Proceedings of the 1998 Winter Simulation Conference
D.J. Medeiros, E.F. Watson, J.S. Carson and M.S. Manivannan, eds.

STOPPING CRITERION FOR A SIMULATION-BASED OPTIMIZATION METHOD

Sigurdur Ólafsson
Leyuan Shi

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

ABSTRACT

We consider a new simulation-based optimization method called the Nested Partitions (NP) method. This method generates a Markov chain and solving the optimization problem is equivalent to maximizing the stationary distribution of this Markov chain over certain states. The method may therefore be considered a Monte Carlo sampler that samples from the stationary distribution. We show that the Markov chain converges geometrically fast to the true stationary distribution, and use these results to derive a stopping criterion for the method.

1 INTRODUCTION

In system optimization it is often desirable to optimize the performance of a system where the design 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 relatively little attention in the research literature, although some important advances have been made in the last few years.

Recent methods proposed for this problem include 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 Ólafsson, 1997;1998). Under certain conditions, all of these methods have been shown to converge almost surely to an optimal solution, but a common difficulty is obtaining good stopping rules. This is of paramount importance in practice, because although asymptotic convergence is of much interest, such results do not indicate how long a method must be applied for satisfactory results, or how and when the method should be terminated. In this paper we propose a stopping rule for the NP method.

The basic idea of the NP method is very simple. The method systematically partitions the feasible region into smaller and smaller subregions, and moves iteratively from one region to another in order to concentrate the search effort in viable regions. This procedure has been shown to generate a Markov chain, which we refer to as the NP Markov chain. It can be shown that the NP Markov chain has a unique stationary distribution and given certain regularity assumptions the maximum stationary probability corresponds to a global optimum. Hence, the NP method converges with probability one to a global optimum (Shi and Ólafsson, 1998). In this paper we consider a stopping rule for this method. Since the stationary distribution of the NP Markov chain is used for inference, the convergence of the NP method can be translated into convergence of the NP Markov chain. The NP method may hence be considered a Monte Carlo sampler that samples from the stationary distribution of the NP Markov chain, and its convergence properties may therefore be inferred by considering its efficiency as an Markov Chain Monte Carlo (MCMC) sampler.

The use of MCMC methods has become increasingly popular in recent years, and many theoretical advancements have also been made (Rosenthal, 1995). However, there is still a considerable gap between theory and practice in the field (Brooks and Roberts, 1998). Although many nice theoretical bounds have been derived for the convergence rate of Markov chains, such bounds normally contain constants that are not available a priori except for special cases. Therefore, such bounds are not directly applicable in practice. In fact, practitioners most often use convergence diagnostics to assess the convergence of a Markov chain simulation (Cowles and Carlin, 1996; Brooks and Roberts,

1998). Such convergence diagnostics usually consider some summary statistic of the Markov chain, and declare convergence when it appears to have settled down in steady state. The problems with such diagnostics are well known. Most seriously, it is inevitable that such methods will sometime declare convergence too soon, since the Markov chain 'settling down' is only a necessary condition for convergence, but not sufficient.

Another method of bridging the gap between theory and practice is to use auxiliary simulation to estimate the required constants (Cowles and Rosenthal, 1996). This is the approach taken here. We use theoretical bounds as a basis for a stopping criteria, and then use auxiliary simulation to estimate a theoretical constant that is required for this bound. By combining these two ingredients we obtain an approximate stopping criterion that can be applied during simulation optimization.

The remainder of this paper is organized as follows. In Section 2 we briefly review the Nested Partitions (NP) method for simulation-based optimization. This algorithm generates a Markov chain and in Section 3 we derive bounds that establish that it converges geometrically fast to its stationary distribution. In Section 4 we show how these bounds may be estimated using auxiliary simulation, and in Section 5 we use these results to derive a stopping criterion for the method. In Section 6 we present a numerical example to illustrate the auxiliary simulation approach and the stopping criterion, and finally, Section 7 contains some concluding remarks.

2 THE NP METHOD

We are interested in solving the general stochastic discrete optimization problem, which can be stated as follows.

$$\min_{\theta\in\Theta} E[L(\theta,\omega)], \tag{1}$$

where Θ is finite, $L(\theta,\omega)$ is the sample performance, and $\omega \in \Omega$ is a point in the sample space of an underlying probability triple $(\Omega, \mathcal{F}, \mathcal{P})$. To keep the notation simple we do not refer explicitly to the sample space in this paper. We assume that the expected performance function cannot be evaluated analytically and is estimated using discrete-event simulation. We now review briefly how the nested partitions (NP) method can be used to solve this problem.

