

**Reference Point Based Multi-Objective Algorithms for
Group Decisions**

Jella Pfeiffer, Uli Golle and Franz Rothlauf

Working paper 2/2008
January 2008

**Working Papers in Information Systems
and Business Administration**

Johannes Gutenberg-University Mainz

Department of Information Systems and Business Administration

D-55128 Mainz/Germany

Phone +49 6131 39 22734, Fax +49 6131 39 22185

E-Mail: sekretariat[at]wi.bwl.uni-mainz.de

Internet: <http://wi.bwl.uni-mainz.de>

Reference Point Based Multi-Objective Algorithms for Group Decisions

Jella Pfeiffer
Dept. of Information Systems
and Business Administration
University of Mainz
jella.pfeiffer@uni-
mainz.de

Uli Golle
Dept. of Information Systems
and Business Administration
University of Mainz
golle@uni-mainz.de

Franz Rothlauf
Dept. of Information Systems
and Business Administration
University of Mainz
rothlauf@uni-mainz.de

ABSTRACT

While in the past decades research on multi-objective evolutionary algorithms (MOEA) has aimed at finding the whole set of Pareto optimal solutions, current approaches focus on only those parts of the Pareto front which satisfy the preferences of the decision maker (DM). Therefore, they integrate the DM early on in the optimization process instead of leaving him/her alone with the final choice of one solution among the whole Pareto optimal set. In this paper, we address an aspect which has been neglected so far in the research on integrating preferences: in most real-world problems, there is not only one DM, but a group of DMs trying to find one consensus decision all participants are willed to agree to. Therefore, our aim is to introduce methods which focus on the part of the Pareto front which satisfies the preferences of several DMs concurrently. We assume that the DMs have some vague notion of their preferences a priori the search in form of a reference point or goal. Thus, we present and compare several reference point based approaches for group decisions and evaluate them on three ZDT and two flow shop problems.

Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search

General Terms

Algorithms, Design

Keywords

Multi-objective optimization, preference-based optimization, decision making, group decisions, reference points.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'08, July 12–16, 2008, Atlanta, Georgia, USA.
Copyright 2008 ACM 978-1-60558-131-6/08/07 ...\$5.00.

1. INTRODUCTION

In many real-world optimization problems, several conflicting objectives have to be optimized simultaneously. In business problems, for instance, there is usually a trade-off between the objectives time, cost, and quality. Since no single optimal solution can be found when objectives are conflicting, multi-objective optimization methods try to find a set of equivalent solutions. Solutions are considered as equivalent when they are Pareto optimal. Evolutionary algorithms (EA) have been applied successfully to multi-objective problems as the population based concept is particularly suitable for finding a set of solutions which approximate the Pareto front in a single run while ensuring a certain diversity of solutions along this front. The advantage of approximating the whole Pareto front is that no other information about the decision maker's (DM) preference is needed prior to the search process [14, 5]. Therefore, most multi-objective evolutionary algorithms (MOEA) work with minimum information about preferences and postpone the decision process after the search process (a posteriori approach).

This a posteriori approach of first computing the complete Pareto front and then letting the DM chose the preferred solution out of the Pareto optimal set has two main disadvantages. First of all, especially for problems with many objectives, the set of solutions representing the Pareto front becomes large, which impedes the choice for the DM. In the MOEA literature, this a posteriori decision process is usually not addressed (for one exception see [9]). Secondly, in most cases the DM has some ideas about his/her preferences prior to the optimization process [3]. Consequently, computing the whole Pareto front is a waste of resources if an articulation of preferences a priori or during the optimization process can help to focus the search on relevant parts of the Pareto front.

Recently, several approaches have been presented that use preferences to focus on parts of the Pareto front early on in the search process [6]. Although this recent development in the research community is very promising, a crucial fact has been ignored until now in MOEA research: many decisions in real-world problems do not only depend on the preferences of one single DM but on several DMs [26]. For instance, in flow shop problems, the DMs representing different departments of a company attach a different importance to each objective like makespan, tardiness, etc. While the production division prefers to maximize their output and therefore wants to minimize the makespan, the sales division aims at

a high customer satisfaction and thus wants to minimize the tardiness. Another example is the optimization of supply chains, where the different partners in the chain have different conceptions of which performance criteria are more important. Although there is a lack of research in MOEA literature on group decisions, research on group decision support systems has addressed this topic profoundly. However, this research mostly assumes the following two points which do not hold for many real-world problems. First, in many cases it is assumed that there is only a small amount of possible alternatives among which a group consensus must be found. Optimization in a huge alternative space is less considered. Second, usually the group is assumed to be homogeneous so that the aggregation into one common preference function is possible.

