Proceedings of the 2000 Winter Simulation Conference
J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, eds.

OPTIMAL SELECTION PROBABILITY IN THE TWO-STAGE NESTED PARTITIONS METHOD FOR SIMULATION-BASED OPTIMIZATION

Sigurdur Ólafsson
Nithin Gopinath

Department of Industrial and Manufacturing Systems Engineering
Iowa State University
2019 Black Engineering
Ames, IA 50011, U.S.A.

ABSTRACT

We investigate a new algorithm for simulation-based optimization where the number of alternatives is finite but very large. This algorithm draws on recent work in adaptive random search and from ranking-and-selection. We show how the ranking-and-selection approach can significantly improve performance of the random search and demonstrate the importance of the probability of correct selection.

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 include 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 (2000) 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 2000). Furthermore, as was first suggested 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 Ólafsson (1999), this method is further improved by drawing on ideas from statistical sampling techniques that have proven useful in simulation in the past, namely ranking-and-selection methods. Thus, the resulting algorithm combines statistical sampling techniques traditionally used for comparing a few alternatives with a global optimization framework aimed at large-scale optimization problems. Here we analyze this new method and in particular focus on the role and importance of the probability of correct selection.

The paper is organized as follows. In Section 2 we discuss the new algorithm and explain its relation to previous work. In Section 3 we analyze the role of the ranking-and-selection probability of correct selection. Section 4 presents some simulation results, and finally, Section 5 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, $f(\theta) = E[L(\theta)]$. 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 are 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.*

Definition 3 *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

$$I(\eta) = \min_{\theta \in \eta} f(\theta), \ \ \forall \eta \in \Sigma, \tag{2}$$

that is, the best solution in a region represents this region. We refer the interested reader to Shi and Ólafsson (1998) 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 hence, is 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 two-stage sampling procedure of Ólafsson (1999) addresses this by incorporating ranking-and-selection into the NP framework, and we will describe this next.

2.2 Two-Stage Sampling

In Ólafsson (1999) it is proposed to use Rinott's two-stage ranking-and-selection procedure for selecting the best subregion (Rinott 1978). 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 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 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:

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

$$X_{ij}(k) = \min_{\theta \in \mathcal{D}_{ij}(k)} L(\theta), \tag{3}$$

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

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 variances

$$\bar{X}_j^{(1)}(k) = \frac{1}{n_0} \sum_{i=1}^{n_0} X_{ij}(k), \tag{4}$$

and

$$S_j^2(k) = \frac{\sum_{i=1}^{n_0} \left[X_{ij}(k) - \bar{X}_j^{(1)}(k) \right]^2}{n_0 - 1}, \tag{5}$$

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

Step 7. Compute the total sample size

$$N_j(k) = \max \left\{ n_0 + 1, \left\lceil \frac{h^2 S_j^2(k)}{\epsilon^2} \right\rceil \right\}, \tag{6}$$

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 overall sample mean be the promising index for each region,

$$\hat{I}\left(\sigma_j(k)\right) = \bar{X}_j(k) = \frac{\sum_{i=1}^{N_j(k)} X_{ij}(k)}{N_j(k)}, \tag{7}$$

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

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

$$\hat{j}_k \in \arg \min_{j=1,...,M+1} \hat{I}(\sigma_j). \tag{8}$$

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

$$\sigma(k+1) = \begin{cases} \sigma_{\hat{i}_k}(k), & \text{if } \hat{i}_k < M+1, \\ s\left(\sigma(k)\right), & \text{otherwise.} \end{cases} \tag{9}$$

Step 11. If $\sigma(k) \in \Sigma_0$ STOP; otherwise go back to Step 1.

Ólafsson (1999) shows that when this algorithm is applied the probability of terminating correctly is bounded by the following inequality:

$$P\left[\hat{\sigma} = \sigma^*_{opt}\right] \geq (P^*)^{d^*} \frac{(1-P^*)^{d^*} - (P^*)^{d^*}}{(1-P^*)^{2d^*} - (P^*)^{2d^*}} \tag{10}$$

where it is assumed that $P^* > \frac{1}{2}$, the indifference zone is selected sufficiently small, and that the sampling is from a normal population. Note that since we can prescribe the value of P^* then we can calculate this probability *a priori* for a given d^*, which simply measures the problem size.

3 PROBABILITY OF CORRECT SELECTION

To implement the NP/Rinott algorithm, several parameters must be selected that can significantly affect the performance of the algorithm. These are the following:

- 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 ($\epsilon > 0$).

First we note that there is clearly a direct 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. Secondly, the indifference zone ϵ depends on how the performance function is scaled and is therefore problem dependent.

