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

OPTIMIZATION VIA ADAPTIVE SAMPLING AND REGENERATIVE SIMULATION

Sigurdur Ólafsson

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

Leyuan Shi

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

ABSTRACT

We investigate a new approach for simulation-based optimization that draws on two recent stochastic optimization methods: an adaptive sampling approach called the nested partitions method and ordinal optimization. An ordinal comparison perspective is used to show that the nested partitions method converges globally under weak conditions. Furthermore, we use those results to determine a lower bound for the required sampling effort in each iteration, and show that global convergence requires relatively little simulation effort in each iteration.

1 INTRODUCTION

In system optimization it is often desirable to optimize the performance of a system where the solution parameters are discrete and the outcomes are uncertain. This means that there is no analytical expression relating the discrete decision parameters to the corresponding expected performance of the system. Such stochastic discrete optimization problems have received considerable attention in recent years, and methods proposed for this problem include, for example, the *stochastic ruler* method (Yan and Mukai, 1992; Alrefaei and Andradóttir, 1997), the method of Andradóttir (1995), the *stochastic comparison* method (Gong, Ho, and Zhai, 1992), *ordinal optimization* (Ho, Sreenivas, and Vakili, 1992), the *stochastic branch-and-bound* (Norkin, Pflug, and Ruszczyński, 1996), and the *nested partitions* (NP) method (Shi and Ólafs-son, 1998a,b). Under certain conditions, all of these methods have been shown to converge almost surely to an optimal solution. For recent reviews of simu-lation-based optimization methods the reader can for example consult Carson and Maria (1997) and Andradóttir (1998).

In this paper we investigate the NP method from the perspective of ordinal comparison, which enables us to gain insights into the convergence of the NP method and proves that the ordinal nature of the method is indeed beneficial. Furthermore, we derive conditions for asymptotic convergence of the algorithm and provide practical guidelines for how to conduct the adaptive sampling in terms of the computational effort needed for each iteration.

The remainder of this paper is organized as follows. In Section 2 we define the problem and discuss the optimization methodology applied. In Section 3 we present convergence results for this method, and finally, Section 4 contains some concluding remarks.

2 OPTIMIZATION METHODOLOGY

In this paper we are concerned with optimizing a performance function $J : \Theta \rightarrow \mathbf{R}$ over a finite feasible region Θ; that is,

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

where $|\Theta| < \infty$. For simplicity of presentation, we assume that there is some unique solution $\theta_{opt} \in \Theta$ that solves this problem, that is, $J\left(\theta_{opt}\right) < J(\theta)$, for all $\theta \in \Theta \setminus \left\{\theta_{opt}\right\}$.

In practice $J(\theta)$ is often the expected performance of a complex system given some underlying solution parameters θ, and there may be no analytical expression available to relate this expected performance to the solution parameters. In such situations, $J(\theta)$ must be estimated from a simulation sample performance $L_t(\theta)$, where t is the simulation time. We assume that regenerative simulation is used to estimate the system performance, that is, $\{L_t\}$ is a regenerative process and the problem is that of simulation-based optimization with discrete decision parameters. Such simulation-based optimization has numerous applications. However, in practice it is very computationally expensive to obtain accurate steady-state simulation estimates of the performance of a complex system, and it is hence often necessary to content with a short simulation that results in noisier estimates; that is, t may be very small and $L_t(\theta)$ may hence only be a rough estimate of $J(\theta)$ for each $\theta \in \Theta$. It has been observed that when dealing with such noisy estimates it is beneficial to focus on the ordinal rather than

cardinal values of the solutions (Ho, Sreenivas, and Vakili, 1992), and we show that due to the ordinal nature of the NP method it may indeed converge despite very noisy simulation estimates. First, however, we describe the NP method itself.

2.1 The Nested Partitions Method

The basic idea of the NP method is simple. In the k-th iteration there is a region $\sigma(k) \subseteq \Theta$ that is considered the *most promising*. In the first iteration nothing is assumed to be known about where good solutions are, so the entire solution space $\sigma(1) = \Theta$ is taken as the most promising region. The most promising region is then *partitioned* into $M_{\sigma(k)}$ subregions, where $M_{\sigma(k)}$ may depend on the subset $\sigma(k)$ but not the iteration. What remains of the solution space, $\Theta \setminus \sigma(k)$, is aggregated into one region called the surrounding region. Clearly the NP method shifts the focus from individual solutions to sets of solutions, and the following definitions, that identify the most important classes of such sets, will be convenient throughout the analysis.