Against this background, the aim of this paper is to address the lack of research in MOEA for group decision problems. We suggest several techniques for integrating group preferences into a MOEA, and show how they support the group in finding a consensus. Thereby, we assume that the group is heterogeneous and the group members have some vague notion about their preferences in form of reference points or goals. We focus the search on a small set of Pareto optimal solutions which are preferred by the DMs and test the proposed approaches on ZDT problems and flow shop problems. Besides our aim to present possible methods and evaluate them in experiments, we hope to arouse the interest of other researchers in the field of MOEAs to address multi-objective group decision problems.

The following section provides a literature review on existing MOEA approaches for incorporating preferences of a DM into a MOEA. Furthermore, section 3 discusses how and when in the search process of a MOEA preferences of groups can be incorporated. Thereafter, we propose four different ways of finding consensus decisions with a MOEA in a reference point based approach. In the experiments in section 5 we evaluate the proposed methods with different reference points and parameter settings on several test problems. The paper closes with concluding remarks.

2. PREFERENCES IN MOEA

Recently, several studies in the field of MOEA have been published which do not aim at approximating the whole Pareto front. Instead, they restrict their search to those Pareto optimal solutions which are preferred by the DM. These approaches often combine methods to elicit preferences from multi-criteria decision making by integrating the DM a priori or during the search process.

One of the first attempts in MOEA allows the DM to interactively choose satisfactory and unsatisfactory solutions, i.e. define goals as well as worst acceptable levels for each objective [21, 22]. A goal or reference point is some kind of preferred point in the solution space which defines an aspiration level for each objective. A similar approach was presented at about the same time by Fonseca and Fleming [14]. They propose a ranking scheme which gives a greater importance to objectives which do not satisfy a goal. The goal can be adjusted and refined interactively by the DM such that more and more solutions are excluded from the search. An a priori MOEA motivated by goal programming was proposed by Deb [8]. Deb takes the absolute deviations from solutions to the goal for each criteria as objective functions and minimizes them with NSGA-II [11].

Other studies aim at estimating the DM's preferences interactively, for example by presenting the DM regularly with a number of solutions which he/she has to evaluate [25]. This information is then used for training an Artificial Neural Network (ANN) which approximates the utility function of the DM. This approximation is used to select a subset of the so-called Pareto population which guides the search process towards the preferred regions of the search space. However, the authors report that the ANN is not truly reflecting the DM preferences due to early convergence of the ANN and inconsistencies of the provided preference information. A further idea in this direction was presented by Phelps and Köksalan [16]. Several times in the course of the search, promising solutions are presented to the DM for pairwise comparison. This information is then used to estimate the utility function by applying the middlemost weights technique [15]. The population is ranked using this estimated function and selection is partially influenced by the ranking.

Other approaches explicitly redefine the concept of Pareto dominance. Branke et al. [4] construct a minimal/maximal linear utility function by asking the DM about their minimal/maximal trade-offs ("how many units of A would you at least/most be willed to trade-off against one unit of B"). The so retrieved information is integrated into a new, so-called *guided dominance principle* where the slope of the borders of the dominated area corresponds to the slope of the maximal/minimal utility function. As a consequence, solutions dominate a larger set in the solution space compared to the original definition of dominance and therefore exclude some solutions from the original Pareto front. Since one disadvantage of the method is that in an m -objective space the DM has to specify $m^2 - m$ trade-offs, Cvetkovic and Parmee [7] focus on the minimization of the cognitive overload for the DM. They aim at minimizing the number of questions posed to the DM while allowing him/her to express their preferences in a fuzzy manner. The weights interfered from the fuzzy preferences are then used to define a weighted dominance relation. However, as [3] mention, the approach only biases the search in a coarse manner since the new dominance scheme only considers whether a solution is better than another one, and not by how much it is better.

