Ólafsson (1998a) introduced such an optimization method, the *nested partitions* (NP) method, for global optimization when the objective function is deterministic. In this context, the method has been found to be quite efficient for combinatorial optimization (Ólafsson and Shi 1998). Furthermore, as is discussed by Shi and Ólafsson (1997), this method can also be applied to stochastic problems, where no analytical expression exists for the objective function and it must be evaluated using simulation. In this paper we suggest a new algorithm that builds on this earlier development in the sense that it also falls within the broad NP method framework. This new algorithm also builds on ideas from statistical sampling techniques that have proven useful in simulation in the past, namely ranking-and-selection methods, and hence combines statistical sampling techniques traditionally used for comparing a few alternatives with a global optimization framework aimed at large-scale optimization problems.

The paper is organized as follows. In Section 2 we introduce the new algorithm and explain its relation to previous work. In Section 3 we prove that the algorithm converges asymptotically under fairly mild conditions to a global optimum. Finally, Section 4 contains some concluding remarks.

2 ALGORITHM DEVELOPMENT

In mathematical notation, we want to solve the problem

$$\min_{\theta \in \Theta} f(\theta), \tag{1}$$

where Θ is a finite feasible region, and $f : \Theta \to \mathbf{R}$ is a performance function that is subject to noise. In other words, for any feasible point $\theta \in \Theta$, $f(\theta)$ cannot be evaluated analytically. Often $f(\theta)$ is an expectation of some random estimate of the performance of a complex stochastic system given a parameter θ, that is, $J(\theta) = E[L(\theta)]$. Here $L(\theta)$ is a random variable which depends on the parameter $\theta \in \Theta$. We assume that $L(\theta)$ is a discrete event simulation estimate of the true performance, and refer to it as the sample

performance. Also, to simplify the analysis, we assume that there exists a unique solution θ_{opt} to problem (1) above.

2.1 The NP Methodology

As we stated in the introduction, the development in this paper builds on a recently proposed method for simulation-based optimization: the NP method. Here we first introduce the basic idea of the NP method, and then show how it may be improved by using statistical selection methodologies. In the k-th iteration of the NP method there is always a region $\sigma(k) \subseteq \Theta$ that is considered the most promising, and as nothing is assumed to be known about location of good solutions before the search is started, $\sigma(0) = \Theta$. The most promising region is then partitioned into M subregions, and what remains of the feasible region, $\Theta \setminus \sigma(k)$, is aggregated into one region called the surrounding region. Therefore, at the k-th iteration $M+1$ disjoint subsets that cover the feasible region are considered. Each of these regions is sampled using some random sampling scheme, and the samples used to estimate the promising index for each region. This index is a set performance function that determines which region becomes the most promising region in the next iteration. If one of the subregions is found to be best, this region becomes the most promising region. If the surrounding region is found to be best, the method backtracks to a larger region. The new most promising region is partitioned and sampled in a similar fashion. This generates a sequence of set partitions, with each partition nested within the last. The partitioning is continued until eventually all the points in the feasible region correspond to a singleton region. The following definitions will be used throughout the analysis.

Definition 1 *A region constructed using a fixed partitioning scheme is called a* valid region *given the fixed partition. The collection of all valid regions is denoted by* Σ. *Singleton regions are of special interest, and* $\Sigma_0 \subset \Sigma$ *denotes the collection of all such valid regions.*

Definition 2 *The singleton regions in* Σ_0, *are called regions of* maximum depth. *More generally, we define the* depth, $d : \Sigma \to \mathbf{N_0}$, *of any valid region iteratively with* Θ *having depth zero, subregions of* Θ *having depth one, and so forth. Since they cannot be partitioned further, we call the singleton regions in* Σ_0 *regions of maximum depth.*

Ultimately only the maximum depth regions are of interest, that is, we want to find a region that contains only one point. Therefore, in each iteration, the estimated best region is the maximum depth region $\hat{\sigma}_{opt}(k) \in \Sigma_0$, that has been most frequently considered the most promising region. Consequently, the method must keep track of the number of times, $\mathcal{N}_k(\sigma)$, a region $\sigma \in \Sigma_0$ has been visited by the k-th iteration. Note that it suffices to keep track of this for regions that have been visited at least once.

