European Journal of Operational Research 242 (2015) 568–579

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Ranking and selection for multiple performance measures using incomplete preference information

V. Mattila*, K. Virtanen

Systems Analysis Laboratory, Department of Mathematics and Systems Analysis, Aalto University School of Science, P.O. Box 11000, FI-00076 Aalto, Finland

ARTICLE INFO

Article history:
Received 10 December 2013
Accepted 14 October 2014
Available online 23 October 2014

Keywords:
Simulation
Ranking and selection
Multi-attribute utility function
Incomplete preference information
Optimal computing budget allocation

ABSTRACT

This paper presents two new procedures for ranking and selection (R&S) problems where the best system designs are selected from a set of competing ones based on multiple performance measures evaluated through stochastic simulation. In the procedures, the performance measures are aggregated with a multi-attribute utility function, and incomplete preference information regarding the weights that reflect the relative importance of the measures is taken into account. A set of feasible weights is determined according to preference statements that are linear constraints on the weights given by a decision-maker. Non-dominated designs are selected using two dominance relations referred to as pairwise and absolute dominance based on estimates for the expected utilities of the designs over the feasible weights. The procedures allocate a limited number of simulation replications among the designs such that the probabilities of correctly selecting the pairwise and absolutely non-dominated designs are maximized.

The new procedures offer ease of eliciting the weights compared with existing R&S procedures that aggregate the performance measures using unique weights. Moreover, computational advantages are provided over existing procedures that identify non-dominated designs based on the expected values of the performance measures. The new procedures allow to obtain a smaller number of non-dominated designs. They also identify these designs correctly with a higher probability or require a smaller number of replications for correct selection. Finally, the new procedures allocate a larger number of replications to the non-dominated designs that are therefore evaluated with greater accuracy. These computational advantages are illustrated through several numerical experiments.

1. Introduction

Ranking and selection (R&S) procedures in simulation-optimization are methods for selecting the best system designs from a set of competing ones (e.g., Goldsman & Nelson, 1998). The designs are compared based on performance measures that are estimated through stochastic simulation. The procedures allocate simulation replications among the designs such that more replications are allocated to promising designs and less to non-promising ones. Thus, the performance of the promising designs is evaluated with greater accuracy, and the best designs can be selected efficiently as well as with a high level of confidence.

In this paper, two new procedures are presented for R&S problems where the best designs are selected based on multiple performance measures. The procedures extend the existing ones by aggregating the performance measures through a multi-attribute utility (MAU) function (e.g., Keeney & Raiffa, 1976) in which a decision-maker's (DM's) preference information regarding the weights that reflect the relative importance of the measures is incomplete (e.g., Hannan, 1981; Hazen, 1986; Kirkwood & Sarin, 1985; Salo & Hämäläinen, 1992; Sarin, 1977; Weber, 1987; White, Sage, & Dozono, 1984). The procedures maximize the probability of correctly selecting non-dominated designs under such information. They apply optimal computing budget allocation (OCBA) (Chen, Lin, Yücesan, & Chick, 2000) as well as multi-objective OCBA (MOCBA) (Lee, Chew, Teng, & Goldsman, 2004) which are originally designed for selecting the best design based on a single performance measure and the non-dominated designs based on multiple performance measures with a limited computing budget in terms of the number of simulation replications. OCBA and MOCBA have been previously extended in several ways (Lee, Chen, et al., 2010) but not in the manner described in this paper.

Existing R&S procedures mostly deal with a single performance measure. Surveys of such procedures appear in Bechhofer, Santner, and Goldsman (1995), Goldsman and Nelson (1998), Swisher, Jacobson, and Yücesan (2003), Kim and Nelson (2006), Kim and Nelson (2007). Surveys of simulation-optimization methods for multiple

* Corresponding author. Tel.: +358 50 560 2712.
E-mail addresses: Ville.A.Mattila@aalto.fi (V. Mattila), Kai.Virtanen@aalto.fi (K. Virtanen).

http://dx.doi.org/10.1016/j.ejor.2014.10.028
0377-2217/

performance measures (Evans, Stuckman, & Mollaghasemi, 1991; Rosen, Harmonosky, & Traband, 2008), on the other hand, acknowledge only few R&S procedures. An approach is to combine the multiple performance measures into a single one based on preference statements given by the DM and apply the procedures for a single performance measure. For instance, in Morrice, Butler, and Mullarkey (1998), and Butler, Morrice, and Mullarkey (2001), the performance measures are aggregated through a MAU function.

An alternative to the aggregation of performance measures is to select non-dominated designs without taking into account the preference information of the DM. A design is dominated if its performance is no better than the performance of another design with respect to all measures and worse with respect to at least one measure. The R&S procedures currently available for selecting non-dominated designs are variants of MOCBA (Chen & Lee, 2010, 2009; Lee, Chew, & Teng, 2007; Lee et al., 2004; Lee, Chew, Teng, & Goldsman, 2010; Teng, Lee, & Chew, 2010). In addition, one line of research considers R&S procedures for maximizing or minimizing a primary performance measure under constraints on a number of secondary performance measures (Andradottir, Goldsman, & Kim, 2005; Andradottir & Kim, 2010; Batur & Kim, 2010; Hunter & Pasupathy, 2013; Morrice & Butler, 2006).