Definition 1 *A set constructed using a fixed partitioning strategy is called a* valid region. *The collection of all valid regions is denoted by* Σ. *Singleton regions are of special interest, and we let* $\Sigma_0 \subset \Sigma$ *denote the collection of all such valid regions. Finally, we let* $\Sigma_g \subset \Sigma$ *denote all the 'good' subregions, that is,* $\sigma \in \Sigma_g$ *if and only if* $\theta_{opt} \in \sigma$.

It will also be convenient to be able to identify the valid region which was partitioned to obtain the current most promising region, which motivates the next two definitions.

Definition 2 *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.

Definition 3 *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* $s(\Theta) = \Theta$.

It will also be necessary to keep track of distance between valid regions, so we define two concepts, the depth of a region, which essentially is its distance from the entire solution space Θ, and a metric that defines the distance between any two valid regions.

Definition 4 *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.*

Definition 5 *We let* $m(\cdot,\cdot)_\Sigma : \Sigma \times \Sigma \to \mathbf{R}$ *be a metric given the partitioning strategy* Σ, *defined by*

$$m(\eta_1, \eta_2)_\Sigma = \min_{\substack{\eta \in \Sigma \\ \eta_1 \subseteq \eta,\ \eta_2 \subseteq \eta}} (d(\eta_1) - d(\eta)) + (d(\eta_2) - d(\eta)), \tag{2}$$

and call it the partitioning metric.

We note that the depth of a region $\eta \in \Sigma$ is its distance from the entire feasible region Θ, that is, $d(\eta) = m(\Theta, \eta)_\Sigma$. Furthermore, the performance of the NP method turns out to depend on how the partitioning is performed, and we can use this metric to define the ideal case.

Definition 6 *A partitioning strategy* Σ *is called* optimal *if and only if the global optimum* σ_{opt} *has the following property: For all* $\eta_1, \eta_2 \in \Sigma$ *such that* $d(\eta_1) = d(\eta_2)$ *and* $m(\sigma_{opt}, \eta_1)_\Sigma < m(\sigma_{opt}, \eta_2)_\Sigma$, *then*

$$J(\theta) < J(\phi),\ \forall \theta \in \eta_1,\ \forall \phi \in \eta_2. \tag{3}$$

Returning to the procedure of the NP method, then given a partitioning of $\sigma(k)$, at the k-th iteration $M_{\sigma(k)}+1$ disjoint subsets that cover the feasible region are considered. Each of these regions is sampled using some *random sampling* scheme, resulting in a set $\mathcal{D}_{\sigma_j(k)}$ of sample points. The samples are then used to estimate the *promising index* for each region. This index is a set performance function $I : \Sigma \to \mathbf{R}$, that determines which region becomes the most promising region in the next iteration and the estimate $\hat{I}\left(\sigma_j(k)\right) = \hat{I}\left(\mathcal{D}_{\sigma_j(k)}\right)$ depends only on the set of sample points. 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. To choose this larger region a fixed backtracking rule is used.

Definition 7 *Let* $\hat{j}_k$ *be the index corresponding to the best region found in the* k*-th iteration.*

$$\hat{j}_k = \arg\min_j \hat{I}(\sigma_j(k)) \tag{4}$$

Based on $\hat{j}_k$, *either move to a subregion or backtrack to the superregion of the current most promising region. That is, let*

$$\sigma(k+1) = \begin{cases} \sigma_{\hat{j}_k}(k), & \text{if } \hat{j}_k < M_{\sigma(k)} + 1, \\ s\left(\sigma(k)\right), & \text{otherwise.} \end{cases} \tag{5}$$

where the function $s : \Sigma \to \Sigma$ *is as in Definition 3 above.*

The new most promising region $\sigma(k+1)$ is then partitioned and sampled in a similar fashion. This generates a sequence of set partitions, with each partition nested within the last. We assume that the partitioning is continued until eventually all the points in the feasible region correspond to a singleton region, and we let the estimate of the best

solution be the singleton region that has been considered the most promising the most often.

Definition 8 *Let $\mathcal{N}_k(\sigma)$ be the number of times region $\sigma \in \Sigma$ has been considered the most promising region by the k-th iteration. The estimate of the best solution is*

$$\hat{\sigma}_{opt}(k) = \arg\max_{\sigma \in \Sigma_0} \mathcal{N}_k(\sigma), \tag{6}$$

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

We note again that the basic idea of the NP algorithm is to shift the focus from the solution space itself to a sequence of subsets of the solution space. These subsets are sampled with variable density and a promising index for each subset is estimated. The ordinal values of these estimates determine how the algorithm proceeds in the next step. It is clear from equation (4) that accurately estimating the promising index is not critical. Only the ordinal values affect how the NP algorithm proceeds. If subregion $\sigma_{j_{opt}} \in \Sigma_g$ contains the true global optimum, then it is sufficient that $\hat{I}\left(\sigma_{j_{opt}}(k)\right) < \hat{I}\left(\sigma_j(k)\right),\ \forall j \neq j_{opt}$. If this holds then the subregion containing the global optimum is identified. We conclude that if the rank is preserved then nothing is gained from more accurate estimates.