Definition 3 *We denote the unique singleton region that corresponds to the optimal solution as* $\sigma_{opt} = \{\theta_{opt}\}$, *and we let the estimate of this best region be*

$$\hat{\sigma}_{opt}(k) = \arg\max_{\sigma \in \Sigma_0} \mathcal{N}_k(\sigma), \tag{2}$$

the most frequently visited singleton region by the k*-th iteration.*

Definition 4 *If a valid region* $\sigma \in \Sigma$ *is formed by partitioning a valid region* $\eta \in \Sigma$, *then* σ *is called a* subregion *of region* η, *and region* η *is called a* superregion *of region* σ. *We define the* superregion function $s : \Sigma \to \Sigma$ *as follows. Let* $\sigma \in \Sigma \setminus \Theta$. *Define* $s(\sigma) = \eta \in \Sigma$, *if and only if* $\sigma \subset \eta$ *and if* $\sigma \subseteq \xi \subseteq \eta$ *then* $\xi = \eta$ *or* $\xi = \sigma$. *For completeness we define* $s(\Theta) = \Theta$.

The NP method shifts the focus from specific points in the feasible region Θ to a space of subsets; namely the space of all valid regions. Consequently, a set performance function $I : \Sigma \to \mathbf{R}$ is needed. This set function can then be used to select the most promising region and is therefore called the promising index of the region. In this paper we let

$$I(\eta) = \min_{\theta \in \eta} f(\theta), \ \forall \eta \in \Sigma, \tag{3}$$

that is, the best solution in a region represents this region. We refer the interested reader to Shi and Ólafsson (1998b) for a comprehensive discussion and analysis of this algorithm, and restrict our attention to the elements that are relevant to our present development.

It is clear that the NP method samples from the entire feasible region in an adaptive fashion, and concentrates the sampling effort by systematically partitioning the feasible region. Thus, in each iteration it selects a most promising region, that is, the subregion that is considered the most likely to contain the global optimum. This selection can be considered a success if the region selected contains the true global optimum, and it would clearly be of practical interest if a minimum probability of success could be guaranteed in each iteration. In the pure NP algorithm described above there is no such assurance.

Also note that when applying the NP method to a stochastic problem there are two sources of randomness that complicate the selection of the correct subregion. First, there is a sampling error due to a relatively small sample being used to estimate the performance of an often large set. Secondly, the performance of each sample points is estimated using simulation and is hence noisy. It is important to observe that the former of these elements implies that the variation within a subregion differs greatly from one region to the next. As an extreme case consider a singleton region that is being compared to the entire surrounding region. That is, a region containing only one solution being compared to a region containing all of the other solutions. Clearly the

first source of randomness has been completely eliminated in the singleton region, whereas it probably accounts for almost all of the randomness in the surrounding region. This implies that to make better use of the sampling effort the number of sample points from each region should be variable and dependent on the variation within the region. The pure NP algorithm does not provide any guidelines or restrictions on how this may be accomplished.

2.2 Two-Stage Sampling

The discussion at the end of the last subsection identified two shortcomings of the pure NP method: the success probability in each iteration cannot be guaranteed, and there may considerable waste involved in the allocation of sample points. We address this by using statistical selection methods and two-stage sampling to compare the subregions as if they were alternative systems, and hence *combine the benefits of global random search and statistical selection*. In particular, we use Rinott's two-stage ranking-and-selection procedure for selecting the best subregion (Rinott 1978). Since ranking-and-selection is applied in each iteration, the new algorithm may be considered an iterative ranking-and-selection algorithm.

To state this approach rigorously, we let $\mathcal{D}_{ij}(k)$ be the i-th set of points selected from the region $\sigma_j(k)$ using a uniform random sampling procedure, $i \geq 1$, $j = 1, 2, ..., M+1$ in the k-th iteration. We let $N = |\mathcal{D}_{ij}(k)|$ denote the number of sample points, which is assumed to be constant. We let $\theta \in \mathcal{D}_{ij}(k)$ denote a point in this set and let $L(\theta)$ be a simulation estimate of the performance of this point. Then in the k-th iteration,

$$X_{ij}(k) = \min_{\theta \in \mathcal{D}_{ij}(k)} L(\theta),$$

is an estimate of the performance (3) of the region σ_j, which we can now also refer to as the i-th system performance for the j-th system, $i \geq 1$, $j = 1, 2, ..., M+1$. The two-stage ranking-and-selection procedure first obtains n_0 such system estimates, and then uses that information to determine the total number N_j of system estimates needed from the the j-th system, that is, subregion $\sigma_j(k)$. This number is selected to be sufficiently large so that the correct subregion is selected with probability at least P^*, subject to an indifference zone of $\varepsilon > 0$.

