FINDING THE PARETO SET FOR MULTI-OBJECTIVE SIMULATION MODELS BY MINIMIZATION OF EXPECTED OPPORTUNITY COST

Loo Hay Lee
Ek Peng Chew
Suyan Teng

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

ABSTRACT

In this study, we mainly explore how to optimally allocate the computing budget for a multi-objective ranking and selection (MORS) problem when the measure of selection quality is the expected opportunity cost (OC). We define OC incurred to both the observed Pareto and non-Pareto set, and present a sequential procedure to allocate the replications among the designs according to some asymptotic allocation rules. Numerical analysis shows that the proposed solution framework works well when compared with other algorithms in terms of its capability of identifying the true Pareto set.

1 INTRODUCTION

Due to the complexity of most real-world systems, it is very hard or even impossible to analytically evaluate the performance measures of these systems and find their optimal decision variables. Therefore, simulation is commonly used to evaluate such complex systems and optimize their responses. When the number of system designs to be compared is relatively small and the input parameters to the systems are discrete, ranking and selection (R&S) techniques are commonly applied to determine the best system with optimal settings for the input parameters, where "best" is defined in terms of the maximum (or minimum) expected value of some function of the simulation output.

There are quite a number of review papers available in this field (Goldsman and Nelson 1998; Kim and Nelson 2003; Swisher, Jacobson and Yücesan 2003). One common issue critical to the R&S problem is how to measure the evidence of correct selection. In most studies, the selection quality is measured by the probability of correct selection (PCS), which is the probability that the selected system is the best system. Under this measure of selection quality, two main distinct lines of research exist for solving the

R&S problem. One is the indifference-zone procedure (Rinott, 1978) based on a least-favorable configuration (LFC) formulation to allocate additional replications. Recent research in this area has resulted in improvements over Rinott's original two-stage procedure. These include the two-stage procedure with elimination at the first stage (Nelson et al. 2001), two-stage procedure with subset selection at the first stage (Alrefaei and Alawneh 2004), and the sequential procedure incorporating sample mean information from the previous stages (Chen and Kelton 2004). Another line of research is based on an average case analysis to allocate additional replications. Typical works are the optimal computing budget allocation (OCBA) procedures developed in Chen et al. (2000) and Chen, Chen and Yücesan (2000). The OCBA framework follows a Bayesian methodology and allocates additional replications by solving the problem as an optimization problem, in which PCS is maximized with a given total computing budget available. Lee et al. (2006) extended both the OCBA procedure and the PCS concept to address the multi-objective ranking and selection (MORS) problem, the R&S problem with systems being evaluated in terms of more than one performance measure. In their solution framework, they incorporate the concept of Pareto optimality into the OCBA procedure and try to find all non-dominated Pareto set of solutions for the decision makers. To evaluate the quality of the Pareto set, the concept of PCS is generalized into Type II/I error corresponding to the Pareto/non-Pareto set, which measures the probability that at least one design from the Pareto/non-Pareto set is dominated/non-dominated. The paper derived asymptotic allocation rules for simulation replications and developed a sequential procedure (MOCBA) to iteratively allocate the simulation replications (according to the asymptotic rules) to identify the non-dominated Pareto set of solutions.

Another measure of selection quality distinct from PCS is the expected opportunity cost (OC), which is defined as the difference in means between the best and the

selected system. While PCS fails to reflect how poor a potential incorrect selection might be, OC does intuitively describe how far away the selected system is from the true best system. The expected OC as a measure of selection quality is applied in some Bayesian decision theoretic methods, in which additional replications are also allocated based on average case analysis but in a way that the expected value of information gained from those replications is maximized (Chick and Inoue 2001). Chick (2003) presented a selection procedure which can provide an upper bound for the expected OC of a potentially incorrect selection in a frequentist sense. The expected OC guarantee is applicable to all configurations of the mean rather than the LFC required in the indifference-zone approach. The paper bridged a gap between the indifference-zone approach (with frequentist guarantees) and the Bayesian approach to selection procedures (where the expected OC is considered). Within the OCBA framework, the expected OC as the measure of selection quality is applied and studied in He, Chick and Chen (2006). The paper extended the original OCBA approach based on maximization of the PCS to a new OCBA procedure seeking minimization of the expected OC. In this study, our main purpose is to study how the expected OC can be applied to measure the quality of the Pareto set in solving the MORS problem, and how the simulation replications should be allocated among the designs when measure of selection quality is changed from type II/I error of the Pareto/non-Pareto set to the OC of those sets.

We now describe the problem considered in this study. Suppose we have a set of *n* designs, each of which is evaluated in terms of *H* performance measures through simulation, we consider a MORS problem which is to find an optimal allocation of the simulation replications to the designs so that the non-dominated set of designs can be identified with the lowest expected OC at the least expense in terms of simulation replications. The organization of this paper is as follows. In Section 2, we first define the expected OC of selecting certain designs into the Pareto set, then introduce the asymptotic allocation rules and the sequential multi-objective computing budget allocation (MOCBA) procedure developed for solving the MORS problem. Some numerical analyses are carried out in Section 3. Finally some conclusions and future research directions are summarized in Section 4.

2 THE MOCBA SOLUTION FRAMEWORK

Without loss of generality, we assume that minimization of the objectives is our goal in this study. Also we assume that the random variables considered follow continuous distributions.

2.1 A Bayesian Framework and the Concept of Dominance

We establish the following notation.

S : The design space, i.e., $|S| = n$.

S_p : True Pareto set.

$\tilde{S}_p$: Observed Pareto set.

S_p^c : True non-Pareto set.