The choice of P^*, on the other hand, 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 expend too much computational effort in each iteration. From Ólafsson (1999) we know that $P^* > 0.5$ is needed to guarantee asymptotic convergence, but it should not be selected too large because then too much effort is spent in each iteration. Indeed, whereas the expected number $E[T]$ of iterations needed to find the global optimum can be shown to decrease exponentially fast in P^* (Ólafsson, 1999), the computational effort dictated by equation (6) increases exponentially in P^* due to constant h (see Figure 1). The optimal value thus lies somewhere between 0.5 and 1, but due to the fact that equation (6) also depends on the sample variance, an *a priori* theoretical prediction of the optimal value does not appear to be possible.

One approach to determining P^* would be to use equation (10) as a guide, that is, select the probability by which we want to terminate correctly and use that to determine what value of P^* to use. On the other hand, we would also like to know how the performance of the algorithm itself, rather than just the stopping rule, depends on how P^* is selected. As indicated above, a theoretical analysis is not likely be be successful, so instead an empirical simulation study can be conducted.

4 SIMULATION RESULTS

In this section we will empirically determine the best value of P^* for a single machine tardiness problem. The experiment was then run for two job-sequencing problems, one with 5 jobs and one with 7 jobs . In each case the program was run for three sets of simulation estimates, short, medium and long. The three different sets of simulation estimates were chosen in order to study the performance at each level of simulation estimates. For each set of simulation estimates, twenty different values of total tardiness was determined and tabulated. To get a good idea of the behavior of the algorithm as a function of P^* on the interval from 0.5 to 1, the experiments were performed for $P^* \in \{0.55, 0.65, 0.75, 0.85, 0.95\}$. Twenty replications were used for each experiment.

The average tardiness results for the 5-job problem are shown in Figure 2 as a function of P^* and for short, medium, and long runs of the algorithm. As should be expected the algorithm continues to make improvements, so as the number of functional evaluation increases we keep getting a better solution. For a short set of functional evaluation the best value is 13.8. For a medium set of functional evaluation the tardiness value obtained is 13.3. In the case of a longer set of functional evaluation the tardiness value obtained is 11.3. Note, however, that in all cases the probability P* = 0.65 gives the best value of tardiness. Furthermore, we note that there is a considerable difference in tardiness calculated when the values of P* = 0.65 and P* = 0.95 are used. The first half of Table 1 highlights these improvements.

From the fourth column of Table 1 we observed that the improvement in the tardiness value is close to 15.9% for a short period of computation time and up to 27% in the case of a longer period of computation. We conclude that by choosing the best P^* value we can improve our solution considerably for the same amount of computation time.

Another important observation is that at higher probabilities there is not a considerable difference in the tardiness values calculated between the different run lengths of the algorithm. On the other hand, for short, medium and long sets of simulation estimates the tardiness values calculated for lower probabilities are considerably different. For example, there is a difference of 3.2 for the probability $P^* = 0.55$ between short and long runs of the algorithm. This should be contrasted with the case of $P^* = 0.95$ where the percentage improvement obtained by increasing the run of the algorithm is only 5.5% (see the second column of Table 2).

An intuitive explanation for this behavior may be that for lower values of P^* the algorithm moves around more freely and can thus explore a greater part of the feasible region. Again this indicates that significant benefits in terms of computation time can be derived if the P^* value is chosen well.

The same simulation experiments were repeated for a larger 7-job sequencing problem. The results can be found in Figure 3, and as before it is clear that $P^* = 0.65$ is the best value. Again, we observe that there is a considerable difference in tardiness calculated when the values of $P^* = 0.65$ and $P^* = 0.95$ are used, and this is highlighted in the last column of Table 1, which shows