More precisely, the procedure is as follows:

2.2.1 Algorithm NP/Rinott

Step 1. Given the current most promising region $\sigma(k)$, partition $\sigma(k)$ into M subregions $\sigma_1(k), ..., \sigma_M(k)$, and aggregate the surrounding region $\Theta \setminus \sigma(k)$ into one region $\sigma_{M+1}(k)$.

Step 2. Let $i = 1$.

Step 3. Use uniform sampling to obtain a set $\mathcal{D}_{ij}(k)$ of N sample points from region $j = 1, 2, ..., M+1$.

Step 4. Use discrete event simulation of the system to obtain a sample performance $L(\theta)$ for every $\theta \in \mathcal{D}_{ij}(k)$ and estimate the performance of the region as

$$X_{ij}(k) = \min_{\theta \in \mathcal{D}_{ij}(k)} L(\theta), \tag{4}$$

$j = 1, 2, ..., M+1$.

Step 5. If $i = n_0$ continue to Step 6. Otherwise let $i = i+1$ and go back to Step 3.

Step 6. Calculate the first-stage sample means and variance

$$\bar{X}_j^{(1)}(k) = \frac{1}{n_0} \sum_{i=1}^{n_0} X_{ij}(k), \tag{5}$$

and

$$S_j^2(k) = \frac{\sum_{i=1}^{n_0} \left[X_{ij}(k) - \bar{X}_j^{(1)}(k) \right]^2}{n_0 - 1}, \tag{6}$$

for $j = 1, 2, ..., M+1$.

Step 7. Compute the total sample size

$$N_j(k) = \max \left\{ n_0 + 1, \left\lceil \frac{h^2 S_j^2(k)}{\epsilon^2} \right\rceil \right\}, \tag{7}$$

where ϵ is the indifference zone and h is a constant that is determined by n_0 and the minimum probability P^* of correct selection (Rinott 1978).

Step 8. Obtain $N_j(k) - n_0$ more simulation estimates of the system performance as in Step 2 - Step 5 above, that is $(N_j(k) - n_0) \cdot N$ more sample points.

Step 9. Let the over all sample mean be the promising index for each region,

$$\hat{I}\left(\sigma_j(k)\right) = \bar{X}_j(k) = \frac{\sum_{i=1}^{N_j(k)} X_{ij}(k)}{N_j(k)}, \tag{8}$$

$j = 1, 2, ..., M+1$.

Step 10. Select the index of the region with the best promising index.

$$\hat{j}_k \in \arg \min_{j=1,...,M+1} \hat{I}(\sigma_j). \tag{9}$$

If more than one region is equally promising, the tie can be broken arbitrarily. If this index corresponds to a region that is a subregion of $\sigma(k)$, then let this be the most promising region in the next iteration. Otherwise, if the index corresponds to

the surrounding region, backtrack to a larger region containing the current most promising region. That is, let

$$\sigma(k+1) = \begin{cases} \sigma_{\hat{i}_k}(k), & \text{if } \hat{i}_k < M+1, \\ s\left(\sigma(k)\right), & \text{otherwise.} \end{cases} \tag{10}$$

Step 11. Update the counters $\{\mathcal{N}_k(\sigma)\}_{\sigma\in\Sigma}$ for the number of times each region has been the most promising region, and if necessary the maximum depth region $\hat{\sigma}_{opt}(k)$ that have been most frequently the most prom-ising region. If $\sigma = \sigma(k+1)$, then let $\mathcal{N}_{k+1}(\sigma) = \mathcal{N}_k(\sigma) + 1$, and otherwise let $\mathcal{N}_{k+1}(\sigma) = \mathcal{N}_k(\sigma)$ for all $\sigma \neq \sigma(k+1)$. If there exists $\sigma \in \Sigma_0$ such that $\mathcal{N}_{k+1}(\sigma) > \mathcal{N}_{k+1}(\hat{\sigma}_{opt}(k))$, then let $\hat{\sigma}_{opt}(k+1) = \sigma$, and otherwise let $\hat{\sigma}_{opt}(k+1) = \hat{\sigma}_{opt}(k)$.

Step 12. If stopping rule is not satisfied let $k = k+1$ and go back to Step 1.

We note that for $n_0 = 1$, with Steps 6-8 omitted, and $\hat{I}(\sigma_i(k)) = X_{1i}$ replacing equation (8), this new iterative ranking-and-selection algorithm reduces to the pure NP algorithm described in Shi and Ólafsson (1998b). On the other hand, by selecting $M = |\Theta|$, the algorithm reduces to a pure Rinott's two-stage ranking-and-selection procedure.

The following parameters must be selected for the new algorithm: the number of sample points used for each system estimate (N), the number of system estimates in the first stage ($n_0 \geq 2$), the probability of correct selection ($P^* \geq 0.5$), and the indifference zone ($\epsilon > 0$). There is clearly a tradeoff between N and n_0 in that $N \cdot n_0$ is the total first stage sample effort, and if we fix $N \cdot n_0$ then increasing N decreases n_0 and vice versa. The choice of P^* deserves special attention. In the pure Rinott procedure, as well as in other ranking-and-selection procedures, this probability is usually selected to be rather large, say $P^* = 0.90$ or $P^* = 0.99$. Here, however, the ranking-and-selection is done iteratively so it is not feasible in practice to expand too much computational effort in each iteration. In the next section we will see that $P^* \geq 0.5$ is needed to guarantee convergence, but it should not be selected too large because then too much effort is spent in each iteration. Finally, the indifference zone ϵ depends on how the performance function is scaled and is therefore problem dependent.

3 CONVERGENCE ANALYSIS

It is straightforward to see that Algorithm NP/Rinott generates a Markov chain and the stationary distribution of this chain can be used for inference about the convergence of the algorithm. To state this precisely, we need the following technical assumption, that can be made without loss of generality.

Assumption 1 *Assume that* $\forall \eta \in \Sigma, \exists \theta \in \eta, \xi \in \Theta \setminus \eta$, *such that* $P\left[L(\theta) < L(\xi)\right] < 1$. □

With this assumption the following proposition follows.

Theorem 1 *If Assumption 1 holds then Algorithm NP/Rinott generates an irreducible recurrent Markov chain* $\{\sigma(k)\}_{k=0}^{\infty}$ *on* Σ, *such that its unique stationary distribution* π *satisfies,*

$$\lim_{k\to\infty} \pi\left(\hat{\sigma}_{opt}(k)\right) > \pi(\eta), \ \forall \eta \in \Sigma_0 \setminus \left\{\sigma_{opt}\right\}, \ w.p.1. \tag{11}$$

In words, the algorithm converges to a maximum of the stationary distribution over all singleton regions.

Proof: This theorem is proven for the pure NP algorithm in Shi and Ólafsson (1998b) in a slightly more general setting, and since that proof also holds for Algorithm NP/Rinott we will only sketch it here. It is clear that $\{\sigma(k)\}_{k=0}^{\infty}$ is a Markov chain, and it is irreducible by Assumption 1. Since Σ is finite, the Markov chain is then positive recurrent with a unique stationary distribution π. Furthermore, it is well known that

$$\lim_{k\to\infty} \frac{\mathcal{N}_k(\eta)}{k} = \pi(\eta),$$

which implies that, in the limit, the most frequently visited region maximizes the stationary distribution. Since $\hat{\sigma}_{opt}(k) = \arg\max_{\sigma\in\Sigma_0} \mathcal{N}_k(\sigma)$ the theorem follows. □

To state the main convergence theorem we need the usual assumption of ranking-and-selection methods, namely that the observations are normally distributed.

Assumption 2 *Assume that* $X_{ij} \sim \mathcal{N}(\mu_j, v_j^2)$, *is normally distributed with mean* μ_j *and variance* v_j^2 *for all* $j \in \{1, 2, ..., M+1\}$, *and* $i \in \{1, 2, ..., N_j(k)\}$, $k \geq 1$. □

We also need to be able to distinguish between the optimum and other solutions.