$\tilde{S}_p^c$: Observed non-Pareto set.

N_i : Number of replications allocated to design *i*.

H : Number of performance measures for each design.

$\boldsymbol{\mu}_i$: Vector of true performance measures of design *i*.

$\tilde{\boldsymbol{\mu}}_i$: Vector of random variables representing the true performance measures of design *i*.

$\hat{\boldsymbol{\mu}}_i$: $H \times l$ matrix representing *l* independent simulation observations for *H* performance measures of design *i*.

$\overline{\mu}_{ik}$: Sample mean of the simulation output for the *k*th objective of design *i*;

$$\overline{\mu}_{ik} = \frac{1}{l}\sum_{s=1}^{l} \hat{\mu}_{ik}^s.$$

σ_{ik}^2 : Variance of the simulation output for the *k*th objective of design *i*, which is to be estimated by sample variance

$$\sigma_{ik}^2 = \frac{1}{l-1}\sum_{s=1}^{l}\left(\hat{\mu}_{ik}^s - \overline{\mu}_{ik}\right)^2.$$

In this study, to make the problem more tractable, we assume that performance measures of the designs are independent of each other. We also assume that simulation outputs obtained from the following situations are independently distributed: 1) different replications; 2) different designs; 3) different performance measures of the same design. We apply a Bayesian framework in which the true performance measure $\boldsymbol{\mu}_i$ is regarded as a random variable and to be estimated by simulation observations $\hat{\boldsymbol{\mu}}_i$. Assume that each unknown performance measure μ_{ik} has a normal distribution as its prior distribution, and no prior knowledge on the performance of any design is available before conducting simulation. Given that $\hat{\mu}_{ik}^1, \hat{\mu}_{ik}^2, ..., \hat{\mu}_{ik}^l$ are *l* independent simulation observations for the *k*th objective of design *i* and σ_{ik}^2 is the known variance of the *k*th objective of design *i*, then according to DeGroot (1970), the true performance measure μ_{ik} can be described by its posterior distribution given below.

$$\tilde{\mu}_{ik} \sim N\left(\overline{\mu}_{ik}, \sigma_{ik}^2/N_i\right)$$

In this way, to compare unknown true performance measures μ_{ik} and μ_{jk} is to compare random variables $\tilde{\mu}_{ik}$ and $\tilde{\mu}_{jk}$ following the above posterior distribution derived from currently available simulation output.

In a noise-free situation, design j dominates design i, denoted by $j \prec i$, if the following condition holds with at least one inequality being strict:

$$\mu_{jk} \leq \mu_{ik} \quad \text{for} \ \ k = 1, 2, ..., H. \tag{1}$$

Thus, if design i is dominated, there exists at least one design j which satisfies the above condition. On the other hand, if design i is non-dominated, then for any other design j, design i has at least one objective better than its corresponding objective.

When the performance measures subject to noise, we need to consider the relationship among the designs in terms of probability or expected value as detailed in the following sections.

2.2 Construction of the Observed Pareto Set

During the computing budget allocation process, the Pareto set is constructed based on observed performance, named as the observed Pareto set. We define design j dominates design i by observation, denoted by $j \stackrel{\sim}{\prec} i$, if the following condition holds with at least one inequality being strict:

$$\bar{\mu}_{jk} \leq \bar{\mu}_{ik} \ \text{for} \ k = 1, 2, ..., H.$$

Then the observed Pareto set $\tilde{S}_p$ can be constructed by putting those designs i into $\tilde{S}_p$ if no design j exists such that $j \stackrel{\sim}{\prec} i$. The rest of the designs are then put into the observed non-Pareto set $\tilde{S}_p^c$.

2.3 Quality of the Observed Pareto Set in terms of Expected Opportunity Cost

In this study, we evaluate the quality of the observed Pareto set in terms of expected opportunity cost (OC) which may incur in the following two cases: 1) a design is non-dominated but it is selected into the non-Pareto set; 2) a design is dominated but it is selected into the Pareto set.

In the single objective problem, suppose the selected best design is denoted as i, and the true best design is denoted as $d*$. Then the OC is defined as:

$$OC = \mu_i - \mu_{d*} = \begin{cases} 0 & \text{if} \ i = d* \\ \geq 0 & \text{if} \ i \neq d* \end{cases}.$$

Here the OC is the same in terms of both selecting the selected best design i and missing the true best design $d*$.

However, in the multi-objective case, both the selected best and the true best are not one single design but a set of designs. This gives rise to several issues we need to address. Firstly, the OC of selecting the not truly non-dominated designs into the observed Pareto set $\tilde{S}_p$ is not the same as the OC of having those truly non-dominated designs in the observed non-Pareto set $\tilde{S}_p^c$. Secondly, to define the OC of design i, which design should design i compare with (serving the purpose of design $d*$ in the single objective case) for calculating its OC? Thirdly, as the OC is also with more than one performance measure, how to evaluate the OC when there is more than one objective?

2.3.1 The Opportunity Cost of the Observed Pareto and Non-Pareto Set

We begin to address the first two issues in this section. For the first issue, we can define the OC for both $\tilde{S}_p$ and $\tilde{S}_p^c$ to handle the two types of OC incurred due to including the dominated designs into the observed Pareto set while missing the truly non-dominated designs.

For the second issue, we realize that there is no unique way to define the OC for the two sets. We intend to borrow the idea from the single objective case and try to find a reference design based on which to compute the OC. We first consider how to define OC for the observed Pareto set $\tilde{S}_p$.

