



Computing budget allocation rules for multi-objective simulation models based on different measures of selection quality[☆]

Loo Hay Lee^{a,*}, Ek Peng Chew^a, Suyan Teng^b

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

^b School of Engineering, Republic Polytechnic, 9 Woodlands Avenue 9, Singapore 738964, Singapore

ARTICLE INFO

Article history:

Received 1 May 2009

Received in revised form

4 July 2010

Accepted 12 July 2010

Available online 8 October 2010

Keywords:

Allocation rules

Multi-objective

Computing budget allocation

Lagrange relaxation

Measure of selection quality

ABSTRACT

In an optimal computing budget allocation problem, different measures of selection quality determine how the best set of designs can be identified and how the simulation budget should be allocated among the designs. In this paper, we look at several measures of selection quality and derive respective allocation rules for the multi-objective computing budget allocation problem. Some computational experiments are carried out to compare the performance of the allocation rules and to identify the suitable ones in certain scenarios.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

Simulation plays a crucial role in analyzing discrete-event systems, in particular, when comparing alternative system designs with a view to optimizing system performance (Chen, Chen, Yücesan & Dai, 1998). However, due to the slow convergence of performance, simulation efficiency is still a major concern especially when the number of competing designs to be compared is large. This explains the increasing popularity of research in ranking and selection (R&S): techniques that determine the number of simulation replications required for each design so that a certain measure of selection quality is guaranteed at a pre-specified level with the least possible computational expense. For a comprehensive review of this field, see Branke, Chick, and Schmidt (2007), Kim and Nelson (2003, 2007) and Swisher, Jacobson, and Yücesan (2003).

To identify the best system from a set of alternatives, one first needs to specify by which criterion we compare the systems and select the best one, i.e., 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

defined as the probability that the selected best is the true best. Solution approaches using PCS as the measure of selection quality include the indifference-zone (IZ) ranking and selection (Nelson, Swann, Goldsman, & Song, 2001; Pritsker, 1986), the decision theoretic methods (Chick, 1997), and the optimal computing budget allocation (OCBA). The IZ-based procedure allocates additional replications based on a least-favorable configuration (LFC) formulation, and causes the allocation efficiency to be more conservative. To improve on this, 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. The numerical results show that the OCBA is able to give better results in terms of using the least computing budget for a given level of PCS. For more details, refer to Chen, Chen, and Dai (1996), Chen, Chen, and Yücesan (2000), Chen, Dai, Chen, and Yücesan (1997), Chen, Donohue, Yücesan, and Lin (2003) and Chen, Lin, Yücesan and Chick (2000). Another measure of selection quality, known as the expected opportunity cost (EOC), aims to quantify how far away the selected system is from the true best system. This measure is important as it not only maximizes the chance of selecting the best design, but also if it fails to find the best design, by optimizing this measure, it will ensure the selected design will not be very far off from the best design. EOC is defined as the difference in means between the selected system and the true best. It has been applied in some Bayesian decision theoretic methods, in which additional replications are allocated in a way that the expected value of information gained from the replications is maximized

[☆] This paper was not presented at any conference. This paper was recommended for publication in revised form under the direction of Editor Berç Rüstem.

* Corresponding author. Tel.: +65 6874 2895; fax: +65 6777 1434.

E-mail addresses: iseleelh@nus.edu.sg (L.H. Lee), isecep@nus.edu.sg (E.P. Chew), teng_suyan@rp.sg (S. Teng).

(Chick & Inoue, 2001). In Chick (2003), a selection procedure which can provide an upper bound for the EOC in a frequentist sense was proposed. The paper bridged a gap between the IZ approach (with frequentist guarantees) and the Bayesian approach (where the EOC is considered) to the selection procedures. Within the OCBA framework, EOC is applied and studied in He, Chick, and Chen (2007), where the original OCBA based on maximization of PCS is extended to a new OCBA seeking minimization of EOC. The expected Net Present Value (NPV), as a measure of selection quality, was proposed in Chick and Gans (2006). NPV concerns not only the economic benefit from implementing the selected system, but also the marginal costs of simulation runs and discounting due to simulation analysis time. NPV is to be maximized to underline the financial significance of the selected systems when the system performance and simulation results are themselves financial measures.

When the systems to be compared are evaluated in terms of more than one performance measure, the R&S problem becomes a multi-objective Ranking and Selection (MORS) problem. The MORS problem has not been as well studied as its single objective counterpart. In the case where a single best solution is pursued, the problem is often transformed into a single-objective problem and PCS is often used as the measure of selection quality (Butler, Morrice, & Mullarkey, 2001). Another line of research treats the multi-objective problem as it is and applies the concept of Pareto Optimality to find all non-dominated solutions. In Lee, Chew, Teng, and Goldsman (2004, 2010), Multi-objective Computing Budget Allocation (MOCBA) frameworks, extensions of OCBA in the single objective case, were proposed. In Lee et al. (2004), the selection measure is a performance index defined as the cumulative probability that a design is dominated by all the other designs. With the assumption that the number of non-dominated designs is known, non-dominated designs can be identified as those with performance indices approaching zero. Lee et al. (2010) improved the measure of selection quality by defining it as two types of errors associated with the Pareto and non-Pareto sets. The two types of errors are a generalization of the PCS concept; when both of them approach zero, the true Pareto set is found with no need of assuming the known number of non-dominated designs.

