European Journal of Operational Research 000 (2014) 1–11

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Refined ranking relations for selection of solutions in multi objective metaheuristics

Ruby L. V. Moritz, Enrico Reich, Maik Schwarz, Matthias Bernt, Martin Middendorf*

Department of Computer Science, University of Leipzig, Leipzig, Germany

ARTICLE INFO

Article history:
Received 17 September 2013
Accepted 21 October 2014
Available online xxx

Keywords:
Multi objective optimization
Ranking relations
Ant colony optimization
Genetic algorithms
Flow shop scheduling problem

ABSTRACT

Two methods for ranking of solutions of multi objective optimization problems are proposed in this paper. The methods can be used, e.g. by metaheuristics to select good solutions from a set of non dominated solutions. They are suitable for population based metaheuristics to limit the size of the population. It is shown theoretically that the ranking methods possess some interesting properties for such applications. In particular, it is shown that both methods form a total preorder and are both refinements of the Pareto dominance relation. An experimental investigation for a multi objective flow shop problem shows that the use of the new ranking methods in a Population-based Ant Colony Optimization algorithm and in a genetic algorithm leads to good results when compared to other methods.

1. Introduction

Optimization problems with several objectives are common in industrial, engineering, or scientific contexts. Abstractly formulated, a *multi objective optimization problem* asks for solutions from a solution space X (also called search space) that are optimal with respect to $d > 1$ objectives. Typically, the objectives are conflicting and it is only possible to optimize a (small) subset of the objectives simultaneously. This is reflected in the concept of the *Pareto dominance* relation. A solution dominates another if it is better in at least one objective and not worse in all other objectives. Based on this relation a solution is called *Pareto optimal* if it is not dominated by any other solution from X. The goal in multi objective optimization is to find the set of Pareto optimal solutions, called Pareto set, or at least a subset of it. Unfortunately, many discrete multi objective optimization problems are NP-hard, i.e. it is not possible to solve them in polynomial time (if $P \neq NP$). For continuous multi objective optimization problems the optimization function is often a black box or algebraically too complicated, i.e. it is impossible to solve them analytically. In these cases, usually the goal is to find a set of non dominated solutions that are close to the Pareto optimal solutions. Another common selective feature is the diversity of the chosen set of solutions.

To explore the solution space the use of algorithms maintaining a population of solutions seems beneficial. Therefore, multi objective variants of various population-based metaheuristics have been developed in recent years. Among the earliest multi objective variants of genetic algorithms is, e.g. the Vector Evaluation Genetic Algorithm (VEGA) from Schaffer (1984) (see overviews in Coello Coello, 2009; Coello Coello, Pulido, & Montes, 2005; Fonseca & Fleming, 1995). Also the ant colony optimization metaheuristic (ACO) was extended to solve multi objective optimization problems, e.g. by Iredi, Merkle, and Middendorf (2001) and Doerner, Hartl, and Reimann (2001) (for an overview see Angus & Woodward, 2009 or Leguizamón & Coello Coello, 2011, Chapter 3). For problems on continuous solution spaces extended particle swarm approaches (PSO) have been proposed, e.g. by Coello Coello and Salazar Lechuga (2002) and Hu and Eberhart (2002) (for an overview see Reyes-Sierra & Coello Coello, 2006).

The Pareto dominance relation is often used in population-based metaheuristics to rank the elements of a set of solutions. But for a set of non dominated solutions this relation is insufficient to guide heuristics into favourable regions of the search space. This is particularly relevant, when the number of objectives is large since the number of non dominated solutions increases along with the number of objectives. This effect has been called a “curse of dimensionality” by Kukkonen and Lampinen (2007). Therefore, several methods to compare and rank elements of sets of (non dominated) solutions have been discussed in the literature. Some of these ranking methods are used for post-Pareto optimality, i.e. to select solutions from the final set of non dominated solutions computed by some metaheuristic.

* Corresponding author. Tel.: +49 341 9732275.
E-mail addresses: rmoritz@informatik.uni-leipzig.de (R. L. V. Moritz), enrico.reich@gmx.de (E. Reich), bernt@informatik.uni-leipzig.de (M. Bernt), middendorf@informatik.uni-leipzig.de (M. Middendorf).

http://dx.doi.org/10.1016/j.ejor.2014.10.044

Aside from ranking the found solutions, there are other methods to prune a set of (non dominated) solutions. One possibility is to apply clustering methods to the set of solutions in order to select a representative solution from each cluster. One potential aim of using clustering methods is to obtain a large diversity in the set of selected solutions (Kukkonen & Deb, 2006).

In this paper we are interested in using ranking methods for meta-heuristics by applying them in each iteration to select only a few solutions from the set of already found (non dominated) solutions. Hence, the ranking methods are used to prune a potentially large set of solutions to a small subset of good solutions. This small set forms the population which is then used in the next iteration to create new and hopefully better solutions. Maintaining a small population has advantages with respect to the memory and time requirements of an algorithm. This is particularly important in restrictive computational environments, e.g. when solutions need to be delivered in real time or the algorithm runs on specific hardware, e.g. as part of an embedded system. Another important reason for the use of small populations is objective functions that are expensive to evaluate. Clearly, the choice of which solutions are to be kept in the population is particularly important for small population sizes.

The simplest kind of ranking schemes is aggregation methods. They use, for instance, an aggregation function calculating a weighted sum of the objectives (Jakob, Gorges-Schleuter, & Blume, 1992) or the distance to a target vector of objective values (Wienke, Lucasius, & Kateman, 1992). Furthermore, there are also objective based ranking methods. One such method is to use given priorities for a lexicographic sorting of the objectives or performing a separate optimization of the single objectives in an order that is compliant with the priorities (Cvetkovic & Parmee, 2002). Weights, target vectors, and priorities should be given by the user. But for many applications this might be difficult or impossible for the user because, for instance, no reasonable weights are known.