When design i is selected into $\tilde{S}_p$ but it is dominated, its OC should be defined in terms of the minimum amount design i needs to "move" in each objective such that design i becomes non-dominated. This measures how much better design i could have been if design i were non-dominated and no OC incurred to it. Suppose we use k_{ij}^* to represent the best objective of design i in reference to design j, i.e.,

$$k_{ij}^* = \arg \min_{k=1,...,H} \left(\mu_{ik} - \mu_{jk} \right).$$

Then, if for each design $j \in S, j \neq i$,

$$\min_{k=1,...,H} \left(\mu_{ik} - \mu_{jk} \right) < 0$$

holds, no design is better than design i at its best objective, therefore design i becomes non-dominated. The reference design, denoted as d_i^*, should be the one furthest away from design i at its best objective, i.e.,

$$d_i^* \equiv \arg \max_{j \in S, j \neq i} \min_{k=1,2,...,H} \left(\mu_{ik} - \mu_{jk} \right).$$

Here the $\min\limits_{k=1,2,...,H}\left(\mu_{ik}-\mu_{jk}\right)$ is to find out the best objective of design i in reference to design j. Once this objective of design i is good enough, design i is not dominated by design j. However, design i can still be dominated by other designs. So the Max operator is to find out the design which dominates design i to the greatest extent.

Then the OC of selecting design i into $\tilde{S}_p$ can be defined as follows:

$$OC_i = \sum_{k=1}^{H}\left(\mu_{ik}-\mu_{d_i^*k}\right)I\left(\mu_{i1}\geq\mu_{d_i^*1},...,\mu_{iH}\geq\mu_{d_i^*H}\right).$$

where $I\left(\mu_{i1}\geq\mu_{d_i^*1},...,\mu_{iH}\geq\mu_{d_i^*H}\right)$ is an indicator function defined as:

$$I\left(\mu_{i1}\geq\mu_{d_i^*1},...,\mu_{iH}\geq\mu_{d_i^*H}\right)=\begin{cases}1 \text{ if } \mu_{ik}\geq\mu_{d_i^*k} \text{ for } k=1,...,H.\\ 0 \text{ otherwise.}\end{cases}$$

Note that, in the deterministic case, since design d_i^* dominates design i for sure, the indicator function can be omitted in this case.

When the performance measures subject to noise, we define design d_i^* in terms of expected value as follows:

$$d_i^* \equiv \arg\max_{j\in S, j\neq i}\min_{k=1,2,...,H}E\left[\left(\tilde{\mu}_{ik}-\tilde{\mu}_{jk}\right)I\left(\tilde{\mu}_{ik}-\tilde{\mu}_{jk}\geq 0\right)\right]. \quad (2)$$

Suppose $\tilde{\mu}_{ijk}$ is the random variable representing the difference in means of the kth objective between design i and design j. Then $\tilde{\mu}_{ijk}$ follows normal distribution with mean $\delta_{ijk}=\bar{f}_{ik}-\bar{f}_{jk}$ and variance $\sigma_{ijk}^2=\sigma_{ik}^2/N_i+\sigma_{jk}^2/N_j$. So

$$\begin{aligned}&E\left[\left(\tilde{\mu}_{ik}-\tilde{\mu}_{jk}\right)I\left(\tilde{\mu}_{ik}-\tilde{\mu}_{jk}\geq 0\right)\right]\\ &=E\left[\tilde{\mu}_{ik}-\tilde{\mu}_{jk}\,\middle|\,\tilde{\mu}_{ik}\geq\tilde{\mu}_{jk}\right]P\left(\tilde{\mu}_{ik}\geq\tilde{\mu}_{jk}\right)\\ &=\int_0^{\infty}xf_{ijk}\left(x\right)dx\\ &=\sigma_{ijk}f\left(Z_{ijk}\right)+\delta_{ijk}\Phi\left(-Z_{ijk}\right).\end{aligned} \quad (3)$$

where $f_{ijk}\left(x\right)$ and $\Phi\left(x\right)$ are the PDF and CDF of standard normal distribution, $Z_{ijk}=-\delta_{ijk}/\sigma_{ijk}$ is the standardized statistic. From (2) and (3) we have,

$$d_i^* \equiv \arg\max_{j\in S, j\neq i}\min_{k=1,2,...,H}\left[\sigma_{ijk}f\left(Z_{ijk}\right)+\delta_{ijk}\Phi\left(-Z_{ijk}\right)\right].$$

Now we can define the expected OC_i in the presence of noise as follows:

$$\begin{aligned}&E[OC_i]=E\left[\sum_{k=1}^{H}\left(\tilde{\mu}_{ik}-\tilde{\mu}_{d_i^*k}\right)I\left(\tilde{\mu}_{i1}\geq\tilde{\mu}_{d_i^*1},...,\tilde{\mu}_{iH}\geq\tilde{\mu}_{d_i^*H}\right)\right]\\ &=\sum_{k=1}^{H}E\left[\left(\tilde{\mu}_{ik}-\tilde{\mu}_{d_i^*k}\right)I\left(\tilde{\mu}_{i1}\geq\tilde{\mu}_{d_i^*1},...,\tilde{\mu}_{iH}\geq\tilde{\mu}_{d_i^*H}\right)\right]\\ &=\sum_{k=1}^{H}E\left[\tilde{\mu}_{ik}-\tilde{\mu}_{d_i^*k}\,\middle|\,\tilde{\mu}_{i1}\geq\tilde{\mu}_{d_i^*1},...,\tilde{\mu}_{iH}\geq\tilde{\mu}_{d_i^*H}\right]P\left(\tilde{\mu}_{i1}\geq\tilde{\mu}_{d_i^*1},...,\tilde{\mu}_{iH}\geq\tilde{\mu}_{d_i^*H}\right)\end{aligned}$$

Due to the independence assumption on the performance measures, we have

$$\begin{aligned}&E[OC_i]\\ &=\sum_{k=1}^{H}E\left[\tilde{\mu}_{ik}-\tilde{\mu}_{d_i^*k}\,\middle|\,\tilde{\mu}_{ik}\geq\tilde{\mu}_{d_i^*k}\right]P\left(\tilde{\mu}_{ik}\geq\tilde{\mu}_{d_i^*k}\right)\prod_{\substack{m=1\\m\neq k}}^{H}P\left(\tilde{\mu}_{im}\geq\tilde{\mu}_{d_i^*m}\right)\\ &=\sum_{k=1}^{H}\int_0^{\infty}xf_{id_i^*k}\left(x\right)dx\prod_{\substack{m=1\\m\neq k}}^{H}P\left(\tilde{\mu}_{id_i^*m}\geq 0\right)\\ &=\sum_{k=1}^{H}\left[\sigma_{id_i^*k}f\left(Z_{id_i^*k}\right)+\delta_{id_i^*k}\Phi\left(-Z_{id_i^*k}\right)\right]\prod_{\substack{m=1\\m\neq k}}^{H}P\left(\tilde{\mu}_{id_i^*m}\geq 0\right)\end{aligned}$$

For all designs selected into the observed Pareto set $\tilde{S}_p$, the average approximated OC is:

$$\begin{aligned}&E[OC]_{\tilde{S}_p}\\ &=\frac{1}{\left|\tilde{S}_p\right|}\sum_{i\in S_p}\sum_{k=1}^{H}\left[\sigma_{id_i^*k}f\left(Z_{id_i^*k}\right)+\delta_{id_i^*k}\Phi\left(-Z_{id_i^*k}\right)\right]\prod_{\substack{m=1\\m\neq k}}^{H}P\left(\tilde{\mu}_{id_i^*m}\geq 0\right)\end{aligned} \quad (4)$$

Now we consider how to define OC for the observed non-Pareto set. For a truly non-dominated design i, if it is selected into the observed non-Pareto set $\tilde{S}_p^c$, the OC of missing design i in the observed Pareto set $\tilde{S}_p$ incurs. Again we propose to define OC_i in terms of the minimum amount design i needs to "move" in each objective such that design i becomes dominated. Given that k_{ij}^* is the best objective of design i in reference to design j, then $k_{ij}^*=\arg\max\limits_{k=1,...,H}\left(\mu_{jk}-\mu_{ik}\right)$ in this case. As design i is non-dominated, there should be no design which is better than it at its best objective. Once we find a design $j\in S, j\neq i$, which satisfies $\max\limits_{k=1,...,H}\left(\mu_{jk}-\mu_{ik}\right)<0$, design i becomes worse than design j even at its best objective. Therefore design i becomes dominated. The reference design, denoted as g_i^*, should be the design that is closest to design i in terms of its best objective, i.e.,

$$g_i^* \equiv \arg\min_{j\in S, j\neq i}\max_{k\in\{1,2,...,H\}}\left(\mu_{jk}-\mu_{ik}\right).$$

Here $\max\limits_{k=1,2,...,H}\left(\mu_{jk}-\mu_{ik}\right)$ is to find out the objective that design i is the best in reference to design j. Once this objective of design i is not as good as design j, design i becomes dominated by design j. The Min operator is to identify the

design to which design i needs to "move" with the least effort to become dominated.

Accordingly, we define OC_i as:

$$OC_i = \sum_{k=1}^{H} \left(\mu_{g_i^* k} - \mu_{ik} \right) I \left(\mu_{g_i^* k} \geq \mu_{ik} \right).$$

Similarly when noise is involved in the performance measures, design g_i^* can be any one design with certain probability, we redefine it in terms of expected value as follows:

$$g_i^* \equiv \arg \min_{j \in S, j \neq i} \max_{k \in \{1,2,...,H\}} E\left[\left(\mu_{jk} - \mu_{ik} \right) I \left(\mu_{jk} \geq \mu_{ik} \right) \right]. \quad (5)$$

Suppose $\tilde{\mu}_{jik}$ is the random variable representing the difference in means of the kth objective between design j and design i. Then $\tilde{\mu}_{jik}$ follows normal distribution with mean $\delta_{jik} = \overline{f}_{jk} - \overline{f}_{ik}$ and variance $\sigma_{jik}^2 = \sigma_{jk}^2 / N_j + \sigma_{ik}^2 / N_i$. By doing similar manipulations as in (3), we can arrive at,

$$g_i^* \equiv \arg \min_{j \in S, j \neq i} \max_{k \in \{1,2,...,H\}} \left[\sigma_{jik} f \left(Z_{jik} \right) + \delta_{jik} \Phi \left(-Z_{jik} \right) \right].$$

So the expected OC can be defined as:

$$E[OC_i] = E\left[\sum_{k=1}^{H} \left(\tilde{\mu}_{g_i^* k} - \tilde{\mu}_{ik} \right) I \left(\tilde{\mu}_{g_i^* k} \geq \tilde{\mu}_{ik} \right) \right]$$

$$= \sum_{k=1}^{H} E\left[\tilde{\mu}_{g_i^* k} - \tilde{\mu}_{ik} \mid \tilde{\mu}_{g_i^* k} \geq \tilde{\mu}_{ik} \right] P\left(\tilde{\mu}_{g_i^* k} \geq \tilde{\mu}_{ik} \right)$$

$$= \sum_{k=1}^{H} \left[\sigma_{g_i^* ik} f \left(Z_{g_i^* ik} \right) + \delta_{g_i^* ik} \Phi \left(-Z_{g_i^* ik} \right) \right].$$

When considering all designs in $\tilde{S}_p^c$, we have the following average expected OC:

$$E[OC]_{\tilde{S}_p^c} = \frac{1}{\left| \tilde{S}_p^c \right|} \sum_{i \in \tilde{S}_p^c} \sum_{k=1}^{H} \left[\sigma_{g_i^* ik} f \left(Z_{g_i^* ik} \right) + \delta_{g_i^* ik} \Phi \left(-Z_{g_i^* ik} \right) \right]. \quad (6)$$

2.3.2 Normalization of the Performance Measures

For the third issue, to handle the multi-objective nature of the OC, we use the sum of the OC incurred in all objectives of design i as its OC. One problem here is that different objectives may not be on the same scale and therefore incomparable. For example, in real life problems, objectives may be very different with one representing investment in millions of dollars, while the other representing customer service level with values within 0 and 1. This renders the summation of the OC to be meaningless. To overcome this difficulty, we can do some preprocessing on the simulation output already obtained: normalize the simulation output of each performance measure so that they are all within interval [0, 1]. Given that $\overline{\mu}_{ik}$ and σ_{ik}^2 are the sample mean and variance of the kth objective of design i $(i = 1,2,...,n)$, then the normalized sample mean

and variance of the kth objective of design i denoted as $\overline{\mu}_{ik}'$ and $\sigma_{ik}^{2'}$ can be calculated as:

$$\overline{\mu}_{ik}' = \frac{\overline{\mu}_{ik} - \min_{i=1,2,...,n} \overline{\mu}_{ik}}{\max_{i=1,2,...,n} \overline{\mu}_{ik} - \min_{i=1,2,...,n} \overline{\mu}_{ik}} \quad (7)$$

$$\sigma_{ik}^{2'} = \frac{\sigma_{ik}^2}{\left(\max_{i=1,2,...,n} \overline{\mu}_{ik} - \min_{i=1,2,...,n} \overline{\mu}_{ik} \right)^2}. \quad (8)$$

2.4 A Sequential Solution Procedure and the Asymptotic Allocation Rules

In solving the MORS problem, we are trying to get the true Pareto set by minimizing expected OC associated with both $\tilde{S}_p$ and $\tilde{S}_p^c$. To this end, we can keep on allocating more replications to certain designs until some termination conditions are met. Two examples of the termination conditions can be: (a) both $E[OC]_{\tilde{S}_p}$ and $E[OC]_{\tilde{S}_p^c}$ are within a preset limit $\varepsilon*$; (b) a total computing budget N_{max} is exhausted. We will focus on the latter one in the numerical analysis of this study.

The general idea is to perform N_0 (relatively small so that it would be generally less than N_i) replications for each design $(i = 1,2,...,n)$ initially. Then the simulation budget can be allocated in a sequential manner as follows. At each step, given the current observed Pareto set (constructed based on the most recent sample mean information) and a new N_{max} (incremented by an additional number of replications) to be allocated, we allocate them among the designs according to allocation rules as stated in Rule 1 and Rule 2 depending on which expected OC ($E[OC]_{\tilde{S}_p}$ and $E[OC]_{\tilde{S}_p^c}$) is larger. Next, we update the expected OC according to new sample mean and variance information. If the termination condition is not met, a new observed Pareto set is constructed and the allocation rules are applied again to do the allocation for a newly incremented N_{max}. The procedure is iterated until the adopted termination condition is satisfied.

To derive the allocation rules, the problem being considered is formulated as an optimization problem with an objective of minimizing the two types of expected OC and a constraint on total computing budget available to be allocated among the designs. Under the condition that total computing budget is infinite, we apply the Lagrange Relaxation method, and derive some asymptotic allocation rules. The rules are now stated as follows.

Rule 1: As $N_{max} \to \infty$, the upper bound of expected OC for $\tilde{S}_p^c$ can be asymptotically minimized when

For a design $l \in \tilde{S}_p^c$,

$$\alpha_l = \frac{\beta_l}{\sum_{l \in \tilde{S}_p^c} \beta_l + \sum_{d \in \tilde{S}_p} \beta_d}.$$

For a design $d \in \tilde{S}_p$,

$$\alpha_d = \frac{\beta_d}{\sum_{l \in \tilde{S}_p^c} \beta_l + \sum_{d \in \tilde{S}_p} \beta_d},$$

where

$$\beta_l = \frac{\alpha_l}{\alpha_m} = \frac{\left(\sigma_{lk_{j_l}^l}^2 + \sigma_{j_l k_{j_l}^l}^2 \big/ \rho_l\right) \big/ \delta_{lj_l k_{j_l}^l}^2}{\left(\sigma_{mk_{j_m}^m}^2 + \sigma_{j_m k_{j_m}^m}^2 \big/ \rho_m\right) \big/ \delta_{mj_m k_{j_m}^m}^2},$$

given that m is any fixed design in $\tilde{S}_p^c$;

$$\beta_d = \frac{\alpha_d}{\alpha_m} = \sqrt{\sum_{i \in \Omega_d} \frac{\sigma_{dk_d^i}^2}{\sigma_{ik_d^i}^2} \beta_i^2},$$

where α_i is the fraction of N_{max} to be allocated to design i; $\delta_{jik} = \overline{\mu}_{jk} - \overline{\mu}_{ik}$;

$$j_i \equiv \arg \min_{\substack{j \in S \\ j \neq i}} \max_{k \in \{1,...,H\}} \left[\sigma_{jik} f\left(Z_{jik}\right) + \delta_{jik} \Phi\left(-Z_{jik}\right)\right],$$