However, the solution framework proposed in Lee et al. (2010) still has several potential limitations. First, the allocation rules are derived separately – one rule for minimizing one type of errors; and the two rules are used iteratively during the allocation. This may not be optimal when it is compared to the allocation rules derived from maximizing the PCS of the Pareto set. Secondly, it fails to reflect how poor a potential incorrect selection might be. For instance, when a dominated solution is selected into the Pareto set, MOCBA cannot tell how much the dominated solution needs to be improved so that it becomes non-dominated. This difficulty arises because quality of the Pareto set is evaluated in terms of probability, and it cannot be overcome unless we employ EOC as a measure of selection quality. In this paper, our main purpose is to address the above issues by applying different measures of selection quality to solve the MORS problem. We will derive and compare allocation rules which optimize different measures of selection quality. The organization of this paper is given below. In Section 2, we present a general framework for solving the MORS problem. Allocation rules derived from minimization of EOC and maximization of PCS are given in Sections 3 and 4, respectively. Then the allocation rules are tested and compared in Section 5. Finally some conclusions and future research directions are summarized in Section 6.

2. A general solution framework for the MORS problem

Without loss of generality, we assume that minimization of the objectives is our goal in this study.

2.1. A Bayesian framework and the concept of dominance

We first establish the following notation:

S : The design space containing all n designs, $|S| = n$.

S_p : The observed Pareto set, i.e., the Pareto set constructed based on observed performances.

\bar{S}_p : The observed non-Pareto set, i.e., $\bar{S}_p = S \setminus S_p$.

N_i : The number of replications allocated to design i .

H : The number of performance measures for each design.

μ_i : The vector of true performance measures of design i ; $\mu_i = (\mu_{i1}, \dots, \mu_{iH})$

$\tilde{\mu}_i$: A vector of random variables having the posterior distribution of the true performance measures of design i ; $\tilde{\mu}_i = (\tilde{\mu}_{i1}, \dots, \tilde{\mu}_{iH})$.

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

$\bar{\mu}_{ik}$: The sample mean of simulation output for the k th objective of design i ; $\bar{\mu}_{ik} = \frac{1}{l} \sum_{s=1}^l \hat{\mu}_{ik}^s$.

σ_{ik}^2 : The variance of simulation output for the k th objective of design i , which is to be estimated by sample variance $\hat{\sigma}_{ik}^2 = \frac{1}{l-1} \sum_{s=1}^l (\hat{\mu}_{ik}^s - \bar{\mu}_{ik})^2$.

In this study, to make the problem more tractable, we assume that simulation outputs are independent across: (1) different replications; (2) different designs; (3) different performance measures of the same design. Therefore $\hat{\mu}_{ik}^s$, the s th simulation observation for the k th objective of design i is independent of all other $\hat{\mu}_{ik}^s$'s.

Remark. We realize that in real-life problems, the performance measures of the same design are often dependent, for example, the cost and the quality of products in product-design optimization. The purpose of the independence assumption is to make it easier to derive the allocation rules which are indeed also applicable to the case when performance measures are dependent. For a more detailed explanation and numerical illustration, refer to Lee et al. (2010).

We address the MORS problem within a Bayesian framework. For any design i , its true performance measures μ_i are unknown, and are to be estimated by observing the performance measures $\tilde{\mu}_i$ through simulation. Assume that each unknown performance measure μ_{ik} , for $k = 1, 2, \dots, H$, has a normal prior distribution, and no prior knowledge on the performance of any design is available before conducting the simulation. Given that $\hat{\mu}_{ik}^1, \hat{\mu}_{ik}^2, \dots, \hat{\mu}_{ik}^l$ 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 unknown true performance measure μ_{ik} can be described by its posterior distribution,

$$\tilde{\mu}_{ik} \sim N(\bar{\mu}_{ik}, \sigma_{ik}^2/l) \quad (1)$$

where $\bar{\mu}_{ik} = \frac{1}{l} \sum_{s=1}^l \hat{\mu}_{ik}^s$ is the sample mean of the simulation output, and σ_{ik}^2 is approximated by the sample variance $\hat{\sigma}_{ik}^2 = \frac{1}{l-1} \sum_{s=1}^l (\hat{\mu}_{ik}^s - \bar{\mu}_{ik})^2$. In this way, the comparison of two performance measures μ_{ik} and μ_{jk} becomes the comparison of two random variables $\tilde{\mu}_{ik}$ and $\tilde{\mu}_{jk}$ following the posterior distribution (1) derived from the most-recently available simulation output.

In addition to using the Bayesian framework to develop a posterior distribution for performance measure $\tilde{\mu}_{ik}$ through observation of simulation output, we also use it to approximate the predicted distribution of $\tilde{\mu}_{ik}$ if additional replications were to be allocated. We assume that, in the predicted distribution, as the simulation budget increases, $\bar{\mu}_{ik}$ and $\hat{\sigma}_{ik}^2$ do not change (Chen et al.,

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات