Proceedings of the 2005 Winter Simulation Conference
M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, eds.

APPLICATION OF MULTI-OBJECTIVE SIMULATION-OPTIMIZATION TECHNIQUES TO INVENTORY MANAGEMENT PROBLEMS

Loo Hay Lee
Suyan Teng
Ek Peng Chew

Dept. of Industrial and Systems Engineering
National University of Singapore
10 Kent Ridge Crescent, 119260, SINGAPORE

Kong Wei Lye
Peter Lendermann

D-SIMLAB Programme
Singapore Institute of Manufacturing Technology
71 Nanyang Drive, 638075, SINGAPORE .

I. A. Karimi

Dept. of Chemical & Biomolecular Engineering
National University of Singapore
10 Kent Ridge Crescent, 119260, SINGAPORE

Yankai Chen
Choon Hwee Koh

Dept. of Industrial and Systems Engineering
National University of Singapore
10 Kent Ridge Crescent, 119260, SINGAPORE

ABSTRACT

In this paper, we present how a solution framework developed for (a special case of) the multi-objective simulation-optimization problems can be applied to evaluate and optimally select the non-dominated set of inventory policies for two case study problems. Based on the concept of Pareto optimality, the solution framework mainly includes how to evaluate the quality of the selected Pareto set by two types of errors, and how to allocate the simulation replications according to some asymptotic allocation rules. Given a fixed set of inventory policies for both case study problems, the proposed solution method is applied to allocate the simulation replications. Results show that the solution framework is efficient and robust in terms of the total number of simulation replications needed to find the non-dominated Pareto set of inventory policies.

1 INTRODUCTION

Today, companies operate in a fast changing business environment. To stay competitive, adaptations and enhancements of manufacturing and service operations and the associated business processes need to take place constantly. One important characteristic of today's high-tech companies is that, they operate in global networks that often involve contract manufacturers and third party logistics providers, driving the underlying systems towards mega-networks. Complex business processes across the mega-networks have become more critical and raised new operational challenges such as how to optimize the collaboration between a line maintenance service company for commercial aircraft and a third party logistics provider to minimize the inventory cost for critical spare components. Most of the planning and scheduling problems associated with these networks are very difficult to solve due to the high variability, underlying non-linear dynamics, large problem size, and possibly multiple objectives.

This paper describes the parts of the results obtained in one out of ten pilot programmes under the Integrated Manufacturing and Service Systems (IMSS) initiative pursued by the Agency for Science, Technology and Research (A*STAR) in Singapore (Lendermann et al. 2005). The objective of this particular programme is to investigate how design, analysis, enhancement and implementation of critical business processes in a manufacturing and service network can be realised using one single simulation/application framework. The overall architecture of the framework outlines how commercial simulation packages and web-service-based business process application components would have to be connected through a commercial application framework to achieve maximum leverage and re-usability of the applications involved. In the pilot phase of this programme, research issues were also addressed with regard to mechanisms for interoperation between commercial simulation packages, symbiotic interaction between simulation-based decision support components and physical systems, and speed-up of simulation analysis by making use of a grid infrastructure.

In this paper we describe how multi-objective simulation-optimization techniques can also help speed-up simulation analysis exercises. This is studied and illustrated through two complex inventory management problems with difficult problem features mentioned above.

The first one is the differentiated service inventory problem, which is to determine an inventory policy to differentiate customer groups and offer different services to different customers. Unlike the traditional inventory system, where all customer demands are treated equally and served on a first-come-first-served basis, here customers are classified into different groups according to their importance to the decision makers. This is because, for some customers, the stock-out cost is so high that they are willing to pay at a higher price for timely replenishment. Obviously, these customers are more valuable and therefore they should have higher priority and be provided with higher service level. In this setting, the problem is to evaluate the cost and service level for customers from different classes and obtain an optimal inventory policy with the lowest cost and highest service level.

The second case study is an aircraft spare parts inventory problem. When a repairable item on an aircraft becomes defective, it is removed and replaced by another item from the spare stock. The defective part then goes into some repair cycle. If the airport does not have the spare part in stock, the aircraft will be grounded and delayed until an incoming flight brings a replacement part from the Central Repair Depot or from a neighboring airport. To decrease departure delays due to unanticipated failures, airlines need to keep inventory of spare parts at the associated airports. The problem is to determine the number of spare parts to be stored at destination airports; and the replacement policy (where to get a replacement part) upon the occurrence of a part failure, so that the average cost involved is minimized and the fill-rate (percentage of failures serviced) is maximized.