$$k_{j_i}^i \equiv \arg \max_{k \in \{1,...,H\}} \left[\sigma_{g_i^* ik} f\left(Z_{g_i^* ik}\right) + \delta_{g_i^* ik} \Phi\left(-Z_{g_i^* ik}\right)\right]$$

with $\tilde{\mu}_{jik} \sim N\left(\delta_{jik}, \sigma_{jik}^2\right)$, where $f(x)$ and $\Phi(x)$ are PDF and CDF of standard normal distribution. Here α_i' is the fraction of N_{max} allocated to design i at the immediate previous iteration. $\Omega_d = \{\text{design } i \mid i \in \tilde{S}_p^c, j_i = d\}$; $\rho_i = \alpha_{j_i}' / \alpha_i'$ initially and after getting α_i and α_{j_i}, it is determined iteratively until it converges.

Rule 2: As $N_{max} \to \infty$, the upper bound of expected OC for $\tilde{S}_p$ can be asymptotically minimized when

For a design $l \in S_p^A$,

$$\alpha_l = \frac{\beta_l}{\sum_{l \in S_p^A} \beta_l + \sum_{u \in S_p^B} \beta_u + \sum_{d \in \tilde{S}_p^c} \beta_d}.$$

For a design $u \in S_p^B$,

$$\alpha_u = \frac{\beta_u}{\sum_{l \in S_p^A} \beta_l + \sum_{u \in S_p^B} \beta_u + \sum_{d \in \tilde{S}_p^c} \beta_d}.$$

For a design $d \in \tilde{S}_p^c$,

$$\alpha_d = \frac{\beta_d}{\sum_{l \in S_p^A} \beta_l + \sum_{u \in S_p^B} \beta_u + \sum_{d \in \tilde{S}_p^c} \beta_d}$$

with

$$\beta_l = \frac{\alpha_l}{\alpha_m} = \frac{\left(\sigma_{lk_{j_l}^l}^2 + \sigma_{j_l k_{j_l}^l}^2 \big/ \rho_l\right) \big/ \delta_{lj_l k_{j_l}^l}^2}{\left(\sigma_{mk_{j_m}^m}^2 + \sigma_{j_m k_{j_m}^m}^2 \big/ \rho_m\right) \big/ \delta_{mj_m k_{j_m}^m}^2},$$

given m is any fixed design in S_p^A;

$$\beta_u = \frac{\alpha_u}{\alpha_m} = \sqrt{\sum_{i \in \Theta_u^*} \frac{\sigma_{uk_u^i}^2}{\sigma_{ik_u^i}^2} \beta_i^2} \quad \text{and} \quad \beta_d = \frac{\alpha_d}{\alpha_m} = \sqrt{\sum_{i \in \Theta_d^*} \frac{\sigma_{dk_d^i}^2}{\sigma_{ik_d^i}^2} \beta_i^2},$$

where, in addition to notations given in Rule 1,

$$j_i \equiv \arg \max_{\substack{j \in S \\ j \neq i}} \min_{k \in \{1,...,H\}} \left[\sigma_{ijk} f\left(Z_{ijk}\right) + \delta_{ijk} \Phi\left(-Z_{ijk}\right)\right],$$

$$k_{j_i}^i \equiv \arg \min_{k \in \{1,...,H\}} P\left(\mu_{ij_i k} \geq 0\right),$$

$$\Theta_d \equiv \{\text{design } i \mid i \in \tilde{S}_p, j_i = d\},$$

$$S_p^A \equiv \left\{\text{design } l \in S_p \left| \frac{\delta_{lj_l k_{j_l}^l}^2}{\sigma_{lk_{j_l}^l}^2 \big/ \alpha_l + \sigma_{j_l k_{j_l}^l}^2 \big/ \alpha_{j_l}} < \min_{i \in \Theta_l} \frac{\delta_{ilk_l^i}^2}{\sigma_{ik_l^i}^2 \big/ \alpha_i + \sigma_{lk_l^i}^2 \big/ \alpha_l} \right.\right\}$$

with α_i being approximated by α_i'; $S_p^B = \tilde{S}_p \setminus S_P^A$; $\Theta_d^* \equiv \{\Theta_d \cap S_P^A\}$.

The sequential procedure applying the above allocation rules, known as the MOCBA_OC algorithm, is outlined as follows.

MOCBA_OC 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 observed Pareto set $\tilde{S}_p$ as stated in Section 2.2, and calculate expected OC ($E[OC]_{\tilde{S}_p}$ and $E[OC]_{\tilde{S}_p^c}$) according to equations (4) and (6).

Step 2: If the termination condition is met, go to Step 6. Otherwise, increase the total computing budget by a certain amount Δ; if $E[OC]_{\tilde{S}_p^c} \leq E[OC]_{\tilde{S}_p}$, go to Step 4.

Step 3: Calculate the new allocation $N_1^{v+1}, N_2^{v+1}, ..., N_n^{v+1}$ according to the allocation rules stated in Rule 1. Go to Step 5.

Step 4: Calculate the new allocation $N_1^{v+1}, N_2^{v+1}, ..., N_n^{v+1}$ according to allocation rules stated in Rule 2.

Step 5: Perform additional $\min(\delta, \max(0, N_i^{v+1} - N_i^v))$ replications for design i $(i = 1, 2, ..., n)$. Set $v = v + 1$ and go to Step 1.

Step 6: Output designs in the observed Pareto set ($\tilde{S}_p$).