In a paper by Branke and Deb [3], the guided dominance principle is compared with the idea of biasing the crowding distance measure. The crowding distance is a concept in NSGA-II which measures the distance between solutions and thereby ensures the diversification of solutions along the Pareto front. The authors improve an approach suggested earlier by Deb [9] where he biases the crowding distance with a weight for each objective provided a priori by the DM. While with that approach some regions of the Pareto front cannot be focused, the improved version in [3] allows a better control of the location and the expansion of the preferred region. The authors project the solutions of the Pareto optimal front onto a linear preference function. The new biased crowding distance sets the distance of projected points (d') in relation to their original distance on the Pareto front (d). Solutions which are in a region on the Pareto front having a similar slope to the utility function and therefore being preferred by the DM, produce a d' similar to d whereas for solutions in regions with a large difference in slope d' is much smaller than d . Hence, regions out of interest are artificially crowded and less solutions are allowed to be located there. For a more detailed overview of incorporating preferences

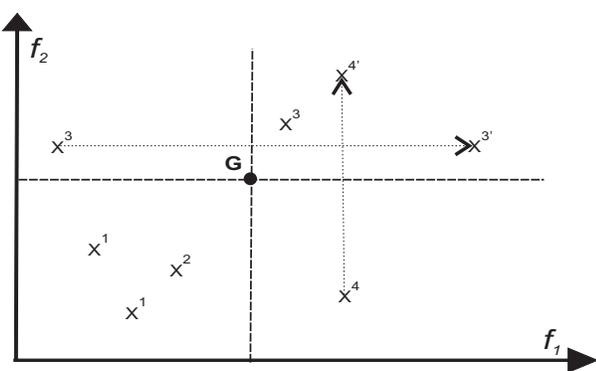


Figure 1: The two-stage Pareto ranking scheme by Tan et al. (1999).

into MOEA, the reader is referred to [6].

In all presented studies, it is assumed that one single preference is given, for instance one linear function or one goal. The question remains how to deal with preferences of several DM which all have to be satisfied at the same time. Lately, there have been a few studies which at least mention that they easily can be extended to the specification of several preferences. However, they assume only one DM and therefore do not deal with the problem of how to support several DMs in finding a consensus solution. Thiele et al. [23] use a reference point method and combine this information with the fitness function in an indicator-based evolutionary algorithm [32]. They mention that with minor changes their method can be adapted to search in regions belonging to several reference points. A similar remark is made by [10] who point out that their proposed reference direction method can be easily extended to handle multiple reference directions simultaneously.

To our best knowledge, there are only two MOEAs which explicitly deal with several preferences. The two-stage Pareto ranking scheme [20] considers both the case where a consensus among goals must be found as well as the case where all regions close to one goal are focused separately. Although the authors do not address explicitly group decisions, their approach could be applied to that case. Yet, as we will see in section 4, this approach is very coarse and therefore fails to focus the search for certain combinations of goals. In a recent paper, Deb and Sundar [12] present a reference point based method which, however, does only find several subsets of Pareto optimal solutions close to each supplied reference point and no consensus. In section 4, we will elaborate on both approaches in more detail. We will show how the reference point based algorithm can be modified and extended such that it meets our purpose, and we will indicate the deficits of the two-stage Pareto ranking scheme.

3. PREFERENCES IN GROUP DECISIONS

In analogy to scenarios with one DM, methods to integrate preferences of a whole group of DM can be classified according to the time when the preferences are articulated [27]: a priori, progressive, or a posteriori the search. In contrast to problems with one single DM, in problems with several DM there must be some kind of aggregation of preferences of all participants so that, finally, one solution can be chosen. This difference has a huge impact on the time

of the preferences' incorporation. In the a priori approach, for instance, the decision maker could agree through discussion on one common preference function. That would be a quite comfortable way from the viewpoint of the optimization process since all approaches considering only one DM (see section 2) could be applied. Furthermore, many techniques for supporting this process of synthesizing individual decisions have been established, like the nominal group technique, Delphi techniques, the dialectical approach, and the analytic hierarchy process (AHP) [29]. Yet, the agreement to one common preference function often is unrealistic. As Zahir et al. [30] point out "in an intermediate-sized group or in a large group, this homogeneity can be neither guaranteed nor achieved".