The new procedures presented in this paper fall between the procedures that aggregate performance measures as well as the procedures that select non-dominated designs. In the new procedures, an additive MAU function is utilized. Thus, the utility reflecting the preferability of a design is calculated as the weighted sum of single-attribute utility functions describing the DM's preference for the values of individual performance measures. The weights reflect the relative importance of the measures and represent the contribution of each performance measure into the utility. The new procedures additionally allow incomplete preference information on the weights. A set of feasible weights is obtained from preference statements that are given by the DM in terms of linear constraints on the weights. The designs are compared through dominance relations established based on the estimates for the expected utilities over the feasible weights. The first procedure, referred to as MOCBA-p, identifies the non-dominated designs according to pairwise dominance (Weber, 1987). Then, a design is dominated if its expected utility is lower than the expected utility of another design with all feasible weights. In this procedure, MOCBA is applied for maximizing the probability of correctly selecting the set of pairwise non-dominated designs with a limited computing budget. The second procedure, referred to as OCBA-a, identifies the designs that are non-dominated according to absolute dominance (Weber, 1987). Then, a design is dominated if its maximal expected utility over the feasible weights is lower than the minimal expected utility of another design. OCBA is used for maximizing the probability of correctly selecting the set of absolutely non-dominated designs with a limited computing budget. MOCBA-p and OCBA-a are considered because both procedures have their advantages. With MOCBA-p, a smaller set of designs is obtained since any pairwise non-dominated design is also absolutely non-dominated. OCBA-a, on the other hand, is a somewhat simpler procedure and more straightforward to implement. Moreover, OCBA-a can be extended such that a non-additive MAU function is used.

MOCBA-p and OCBA-a do not require as strict preference statements as the existing procedures that aggregate performance measures using unique weights. The DM may not be able to provide the unique weights for the MAU function due to time pressure, lack of knowledge, fear of commitment or other reasons (Weber, 1987). Then, the existing procedures are not applicable whereas the new procedures are.

MOCBA-p and OCBA-a also provide several computational advantages over the existing procedures that do not aggregate performance measures. First, by selecting pairwise and absolutely non-dominated designs, the new procedures allow to obtain a smaller number of designs that remain after the simulations than the existing procedures. Therefore, a smaller set of designs remains to be compared by the DM which makes the selection of the most preferred design easier. Second, the procedures identify pairwise and absolutely non-dominated designs correctly with a higher probability or with a smaller computing budget compared with an approach in which non-dominated designs are identified first and preference information is utilized after the simulations. Thus, MOCBA-p and OCBA-a either increase the confidence in correct selection or provide computational savings. Third, because MOCBA-p and OCBA-a concentrate on a smaller set of designs, they allocate a larger number of simulation replications to the designs that remain after the simulations and evaluate such designs more accurately than the existing procedures. These computational advantages are illustrated through numerical experiments in this paper. Moreover, the effect of the set of feasible weights as well as the effect of the number of performance measures on the advantages is examined in the numerical experiments in order to assess the applicability of the procedures.

Two earlier R&S procedures utilizing MAU functions take into account incomplete preference information regarding weights using probability distributions (Branke & Gamer, 2007; Frazier & Kazachkov, 2011). In these procedures, the estimate for the expected utility of each design is calculated uniquely which leads to an R&S problem with a single performance measure. The weights of the MAU function are interpreted as the contribution of each performance measure into the utility of a design. However, there is no such interpretation for the probability distributions. Therefore, these distributions may be challenging to provide. In turn, constraints for the weights utilized in this paper can be determined based on this interpretation. Thus, they provide a more transparent representation of the incomplete preference information.

The rest of the paper is organized as follows. Section 2 discusses existing procedures for R&S with multiple performance measures which are based on optimal computing budget allocation. This discussion helps to introduce the use of MOCBA and OCBA as well as the use of MAU functions in the new procedures. Section 3 describes the MAU function with incomplete preference information on weights as well as pairwise and absolute dominance relations. In Section 4, the new MOCBA-p and OCBA-a procedures are presented. The numerical experiments illustrating the computational advantages of the procedures are presented in Section 5. Finally, concluding remarks are given in Section 6.

2. Existing procedures based on optimal computing budget allocation

In this section, the MOCBA (Chen & Lee, 2010) procedure is summarized first. Second, a procedure in which performance measures are aggregated with a MAU function and the design with the highest expected utility is selected using OCBA is described (Chen & Lee, 2010).

The R&S problem with multiple performance measures is formally described as follows. There are K designs and n performance measures. $\mathbf{X}_k = (X_{k1}, \ldots, X_{kn})$ denotes a vector of independent random variables where X_{ki} represents the performance of the kth design with respect to the ith performance measure. Now, the designs that minimize the expected values $E[X_{ki}]$ of the performance measures are selected, i.e.,

$$\min_{k \in \{1,\ldots,K\}} \left(E[X_{k1}], \ldots, E[X_{kn}]\right). \tag{1}$$

Moreover, $E[X_{ki}]$ are estimated by generating realizations of the measures through replications of a stochastic simulation model. Thus, the replications must be allocated among the designs such that the ones minimizing $E[X_{ki}]$ are selected correctly with high confidence and as few replications as possible are performed.

Which designs are the solution of Problem (1) depends on how the multiple performance measures are handled. When using MOCBA, the