3 NUMERICAL ANALYSIS

In this section, we examine the performance of the proposed MOCBA_OC solution framework. Here the performance is evaluated in terms of estimated true opportunity cost — OC_{S_p} and $OC_{S_p^c}$ (OC incurred to both the Pareto and non-Pareto set averaged over total independent experiments conducted), estimated true Type II and I errors (percentage of experiments that misclassify designs into the Pareto and non-Pareto sets), and the estimated true probability of correct selection — PCS (percentage of experiments that can identify the true Pareto set). The proposed MOCBA_OC algorithm is compared with the uniform computing budget allocation (UCBA) algorithm. In the following computational experiment, we set the parameters as follows: $N_0 = 5$, $\Delta = 10$ and $\delta = 5$.

For the convenience of comparison, we use the same test problem presented in Lee et al. (2006) except that we multiply 10 to the mean and standard deviation of the third performance measure. The purpose of the multiplication is to make the third objective apparently in different scale from the other two objectives, so that we can study whether difference in objectives will affect the performance of the algorithm and whether normalization is an apt tool to handle the multi-objective nature of the OC. For this test problem, there are 25 design alternatives, each is evaluated in terms of three objectives. We assume that the true performance measures of the designs are known, therefore we know the non-dominated set of designs is 1, 2, 4, 5, and 9. When the performance measures are unknown, we assume that they can be evaluated through simulation whose output follows normal distribution with given means and standard deviations. The means are assumed to be the same as the true performance measures and the standard deviation is 8 for each performance measure of the designs.

To study on average how the algorithms perform as computing budget increases, we conduct 2000 independent experiments each with total computing budget N_{max} set as: 200, 400, …, 2000. Results are shown in Table 1 and Figures 1 to 4.
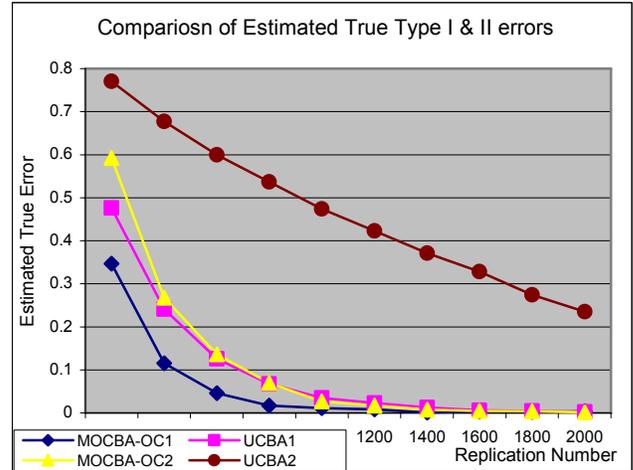
Figure 1: Estimated true errors of the procedures.

Figure 1 illustrates the estimated true Type I and II errors (labeled as "1" and "2" respectively in the legend) for the two algorithms. It is very clear that MOCBA_OC performs much better than UCBA for both types of errors. Moreover, it produces a more balanced minimization of the two types of errors: the difference between Type II error and Type I error is much smaller than that of the UCBA). Now we compare their capabilities of identifying the final Pareto set. Figure 2 displays how PCS differs between the two procedures as computing budget changes.
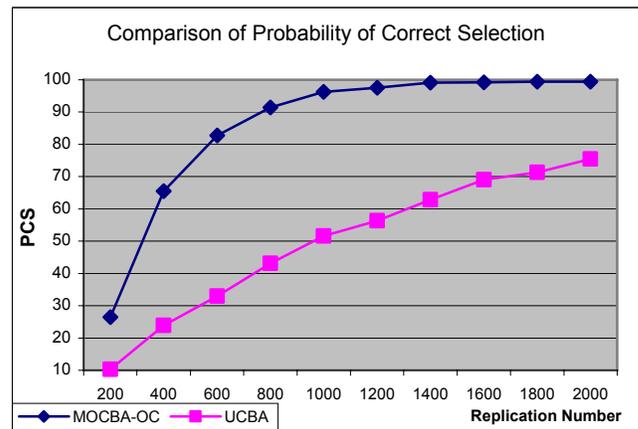
Figure 2: Comparison of PCS among the algorithms.

We can observe from Figure 2 that, MOCBA-OC performs substantially better than UCBA. The PCS produced by MOCBA-OC is generally about 40% better than that given by UCBA. Now we study when compared with UCBA, how MOCBA-OC speed up the simulation process in terms of the savings in computing budget. Table 1 lists out the results when the required PCS is set as: 99%, 95%, …, and 80%.

Table 1: Speedup of MOCBA-OC over UCBA in terms of simulation replications.

PCS	UCBA	MOCBA-OC	
		Rep. No.	Ratio
99%	6860	1400	4.9
95%	4500	930	4.84
90%	3650	750	4.87
85%	2800	650	4.31
80%	2400	590	4.07

The results in Table 1 indicate that UCBA takes about 4 to 5 times of the computing budget that MOCBA-OC takes. The worse performance of UCBA highlights the point that we should allocate computing budget among the designs according to their performance rather than uniformly. Figure 3 depicts how MOCBA-OC and UCBA allocate the simulation budget when N_{max} = 1600. When compared with UCBA, MOCBA-OC exhibits different patterns in terms of distribution of the replications among the designs — only the following designs are allocated more replications: (a) those designs should be in the Pareto set, and (b) those designs whose performances are very close to designs in (a). As MOCBA-OC has directed most of its efforts in singling out those non-dominated designs from those designs with close performances, this explains why it can outperform UCBA and result in great speedup of the simulation process.
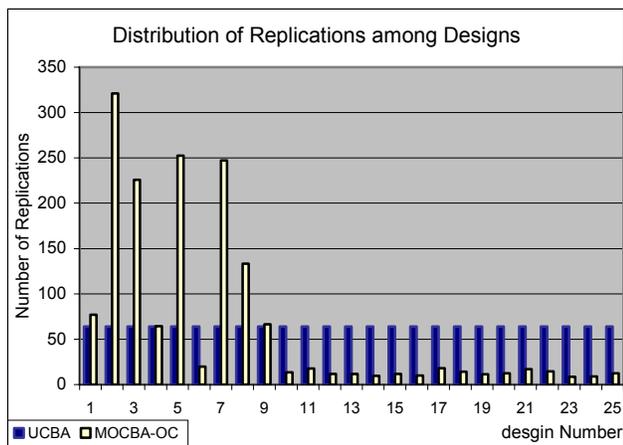
Figure 3: Distribution of the computing budget among the designs.

We now present how expected OC changes for MOCBA_OC and UCBA as total computing budget N_{max} increases from 200 to 2000. Results are shown in Figure 4.
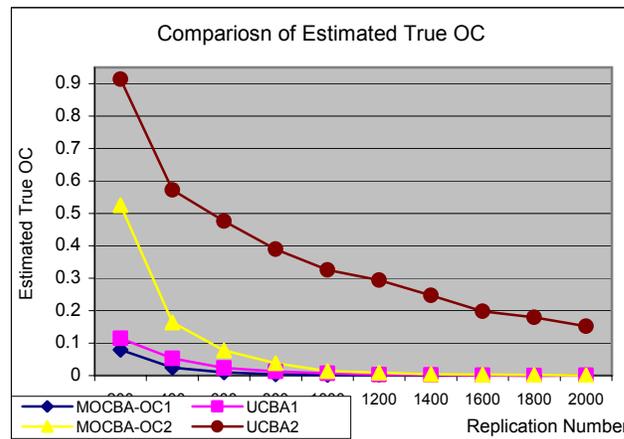
Figure 4: Comparison of expected OC.

Here the expected OC ("1" and "2" correspond to the observed non-Pareto and Pareto set respectively) is calculated based on normalized performance measures as shown in (7). Again MOCBA-OC achieves much lower expected OC than UCBA for both observed Pareto and non-Pareto set. Though UCBA can achieve relatively lower OC of the non-Pareto set, it results in very high OC for the Pareto set. MOCBA-OC alleviates this problem by allocating in a different way (Figure 3). However, especially when N_{max} is low, the two types of OC still deviate highly from each other. This implies that there is still room to improve the performance of the MOCBA-OC by balancing the minimization of OC associated with the Pareto and non-Pareto set.

4 CONCLUSIONS

In this study, we present a solution framework for solving the multi-objective ranking and selection (MORS) problem based on minimization of opportunity cost (OC). A sequential solution procedure is proposed to iteratively identify the Pareto set by allocating the simulation replications among the designs according to some asymptotic allocation rules. Numerical analysis shows that, the proposed framework employing OC as the measure of selection quality works substantially better than Uniform Computing Budget Allocation (UCBA) in terms of both probability of correct selection and expected OC. In future research, we need to do more numerical analysis to study the performance of MOCBA_OC. For example, we can compare it with the procedure proposed in Lee et al. (2006) (a framework employs Type I & II errors to evaluate the quality of the Pareto set) to see how different measures of selection quality affect the quality of the final Pareto set and the distribution of the simulation replications among the designs. We may also study how correlation among the objectives would affect the allocation rules and therefore the performance of the solution framework.

REFERENCES

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

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

Chen, E.J. and Kelton, W.D. (2004) Sequential selection procedures: Using sample means to improve efficiency. *European Journal of Operational Research*,

Chen, H.C., Chen C.H. and Yücesan, E. (2000) Computing Efforts Allocation for Ordinal Optimization and Discrete Event Simulation. *IEEE Transactions on Automatic Control*, 45 (5), 960-964.

Chick, S.E. (2003) Expected opportunity cost guarantees and indifference-zone selection procedures, in *Proceedings of the 2003 Winter Simulation Conference*, IEEE Press, Piscataway, NJ, pp. 465–473.

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

DeGroot, M. H. 1970. *Optimal Statistical Decisions*. McGraw-Hill, Inc.

Goldsman, D. and Nelson, B.L. (1998) Statistical screening, selection, and multiple comparison procedures in computer simulation, in *Proceedings of the 1998 Winter Simulation Conference*, IEEE Press, Piscataway, NJ, pp. 159–166.

He, D., Chick, S. E., and Chen, C. H. (2006). The opportunity cost and OCBA selection procedures in ordinal optimization. Submitted for publication.

Kim, S.H. and Nelson, B.L. (2003) Selecting the best system: theory and methods, in *Proceedings of the 2003 Winter Simulation Conference*, IEEE Press, Piscataway, NJ, pp. 101-112.

Lee, L.H., Chew, E.P., Teng, S.Y. and Goldsman, D. (2006). Optimal computing budget allocation for multi-objective simulation models. Submitted to *IIE Transactions*.

Nelson, B.L., Swann, J., Goldsman, D. and Song W.M. (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.

Swisher, J.R., Jacobson, S.H. and Yücesan, E. (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.

AUTHOR BIOGRAPHIES

LOO HAY LEE is an Associate 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, committee member 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 web address is <www.ise.nus.edu.sg/staff/leelh/>.

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 <isecep@nus.edu.sg> and his web address is <www.ise.nus.edu.sg/staff/chewep/>.

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 <isets@nus.edu.sg>.