Assumption 3 *Assume that the indifference zone* ϵ *satisfies* $\epsilon \leq \min_{\theta\in\Theta\setminus\theta_{opt}} f(\theta) - f(\theta_{opt})$. □

We now have the following main convergence theorem for this algorithm.

Theorem 2 *If Assumption 1-3 hold and* $P^* \geq 0.5$, *then Algorithm NP/Rinott converges with probability one to a global optimum, that is,*

$$\lim_{k\to\infty} \hat{\sigma}_{opt}(k) = \sigma_{opt}, \ \ w.p.1. \tag{12}$$

Proof: By Theorem 1 the algorithm converges to the singleton region that has the largest stationary probability, so we only need to show that $\pi(\sigma_{opt}) \geq \pi(\eta)$ for all $\eta \in \Sigma_0$. Hence, let $\eta \in \Sigma_0$ be an arbitrary singleton region. Now let η^* be the smallest region that contains both η and σ_{opt}, that is

$$\eta^* = \min\{\sigma \in \Sigma : \eta \subseteq \sigma, \sigma_{opt} \subseteq \sigma\}.$$

Furthermore, let $\eta_1, \eta_2, ..., \eta_{d(\eta^*)}$ be the sequence of regions that satisfies $\sigma_{opt} = \eta_1 \subset \eta_2 \subset ... \subset \eta_{d(\eta^*)} = \eta^*$, that is this is the sequence of regions the algorithm must traverse to get from σ_{opt} to η^* and vice versa. Furthermore, let $\eta_{d(\eta^*)+1}$, $\eta_{d(\eta^*)+2}$,... ,$\eta_{2d(\eta^*)}$ be the same type of sequence with $\eta_{2d(\eta^*)} = \eta$, that is, the sequence needed to get from η^* to η and vice versa. Thus, the sequence $\eta_1, \eta_2, ..., \eta_{2d(\eta^*)}$ represent the shortest path from σ_{opt} to η and back. Now for any $i \in \{1, 2, ..., 2d(\eta^*) - 1\}$ it is clear that since the Markov chain is clearly reversible (note the tree structure) then

$$P(\eta_i, \eta_{i+1})\pi_{\eta_i} = P(\eta_{i+1}, \eta_i)\pi_{\eta_{i+1}},$$

and similarly

$$P(\eta_{i+1}, \eta_{i+2})\pi_{\eta_{i+1}} = P(\eta_{i+2}, \eta_{i+1})\pi_{\eta_{i+2}}.$$

These equations can of course be verified by setting up the full balance equations. Thus we have

$$\frac{P(\eta_i, \eta_{i+1})}{P(\eta_{i+1}, \eta_i)}\pi_{\eta_i} = \pi_{\eta_{i+1}} = \frac{P(\eta_{i+2}, \eta_{i+1})}{P(\eta_{i+1}, \eta_{i+2})}\pi_{\eta_{i+2}},$$

or

$$\pi_{\eta_i} = \frac{P(\eta_{i+1}, \eta_i) \cdot P(\eta_{i+2}, \eta_{i+1})}{P(\eta_i, \eta_{i+1}) \cdot P(\eta_{i+1}, \eta_{i+2})} \cdot \pi_{\eta_{i+2}}.$$

By induction we get

$$\pi_{\eta_1} = \frac{P(\eta_2, \eta_1) \cdot ... \cdot P(\eta_{2d(\eta^*)}, \eta_{2d(\eta^*)-1})}{P(\eta_1, \eta_2) \cdot ... \cdot P(\eta_{2d(\eta^*)-1}, \eta_{2d(\eta^*)})} \cdot \pi_{\eta_{2d(\eta^*)}}. \quad (13)$$

We also know that $P(\eta_{i+1}, \eta_i)$ is the probability of correct selection, that is, moving towards the true optimum, so by Assumptions 2-3 and Proposition 1 in Rinott (1978) we have $P(\eta_{i+1}, \eta_i) \geq P^* \geq \frac{1}{2}$. This implies that

$$P(\eta_{i+1}, \eta_i) \geq P(\eta_{i+1}, \eta_{i+2}), \ \forall i = 1, 2, ..., 2d(\eta^*) - 2,$$

which together with equation (13) says that

$$\pi_{\eta_1} \geq \frac{P(\eta_{2d(\eta^*)}, \eta_{2d(\eta^*)-1})}{P(\eta_1, \eta_2)}\pi_{\eta_{2d(\eta^*)}}. \quad (14)$$

On the other hand, it is clear that moving from $\eta_{2d(\eta^*)}$ to $\eta_{2d(\eta^*)-1}$ is the correct selection so

$$P(\eta_{2d(\eta^*)}, \eta_{2d(\eta^*)-1}) \geq P^* \geq \frac{1}{2},$$

and vice versa moving from η_1 to η_2 is incorrect selection so $P(\eta_1, \eta_2) \leq \frac{1}{2}$. Thus equation (14) reduces to

$$\pi_{\sigma_{opt}} = \pi_{\eta_1} \geq \pi_{\eta_{2d(\eta^*)}} = \pi_\eta,$$

which is precisely what is needed to prove the theorem. □

The last theorem shows that Algorithm NP/Rinott converges asymptotically. Next we consider how fast it converges. By Definition 3, for the algorithm to correctly consider the optimum σ_{opt} as the best solution, this state must be visited at least once. Hence it is of interest to look at the expected time until the algorithm will visit this state for the first time. Clearly we would like this to be as small as possible, and the next theorem provides an upper bound for this expected time.

Theorem 3 *Let Assumption 1-3 hold and assume that $P^* > 0.5$. Let T_1 denote the first time Algorithm NP/Rinott visits the optimal solution. Then*

$$E[T_1] \leq \frac{d^*}{2P^* - 1}. \quad (15)$$

Proof: Recall that the Markov chain $\{\sigma(k)\}_{k=1}^{\infty}$ has a minimum success probability of P^* given its current state $\sigma(k) \in \Sigma$, that is, with probability at at least P^*, $\sigma(k+1)$ will be closer to σ_{opt} than $\sigma(k)$ in terms of the number of transitions required to move between the regions. Now imagine a Markov chain that is identical to $\{\sigma(k)\}$ except that this success probability is even and equal to P^* for every state $\sigma \in \Sigma$. Now note that since the success probability is constant, the exact state is not of any consequence, but rather the number of transitions it takes to move from the current state $\sigma(k)$ to the optimum. The maximum such distance is $2d^*$, and we can therefore, without losing any information, reduce the state space to $\mathcal{S} = \{0, 1, 2, ..., 2d^*\}$. With this representation the entire feasible region Θ corresponds to state d^*, and we can let the global optimum correspond to state zero. Given a state $x \in \mathcal{S}$ the probability of moving to $x - 1$ is fixed and equal to P^*, and the probability of moving to $x + 1$ is equal to $1 - P^*$, regardless of the state. Therefore, the new Markov chain is a simple random walk. Furthermore, it is clear that $E[T_1] \leq E\left[T_1'\right]$, where T_1' is the first time the random walk visits σ_{opt} if it starts in state d^*, which corresponds to Θ, the starting state of Algorithm NP/Rinott.

Hence, if we calculate the expected hitting time for the random walk this automatically gives us an upper bound for the original Markov chain. Furthermore, since we are only interested in the time the global optimum is found for the first time, we can assume 0 is an absorbing barrier and look at the time of absorption. Note also that $2d^*$ is a reflective barrier. Then it is known that the expected time T of absorption when starting in state u is (Weesakul, 1961)

$$E_u[T] = \frac{u}{2P^* - 1} + \frac{(1 - P^*)^{2d^*+1}}{(P^*)^{2d^*}(2P^* - 1)^2}\left(1 - \left(\frac{P^*}{1 - P^*}\right)^u\right). \quad (16)$$

Thus, for $u = d^*$, that is, when the algorithm starts in state $\sigma(0) = \Theta$, we have

$$E_{d^*}[T] = \frac{d^*}{2P^* - 1} + \frac{(1-P^*)^{2d^*+1}}{(P^*)^{2d^*}(2P^*-1)^2} \left(1 - \left(\frac{P^*}{1-P^*}\right)^{d^*}\right). \quad (17)$$