The difficult features related to the above problems often make them mathematically intractable. For optimization-based approaches to be applicable and effective, they often require too many assumptions and simplifications made on the problems. On the other hand, simulation-based techniques are not constrained by analytical assumptions and simplifications and can give reasonable solutions within acceptable time. However, simulation may be limited by human imagination of possible alternatives, and also it can be both expensive and time consuming. Therefore, it is important to improve the performance of the simulation-based techniques through optimization approaches. This area of research, known as Simulation Optimization, has become a hot and important topic recently.

A very general formulation of the above simulation optimization problem is to minimize the expected value of the objective function with respect to its constraint set as:

$$\min_{\theta\in\Theta} J(\theta).$$

where $J(\theta)=E[L(\theta,\varepsilon)]$ is the performance measure of the problem, $L(\theta,\varepsilon)$ is the sample performance, ε represents the stochastic effects in the system, θ is a p-vector of discrete controllable factors and Θ is the discrete constraint set on θ. If $J(\theta)$ is a scalar function, the problem is single objective optimization; whereas if it is a vector, the problem becomes multi-objective.

The above problem can be very challenging, both analytically and computationally, due to three kinds of difficulties inherent to the problem considered: lack of structure of the solution space (Θ), huge size of the solution space (Θ) and large uncertainties (ε) in the performance measures. In this study, we focus on a simplified version of the problem where Θ is defined as a finite set with a number of alternatives. The simplified problem is usually known as Ranking and Selection (R&S) problem (Swisher, Jacobson, and Yücesan 2003) in the literature. Also, we will consider the case where the objective function $J(\theta)$ is a vector of performance measures. So, the problem considered in this study is a multi-objective ranking and selection (MORS) problem (Lee et al. 2004).

For the single objective R&S problem, several different approaches have been proposed: indifference-zone ranking and selection (Rinott 1978; Nelson et al. 2001; Alrefaei and Alawneh 2004), optimal computing budget allocation (Chen et al. 1997; Chen et al. 2000), decision theoretic methods (Chick 1997; Chick and Inoue 2001), subset selection (Gupta 1956; Nelson et al. 2001), and multiple comparisons procedures (Fu 1994; Hsu 1996). Among them, the optimal computing budget allocation (OCBA) method is relatively recently developed and more efficient in terms of total number of replications needed to find the best alternative. Moreover, OCBA is capable of solving problems with relatively large number of alternatives. Specifically, OCBA follows a Bayesian methodology, making use of information on both sample means and sample variances. The rationale here is to only simulate likely competitors for the “best”. This is done by developing lower bounds for the probability of correct selection P(CS), and solving the problem as an optimization problem, in which P(CS) is maximized with a given total computing budget available. In Chen et al. (1997), additional replications are allocated based on gradient information from estimated P(CS). In Chen et al. (2000) and Chen et al. (2003), simpler and more efficient asymptotic allocation rules are derived when an infinite computing budget is assumed to be available.

For the multi-objective R&S (MORS) problem, one common solution framework is to weight several parameters of interest to form a single measure of effectiveness by applying multiple attribute utility (MAU) theory (Butler, Morrice, and Mullarkey 2001; Morrice, Butler, and Mullarkey 1998; Swisher and Jacobson 2002). The problem reduces to a single-objective model, and existing methods can be applied. One disadvantage of this method is that, the decision maker needs to specify the relative importance of the performance measurers before optimization is done. As a result, the best solution selected would be strongly dependent on these preferences. Once the preferences of the decision makers are changed, the solution may become inferior. Therefore, ideally, all non-dominated solutions should be provided to the decision makers, so that they can choose their favorite solution under specific circumstances. The complete set of non-dominated solutions is referred to as the Pareto set of solutions. They represent the "best" designs and are characterized by the definition that no other solution exists that is superior in all the objectives. In Lee et al. (2004), the authors incorporated the concept of Pareto optimality into the R&S scheme, and developed a simple sequential procedure to find the non-dominated Pareto set of designs to the MORS problem. However, in that study, they assume that the number of non-dominated designs in the Pareto set is known in advance. In Lee et al. (2005), to relax the above-mentioned assumption and consider the problem from a more realistic aspect, they presented a different framework for the MORS problem. Computational results show that, the solution framework is robust and efficient in terms of the number of replications needed to find the Pareto set.