The NP method proceeds as follows. In the k-th iteration there is a region $\sigma(k) \subseteq \Theta$ that is considered the most promising. In the first iteration $\sigma(1) = \Theta$, the entire feasible region. In each iteration, this most promising region is partitioned into a fixed number of subregions and the entire surrounding region is aggregated into one region. Therefore, at each iteration, a fixed number of disjoint subsets that cover the feasible region is considered. Each of these regions is sampled using a random sampling scheme, with the only restriction being that each point must be sampled with a positive probability. The estimated performance function values of the randomly selected samples are used to estimate the promising index for each region. This index 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. To choose this larger region a fixed backtracking rule is used. The new most promising region is then partitioned and sampled in a similar fashion. We note that this partitioning creates a tree structure similar to that of the branch-and-bound method.

We assume that Θ is partitioned such that eventually every point corresponds to a singleton region. We call such singletons regions of maximum depth, and more generally talk about the depth of any region. This is defined iteratively with Θ having depth zero and so forth. We let Σ denote the space of all regions that are formed by a fixed partitioning scheme as described above, and we let $d(\sigma)$ denote the depth of a region $\sigma \in \Sigma$. Ultimately, we are interested in finding the best singleton region. Therefore, we let Σ_0 denote all singleton regions, and in the k-th iteration we let the estimated best region, $\sigma_k^{(1)}$, be the singleton region that has been most frequently considered the most promising region. To be able to update $\sigma_k^{(1)}$, we let $N_k(\sigma)$ be the number of times a region $\sigma \in \Sigma_0$ has been visited. We only define these counters for maximum depth regions that have been visited at least once. To keep track of the sequence of set partitions leading to the current most promising region, we define the 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 $s(\Theta) = \Theta$. Therefore, $s(\sigma) \in \Sigma$ is the region that was partitioned to obtain the region $\sigma \in \Sigma$.

Using the notation defined above, the NP method can be implemented as follows.

0. **Initialization.** Set $k = 0$ and $\sigma(k) = \Theta$.

1. **Partitioning.** If the depth is not the maximum depth, partition the most promising region, $\sigma(k)$, into $M_{\sigma(k)}$ subregions $\sigma_1(k), ..., \sigma_{M_{\sigma(k)}}(k)$. If the depth is maximum depth then let $M_{\sigma(k)} = 1$ and $\sigma_1(k) = \sigma(k)$. If the depth is not zero, i.e., $\sigma(k) \neq \Theta$, aggregate the surrounding region $\Theta \setminus \sigma(k)$ into one region $\sigma_{M_{\sigma(k)}+1}(k)$.

2. **Random Sampling.** Use a random sampling procedure to select N_j points $\theta^{j1}, \theta^{j2}, ..., \theta^{jN_j}$ and calculate the corresponding sample performance values

$L(\theta^{j1}), L(\theta^{j2}), ..., L(\theta^{jN_j})$ from each of the regions $\sigma_j(k),\ j = 1, 2, ...,\ M_{\sigma(k)} + 1$.

3. **Estimating the Promising Index.** Given a *promising index* function, $I: \Sigma \to \mathbf{R}$, for each region $\sigma_j,\ j = 1, 2, ..., M_{\sigma(k)} + 1$, calculate the estimated promising index $\hat{I}(\sigma_j)$ of the region. In this paper we select the promising index as $I(\eta) = \min_{\theta \in \eta} J(\theta)$, which can be estimated using

$$\hat{I}(\sigma_j) = \min_{\theta = \theta^{j1}, ..., \theta^{jN_j}} L(\theta), \tag{2}$$

for $j = 1, 2, ..., M_{\sigma(k)} + 1$.

4. **Backtracking.** Calculate the index of the region with the best promising index.

$$\hat{j}_k \in \arg \min_{j=1,...,M_{\sigma(k)}+1} \hat{I}(\sigma_j). \tag{3}$$

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. That is $\sigma(k+1) = \sigma_j(k),\ j \leq M_{\sigma(k)}$. If the index corresponds to the surrounding region, *backtrack* to the superregion of $\sigma(k)$

$$\sigma(k+1) = s(\sigma(k)). \tag{4}$$

5. **Updating Counters.** Update the number of times each maximum depth region has been visited. If $\sigma(k+1) \in \Sigma_0$ let $N_{k+1}(\sigma(k+1)) = N_{k+1}(\sigma(k+1)) + 1$. Let $N_{k+1}(\sigma) = N_{k+1}(\sigma)$ for all other $\sigma \in \Sigma_0$ that have been visited at least once. Update the estimated best region as follows. If $\sigma(k+1) \in \Sigma_0$ and $N_{k+1}(\sigma(k+1)) > N_{k+1}(\sigma_k^{(1)})$, then let $\sigma_{k+1}^{(1)} = \sigma(k+1)$. Otherwise, let $\sigma_{k+1}^{(1)} = \sigma_k^{(1)}$. Let $k = k+1$ and go back to Step 1.