The second approach of progressively integrating preferences demands a high involvement of the DM. Therefore, in the case of a whole group of DMs, this approach is a very demanding and time-consuming one for the participants. In the a posteriori approach, the Pareto-front would be presented to the group which then has to find a consensus again in discussions or with the aid of one of the already mentioned tools, like AHP etc.

4. FINDING CONSENSUS WITH MOEA

In this section, we propose four variations of the reference point based approach by [12] which are able to find a group consensus. The presented approach focuses on relevant parts of the Pareto front and still presents the DMs with a range of several solutions at the end of the optimization process. We think that this procedure addresses the topics discussed in section 3 adequately. On the one hand, it supports the group's decision by improving only solutions which are of interest for the group and, on the other hand, it leaves a certain degree of autonomy and flexibility to the group to make the final decision among several alternatives. This flexibility is important since we do not know whether all preferences are of equal importance or whether one of the DM has more power of decision than others.

Coello [5] points out that when addressing several DMs, one must be aware that the aggregation of individual preferences into a group preference has several negative consequences due to *Arrow's Impossibility Theorem* [1]. Arrow specified several requirements of social voting systems which should be satisfied when constructing a group solution, and showed that no voting system can fulfill all those requirements at the same time. This theorem was originally stated for ordinal preferences. Further works show that the impossibility results also hold for cardinal utility functions if these functions are not interpersonally comparable [17]. Since we propose a method where preferences are not aggregated into one single preference function, we presume that the negative results of the impossibility theorem do not hold for our approach. Nevertheless, this conjecture needs to be proved in future work.

As pointed out in section 2, there are two existing approaches which consider several preferences, the two-stage Pareto ranking scheme and the reference point based approach. We will now first introduce both approaches, then we indicate the deficits of the two-stage ranking scheme and finally we show how the reference point based algorithm can be modified and extended such that it meets our purpose.

4.1 Two-stage Pareto ranking scheme

Tan et al. [20] suggest the following two-stage Pareto ranking scheme computed for each of the m goals $G_k, k = 1, \dots, m$ specified by the DM. A goal is equivalent to a reference point and defines (aspiration) levels for each objective the DM would like to achieve. In the first stage, only solutions satisfying G_k get allocated a Pareto rank. That is, all non-dominated solutions get rank 1 and all others get the rank according to the number of solutions in the population dominating them plus 1. In the second stage, all solutions not satisfying G_k in at least one criteria are ranked. Let F_a^x and F_b^x denote the components of vector F_a and F_b respectively in which F_a does not meet the goal G_k . Then F_a dominates F_b :

$$F_a \prec_{G_k} F_b \text{ if and only if } F_a^x \prec F_b^x \text{ or } |F_a - G_k| \prec |F_b - G_k|$$

The rank value in this second stage starts from the maximum rank of the first stage plus 1. Using the difference $|F_a - G_k|$ can be interpreted as moving a solution x^p to $x^{p'}$ as shown in figure 1. There, all solutions which are not located in the left bottom quadrant do not satisfy the goal G and are therefore transformed in the second ranking stage into the right upper quadrant. Tan et al. propose the following two approaches in case of several goals. They concatenate the ranks of all goals G_k to a solution F ($rank(F, G_k)$) with boolean operations *OR* or *AND*,

$$OR : \min_{k=1, \dots, m} \{rank(F, G_k)\} \quad (1)$$

$$AND : \max_{k=1, \dots, m} \{rank(F, G_k)\}. \quad (2)$$

While the *OR* operation evolves the population toward either of the goals, the *AND* operation tries to minimize the deviation from all goals concurrently and therefore attempts to find a consensus.