2.2 An Ordinal Promising Index

It is clear from the description of the NP method that a critical element is the selection and estimation of a promising index. Indeed, the estimated values of this index determine, in each iteration, how the sampling is concentrated in the next iteration. In its simplest form the estimated promising index can be taken as a summary statistic for the sampling information (Shi and Ólafsson, 1998a). We can for example define the promising index function as

$$I(\sigma) = \min_{\theta \in \sigma} J(\theta),\ \ \sigma \in \Sigma. \tag{7}$$

For a given region $\sigma \in \Sigma$ and a set of sample points $\mathcal{D}_\sigma \subseteq \sigma$, we need to obtain an estimate $\hat{I}(\sigma)$ of the promising index value $I(\sigma)$. This estimate must be based on the sample performance $L_t(\theta)$ for each sample point $\theta \in \mathcal{D}_\sigma$, but the problem is that an accurate estimate of the performance is very expensive. If the performance is estimated using simulation it is well known that the estimate $\hat{J}(\theta)$ converges to $J(\theta)$ at a rate that is at the most $O(\frac{1}{\sqrt{t}})$ in the total simulation time t. This in turn implies that the estimate $\hat{I}(\sigma)$ converges to $I(\sigma)$ at a rate that is at least as slow. However, recall that if it is desirable to move into a good region $\sigma_g \in \Sigma_g$, where Σ_g is as in Definition 1, and this is being compared to another bad region $\sigma_b \in \Sigma \setminus \Sigma_g$, then it is sufficient that

$$\hat{I}(\sigma_g) < \hat{I}(\sigma_b), \tag{8}$$

that is, if the rank is preserved then the correct valid region is selected. The advantage of this being sufficient is that the estimated rank of a random variable may converge to its true rank at an exponential rate even if the cardinal values converge at a much slower rate (Dai, 1996). The implication is that it is not necessary to accurately estimate $J(\theta)$ for each $\theta \in \mathcal{D}_\sigma$ to obtain a sufficiently good estimate of the promising index. Therefore, for every $\sigma \in \Sigma$ and corresponding set of sample points $\mathcal{D}_\sigma$, we let

$$\hat{I}(\sigma) = \min_{\theta \in \mathcal{D}_\sigma} L_t(\theta). \tag{9}$$

Since $L_t(\theta)$ is obtained using regenerative simulation, and such estimates are strongly consistent, we have that

$$\hat{I}(\sigma) \to \min_{\theta \in \mathcal{D}_\sigma} J(\theta), w.p.1.$$

So, in the long-run, if $\min_{\theta \in \mathcal{D}_{\sigma_g}} J(\theta) < \min_{\theta \in \mathcal{D}_{\sigma_b}} J(\theta)$ then σ_g will be selected. However, it is also known that this convergence occurs rather slowly. On the other hand, as we pointed out above we do not need accurate estimates of the cardinal values and we will show that if the estimated promising index (9) is used then, for certain systems, the probability of equation (8) holding converges to a sufficiently large value at an exponential rate.

3 CONVERGENCE ANALYSIS

By noting how the NP algorithm moves from one region in Σ to the next, based only on the current sampling information, it is clear that the algorithm generates a Markov chain $\{\sigma(k)\}_{k=1}^{\infty}$ with state space Σ. Furthermore, it is not difficult to show that this Markov chain has a unique stationary distribution. To prove asymptotic convergence of the method, we show that given certain regularity conditions, the stationary probability of the singleton $\sigma_{opt} = \{\theta_{opt}\}$ is greater than that of any other singleton region and the NP algorithm converges to this maximum stationary probability singleton Shi and Ólafsson (1998a,b).

3.1 Asymptotic Convergence

We begin by stating the asymptotic convergence result precisely.

Theorem 1 *Assume that*

$$P\left[\hat{I}\left(\sigma_g\right) \leq \hat{I}\left(\sigma_b\right)\right] \geq P\left[\hat{I}\left(\sigma_g\right) \geq \hat{I}\left(\sigma_b\right)\right], \tag{10}$$

$\forall \sigma_g \in \Sigma_g, \sigma_b \in \Sigma \setminus \Sigma_g$. *Then the NP method converges with probability one to the global optimum* $\sigma_{opt} = \{\theta_{opt}\}$, *that is, as* $k \to \infty$ *then*

$$\hat{\sigma}_{opt}(k) \to \sigma_{opt}, \ \ w.p.1. \tag{11}$$

Proof: We will only sketch a proof here and refer to Shi and Ólafsson (1998b) for full analysis of the stochastic NP method. We start by observing that $\{\sigma(k)\}_{k=1}^{\infty}$ is an irreducible positive recurrent Markov chain. Therefore, it has a unique stationary distribution π, and it is well known that with probability one, as $k \to \infty$,

$$\frac{\mathcal{N}_k(\sigma)}{k} \to \pi(\sigma), \ \ \forall \sigma \in \Sigma,$$

where $\mathcal{N}_k(\sigma)$ counts, as before, the number of times $\sigma \in \Sigma$ is visited. Since, by Definition 8 the NP method estimates the best solution as

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

it can be seen that with probability one as $k \to \infty$,

$$\hat{\sigma}_{opt}(k) \to \arg\max_{\sigma \in \Sigma_0} \pi(\sigma).$$

Hence, the algorithm converges to the singleton region that maximizes the stationary distribution. Now to show that this singleton region is indeed $\sigma_{opt} = \{\theta_{opt}\}$, first note that the Markov chain is reversible and we hence have that for any $\eta \in \Sigma_0$,

$$P^{\kappa(\eta,\sigma_{opt})}\left(\eta, \sigma_{opt}\right)\pi(\eta) = P^{\kappa(\eta,\sigma_{opt})}\left(\sigma_{opt}, \eta\right)\pi\left(\sigma_{opt}\right),$$

where $\kappa(\eta, \sigma_{opt})$ is the number of transitions it takes to go from η to σ_{opt} and vice versa. Hence, if the $\kappa(\eta, \sigma_{opt})$-step transition probability from η to σ_{opt} is larger than the $\kappa(\eta, \sigma_{opt})$-step transition probability from η to σ_{opt} to η for all $\eta \in \Sigma_0 \setminus \{\sigma_{opt}\}$, then

$$\sigma_{opt} = \arg\max_{\eta \in \Sigma_0} \pi(\eta)$$

and the theorem is proven.

To see why this holds, we look at the superregion of the optimum, $s\left(\sigma_{opt}\right)$. By equation (10) it is clear that the probability of moving to the σ_{opt} is larger than the probability of backtracking to $s\left(s\left(\sigma_{opt}\right)\right)$,

$$\begin{aligned} P\left(s\left(\sigma_{opt}\right), \sigma_{opt}\right) &= P\left[\hat{I}\left(\sigma_{opt}\right) \le \hat{I}\left(\Theta \setminus s\left(\sigma_{opt}\right)\right)\right] \\ &\ge P\left[\hat{I}\left(\Theta \setminus s\left(\sigma_{opt}\right)\right) \le \hat{I}\left(\sigma_{opt}\right)\right] \\ &= P\left(s\left(\sigma_{opt}\right), s\left(s\left(\sigma_{opt}\right)\right)\right), \end{aligned}$$

and in general, the same result holds for any region on the 'path' between σ_{opt} and an arbitrary $\eta \in \Sigma_0 \setminus \{\sigma_{opt}\}$. We conclude that σ_{opt} is a singleton region that maximizes the stationary probability and the theorem holds.

It remains to justify that equation (10) may indeed be satisfied then applying the method in practice, and how (10) relates to the implementation parameters of the method, in particular the partitioning and sampling. We approach this via the perspective of ordinal comparisons.

3.2 Ordinal Comparison

To show analytically that using ordinal comparison is beneficial we use the following theorem from Dai (1996), which shows that (9) converges rapidly when used to estimate the promising index.

Theorem 2 *Let* $\mathcal{D} \subseteq \Theta$ *and let* $\Theta_g = \mathcal{D} \cap \mathcal{G}$ *be the good solutions and let* $\Theta_b = \mathcal{D} \setminus \Theta_b$ *denote the bad solutions. We assume that* $\Theta_g \neq \emptyset$ *and* $\Theta_b \neq \emptyset$. *Then the probability of the estimated best solution in* Θ_g *being better than the estimated best solution in* Θ_b *converges to one at an exponential rate.*

$$P\left[\min_{\theta \in \Theta_g} L_t(\theta) \le \min_{\theta \in \Theta_b} L_t(\theta)\right] = 1 - O\left(e^{-\alpha t}\right), \tag{12}$$

and

$$P\left[\min_{\theta \in \Theta_g} L_t(\theta) > \min_{\theta \in \Theta_b} L_t(\theta)\right] = O\left(e^{-\alpha t}\right). \tag{13}$$

Proof: See Theorem 4.5 in Dai (1996).

We immediately obtain the following theorem.

