Robust Multi-Objective Optimization in High Dimensional Spaces

André Sülflow, Nicole Drechsler, and Rolf Drechsler

Institute of Computer Science
University of Bremen
28359 Bremen, Germany
{suelflow,nd,drechsle}@informatik.uni-bremen.de

Abstract. In most real world optimization problems several optimization goals have to be considered in parallel. For this reason, there has been a growing interest in Multi-Objective Optimization (MOO) in the past years. Several alternative approaches have been proposed to cope with the occurring problems, e.g. how to compare and rank the different elements. The available techniques produce very good results, but they have mainly been studied for problems of "low dimension", i.e. with less than 10 optimization objectives.

In this paper we study MOO for high dimensional spaces. We first review existing techniques and discuss them in our context. The pros and cons are pointed out. A new relation called ϵ-*Preferred* is presented that extends existing approaches and clearly outperforms these for high dimensions. Experimental results are presented for a very complex industrial scheduling problem, i.e. a utilization planning problem for a hospital. This problem is also well known as *nurse rostering*, and in our application has more than 20 optimization targets. It is solved using an evolutionary approach. The new algorithms based on relation ϵ-*Preferred* do not only yield better results regarding quality, but also enhances the robustness significantly.

1 Introduction

To solve complex optimization problems today, it is often not sufficient to only consider a single optimization criteria. In contrast, many real world problems have several – often contradicting – optimization goals. Thus, in the recent past several techniques for *Multi-Objective Optimization* (MOO) have been proposed.

One of the first approaches in this direction was the use of *Pareto-optimal elements*. This has been discussed in the context of *Evolutionary Algorithms* (EAs) in [1]. The goal is to determine elements from the *Pareto set*. To guide this search, there exist several alternative methods (see e.g. [2,3]) where the core is a relation that allows to compare different elements. E.g. the relation *Dominate* proposed in [1] can be applied. These methods are well known and have been studied intensively. But so far these studies mainly consider problems with a small number of optimization criteria, e.g. in [2] comparisons for dimensions two or three are given.

S. Obayashi et al. (Eds.): EMO 2007, LNCS 4403, pp. 715–726, 2007.

For higher dimensional spaces there only exist a few studies (see e.g. [4,5,6,7]). As testcases scalable test functions proposed in [4] are considered. For example, in [3] it is reported that the number of individuals in the Pareto set, i.e. the non-dominated solutions, increase with the number of optimization objectives. Experiments have shown that for 20 objectives the percentage of solutions that cannot be distinguished using relation *Dominates* in random populations is nearly 100%. For this reason, new measures and relations have to be defined that help to automate and guide the optimization process.

In general for higher dimensions *weighted sums* or *aggregation* have been proposed, since they are easy to describe and, on a first sight, scale well. But for high dimensions these techniques reach their limits, since it is hard (or even impossible) to determine good weights or the fitness of the optimal solution is not known in advance, respectively.

In [8,9] an alternative relation called *Preferred* (originally introduced as *Favour*) has been proposed and applied for five dimensions. Experiments have shown that *Preferred* clearly outperformed relation *Dominates* and an approach based on weighted sums. But in all cases described above the dimensions considered are rather small, i.e. less than 10. However, for complex optimization problems, where especially EAs are frequently used, often a higher number of dimensions occur. Of course, the standard algorithms can also be applied in the case of higher dimensions, but it will be shown in this paper by a detailed discussion and also by experimental studies for an industrial application that other techniques should be applied.

In this paper we first discuss the existing techniques and point out their main properties. Then, an experimental study shows the weaknesses of the above techniques for higher dimensions. For the experiments an industrial application where a very complex scheduling problem with many constraints occurs is considered. I.e. the *nurse rostering problem* [10], a well-known problem in mixed integer optimization, where a highly constraint schedule for employees in a hospital is generated. In this problem, 25 optimization goals are considered in parallel. As an additional difficulty, there are different types of constraints. Some can be seen as "hard constraints" that are enforced by state laws, while others are "soft constraints" that should be fulfilled as good as possible. It is demonstrated by experiments that for high dimensions the approach using relation *Preferred* outperforms traditional methods based on non-dominated sorting (relation *Dominates*). But relation *Preferred* is not robust for high dimensions and has to be extended accordingly. Therefore, we propose an extension of *Preferred* that also takes the relative difference over all dimensions into account. It considers environments of radius ϵ, where elements outside this region are "punished". The new relation is called ϵ-*Preferred*. Experimental results show that the new approach results in higher quality and, additionally, gives very robust optimization results.

The paper is structured as follows: In Section 2 previous work is reviewed and properties of the different relations are discussed. An experimental study for a complex scheduling problem is presented in Section 3. This study clearly