In this paper, we study how the solution framework in Lee et al. (2005) can be applied to solve real world problems, and examine its robustness and efficiency through the case study problems. The paper is organized as follows. Section 2 briefly introduces the sequential solution framework for the MORS problem. In Section 3, we present two case study problems and illustrate how efficient and robust the solution framework can be when allocating the simulation replications to the designs. Finally some conclusions and future research directions are summarized in Section 4.

2 THE SOLUTION FRAMEWORK FOR THE MORS PROBLEM

We present a brief description of a solution framework for the MORS problem in this section. The solution method is called the Multi-objective Optimal Computing Budget Allocation (MOCBA) algorithm. For more details about how MOCBA works, please refer to Lee et al. (2005).

Without loss of generality, we assume that minimization of the objectives is our goal throughout this paper. In case some of the objectives are maximization, we would minimize the corresponding negatives of those objectives. Also, we assume that the random variables under study follow continuous distributions.

2.1 A Performance Index to Measure the Non-dominated Designs

When considering Pareto optimality, we are trying to find a complete set of those non-dominated designs. Here, we first need to find a way to measure how non-dominated a design is.

Suppose we have two designs θ_i and θ_j, each of which is evaluated in terms of H performance measures:

$$\begin{aligned}\theta_i &: J_1(\theta_i), J_2(\theta_i), ..., J_H(\theta_i)\\ \theta_j &: J_1(\theta_j), J_2(\theta_j), ..., J_H(\theta_j)\end{aligned}.$$

Here performance measures $J_k(\theta_i)$ and $J_k(\theta_j)$ ($k = 1, 2, ..., H$) are obtained through simulation, they are random variables subject to noise, so we consider the probability that design θ_j dominates design θ_i, as expressed in the following condition with at least one inequality being strict:

$$P(\theta_j \prec \theta_i) = P(J_k(\theta_j) \le J_k(\theta_i) \text{ for } k = 1, 2, ..., H).$$

Under the condition that the performance measures are independent from one another and they follow continuous distributions, we have

$$P(\theta_j \prec \theta_i) = \prod_{k=1}^{H} P(J_k(\theta_j) \le J_k(\theta_i)). \tag{1}$$

Now suppose we have a total of n designs, i.e., $|\Theta| = n$, we introduce the following performance index to measure how non-dominated a design i is:

$$\psi_i = \prod_{j=1, j\ne i}^{n} [1 - P(\theta_j \prec \theta_i)]. \tag{2}$$

This performance index measures the probability that design i is non-dominated by all the other designs. If ψ_i is very close to 1, the probability that design i is non-dominated is very high. Therefore, at the end of simulation, all designs in the Pareto set should have ψ_i close to 1, and those designs outside of the Pareto set should have ψ_i close to 0, because they are dominated.

2.2 Two Types of Errors of the Selected Pareto Set

When the true Pareto set is found, all designs in it should be non-dominated with probability 1, and all designs outside it should be dominated with probability 1. During the allocation process, the Pareto set is constructed based on observed performance. Here we call it the selected Pareto set (S_p). The quality of the selected Pareto set can be evaluated by two types of errors: Type I error (e_1) and Type II error (e_2).

Type I error (e_1) is the probability that at least one design in the selected non-Pareto set ($\bar{S}_p$) is non-dominated; while Type II error (e_2) is the probability that at least one design in the selected Pareto set is dominated by other designs. When both types of errors approach to 0, the true Pareto set is found. The two types of errors can be bounded by the approximated errors ae_1 and ae_2 respectively as follows.

$$e_1 \le ae_1 = \sum_{i \in \bar{S}_p} \psi_i. \tag{3}$$

$$e_2 \le ae_2 = \sum_{i \in S_p} (1-\psi_i). \tag{4}$$

Now based on the non-dominated performance index ψ_i and the approximated type I and type II errors (ae_1 and ae_2), after ranking all the designs in descending order of ψ_i, we can construct the selected Pareto set according to one of the three criteria below.

C1: Assign a maximum number of k designs with the highest ψ_i into the selected Pareto set S_p, so that $ae_2 = \sum_{i=1}^{k}(1-\psi_i) \le \varepsilon$, where ε is a predefined error limit.