In this paper we assume that there is a unique global optimum that corresponds to a singleton region $\sigma_{opt} \in \Sigma_0$. In Shi and Ólafsson (1998) it is shown that the NP algorithm above generates a Markov chain and under certain regularity conditions σ_{opt} is the singleton state that has the largest stationary probability. In each iteration of the NP algorithm an estimate of the stationary distribution is generated,

$$\mu_k(\sigma) = \frac{N_k(\sigma)}{k},\ \sigma \in \Sigma, \tag{5}$$

and as is well known, this estimate converges asymptotically to the true stationary distribution. Therefore, by taking $\sigma_k^{(1)} = \arg\max_{\sigma \in \Sigma_0} N_k(\sigma)$ as the estimate of the global optimum, the NP algorithm converges to the unique global optimum (Shi and Ólafsson, 1998). Because of this the NP algorithm may be considered an MCMC sampler that samples from the stationary distribution. When a sufficiently good estimate of the stationary distribution has been obtained, the algorithm has converged. Here we are therefore concerned with how fast this estimated stationary distribution converges to the true stationary distribution. Perhaps the most important feature of an efficient MCMC sampler is that the Markov chain is geometrically ergodic (Roberts and Rosenthal, 1998), and in the next section we show this holds here.

3 GEOMETRIC CONVERGENCE RATE

As before we let μ_k denote the empirical stationary distribution. We are interested in determining when the distance between μ_k and the true stationary distribution π is within a certain predetermined tolerance. When such sufficient accuracy is obtained the estimated stationary distribution can be used for inference about which singleton region has the largest stationary probability and the NP method has converged. To measure the distance between these two probability distributions we use the *total variation distance* norm.

Definition 1 *Given two probability measures υ_1 and υ_2 the norm*

$$||\upsilon_1 - \upsilon_2||_{var} = \sup_{A \subseteq \Sigma} |\upsilon_1(A) - \upsilon_2(A)|$$

is called the total variation distance. □

We let $1 = \lambda_1 > \lambda_2 \geq \lambda_3 ... \geq \lambda_{|\Sigma|} > -1$ denote the eigenvalues of the transition matrix of the NP Markov chain. It is well known that the rate of convergence to the stationary distribution is determined by the eigenvalue that is the second largest in absolute value $\lambda_{max} = \max\{\lambda_2, |\lambda_{|\Sigma|}|\}$. If all the eigenvalues are positive then $\lambda_{max} = \lambda_2$. Negative eigenvalues correspond to near cyclic behavior and we can always modify the NP algorithm in a trivial manner to ensure that all the eigenvalues are positive. This can for example be achieved by introducing a self-loop with probability $\frac{1}{2}$ for every state, that is, with probability $\frac{1}{2}$ we move to a subregion or backtrack, and with probability $\frac{1}{2}$ we stay in the current most promising region for one more iteration. We can therefore assume from now on that $\lambda_{max} = \lambda_2$.

We also need the following definitions.

Definition 2 *A set $\mathcal{A} \subseteq \Sigma$ is* connected *if for every $\eta_1, \eta_2 \in \mathcal{A}$, there exists a path from η_1 to η_2 such that every state in the path is in $\mathcal{A}$.*

Definition 3 *The* ergodic flow *out of $\mathcal{A}$ is defined as*

$$F(\mathcal{A}) = \sum_{\substack{\eta\in\mathcal{A}\\ \xi\in\overline{\mathcal{A}}}} P(\eta,\xi)\cdot\pi(\eta). \tag{6}$$

The conductance *of the Markov chain defined by (P,Σ) is*

$$\Phi = \min_{0<\pi(\mathcal{A})\le\frac{1}{2}} \Phi(\mathcal{A}), \tag{7}$$

where

$$\Phi(\mathcal{A}) = \frac{F(\mathcal{A})}{\pi(\mathcal{A})}. \tag{8}$$

Our objective is to calculate the conductance of the NP Markov chain. The next two lemmas show how the minimum in equation (7) can be taken over a limited number of sets.

Intuitively, the conductance is a measure of the bottlenecks in the probability 'flow' in the partitioning tree. Intuitively this 'flow' is likely to be small for subsets where there is only one state through which the NP Markov chain can exit the subset. This leads us to connected subsets that are such that the Markov chain can only leave the set through the 'root' node, that is, entire branches. The first lemma shows that the set on which the minimum (7) is realized must be a connected set.