4.2 Reference point based method

The reference point based method proposed by Deb and Sundar [12] seizes the suggestion of the *OR* operation. It rests upon achievement scalarizing functions proposed in [28]. In its simplest version, the achievement scalarizing function minimizes the maximal weighted distance of solutions $f(x)$ towards the reference point r and therefore turns the d -objective problem into a single-objective one.

$$\text{Minimize } \max_{i=1, \dots, d} [w_i (f_i(x) - r_i)]. \quad (3)$$

Based on this approach, two MOEAs were suggested, an interactive [23] and an a priori [12] approach. Due to the arguments pointed out in section 3 we will focus on the a priori approach [12]. In this approach a new way of calculating the crowding distance in NSGA-II and still preserving diversification along the front is proposed. The crowding distance is taken to differentiate between solutions belonging to the same front (the same non-domination rank). In this case, a tournament selection selects the individual with a low crowding distance d_c .¹ The following three steps describe the modifications which are necessary to adopt NSGA-II to the reference point method [12]:

¹In the original NSGA-II crowding distances of higher value are preferred while in the reference point based approach crowding distances of lower value are favored.

- 1. Ranking** For each reference point, a list of the normalized Euclidean distances d_i to each solution of the current front is sorted in an ascending manner and then ranked.

$$d = \sqrt{\sum_{i=1}^M \frac{1}{M} \left(\frac{f_i(x) - r_i}{f_i^{max} - f_i^{min}} \right)^2},$$

where $f_i(x)$ is the i -th objective value of the current solution and f_i^{max} and f_i^{min} are the population maximum and minimum values of the i -th objective. The solution closest to the reference point is assigned rank one. Each solution has therefore several rank values, each corresponding to a different reference point.

- 2. Min** The minimum of the ranks assigned to a solution is taken as its crowding distance $d_c = \min_{k=1, \dots, m} [rank_k]$ where m is the number of reference points. Therefore, a solution closest to one of the reference points is assigned the lowest crowding distance of one.
- 3. Grouping** To achieve a diversification of solutions, all solutions having a sum of normalized difference in objective values of ϵ or less between them are grouped. To allow only one random solution of each group to be kept in the next generation, all others are penalized with a high crowding distance.

From our viewpoint, the two-stage Pareto ranking and the reference point based method incorporate concepts which are valuable when a group consensus must be found. Using the max rank in the *AND* operation helps to find solutions which are acceptable to all group members and should therefore be tested further. However, there are constellations of reference points (goals) where the two-step ranking fails to focus the search on the relevant parts of the search space. Such an example is shown in fig. 2(a) on a slightly modified version ZDT1* of the ZDT1 problem (see section 5.1). The two reference points are set to $[0.1, 0.5]$ and $[0.5, 0.1]$. All solutions of the Pareto front are assigned the same rank, therefore no part of the front is focused. That is because there is no differentiation between solutions which are closer to the reference points. Instead, a pure Pareto dominance concept is used as method for assigning rankings. We think that the Euclidean distances to the reference points should have an effect on the evaluation of solutions as they represent a point in the solutions space where the DM wants solutions to be located. Furthermore, finer graduations of ranks or crowding distances are needed to control the focus on the Pareto front better. Concerning the reference point based method, we will see later that it is easily adoptable to the task of consensus finding. However, in the current version it focuses on each reference point separately (see figure 2(b)) and therefore does not find a set of consensus solutions.

4.3 Proposed Consensus Finding Approaches

We propose four different variations on how the reference point based method as suggested in [12] is applied to group consensus finding. In the *ranking based approach*, only the second step of the reference point method must be adopted, while for the *distance based approach* the first and second step are modified.

Ranking based approach.

In the ranking based approach, the second step (Min) (see section 4.2) can be changed in two different ways. First,

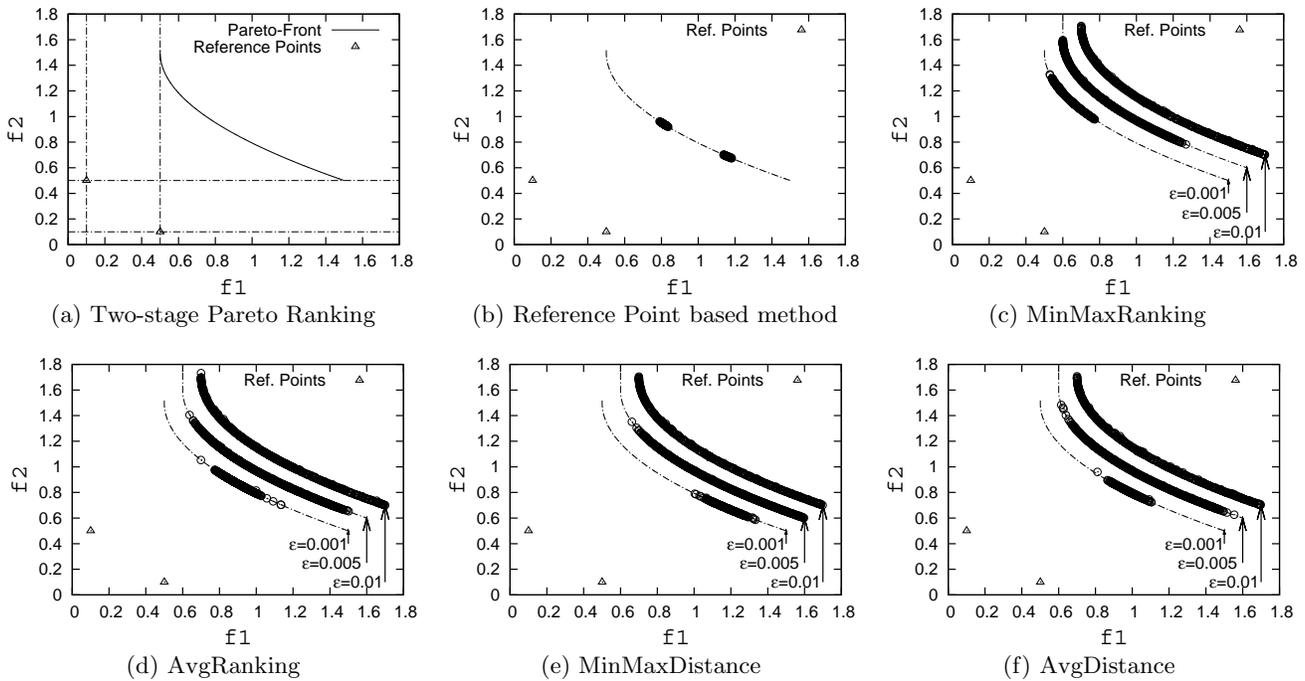


Figure 2: Comparison between "AND" operator of two-stage ranking, reference point based method by Deb and Sundar 2006, and the four newly proposed consensus finding methods with different ϵ values for a modified version of ZDT1.

instead of the minimum of the ranks, the maximum of the ranks is taken as crowding distance (in reference to the *AND* operation). Since in contrast to the original NSGA-II, the crowding distance is supposed to be minimized, we call this approach the *MinMaxRanking*. The second idea is to assign the average rank as crowding distance (*AvgRanking*). For example, if one solutions has the ranks 1,4,10 to each of the three reference points, respectively, the crowding distance after step 2 would be 5.

Distance based approach.

In the distance based approach, the first step (Ranking) is simplified because the ranking is skipped and only the normalized Euclidean distances are used as crowding distances. Then, in the second step, the same changes are made as in the ranking based approach. We consider all solutions and assign as crowding distance either the maximal normalized Euclidean distance they have to a reference point (*MinMaxDistance*) or the average distance (*AvgDistance*). Since for the ranking and the distance based approach, the grouping step is unchanged, the ϵ parameter can control the diversification of the Pareto front.

5. EXPERIMENTS

We present experiments for all four variations (*AvgRanking*, *MinMaxRanking*, *MinMaxDistance*, *AvgDistance*) of the proposed approach on three ZDT test problems and two test problems of the flow shop problem. In the ZDT experiments, we show how the proposed approaches perform for different numbers of reference points, different locations of reference points, and different settings of the grouping parameter ϵ . Afterwards, on the basis of the flow shop problem,

we explain how the method can be applied to a realistic problem setting. Since the visualization of the results is important for the evaluation of the proposed consensus finding method, we limit ourselves to two-dimensional test problems.

5.1 ZDT Problems

For the experiments with the three 30-variable ZDT problems, we use the parameters proposed for the reference point based method [12]. The distribution index is 10 for the simulated binary crossover (SBX) and 20 for the polynomial mutation operator. The mutation probability is set to $\frac{1}{30}$ (1/length of genotype) and the crossover probability to 0.9. Furthermore, the population size is set to 100, and each run is stopped after 50,000 evaluations. The results of each experiment are based on 10 independent runs.