Theorem 3 *Assume that two regions* $\sigma_g \in \Sigma_g$ *and* $\sigma_b \in \Sigma \setminus \Sigma_g$ *are compared, where* σ_g *contains the global optimum but* σ_b *does not. Let* $\mathcal{D}_{\sigma_g}$ *denote the set of sample points from* σ_g, *and similarly* $\mathcal{D}_{\sigma_b}$ *denote the set of sample points from* σ_b. *Then*

$$\begin{aligned} P\left[\hat{I}(\sigma_g) \le \hat{I}(\sigma_b)\right] &= P\left[\min_{\theta \in \mathcal{D}_{\sigma_g}} J(\theta) < \min_{\theta \in \mathcal{D}_{\sigma_b}} J(\theta)\right] \\ &\quad + O\left(e^{-\alpha t}\right), \end{aligned} \tag{14}$$

where t *is the simulation time.*

Proof: By conditioning on the best solution sampled being from the good region, that is,

$$A = \left[\min_{\theta \in \mathcal{D}_{\sigma_g}} J(\theta) < \min_{\theta \in \mathcal{D}_{\sigma_b}} J(\theta)\right],$$

this follows directly from Theorem 4:

$$
\begin{aligned}
P\left[\hat{I}(\sigma_g) \le \hat{I}(\sigma_b)\right] &= P\left[\min_{\theta\in\mathcal{D}_{\sigma_g}} L_t(\theta) \le \min_{\theta\in\mathcal{D}_{\sigma_b}} L_t(\theta)\right] \\
&= P\left[\min_{\theta\in\mathcal{D}_{\sigma_g}} L_t(\theta) \le \min_{\theta\in\mathcal{D}_{\sigma_b}} L_t(\theta)\Big| A\right]\cdot\ P[A] \\
&+\ P\left[\min_{\theta\in\mathcal{D}_{\sigma_g}} L_t(\theta) \le \min_{\theta\in\mathcal{D}_{\sigma_b}} L_t(\theta)\Big| A^c\right]\cdot\ (1-P[A]) \\
&= \left(1-O\left(e^{-\alpha t}\right)\right)\cdot P[A] +\ O\left(e^{-\alpha t}\right)\cdot(1-P[A]) \\
&= P\left[\min_{\theta\in\mathcal{D}_{\sigma_g}} J(\theta) < \min_{\theta\in\mathcal{D}_{\sigma_b}} J(\theta)\right] + O\left(e^{-\alpha t}\right).
\end{aligned}
$$

This proves the theorem.

In the k-th iteration of the NP method exactly one of the subregions sampled, say $\sigma_{j^*}(k) \in \Sigma_g$, contains the global optimum. This subregion is compared with all of the other regions, and will be selected if $\hat{I}(\sigma_{j^*}(k)) \le \hat{I}(\sigma_j(k))$, $\forall k = 1, 2, ..., M(\sigma(k)) + 1$. It follows that the method is inherently ordinal and by Theorem 3 the probability of moving towards $\sigma_{j^*}(k)$ in the next iteration converges exponentially fast to a probability that depends only on which solutions were randomly selected in the current iteration. In other words, if we define the probability of selecting the best solution from the right region as $P^*\left(\sigma_{j^*}(k)\right) = P\left[\min_{\theta\in\mathcal{D}_{\sigma_{j^*}(k)}} J(\theta) < \min_{\theta\in\mathcal{D}_{\sigma_{j}(k)}} J(\theta),\ \forall j \ne j^*\right]$, then Theorem 3 states that

$$
\begin{aligned}
&P\left[\hat{I}(\sigma_{j^*}(k)) \le \hat{I}(\sigma_j(k)),\ \forall k = 1, 2, ..., M(\sigma(k)) + 1\right] \\
&\qquad = P^*(\sigma_{j^*}(k)) + O\left(e^{-\alpha t}\right),
\end{aligned}
$$

where t is as before the simulation time. The probability $P^*(\sigma_{j^*}(k))$ can be made large by partitioning such that many good solutions fall in the same regions or by increasing the sampling effort in each iteration. This probability depends on comparing multiple regions, but to simplify the analysis we assume without loss of generality that we only compare two regions $\sigma_g \in \Sigma_g$ and $\sigma_g \in \Sigma \setminus \Sigma_g$. Accordingly, we define the *success probability*

$$
P^*\left(\sigma_g, \sigma_b\right) = P\left[\min_{\theta\in\mathcal{D}_{\sigma_g}} J(\theta) < \min_{\theta\in\mathcal{D}_{\sigma_b}} J(\theta)\right], \qquad (15)
$$

for all $\sigma_g \in \Sigma_g, \sigma_g \in \Sigma \setminus \Sigma_g$. For the remainder of the paper we focus on how $P^*\left(\sigma_g, \sigma_b\right)$ depends on the partitioning strategy and the sampling effort, and how it can be made sufficiently large.

3.3 Partitioning and Sampling

To better understand the relationship between the partitioning and the required sampling effort, we start by looking at the ideal case.

Theorem 4 *Let the assumptions and definitions of $\sigma_g \in \Sigma_g, \sigma_b \in \Sigma \setminus \Sigma_g$ be as in Theorem 3. If Σ is an optimal partition then*

$$
P\left[\hat{I}(\sigma_g) \le \hat{I}(\sigma_b)\right] = 1 + O\left(e^{-\alpha t}\right), \qquad (16)
$$

where t is the simulation time.

Proof: By Definition 5 we have that $m(\sigma_{opt}, \sigma_g)_\Sigma < m(\sigma_{opt}, \sigma_b)_\Sigma$, so by Definition 6 of an optimal partition

$$
J(\theta) < J(\phi),\ \forall\theta \in \sigma_g, \phi \in \sigma_b.
$$

Therefore,

$$
P\left[\min_{\theta\in\mathcal{D}_{\sigma_g}} J(\theta) < \min_{\theta\in\mathcal{D}_{\sigma_b}} J(\theta)\right] = 1
$$

so the theorem follows directly from Theorem 3 above.

We note that Theorem 3 and Theorem 4 provide new insights into when the NP method converges to the global optimum. In particular, Theorem 3 implies that if $P^*\left(\sigma_g, \sigma_b\right) > \frac{1}{2}$ for all $\sigma_g \in \Sigma_g$, $\sigma_b \in \Sigma \setminus \Sigma_g$, then the global convergence condition (10) will be satisfied at an exponential rate in terms of the simulation effort used for evaluating each solution. By Theorem 6 this clearly holds for optimal partitioning. In practice, however, optimal partitioning is never realized, and it is therefore of interest to determine how good the partitioning needs to be. It is also clear that as the partitioning becomes worse, more sample effort may be needed from each region. To measure the quality of a partitioning strategy Σ we define the *non-overlap set* function, $\Psi : \Sigma_g \to \Theta$ by

$$
\Psi(\sigma_g) = \left\{\theta \in \sigma_g : J(\theta) < J(\psi), \forall\psi \in \Theta \setminus \sigma_g\right\}, \qquad (17)
$$

$\sigma_g \in \Sigma_g$. This function counts, for each good region $\sigma_g \in \Sigma_g$, how many of the solutions in the good region have better expected performance than all of the solutions outside this region, that is, the non-overlap in expected performance. A high value indicates that it may be easy to differentiate between the the good region and other regions, and vice versa for low values. By definition of Σ_g we have that $\theta_{opt} \in \Psi(\sigma_g)$ so $\Psi(\sigma_g) \ne \emptyset$ for all $\sigma_g \in \Sigma_g$. It is also clear that if Σ is optimal then by Definition 6, $J(\theta) < J(\psi)$ for all $\theta \in \sigma_g, \psi \in \Theta \setminus \sigma_g$, so $\Psi(\sigma_g) = \sigma_g$ for all $\sigma_g \in \Sigma_g$. Therefore, the size of this set $|\Psi(\sigma_g)| \in \{1, 2, ..., |\sigma_g|\}$ for all $\sigma \in \Sigma_g$ is a measure of the quality Σ. We now obtain the following theorem.

Theorem 5 *Let the assumptions and definitions of $\sigma_g \in \Sigma_g, \sigma_b \in \Sigma \setminus \Sigma_g$ be as in Theorem 5. Let $n(\sigma_g) = |\mathcal{D}_{\sigma_g}|$ be the number of sample points from $\sigma_g \in \Sigma_g$. Define $r(\sigma_g) = \frac{|\sigma_g| - |\Psi(\sigma_g)|}{|\sigma_g|}$ to be the percentage overlap, and assume that*

$$n(\sigma_g) \geq \frac{\log(\frac{1}{2})}{\log(r(\sigma_g))}, \tag{18}$$

and that uniform sampling is used. Then the global convergence condition (10) is satisfied at an exponential rate, that is,

$$P\left[\hat{I}(\sigma_g) \leq \hat{I}(\sigma_b)\right] \geq \frac{1}{2} + O\left(e^{-\alpha t}\right), \tag{19}$$

where t is the simulation time.

Proof: It is clear that if one of the solutions in $\Psi(\sigma_g)$ is selected in $\mathcal{D}_{\sigma_g}$ then the best solution in $\mathcal{D}_{\sigma_g}$ is better than the best solution in $\mathcal{D}_{\sigma_b}$. That is,

$$\begin{aligned} &P\left[\min_{\theta \in \mathcal{D}_{\sigma_g}} J(\theta) < \min_{\theta \in \mathcal{D}_{\sigma_b}} J(\theta)\right] \\ &\quad \geq P\left[\Psi(\sigma_g) \cap \mathcal{D}_{\sigma_g} \neq \emptyset\right] \\ &\quad = 1 - P\left[\Psi(\sigma_g) \cap \mathcal{D}_{\sigma_g} = \emptyset\right] \\ &\quad = 1 - \left(\frac{|\sigma_g| - |\Psi(\sigma_g)|}{|\sigma_g|}\right)^{n(\sigma_g)}, \end{aligned}$$

where the last equation follows from the uniform sampling strategy. On the other hand, by the assumption (18) we have

$$\begin{aligned} \left(\frac{|\sigma_g| - |\Psi(\sigma_g)|}{|\sigma_g|}\right)^{n(\sigma_g)} &= r(\sigma_g)^{n(\sigma_g)} \\ &\leq r(\sigma_g)^{\frac{\log(\frac{1}{2})}{\log(r(\sigma_g))}} \\ &= \left(e^{\log(r(\sigma_g))}\right)^{\frac{\log(\frac{1}{2})}{\log(r(\sigma_g))}} \\ &= e^{\log(\frac{1}{2})} \\ &= \frac{1}{2}, \end{aligned}$$

so

$$\left[\min_{\theta \in \mathcal{D}_{\sigma_g}} J(\theta) < \min_{\theta \in \mathcal{D}_{\sigma_b}} J(\theta)\right] \geq \frac{1}{2},$$

which, combined with Theorem 3, proves the theorem.

This theorem illustrates the relationship between the partitioning and the sampling effort needed. If the partitioning is poor, that is $|\Psi(\sigma_g)|$ small for at least some $\sigma_g \in \Sigma_g$, then more sample effort is need, and vice versa. In particular, if $|\Psi(\sigma_g)| \geq \frac{|\sigma_g|}{2}$ for all $\sigma_g \in \Sigma_g$, then (19) is satisfied even if we use only one sample solution from each region. Moreover, Theorem 5 illustrates just how important a good partitioning strategy is, because the lower bound (18) on the number of sample solutions needed converges to one at an exponential rate as $|\Psi(\sigma_g)|$ goes to $\frac{|\sigma_g|}{2}$ from below. This is illustrated in Figure 1 where the minimum required number of sample points to obtain a given success probability $P^*\left(\sigma_g, \sigma_b\right) \in \{0.25, 0.50, 0.75\}$ is plotted against the percentage overlap

$$r(\sigma_g) = \frac{|\sigma_g| - |\Psi(\sigma_g)|}{|\sigma_g|} \in [0.50, 0.95],$$

that is $|\Psi(\sigma_g)| \in [0.05, 0.50]$. The opposite is also true, increasing the sampling effort in each iteration leads to exponential improvement in the success probability as is illustrated in Figure 2 for four different partitioning quality levels $r(\sigma_g) \in \{0.5, 0.7, 0.9, 0.99\}$.

We conclude that when optimizing certain systems using regenerative simulation, ordinal rather than cardinal optimization is indeed beneficial. Furthermore, this translates into weak convergence conditions for the NP algorithm, and relatively little simulation effort being needed in each iteration.

4 CONCLUSIONS

We have analyzed a new simulation-based optimization algorithm that draws from the paradigm of ordinal optimization and a recently proposed adaptive sampling algorithm called the nested partitions (NP) method. The new algorithm falls into the NP method framework, which guarantees global convergence under certain conditions, and the ordinal optimization perspective is used to show that for certain problems the method also has certain exponential convergence rate characteristics. We derived new conditions under