Now since $P* > 0.5$ then $\frac{P^*}{1-P^*} > 1$ so

$$\left(1 - \left(\frac{P^*}{1-P^*}\right)^{d^*}\right) < 0,$$

and hence

$$\frac{(1-P^*)^{2d^*+1}}{(P^*)^{2d^*}(2P^*-1)^2} \left(1 - \left(\frac{P^*}{1-P^*}\right)^{d^*}\right) < 0.$$

Therefore,

$$E[T_1] \leq E_{d^*}[T] \leq \frac{d^*}{2P^* - 1},$$

which proves the theorem. □

To obtain the simple bound in equation (16) we ignored the second negative term in equation (17). It is therefore appropriate to consider how loose this bound is, and we observe that this second term goes to zero as either $d^* \to \infty$ or $P^* \to 1$, and indeed, unless d^* is small, say less than ten, and P^* is close to one half, say $P^* < 0.55$, the first term in equation (16) is much larger than the second term in absolute value.

Now lets consider if an optimal selection probability $P^*(n_0, M)$ and can be found. It is clear that as $P^*(n_0, M)$ increases $E[T_1]$, that is, the expected time until the global optimum is encountered decreases. This occurs, however, at a decreasing rate. On the other hand, as $P^*(n_0, M)$ increases $h(n_0, M, P^*)$ also increases and this occurs at an increasing rate. Therefore, an optimal probability is somewhere between the extreme values of $P^*(n_0, M) = 0.5$ and $P^*(n_0, M) = 1$. However, since the second-stage sample size depends on the sample variance from the first stage sampling and the indifference zone, both of which are clearly problem dependent, so does the optimal value of $P^*(n_0, M)$. It is therefore not possible to give an *a priori* prescription for the optimal probability.

Another quantity of interest when applying the Algorithm NP/Rinott is the probability of the first maximum depth region visited being the one corresponding to the global optimum. If this probability is fairly high then a reasonable stopping rule would be to stop whenever maximum depth is reached. We can again use a random walk analysis, this time for a simple random walk with two absorbing barriers, to calculate this probability.

Table 1: Expected first hitting time of the optimum.

	Maximum Depth (d^*)				
Success Prob.	2	5	10	20	30
55%	20	50	100	200	300
60%	10	25	50	100	150
65%	7	17	33	67	100
70%	5	13	25	50	75
75%	4	10	20	40	60
80%	3	8	17	33	50
85%	3	7	14	29	43
90%	3	6	13	25	38
95%	2	6	11	22	33

Theorem 4 *Let Assumption 1-3 hold and assume that $P^* > 0.5$. Let $\hat{\sigma}$ denote the first maximum depth region visited. Then*

$$P\left[\hat{\sigma} = \sigma^*_{opt}\right] = \begin{cases} (P^*)^{d^*} \frac{(1-P^*)^{d^*} - (P^*)^{d^*}}{(1-P^*)^{2d^*} - (P^*)^{2d^*}} & P^* \neq \frac{1}{2}, \\ \frac{1}{2} & P^* = \frac{1}{2}. \end{cases} \quad (18)$$

Proof: Since the success probability is constant we can again consider the random walk with state space $\mathcal{S} = \{0, 1, ..., 2d^*\}$ defined in the proof of Theorem 3 above. Here the only question is thus if state 0 or $2d^*$ will be visited first; that is, the probability that the first maximum depth visited contains the global optimum is equal to the probability that the random walk visits state 0 before it visits state $2d^*$. This probability is thus equal to the absorption probability at zero for a simple random walk with two absorbing barriers, which can for example be found on p. 32 in Cox and Miller (1965). □

For insights into this theorem consider Table 2 which shows the results of equation (18) for $d \in \{2, 5, 10, 20, 30\}$ and P^* ranging from 0.50 to 0.95. From this table we see that unless the problem is small, say $d^* < 10$, then the probability of the first maximum depth region visited corresponding to the optimum is very high even for success probability as low as 55%. For problems as small as $d^* = 5$ it is sufficient to have 75% success probability for it to be virtually certain that the first maximum depth region will correspond to the optimum. Thus, stopping when the algorithm reaches maximum depth is a reasonable strategy for the Algorithm NP/Rinott, and if this stopping rule is applied then equation (18) can be used to calculate the probability of this being a correct termination.