C2: Construct the selected Pareto set by optimizing k so that both ae_1 and ae_2 are minimized.

C3: Select k designs with $\psi_k \ge \psi^*$ into the selected Pareto set, where ψ^* is a predetermined non-dominated probability.

2.3 A Sequential Solution Procedure and the Asymptotic Allocation rules

In solving the MORS problem, we are trying to get the true Pareto set with high probability by minimizing both Type I and Type II errors. This can be done by keeping on allocating more replications to certain designs until both approximated Type I and Type II errors (ae_1 and ae_2) are within an error limit ε^*.

At the beginning of the allocation process, we can perform N_0 replications for each design. Based on the simulation output, we then estimate the performance index ψ_i and construct the selected Pareto set S_p. Then the MORS problem is to determine the optimal allocation of the replications to the designs so that both ae_1 and ae_2 are within error limit ε^*, and the total number of simulation replications is minimized.

In this study, we propose a sequential approach for solving the above problem. The procedure iteratively allocates the simulation replications according to some asymptotic allocation rules (Lee et al. 2005) until both ae_1 and ae_2 are within error limit ε^*.

The sequential procedure, known as the MOCBA algorithm, is outlined as follows.

MOCBA algorithm

Step 0: Perform N_0 replications for each design. Set iteration index v: = 0. $N_1^v = N_2^v = ... = N_n^v = N_0$.

Step 1: Construct the selected Pareto set S_p according to criterion C2 of Section 2.2. Calculate ae_1 and ae_2 according to equations (3) and (4) in Section 2.2. If ($ae_1 < ae_2$), construct S_p according to the criterion C1 of Section 2.2 with $\varepsilon = ae_1$.

Step 2: If (($ae_1 < \varepsilon^*$) and ($ae_2 < \varepsilon^*$)), go to Step 5.

Step 3: Increase the simulation replications by a certain amount Δ, and calculate the new allocation $N_1^{v+1}, N_2^{v+1}, ..., N_n^{v+1}$ according to the asymptotic allocation rules (Lee et al. 2005).

Step 4: Perform additional min(δ, $\max(0, N_i^{v+1} - N_i^v)$) replications for design i (i =1,…,n). Set $v = v+1$ and go to Step 1.

Step 5: Output designs in the selected Pareto set (S_p).

3 TWO INVENTORY MANAGEMENT CASE STUDY PROBLEMS

In this section, we present two case study problems both of which are multi-objective inventory management problems.

To avoid too many assumptions and unnecessary simplifications on the problems, we employ simulation to evaluate the performances of the alternatives rather than using the analytical methods. Therefore, with a set of given design alternatives (inventory policies), the MOCBA algorithm is employed to allocate simulation replications efficiently to obtain the Pareto set of solutions. In both cases, results from MOCBA are compared with those from the Uniform Computing Budget Allocation (UCBA), which is to uniformly allocate the same number of replications to each design. In both cases, the parameter setting of the MOCBA algorithm is as follows: the initial runs $N_0 = 10$, the incremental runs $\Delta = 10$, and additional maximum number of runs $\delta = 5$.

3.1 The Differentiated Service Inventory Problem

The inventory systems with several demand classes can be found in many cases. Several decision policies aiming at differentiating customer groups and offering different services can be found in the literature.

One such decision policy is called Critical Level Policy with $n-1$ critical levels for n customer groups. It is assumed that demand from the customer with the highest priority will always be satisfied. When demand from class m arrives, it will be satisfied if the inventory-on-hand is higher than the jth critical level; otherwise, it would be rejected. Veinott (1965) was the first to consider such inventory policy under periodic review. Dekker, Hill and Klejin (1998) applied it in a lost sales continuous (s,Q) model. Instead of fixing the $n-1$ critical levels, another policy called Dynamic Critical Level Policy also considers the lead times preceding the order arrivals. In this policy, the decision is made based on both the customer class and the lead times. The critical level is thus non-increasing as time approaches to an order arrival. This is an approach similar to airline revenue management, where the number of seats to be sold is controlled based on both the remaining time before departure and the type of customers.