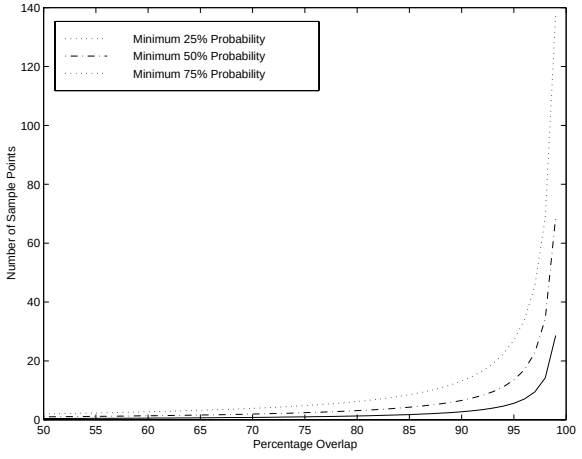

Figure 1: Sample Effort Needed In Each Iteration

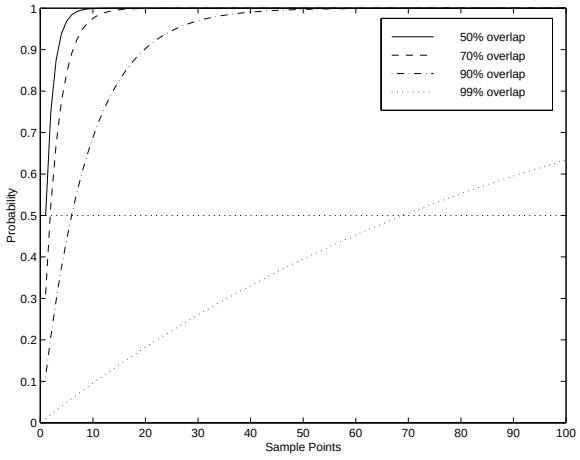

Figure 2: Probability of Correct Selection

which asymptotic convergence holds and provided practical guidelines for determining the sampling effort in each iteration.

5 ACKNOWLEDGEMENT

This research was supported in part by the National Science Foundation under grant DMI-9713647.

REFERENCES

Alrefaei, M.H., and S. Andradóttir. 1997. Accelerating the convergence of the stochastic ruler method. In *Proceedings of the 1997 Winter Simulation Conference*, ed. S. Andradóttir, K.J. Healy, D.H. Withers, and B.L. Nelson, 352-357. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Andradóttir, S. 1995. A method for discrete stochastic optimization. *Management Science* 41: 1946-1961.

Andradóttir, S. 1998. A review of simulation optimization techniques. In *Proceedings of the 1998 Winter Simulation Conference*, ed. D.J. Medeiros, E.F. Watson, J.S. Carson, and M.S. Manivannan, 151-158. Institute of Electrical and Electronics Engineers. Piscataway, New Jersey.

Carson, Y. and A. Maria. 1997. Simulation optimization: methods and applications. In *Proceedings of the 1997 Winter Simulation Conference*, ed. S. Andradóttir, K.J. Healy, D.H. Withers, and B.L. Nelson, 118-126. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Dai, L. 1996. Convergence Properties of Ordinal Comparison in the Simulation of Discrete Event Dynamic Systems. *Journal of Optimization Theory and Applications*, 91, 363-388

Gong, W.-B., Y.-C. Ho, and W. Zhai. 1992. Stochastic Comparison Algorithm for Discrete Optimization with Estimation. In *Proceedings of the 31st IEEE Conference on Decision and Control*, 795-800.

Ho, Y.-C., R.S. Sreenivas, and P. Vakili. 1992. Ordinal Optimization of DEDS. *Discrete Event Dynamic Systems: Theory and Applications*, 2, 61-88.

Ólafsson, S. and L. Shi. 1998. Stopping Criterion for a Simulation-Based Optimization Method. In *Proceedings of the 1998 Winter Simulation Conference*, ed. D.J. Medeiros, E.F. Watson, J.S. Carson, and M.S. Manivannan, 743-750. Institute of Electrical and Electronics Engineers. Piscataway, New Jersey.

Shi, L. and S. Ólafsson. 1997. An integrated framework for deterministic and stochastic optimization. In *Proceedings of the 1997 Winter Simulation Conference*, ed. S. Andradóttir, K.J. Healy, D.H. Withers, and B.L. Nelson, 352-357. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey.

Shi, L and S. Ólafsson. 1998a. Nested Partitions Method for Global Optimization. *Operations Research*, to appear.

Shi, L and S. Ólafsson. 1998b. Nested Partitions Method for Stochastic Optimization. Working Paper, 98-115, Department of Industrial and Manufacturing Systems Engineering, Iowa State University.

Tang, Z.B. 1994. Adaptive Partitioned Random Search to Global Optimization. *IEEE Transactions on Automatic Control*, 39, 2235-2244.

Yan, D. and H. Mukai. 1992. Stochastic Discrete Optimization. *SIAM Journal Control and Optimization*, 30, 594-612.

AUTHOR BIOGRAPHIES

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.

LEYUAN SHI is an Assistant Professor in the Department of Industrial Engineering at the University of Wisconsin-Madison. She holds a B.S. degree in Mathematics from Nanjing Normal University, China (1982), an M.S. degree in Applied Mathematics from Tsinghua University, China (1985), and an M.S. and a Ph.D. degrees in Applied Mathematics from Harvard University (1990,1992). Her research interests include modeling, analysis, and optimization of discrete event systems, discrete-event simulation, and sensitivity analysis.