Lemma 1 *Let $\mathcal{A}\subset\Sigma$ be an unconnected subset. Then there exist proper subsets $\mathcal{A}_1\subset\mathcal{A}$ and $\mathcal{A}_2\subset\mathcal{A}$ such that $\Phi(\mathcal{A})\ge\min\{\Phi(\mathcal{A}_1),\Phi(\mathcal{A}_2)\}$.*

Proof: Since $\mathcal{A}$ is disjoint we can select $\mathcal{A}_1$ and $\mathcal{A}_2$ such that $\mathcal{A}=\mathcal{A}_1\cup\mathcal{A}_2$, $\mathcal{A}_1\cap\mathcal{A}_2=\emptyset$, and $P(\mathcal{A}_1,\mathcal{A}_2)=0$. Then it is clear that $\pi(\mathcal{A}_1)+\pi(\mathcal{A}_2)=\pi(\mathcal{A})$ and $F(\mathcal{A}_1)+F(\mathcal{A}_2)=F(\mathcal{A})$. Therefore,

$$\begin{aligned}
\Phi(\mathcal{A}_1) &= \frac{F(\mathcal{A}_1)}{\pi(\mathcal{A}_1)}\\
&= \frac{F(\mathcal{A})-F(\mathcal{A}_2)}{\pi(\mathcal{A}_1)}\\
&= \frac{\pi(\mathcal{A})\Phi(\mathcal{A})-\pi(\mathcal{A}_2)\Phi(\mathcal{A}_2)}{\pi(\mathcal{A}_1)}\\
&= \frac{(\pi(\mathcal{A}_1)+\pi(\mathcal{A}_2))\Phi(\mathcal{A})-\pi(\mathcal{A}_2)\Phi(\mathcal{A}_2)}{\pi(\mathcal{A}_1)}\\
&= \Phi(\mathcal{A})+\frac{\pi(\mathcal{A}_2)}{\pi(\mathcal{A}_1)}\left(\Phi(\mathcal{A})-\Phi(\mathcal{A}_2)\right).
\end{aligned}$$

Now if $\Phi(\mathcal{A})<\Phi(\mathcal{A}_2)$ then $\Phi(\mathcal{A})>\Phi(\mathcal{A}_1)$, so $\Phi(\mathcal{A})\ge\min\{\Phi(\mathcal{A}_1),\Phi(\mathcal{A}_2)\}$. □

This establishes that equation (7) is realized on a connected subset. The next lemma shows that this subset has even more structure.

Lemma 2 *Let $\mathcal{A}\subset\Sigma$ be a connected set. Let $\mathcal{B}$ be the set of all regions that are a subset of a region in $\mathcal{A}$,*

$$\mathcal{B} = \bigcup_{\xi:\exists\eta\in\mathcal{A},\xi\subset\eta}\xi \tag{9}$$

Let $\mathcal{A}^*=\mathcal{A}\cup\mathcal{B}$. Then

$$\Phi(\mathcal{A}^*)\le\Phi(\mathcal{A}). \tag{10}$$

Proof: Clearly

$$\pi(\mathcal{A}^*)=\pi(\mathcal{A})+\pi(\mathcal{B})\ge\pi(\mathcal{A}),$$

and since the Markov chain can only exist $\mathcal{B}$ by entering $\mathcal{A}$, it is also clear that

$$F(\mathcal{A}^*)\le F(\mathcal{A}).$$

Therefore,

$$\Phi(\mathcal{A}^*)=\frac{F(\mathcal{A}^*)}{\pi(\mathcal{A}^*)}\le\frac{F(\mathcal{A})}{\pi(\mathcal{A})}=\Phi(\mathcal{A})$$

and the lemma is proved. □

We now obtain the first major result, which is a lower bound on the conductance.

Theorem 1 *The conductance Φ of the NP Markov chain is bounded from below by*

$$\Phi\ge\frac{(1-C)P(\sigma_{opt},s(\sigma_{opt}))}{1-C^{d^*}}, \tag{11}$$

where

$$C=\frac{1-P(\sigma_{opt},s(\sigma_{opt}))}{P(\sigma_{opt},s(\sigma_{opt}))}.$$

Proof: By Lemma 1 and Lemma 2 it suffices to consider connected sets $\mathcal{A}$ which contain all the regions that are a subregion of a region in $\mathcal{A}$. Then there exists a state $\sigma_{\mathcal{A}}\in\mathcal{A}$ that contains all the other states $\sigma\subseteq\sigma_{\mathcal{A}}$, for all $\sigma\in\mathcal{A}$. Furthermore, since the minimum in equation (7) is taken of all sets that have less that or equal to $\frac{1}{2}$ in probability mass, it is clear that $\sigma_{\mathcal{A}}\neq\Theta$. Therefore, the flow out of set $\mathcal{A}$ is

$$F(\mathcal{A})=\pi(\sigma_{\mathcal{A}})P(\sigma_{\mathcal{A}},s(\sigma_{\mathcal{A}})).$$

Now we can write down the balance equations for states in $\mathcal{A}$ that are of different but adjacent depth levels as follows,

$$\sum_{\substack{\eta\in\mathcal{A}\\ d(\eta)=k+1}} P(\eta,s(\eta))\pi(\eta)=\sum_{\substack{\eta\in\mathcal{A}\\ d(\eta)=k}}(1-P(\eta,s(\eta)))\pi(\eta),$$

and furthermore, we can assume that for the unique global optimum σ_{opt} we have $P(\sigma_{opt}, s(\sigma_{opt})) \leq P(\eta, s(\eta))$ forall $\eta \in \Sigma$. Therefore,

$$\begin{aligned} P(\sigma_{opt}, s(\sigma_{opt})) \sum_{\substack{\eta\in\mathcal{A} \\ d(\eta)=k+1}} \pi(\eta) & \\ \leq & \sum_{\substack{\eta\in\mathcal{A} \\ d(\eta)=k+1}} P(\eta, s(\eta))\pi(\eta) \\ = & \sum_{\substack{\eta\in\mathcal{A} \\ d(\eta)=k}} (1 - P(\eta, s(\eta)))\pi(\eta) \\ \leq & (1 - P(\sigma_{opt}, s(\sigma_{opt}))) \sum_{\substack{\eta\in\mathcal{A} \\ d(\eta)=k}} \pi(\eta) \end{aligned}$$

Therefore,

$$\sum_{\substack{\eta\in\mathcal{A} \\ d(\eta)=k+1}} \pi(\eta) \leq \frac{1 - P(\sigma_{opt}, s(\sigma_{opt}))}{P(\sigma_{opt}, s(\sigma_{opt}))} \sum_{\substack{\eta\in\mathcal{A} \\ d(\eta)=k}} \pi(\eta).$$

If we denote $C = \frac{1-P(\sigma_{opt}, s(\sigma_{opt}))}{P(\sigma_{opt}, s(\sigma_{opt}))}$ then by induction

$$\sum_{\substack{\eta\in\mathcal{A} \\ d(\eta)=k}} \pi(\eta) \leq C^{k-d(\sigma_{\mathcal{A}})}\pi(\sigma_{\mathcal{A}}),$$

and

$$\begin{aligned} \sum_{\eta\in\mathcal{A}} \pi(\eta) & = \sum_{k=d(\sigma_{\mathcal{A}})}^{d^*} \sum_{\substack{\eta\in\mathcal{A} \\ d(\eta)=k}} \pi(\eta) \\ & \leq \sum_{k=d(\sigma_{\mathcal{A}})}^{d^*} C^{k-d(\sigma_{\mathcal{A}})}\pi(\sigma_{\mathcal{A}}) \\ & = \sum_{k=0}^{d^*-d(\sigma_{\mathcal{A}})} C^{k}\pi(\sigma_{\mathcal{A}}) \\ & = \frac{1 - C^{d^*-d(\sigma_{\mathcal{A}})+1}}{1 - C}\pi(\sigma_{\mathcal{A}}). \end{aligned}$$

Now we can use these results together with the definition of $\Phi(\mathcal{A})$ to obtain the following,

$$\begin{aligned} \Phi(\mathcal{A}) & = \frac{F(\mathcal{A})}{\pi(\mathcal{A})} \\ & \geq \frac{\pi(\sigma_{\mathcal{A}})P(\sigma_{\mathcal{A}}, s(\sigma_{\mathcal{A}}))}{\frac{1-C^{d^*-d(\sigma_{\mathcal{A}})+1}}{1-C}\pi(\sigma_{\mathcal{A}})} \\ & = \frac{(1 - C)P(\sigma_{\mathcal{A}}, s(\sigma_{\mathcal{A}}))}{1 - C^{d^*-d(\sigma_{\mathcal{A}})+1}} \\ & \geq \frac{(1 - C)P(\sigma_{opt}, s(\sigma_{opt}))}{1 - C^{d^*}}. \end{aligned}$$

In the last inequality we used $C > 0$ and $d(\sigma_{\mathcal{A}}) > 1$ so $C^{d^*-d(\sigma_{\mathcal{A}})+1} \leq C^{d^*-1+1} = C^{d^*}$. This completes the proof. □

We will need the following two results.

Proposition 1 *The total variation distance can be bounded as follows,*

$$4 \cdot ||P^n(\Theta, \cdot) - \pi||^2_{var} \leq \frac{1 - \pi(\Theta)}{\pi(\Theta)}\lambda_2^{2n}, \qquad (12)$$

where λ_2 is as before the second largest eigenvalue.

Proof: See Proposition 3 in Diaconis and Stroock (1991). □

Proposition 2 *The second largest eigenvalue can be bounded in terms of the conductance of the underlying graph.*

$$\lambda_2 \leq 1 - \frac{\Phi^2}{2}. \qquad (13)$$

Proof: See Lemma 2.4 in Sinclair (1993). □

Now we obtain an explicit bound on the total variation distance in terms of the stationary distribution and transition probabilities of the NP Markov chain.

Theorem 2 *The total variation distance between the law of the Markov chain and the true stationary distribution is bounded as follows,*

$$\begin{aligned} ||P^k(\Theta, \cdot) - \pi||_{var} & \leq \sqrt{\frac{1 - \pi(\Theta)}{4 \cdot \pi(\Theta)}} \qquad (14) \\ & \cdot \left(1 - \frac{((1-C)P(\sigma_{opt}, s(\sigma_{opt})))^2}{2 \cdot \left(1-C^{d^*}\right)}\right)^k, \end{aligned}$$

where as before

$$C = \frac{1 - P(\sigma_{opt}, s(\sigma_{opt}))}{P(\sigma_{opt}, s(\sigma_{opt}))}.$$

Proof: Follows directly from combining Theorem 1, Proposition 1, and Proposition 2. □

We note that these bounds converge rapidly to zero if the transition probability $P(\sigma_{opt}, s(\sigma_{opt}))$ is large, and the constant in the bound is small if the stationary probability $\pi(\Theta)$ is large.

The bounds in equation (14) show that the law of the NP Markov chain converges geometrically fast to the true stationary distribution, but this is only a qualitative bound since the transition probabilities and stationary probabilities are unknown. To obtain a quantitative bound we need to estimate these probabilities and we will do that in the next section using *auxiliary simulation*.

4 AUXILIARY SIMULATION

Since both $\pi(\Theta)$ and $P(\sigma_{opt}, s(\sigma_{opt}))$ are unknown a priori, the bounds on the conductance (11) and the bounds on the total variation distance (14) are not directly computable. Therefore, we propose using auxiliary simulation to estimate these quantities during simulation-based optimization. To derive these estimates, we need to define one more quantity, namely we let $D_k(\eta)$ denote the number of departures from $\eta \in \Sigma_0$ by the k-th iteration. The stationary quantities of interest can then be estimated as follows.

$$\frac{1-\pi(\Theta)}{\pi(\Theta)} \approx \frac{1-\frac{N_k(\Theta)}{k}}{\frac{N_k(\Theta)}{k}} = \frac{k-N_k(\Theta)}{N_k(\Theta)},$$

$$P(\sigma_{opt}, s(\sigma_{opt})) \approx \frac{D_k(\sigma_k^{(1)})}{N_k(\sigma_k^{(1)})},$$

$$C \approx \frac{1-\frac{D_k(\sigma_k^{(1)})}{N_k(\sigma_k^{(1)})}}{\frac{D_k(\sigma_k^{(1)})}{N_k(\sigma_k^{(1)})}} = \frac{N_k(\sigma_k^{(1)}) - D_k(\sigma_k^{(1)})}{D_k(\sigma_k^{(1)})}.$$

Finally, we get an estimate of the conductance

$$\hat{\Phi}(k) = \frac{2D_k(\sigma_k^{(1)}) - N_k(\sigma_k^{(1)})}{N_k(\sigma_k^{(1)})\left(1-\left(\frac{N_k(\sigma_k^{(1)})-D_k(\sigma_k^{(1)})}{D_k(\sigma_k^{(1)})}\right)^{d^*}\right)} \quad (15)$$

By substituting this into equation (14) we get the following approximate quantitative bound that holds asymptotically,

$$||P^k(\Theta,\cdot) - \pi||_{var} \leq \sqrt{\frac{k-N_k(\Theta)}{4\cdot N_k(\Theta)}}\left(1-\frac{\hat{\Phi}(k)^2}{2}\right)^k. \quad (16)$$

In the next section we use this bound to derive approximate stopping criteria.

5 STOPPING CRITERIA

Although it is reassuring that the NP method converges asymptotically at a geometric rate, it would be of more practical importance to have a sufficient condition such that if this condition is satisfied after k iterations, the NP method can be terminated. In this section we use bounds on the total variation distance between the true and estimated stationary distribution to derive such stopping criterion.

To implement the stopping criterion we need to maintain two singleton regions that have the largest and second largest estimated stationary probability.

Definition 4 *In the k-th iteration, as before let $\sigma_k^{(1)}$ denote the region that is considered to be the best singleton region, and let $\sigma_k^{(2)}$ denote the singleton region that is considered the second best. That is, $\mu_k(\sigma_k^{(2)}) \geq \mu_k(\eta)$ for all $\eta \in \Sigma_0 \setminus \{\sigma_k^{(1)}, \sigma_k^{(2)}\}$, and $\mu_k(\sigma_k^{(1)}) > \mu_k(\sigma_k^{(2)})$.* □

We note that assuming the strict inequality $\mu_k(\sigma_k^{(1)}) > \mu_k(\sigma_k^{(2)})$ imposes no lack of generality.

Using this notation the algorithm may be terminated once the following equation is satisfied.

Theorem 3 *If the following inequality is satisfied for some $k > 0$, then $\sigma_k^{(1)} = \sigma_{opt}$,*

$$||\mu_k - \pi||_{var} \leq \frac{\mu_k\left(\sigma_k^{(1)}\right) - \mu_k\left(\sigma_k^{(2)}\right)}{2}. \quad (17)$$

Proof: See Theorem 13 in Ólafsson (1998).

We note that this theorem does not guarantee the existence of a finite k such that the inequality (17) is satisfied. In fact, for any $k < \infty$ there is a positive probability that it is not satisfied. However, as $\lim_{k\to\infty} \mu_k(\eta) = \pi(\eta)$, $\forall \eta \in \Sigma$, this probability of (17) not being satisfied converges to zero.

We let $\Psi(k) = \frac{1}{2}\left(\mu_k\left(\sigma_k^{(1)}\right) - \mu_k\left(\sigma_k^{(2)}\right)\right)$ denote the right hand side of the inequality (17), and note that this is easily computable after each iteration

$$\Psi(k) = \frac{1}{2k}\left(N_k\left(\sigma_k^{(1)}\right) - N_k\left(\sigma_k^{(2)}\right)\right). \quad (18)$$

We have shown that if $||\mu_k - \pi||_{var}$ is sufficiently small then the algorithm may be terminated. This cannot be bounded directly; however, since μ_k is a sample that is generated according to the law $P^k(\Theta,\cdot)$ of the NP Markov chain, we propose approximating it with $||P^k(\Theta,\cdot) - \pi||_{var}$, that is,

$$||\mu_k - \pi||_{var} \approx ||P^k(\Theta,\cdot) - \pi||_{var}. \quad (19)$$

With this approximation we have an approximation of the left hand side of the inequality (17) that we can bound with inequality (14) and inequality (16) above. Furthermore, the latter of these is easily computable.

We note that the approximation (19) is exact for both $k = 0$ as then both μ_k and the law of the chain place all the probability mass at a single state Θ, and as $k \to \infty$ since both μ_k and $P^k(\Theta,\cdot)$ converge to the stationary distribution. Additional arguments for this approximation may be found in Ólafsson (1998).

With the approximation (19) of $||\mu_k - \pi||_{var}$, and the bound (16) on $||P^k(\Theta,\cdot) - \pi||_{var}$ derived using the auxiliary simulation approach in the last section, we have a computable approximate bound on the left hand side

of inequality (17) in the stopping criterion. Since from equation (18) the right hand side $\Psi(k)$ of the inequality is directly computable, we obtain the following approximate stopping rule that is directly applicable during simulation optimization.

Stopping Criterion. *Terminate the NP algorithm at iteration k if*

$$\sqrt{\frac{k - N_k(\Theta)}{4 \cdot N_k(\Theta)}} \left(1 - \frac{\hat{\Phi}(k)^2}{2}\right)^k \leq \Psi(k). \qquad (20)$$

A potential problem with this stopping rule is that it relies on the estimates of the conductance being sufficiently good. Hence, a warmup period is usually advisable before it is applied.

6 NUMERICAL EXAMPLE

In this section we consider the application of the NP method and the stopping criterion to a scheduling problem in cellular manufacturing systems where the cells are configured in parallel. The objective of this problem is to simultaneously sequence jobs within each cell, and allocate limited flexible resource to the cells, in such a way that the makespan is minimized. The details of the problem formulation, as well as further numerical results, may be found in Ólafsson (1998).

We refer the interested reader to Ólafsson (1998) for the details of the implementation of the NP algorithm, and simply concentrate on the behavior of the stopping criteria. The approximate stopping criteria (20) depends on a sufficiently good estimate of the conductance of the NP Markov chain. A necessary condition for the auxiliary simulation estimate of the conductance to have converged is that it has stabilized or 'settled down'. We therefore start by considering the estimated conductance $\hat{\Phi}(k)$ to verify that it does seem to converge to a limit. This is shown in Figure 1 for a problem where five jobs are to be scheduled in two cells, and there are three flexible resources that can be allocated dynamically to the cells. It is apparent that this auxiliary simulation estimate of the conductance stabilizes quite rapidly. These results indicate that the auxiliary simulation estimate of the conductance is a reasonable estimate.

Since the auxiliary simulation method has been somewhat validated, we consider the performance of the stopping criterion. Figure 2 shows the left hand side (LHS) and right hand side (RHS) of inequality (20) for the same problem as before. For this problem, the NP algorithm found the true optimum, and the LHS of the stopping criterion converges rapidly below the RHS and terminates correctly. This indicates the stopping criterion may be useful.

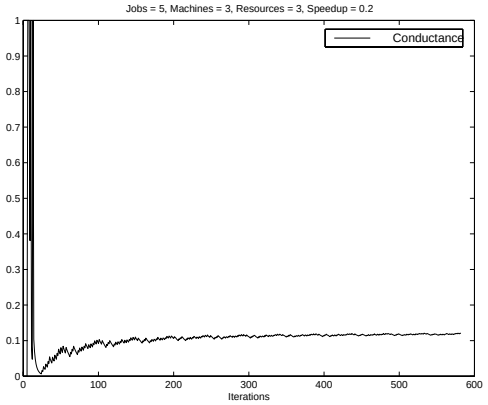

Figure 1: Conductance of the NP Markov Chain

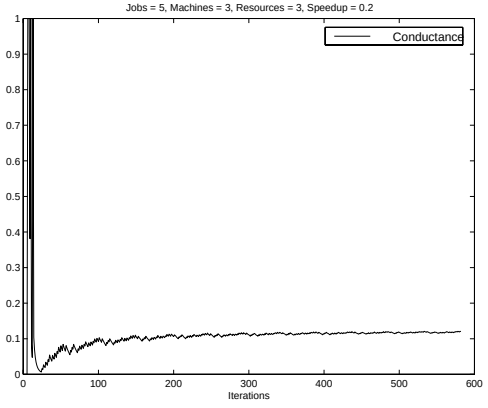

Figure 2: Performance of the Stopping Criterion

7 SUMMARY

We have considered the convergence rate of a new method for simulation-based optimization, the Nested Partitions (NP) method. This method generates a Markov chain and its convergence rate is dependent on how fast this chain converges to its stationary distribution. By bounding the total variation distance between the law of the Markov chain and its stationary distribution, we showed that this occurs at a geometric rate. We also used these bounds to derive an approximate stopping criterion for the method and illustrated our results with a numerical example.

ACKNOWLEDGMENTS

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.

Brooks, S.P., and G.O. Roberts. 1998. Diagnosing convergence of Markov Chain Monte Carlo algorithms. *Statistics and Computing*, to appear.

Cowles, M.K. and B.P. Carlin. 1996. Markov Chain Monte Carlo convergence diagnostics: a comparative review. *Journal of the American Statistical Association* 91: 883-904.

Cowles, M.K., and J.S. Rosenthal. 1996. A simulation approach to convergence rates for Markov Chain Monte Carlo. Technical report, Harvard.

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. *Journal of Discrete Event Dynamic Systems: Theory and Applications* 2:61-88.

Norkin, W.I., G.C. Pflug, and A. Ruszczyński. 1996. A branch and bound method for stochastic global optimization. *Mathematical Programming*, to appear.

Ólafsson, S. 1998. *Convergence Properties of the Nested Partitions Method for Stochastic Optimization*, Ph.D. Thesis, Department of Industrial Engineering, University of Wisconsin - Madison

Roberts, G.O. and J.S. Rosenthal. 1998. Markov Chain Monte Carlo: some practical implications of theoretical results. *Canadian Journal of Statistics*, to appear.

Rosenthal, J.S. 1995. Theoretical convergence rates for Markov chains. *SIAM Review* 37:387-405.

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. 1998. Nested Partitions Method for Stochastic Optimization. Working paper. Department of Industrial Engineering, University of Wisconsin - Madison

Sinclair, A. 1993. *Algorithms for Random Generation & Counting*, Boston, MA: Birkhäuser.

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

AUTHOR BIOGRAPHIES

SIGURDUR ÓLAFSSON is a Ph.D. candidate in the Department of Industrial Engineering at the University of Wisconsin - Madison. He received a B.S. in Mathematics from the University of Iceland in 1995, and an M.S. in Industrial Engineering from the University of Wisconsin - Madison in 1996. His research interests include stochastic optimization and simulation.

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.