In this case study, we examine the Dynamic Critical Level Policy within a (s,Q) inventory model. We simulate the arrival of customers from 2 different demand classes and evaluate the cost and service level of this policy. We assume that customer arrivals follow a Poisson distribution, and the annual demand of both customer classes is 300. The Critical Level Policy is defined in terms of a two dimensional matrix with 2 columns corresponding to 2 customer classes and 600 rows corresponding to 600 different possible remaining lead times. The lead time is set as 0.03 year. For customer class 1, all critical levels are set as 0, because customers from this class have top priority. Figure 1 illustrates how critical level changes with remaining lead time for customer class 2.

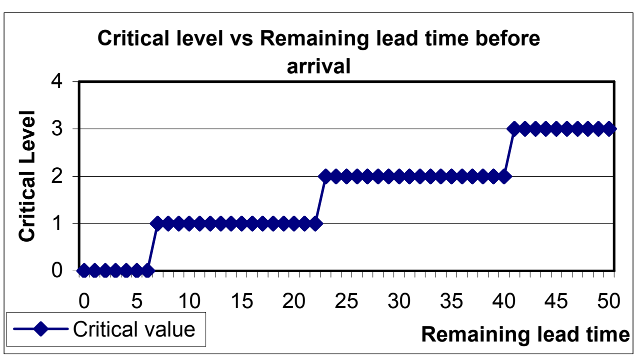

Figure 1 Change of Critical Level with Remaining Time

Here, as time passes, the lead time decreases and the critical level decreases accordingly. The objective is to obtain the optimal (s,Q) with the lowest cost and highest service level for this Dynamic Critical Level Policy. Here, the cost consists of setup cost, inventory holding cost and backorder cost, where the latter two costs are computed based on per customer per unit time. The set up cost is assumed to be 1, and the inventory holding cost is set at 2.5. The backorder penalty cost used for customer class 1 is 100 and customer class 2 is 10. The service level is defined in terms of backorder fill rate of the two classes of customers.

We first generate 25 promising alternatives based on different reorder point s and fixed reorder quantity Q. It is assumed that any improvement of backorder fill rate below 0.01 is insignificant. Therefore, the Pareto sets are the inventory policies with lowest cost and relatively low backorder rate. We apply both MOCBA and UCBA to allocate the simulation replications. The number of replications needed for both MOCBA and UCBA is illustrated in Figure 2.

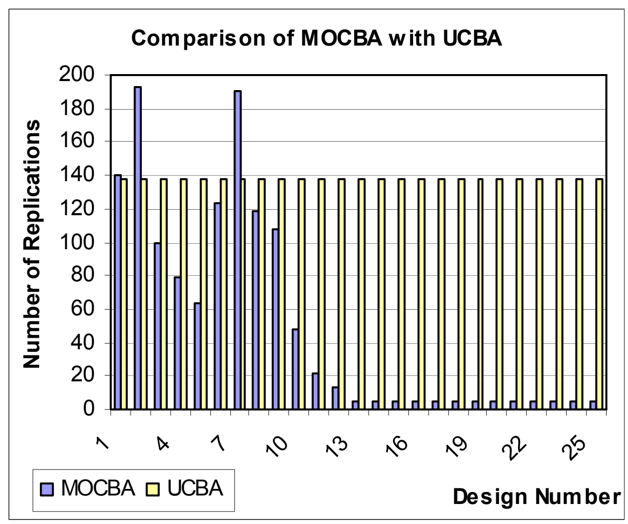

Figure 2 Compare MOCBA and UCBA (Case Study 1)

In this case, with similar error limit ($\varepsilon^* = 0.05$) and the same set of designs (designs 1 to 7) in the Pareto set, MOCBA takes a total of 1262 replications, and UCBA requires 3450 replications. The speedup of MOCBA over UCBA is about 2.73 times. We can observe from Figure 2 that, for MOCBA, the following designs are allocated more replications: (a) those designs that should be in the Pareto set, and (b) those designs whose performances are very close to designs in (a). This indicates that the MOCBA algorithm is effective

3.2 Planning Aircraft Spare Parts Inventory among Airports

In the aircraft spare parts inventory problem, the main decision is to determine how many spare parts to be stored at each destination airport; and the replacement policy (where to get a replacement part, from airport's own inventory, neighboring airport or Central Repair Depot) upon occurrence of a failure.