Table 2: Probability of first maximum depth region being the optimum.

	Maximum Depth (d^*)				
Success Prob.	2	5	10	20	30
55%	0.60	0.73	0.88	0.98	1.00
60%	0.69	0.88	0.98	1.00	1.00
65%	0.78	0.96	1.00	1.00	1.00
70%	0.84	0.99	1.00	1.00	1.00
75%	0.90	1.00	1.00	1.00	1.00
80%	0.94	1.00	1.00	1.00	1.00
85%	0.97	1.00	1.00	1.00	1.00
90%	0.99	1.00	1.00	1.00	1.00
95%	1.00	1.00	1.00	1.00	1.00

4 SUMMARY

We have introduced a new algorithm for optimizing systems where the number of alternatives is very large and the performance of each alternative must be evaluated using simulation. The approach combines an adaptive sampling method called the nested partitions method with traditional ranking-and-selection procedures. We have proved the asymptotic convergence of the algorithm but numerical testing of the algorithm is needed and is currently underway.

Future research directions include further refining of the algorithm, analyzing how fast it converges, and deriving efficient stopping rules. Also of interest would be to incorporate other statistical selection procedures into the method. This could include the optimal computing budging allocation (OCBA) procedure (Chen et al. 1998) or a combined subset selection and Rinott's procedure (Goldsman and Nelson 1998).

REFERENCES

Andradóttir, S. 1998. "A Review of Simulation Optimization Techniques," in *Proceedings of the 1998 Winter Simulation Conference*, 151-158.

Chen, C.-H., E. Yücesan, Y. Yuan, H.-C. Chen, and L. Dai. 1998. "Computing Budget Allocation for Simulation Experiments with Different System Structures," in Proceedings of the 1998 Winter Simulation Conference, 735-741.

Chick, S.E. 1997. "Selecting the Best System: A Decision Theoretic Approach," in *Proceedings of the 1997 Winter Simulation Conference*, 326-333.

Cox, D.R. and H.D. Miller, 1965. *The Theory of Stochastic Processes*, Chapman & Hall, London.

Fu, M.C. 1994. "Stochastic Optimization using Simulation: A review," *Annals of Operations Research*, 53, 199-248.

Goldsman, D. and B.L. Nelson. 1998. "Statistical Screening, Selection, and Multiple Comparison Procedures in Computer Simulation," in *Proceedings of the 1998 Winter Simulation Conference*, 159-166.

Jacobson, S.H. and L.W. Schruben. 1989. "A Review of Techniques for Simulation Optimization," *Operations Research Letters*, 8, 1-9.

Matejcik, F.J. and B.L. Nelson. 1995. "Two-Stage Multiple Comparisons with the Best for Computer Simulation," *Operations Research*, 43, 633-640.

Ólafsson, S. and L. Shi. 1998. "A Method for Scheduling in Parallel Manufacturing Systems with Flexible Resources," *IIE Transactions on Scheduling and Logistics*, accepted for publication.

Rinott, Y. 1978. "On Two-Stage Selection Procedures and Related Probability-in-Equalities," *Communications in Statistics*, A7, 799-811.

Shi, L. and S. Ólafsson. 1997. "An Integrated Framework for Deterministic and Stochastic Optimization," in *Proceedings of the 1997 Winter Simulation Conference*, 358-365.

Shi, L. and S. Ólafsson. 1998a. "Nested Partitions Method for Global Optimization," *Operations Research*, accepted for publication.

Shi, L. and S. Ólafsson. 1998b. "Nested Partitions Method for Stochastic Optimization," Working Paper Series, 98-116, Department of Industrial and Manufacturing Systems Engineering, Iowa State University.

AUTHOR BIOGRAPHY

SIGURDUR ÓLAFSSON is an assistant professor in the Department of Industrial and Manufacturing Systems Engineering at Iowa State University. He received a B.S. in Mathematics from the University of Iceland in 1995, and an M.S. and a Ph.D. in Industrial Engineering from the University of Wisconsin - Madison in 1996 and 1998, respectively. His research interests include applied probability, stochastic optimization, and simulation. He is a member of IIE and INFORMS.