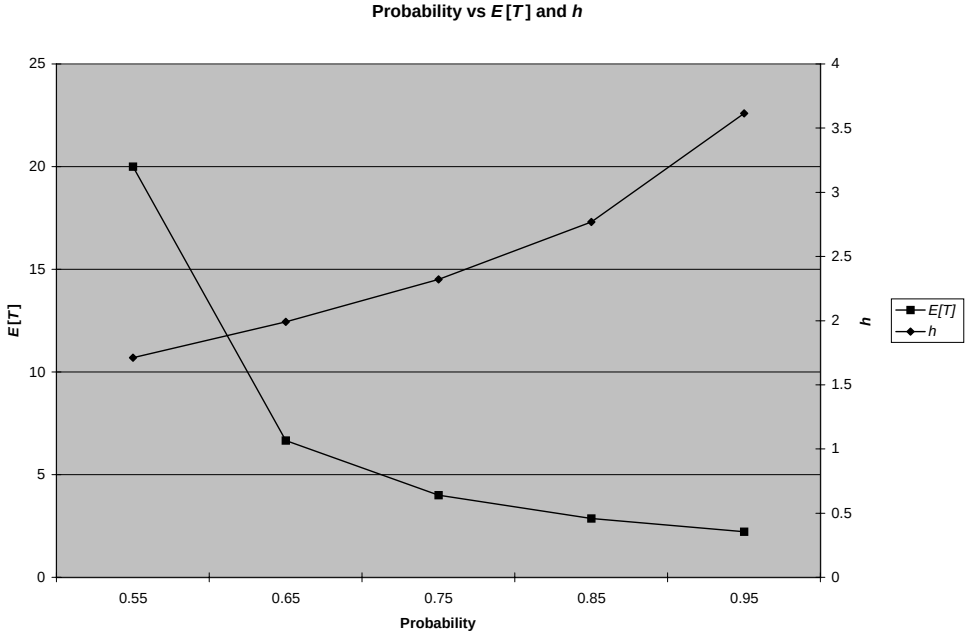

Figure 1: Computation Time Within Each Iteration Versus the Expected Number of Iterations

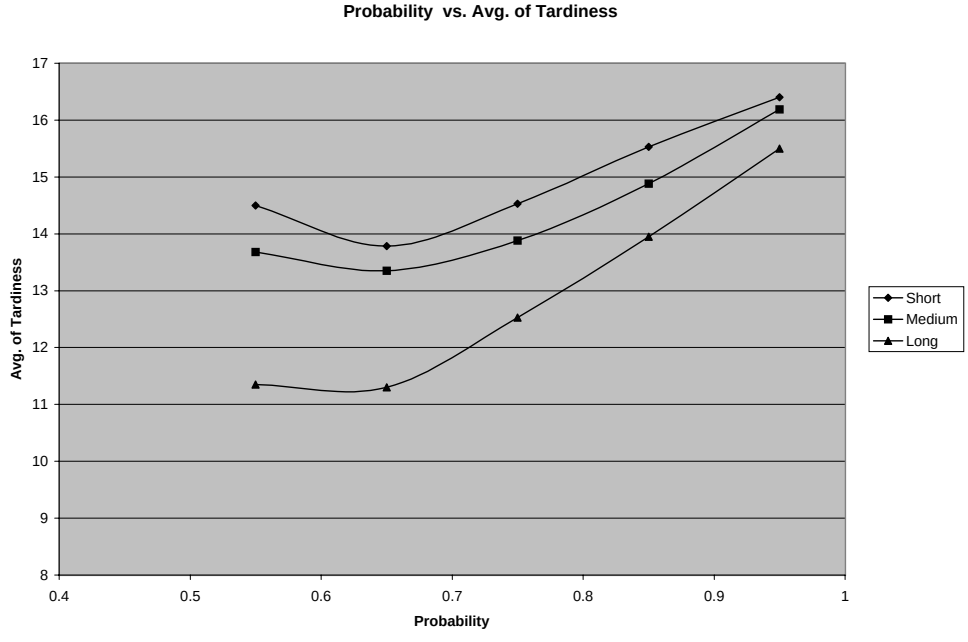

Figure 2: Simulation Results for the 5-Job Problem

Table 1: Improvement in Average Tardiness for Different P^* Values

	5-job Problem			7-job Problem		
Computation Time	$P^* = 0.95$	$P^* = 0.65$	% Improved	$P^* = 0.95$	$P^* = 0.65$	% Improved
Short	16.4	13.8	15.9	94.2	87.9	6.7
Medium	16.2	13.4	17.3	90.6	78.5	13.4
Long	15.5	11.3	27.1	89.4	76.4	14.5

Table 2: Improvement in Tardiness from Short to Long Computation Time

P^*	Percentage Improvement	
	5-job problem	7-job problem
0.55	21.4	16.9
0.65	18.0	13.1
0.75	13.8	11.2
0.85	10.3	7.8
0.95	5.5	5.1