This problem is essentially a two-echelon setting for supplying repairable spares. Work on the multi-echelon technique for recoverable item control (METRIC) by Sherbrooke (1968) forms the main motivation for this line of research. One difference with our problem is that in the METRIC model a defective part can be replaced only by another item available at the station or at the central depot. Batchoun, Ferland and Cléroux (2003) also considered a similar aircraft spare parts inventory problem. They applied Genetic Algorithms in an adaptive search procedure to allocate the initial quantity of spare parts to the airports. In both cases above, the problems are assumed to have single objective which can be estimated by the average cost through analytical methods.

In this case study, we consider an airport network which consists of 20 airports and 1 Central Repair Depot. We assume that only one spare part type is considered, and there are 60 repairable spares in total. At every maintenance check, a failure occurs at the rate of 0.5%. A defective part goes into some repair cycle with repair time being assumed to be uniformly distributed between 12-24 days. We also assume that the repair capacity at the Central Repair Depot is infinite. This assumption can be justified by the very low failure rate (0.5%) of the aircraft repairable parts, as in this case, the occurrence of the failure event would be very rare. When a part failure occurs, re-supply of the spare part comes from either the inventory at the airport, the Central Repair Depot or from the neighboring airports. 2 flights are required to replace a part from either the Central Repair Depot or from the airport's own inventory. This cost is associated with transport of the defective part to and from the Central Repair Depot. 3 flights are required to replace a part from the neighboring airports, two for transporting the defective part to and from the Central Repair Depot, one for obtaining a new spare part. The problem is to determine the optimal inventory policy (allocation of the spare parts among the airports and selection of the replacement policy) so that the expected cost of the inventory policy is minimized and the average service level (fill rate) of the entire network is maximized. Here the cost is defined in terms of inventory cost, and distance traveled to repair a defective part and obtain a new spare part. And "fill rates" is defined in terms of percentage of failures/defectives serviced.

We also generate 25 promising alternatives based on different number of spare parts stocked at each airport. Figure 3 illustrates the number of replications needed for both MOCBA and UCBA.

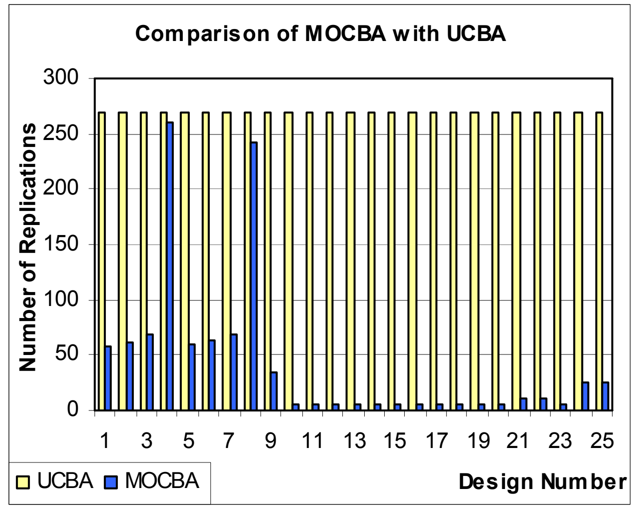

Figure 3 Compare MOCBA and UCBA (Case Study 2)

In this case, with similar error limit ($\varepsilon^* = 0.1$) and the same set of designs in the Pareto set (designs 1 to 8), MOCBA takes a total of 1048 replications, and UCBA requires 6750 replications. The speedup of MOCBA over UCBA is about 6.44 times in this instance. Similar to case study 1, we can observe from Figure 3 that most of the replications are allocated to the competitive designs.

4 CONCLUSIONS

In this paper, two case study problems in the area of inventory management are presented to illustrate the applicability of a multi-objective simulation-optimization solution framework (MOCBA) to address real world complex and difficult problems. The MOCBA is developed for the multi-objective ranking and selection problem (a special case of multi-objective simulation-optimization problem) to find all the non-dominated designs in the Pareto set. Results show that, in comparison with the UCBA (uniform computing budget allocation), MOCBA can more efficiently allocate the simulation replications to the designs. In the first case study, the speedup of MOCBA over UCBA is about 2.73 times, with similar type I and type II errors resulted upon termination of the algorithm; while in the second one, where the uncertainty involved in the problem is much higher, the speedup is about 6.44 times. In this paper, we assume that the solution space of the case study problems is finite and consists of a given set of promising alternatives. In future research, it is important to relax this assumption and consider how to incorporate a search procedure into the solution framework so that promising alternatives can be found through efficient exploration of the solution space. At the same time, we will be exploring how MOCBA techniques can be further leveraged by making use of a Grid infrastructure for simulation execution. This has also been initiated in the above-mentioned pilot programme. Details can be found in Julka et al. (2005).