In contrast to the above ranking methods that need different kinds of user input, several methods have been developed based on the notion of dominance. To rank a solution a within a set of solutions X one could use the number of solutions in X that are dominated by a (*dominance rank*) or the number of solutions in X dominating a (*dominance count*) (see also Fonseca & Fleming, 1993; Zitzler & Thiele, 1999). An alternative is to use the number of times the current set of non dominated solutions has to be removed from the remaining set of solutions until the solution itself becomes non dominated. This measure has been called *dominance depth* by Srinivas and Deb (1994). Additionally to the comparison of single objective vectors also sets of non dominated solutions can be compared. This is commonly done in the indicator function framework, which assigns a real value to a set or sets of non dominated solutions. One example is the binary hypervolume indicator that is defined as the hypervolume of the objective space (i.e. the space of vectors of potential values of the different objectives) dominated by one set of solutions but not by another (we refer the interested reader to Zitzler, Thiele, Laumanns, Fonseca, & da Fonseca, 2003).

Some ranking methods consider how much two solutions differ with respect to the different objectives. Garza-Fabre, Toscano Pulido, and Coello Coello (2009) have proposed three such methods. An example is the *Global Detriment method* which computes for each solution the sum of the differences to other solutions over all those objectives where the solution is worse than the corresponding other solution. A potential disadvantage of these methods is that they typically need some normalization between the different objectives. Therefore, these methods are not considered further in this paper. However, we use one representative of these methods as a comparison method. This method uses the same measure as the Global Detriment method, but only for pairwise comparison, and random weights for the normalization (for details see Section 3.2).

By regarding only the information whether two objective values differ and thereby disregarding the magnitude of the difference of the objective values the decision process can be greatly simplified. The number of objectives for which one solution is better or worse than the other can be helpful for a decision between alternative solutions. A similar reasoning simplifies everyday decisions when comparing different choices by the number of advantages and disadvantages. This basic idea was formally captured by the relation *favour* that prefers a solution over another if it wins, i.e. is better, in more objectives than the other (Drechsler, Drechsler, & Becker, 2001). Several extensions or similar relations have been proposed in the literature. They modify the decision if an objective is counted as won or lost (Laumanns, Thiele, Deb, & Zitzler, 2002; Sülflow, Drechsler, & Drechsler, 2007) or how many won objectives are required for preference (Farina & Amato, 2004; Zou, Chen, Liu, & Kang, 2008). Approaches based on the number of won objectives can also be used to rate a solution with respect to a set of solutions e.g. (Bentley & Wakefield, 1998; Maneeratana, Boonlong, & Chaiyaratana, 2006; Mostaghim & Schmeck, 2008). A detailed description of these methods is given in Section 3.

In this paper we propose two new ranking relations which are based on the relation favour (Section 4). We prove that both relations are refinements of the Pareto dominance relation and are total preorders. The ranking relations can be used in multi objective population based meta-heuristics for selecting the solutions that are included into the population and thus guide the search for better solutions. We compare the different ranking schemes when used in a Population-based Ant Colony Optimization algorithm (P-ACO) and in a genetic algorithm (GA). As test problem a multi objective flow shop scheduling problem is used. The results show that the new ranking schemes are advantageous to select the solutions for the population of the P-ACO and the GA.

The paper starts with the introduction of basic definitions in Section 2. A detailed review of related ranking relations from the literature and a few corresponding theoretical results are presented in Section 3. The new ranking relations are introduced in Section 4. The metaheuristics P-ACO and GA which are used for the experiments are described in Section 5. The test problem and the experiments are described in Section 6. The experimental results are presented in Section 7. A short conclusion is given in Section 8. Note, that this paper is an extension of Moritz, Reich, Schwarz, Bernt, and Middendorf (2013).

2. Basic definitions

Consider a multi objective optimization problem with a set of solutions X and a vector of *objective functions* $\vec{f}(a) = (f_1(a), \ldots, f_d(a))$, where $a \in X$ and $f_i : X \mapsto \mathbb{R}$. The solutions in X can, for example, be vectors of real values in case of continuous optimization problems or vectors of elements from a finite set in case of combinatorial optimization problems. The aim is to find solutions from X that minimize the objectives, i.e.

$$\min_{a \in X} \vec{f}(a) = \min_{a \in X} (f_1(a), \ldots, f_d(a)). \quad (1)$$

Note, that by using $-f_i(a)$ maximization is also possible.

In order to find a minimum in a two or higher dimensional space the Pareto dominance relation ($\prec$) is used. Let $a, b \in X$, then

$$a \prec b \iff \forall i \in [1:d] : f_i(a) \le f_i(b) \wedge \vec{f}(a) \ne \vec{f}(b). \quad (2)$$

Solution a *dominates* b if $a \prec b$. Note that, if $a \prec b$ then there is at least one $i \in [1:d]$ with $f_i(a) < f_i(b)$. Two solutions $a, b \in X$ are called incomparable if $a \nprec b \wedge b \nprec a$ or indifferent in case of $\vec{f}(a) = \vec{f}(b)$. A solution $a \in X$ is called *Pareto optimal* if $\nexists b \in X : b \prec a$. A solution $a \in X$ is called *non dominated* solution with respect to a subset $X' \subseteq X$ if $\nexists b \in X' : b \prec a$. The set of all Pareto optimal solutions from X is called the *Pareto set* and the corresponding set of objective vectors in $\mathbb{R}^d$ is called the *Pareto front* of X.