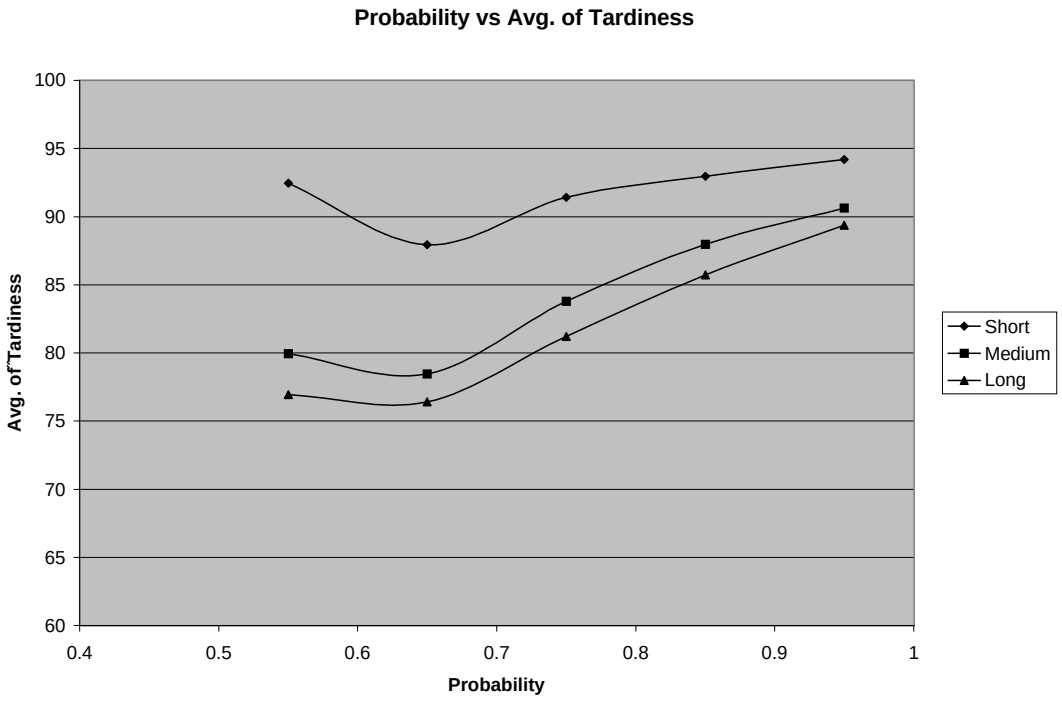

Figure 3: Simulation Results for 7-Job Problem

the percentage improvement in tardiness. We see from this table that the improvement in the tardiness value is close to 6.7% for a short run of the algorithm, and up to 14.5% in the case of longer computation. By choosing the best P^* value we can improve our solution considerably for the same amount of time. Similar to the previous problem, at higher probabilities there is not a considerable difference in tardiness values between short and long runs of the algorithm. There is a difference of 15.5 for the probability $P^* = 0.55$ between short and long runs of the algorithm, but in the case of $P^* = 0.95$ the improvement is only 4.8 time units. A comparison of the percentage improvement further illustrates the advantage of choosing lower probabilities. From the last column of Table 2 we observe that the percentage improvement is 16.9% for $P^* = 0.55$ and in the case of a higher probability, $P^* = 0.95$, the improvement in tardiness is only 5.1%. It is again evident that to get a good improvement in a short period of time we need to choose lower probabilities. We also observe that as in the case of the five job sequencing problem choosing the best value of P^* makes a better use of the computer time. Finally, we observe that our simulation experiments indicate that the best P^* value is insensitive to problem size and for the run length of the algorithm.

5 SUMMARY

This paper has considered how to determine the probability of correct selection in the two-stage NP method proposed by Ólafsson (1999). Our results show that this determination is critical as the performance of the algorithm varies greatly depending on how it is performed. Furthermore, a clear pattern emerged in terms of the performance of the algorithm as a function of the selection probability and, for the test problem, the optimal probability was insensitive to other factors. This indicates that it may be possible, for a given problem, to determine an optimal or a close to optimal selection probability empirically for a relatively small problem instance and that this setting of the algorithm will perform well when other instances of the problem

are solved. Further investigation is required, however, to confirm this observation.

REFERENCES

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. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

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*, ed. D. J. Medeiros, E. F. Watson, J. S. Carson, and M. S. Manivannan, 735–741. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Chick, S.E. 1997. Selecting the best system: A decision theoretic approach. In *Proceedings of the 1997 Winter Simulation Conference*, ed. S. Andradóttir, K. J. Healy, D. H. Withers, and B. L. Nelson, 326–333. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

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*, ed. D. J. Medeiros, E. F. Watson, J. S. Carson, and M. S. Manivannan, 159–166. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

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. 1999. Iterative ranking-and-selection for large-scale optimization. In *Proceedings of the 1999 Winter Simulation Conference*, ed. P. A. Farrington, H. B. Nembhard, D. T. Sturrock, and G. W. Evans, 479–485. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Ólafsson, S. and L. Shi. 2000. A method for scheduling in parallel manufacturing systems with flexible resources. *IIE Transactions on Scheduling and Logistics* 32:135–146.

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*, ed. S. Andradóttir, K. J. Healy, D. H. Withers, and B. L. Nelson, 358-365. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Shi, L. and S. Ólafsson. 1998. Nested partitions method for stochastic optimization. Working Paper 98-116, Department of Industrial and Manufacturing Systems Engineering, Iowa State University.

Shi, L. and S. Ólafsson. 2000. Nested partitions method for global optimization *Operations Research*, 48 (3): 390–407.

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.

NITHIN GOPINATH is a graduate student in the Department of Industrial and Manufacturing Systems Engineering at Iowa State University. He received his B.S. in Mechanical Engineering from the University of Madras in 1998. His research focuses on methodology and applications in simulation-based optimization.