ACKNOWLEDGMENTS

This pilot programme was partly funded by the Agency for Science, Technology and Research (A*STAR) in Singapore under the "Industrial IT – Integrated Manufacturing and Service Systems" initiative.

REFERENCES

Alrefaei, M. H. and A. J. Alawneh. 2004. Selecting the best stochastic system for large scale problems in DEDS. *Mathematics and Computers in Simulation* 64: 237–245.

Batchoun, P., J. A. Ferland, and R. Cléroux. 2003. Allotment of aircraft spare parts using Genetic Algorithms. *Pesquisa Operacional* 23(1):141–159.

Butler, J., D. J. Morrice, and P. W. Mullarkey. 2001. A multiple attribute utility theory approach to ranking and selection. *Management Science* 47(6): 800–816.

Chen, C. H., K. Donohue, E. Yücesan, and J. Lin. 2003. Optimal computing budget allocation for Monte Carlo simulation with application to product design. *Simulation Modelling Practice and Theory* 11: 57–74.

Chen, C. H., J. W. Lin, E. Yücesan, and S. E. Chick. 2000. Simulation budget allocation for further enhancing the efficiency of ordinal optimization. *Discrete Event Dynamic Systems: Theory and Applications* 10: 251–270.

Chen H. C., L. Dai, C. H. Chen, and E. Yücesan. 1997. New development of optimal computing budget allocation for discrete event simulation. In *Proceedings of the 1997 Winter Simulation Conference* (ed. S. Andradóttir, K. J. Healy, D. H. Withers, and B. L. Nelson), 334–341. 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.

Chick, S. E. and K. Inoue. 2001. New two-stage and sequential procedures for selecting the best simulated system. *Operations Research* 49 (5): 732–743.

Dekker, R., R. M. Hill, and M. J. Klejin. 1998. On the $(S-1,S)$ lost sales inventory model with priority demand classes. In *Inventory Modelling in Production and Supply Chains* (ed. Papachristos S. and Ganas I.), 99–103. International Society for Inventory Research (Budapest) and University of Ioannina

Fu, M. 1994. Optimization via simulation: A review. *Annals of Operations Research* 53: 199–247.

Gupta S. S. 1956. On a decision rule for a problem in ranking means. Mimeograph Series No. 150, Institute of Statistics. University of North Carolina, Chapel Hill, Chapel Hill, N.C.

Hsu, J. C. 1996. Multiple Comparisons: Theory and Methods. Chapman & Hall, London, England.

Julka, N., P. Lendermann, W.T. Cai, N.Q. Khanh, and T. Hung. 2005. Grid Computing for Virtual Experimentation and Analysis, *Proceedings of the 2nd Workshop on Grid Computing & Applications*, May 2005, Singapore.

Lee, L. H., E. P. Chew, S. Y. Teng, and D. Goldsman. 2004. Optimal computing budget allocation for multi-objective simulation models. In *Proceedings of the 2004 Winter Simulation Conference* (ed. R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters), 586–594. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Lee, L.H., E.P. Chew, S.Y. Teng, and D. Goldsman. 2005. Finding the non-dominated Pareto set for multi-objective simulation models. Submitted to IIE Transactions.

Lendermann, P., M.Y.H. Low, B.P. Gan, N. Julka, L.P. Chan, S.J. Turner, W.T. Cai, X.G. Wang, L.H. Lee, T. Hung, S.J.E. Taylor, L.F. McGinnis, and S. Buckley. 2005. An integrated and adaptive decision-support framework for high-tech manufacturing and service networks. In *Proceedings of the 2005 Winter Simulation Conference* (ed. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines). Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Morrice, D. J., J. Butler, and P. W. Mullarkey. 1998. An approach to ranking and selection for multiple performance measures. In *Proceedings of the 1998 Winter Simulation Conference* (ed. D. J. Medeiros, E. F. Watson, J. S. Carson, and M. S. Manivannan), 719–725. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.

Nelson, B. L., J. Swann, D. Goldsman, and W. M. Song. 2001. Simple procedures for selecting the best simulated system when the number of alternatives is large. *Operations Research* 49: 950–963.

Rinott, Y. 1978. On two-stage selection procedures and related probability-inequalities. *Communications in Statistics* A7 (8): 799–811.

Sherbrooke, C. C. 1967. METRIC: A Multi-Echelon Technique for Recoverable Item Control. *Operations Research* 16: 122–141.

Swisher, J. R. and S. H. Jacobson. 2002. Evaluating the design of a family practice healthcare clinic using discrete-event simulation. *Health Care Management Science* 5 (2): 75–88.

Swisher, J. R., S. H. Jacobson, and E. Yücesan. 2003. Discrete-event simulation optimization using ranking, selection, and multiple comparison procedures: A survey. *ACM Transactions on Modeling and Computer Simulation* 13 (2): 134–154.

Veinott A. F. 1965. Optimal policy in a dynamic, single product, nonstationary inventory model with several demand classes. *Operations Research* 13: 761–778.

AUTHOR BIOGRAPHIES

LOO HAY LEE is an Assistant Professor in the Department of Industrial and Systems Engineering, National University of Singapore. He received his B.S. (Electrical Engineering) degree from the National Taiwan University in 1992 and his S.M. and Ph.D. degrees in 1994 and 1997 from Harvard University. He is currently a senior member of IEEE, vice president of ORSS, and a member of INFORMS. His research interests include simulation-based optimization, production scheduling and sequencing, logistics and supply chain planning, and vehicle routing. His email address is <iseleelh@nus.edu.sg> and his website is <www.ise.nus.edu.sg/staff/leelh/>.

SUYAN TENG is currently a Research Fellow in the Department of Industrial and Systems Engineering, National University of Singapore. She received her Ph.D. degree from National University of Singapore in 2004. Her research interests include simulation-based optimization and heuristic algorithms development for solving vehicle routing and scheduling problems. Her e-mail address is <isetsy@nus.edu.sg> .

EK PENG CHEW is an Associate Professor and Deputy Head (Academic) in the Department of Industrial and Systems Engineering, National University of Singapore. He received his Ph.D. degree from the Georgia Institute of Technology. His research interests include logistics and inventory management, system modeling and simulation, and system optimization. His email address is <misecep@nus.edu.sg> and his web address is <www.ise.nus.edu.sg/staff/chewep/>.

I. A. KARIMI is a faculty member in the Department of Chemical & Biomolecular Engineering at the National University of Singapore. His research interests include optimization, supply chain management, planning & scheduling, and mathematical modeling. His e-mail address is <cheiak@nus.edu.sg>.

KONG WEI LYE received his B.A.Sc. (Hon) degree in Computer Engineering from the Nanyang Technology University, Singapore in 1996, and the M.S. amd Ph.D. degrees in Electrical and Computer Engineering from Carnegie Mellon University, USA, in 1999 and 2004 respectively. Currently, he is a research engineer with the Singapore Institute of Manufacturing Technology where he helps conduct simulation studies and develop decision support systems. Lye has been a member of the IEEE since 1997. His e-mail address is <kwlye@SIMTech.a-star.edu.sg>.

PETER LENDERMANN is a Senior Scientist and Head of the D-SIMLAB Programme at Singapore Institute of Manufacturing Technology (SIMTech). Previously he was a Managing Consultant with agiplan in Germany where his focus was on the areas of supply chain management and production planning. He also worked as a Research Associate at the European Laboratory for Particle Physics CERN in Geneva (Switzerland) and Nagoya University (Japan). He obtained a Diploma in Physics from the University of Munich (Germany), a Doctorate in Applied Physics from Humboldt-University in Berlin (Germany) and a Master in International Economics and Management from Bocconi-University in Milan (Italy). His research interests include modelling and analysis of manufacturing and logistics systems as well as distributed simulation. His email address is <mpeterl@SIMTech.a-star.edu.sg>.

YANKAI, ALOYSIUS CHEN is currently in the final year of reading an undergraduate degree in Industrial and Systems Engineering at the National University of Singapore. His research interests include simulation-based optimization and systems optimization. His email address is <eng11099@nus.edu.sg>.

CHOON HWEE KOH is currently a final year engineering student in the Department of Industrial Systems Engineering, National University of Singapore. He is currently doing his final year thesis on simulation-based optimization on inventory control. His email address is <eng10610@nus.edu.sg>.