ZDT1*.

ZDT1* is a slightly modified version of ZDT, in which a constant value of 0.5 is added to each of the two objectives of ZDT1. This transformation was necessary to better demonstrate the effects of certain reference points in the two stage Pareto ranking. ZDT1* forms a convex Pareto front ([31]). Figures 2(c)-(f) show the results for all four variations of the consensus approach for different values of ϵ for the 10 independent runs. We assume the same two reference points as in figures 2(a) and 2(b). For a better visualization of $\epsilon = 0.005$ and $\epsilon = 0.01$, we plot the two Pareto fronts with a parallel translation of plus 0.1 and 0.2 respectively in each dimension. The parameter ϵ controls the focus of the NSGA-II. For higher ϵ (0.01 and above) the front is as diverse as the one for the two-stage Pareto ranking because a high ϵ prevents all solutions except one in a quite large neighborhood

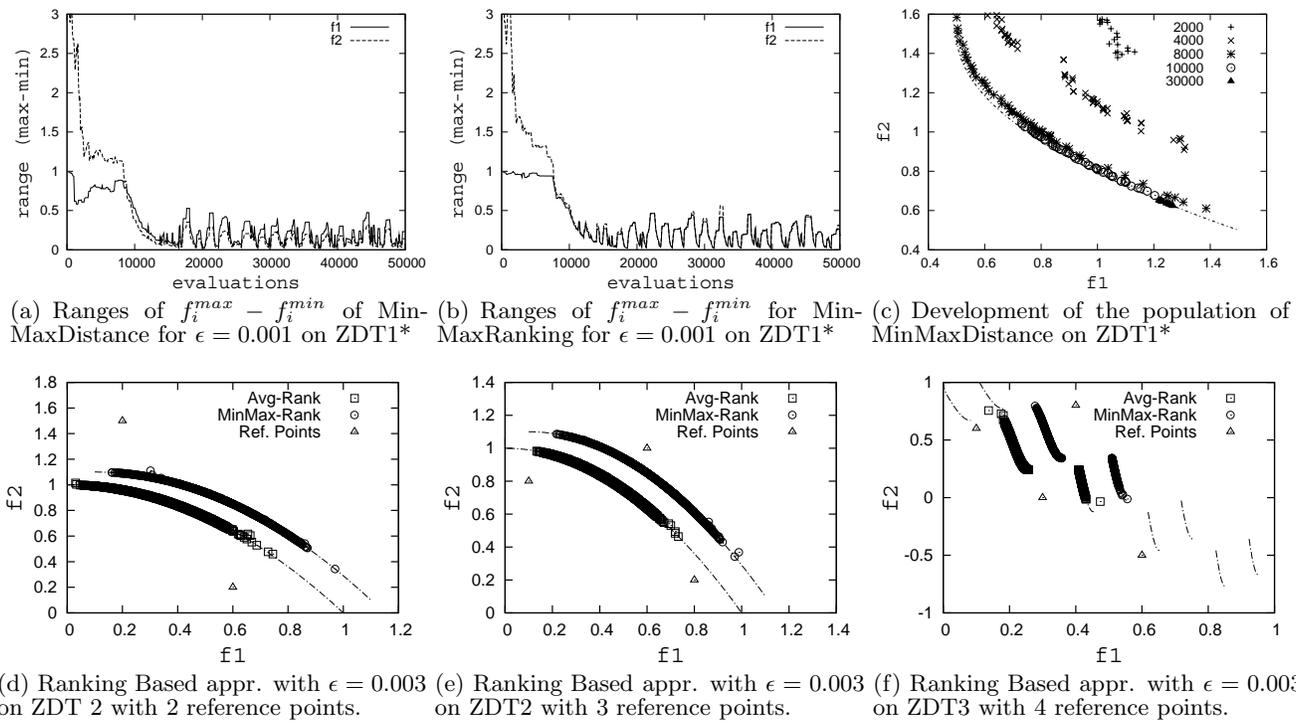


Figure 3: Runs for all three ZDT problems for different combinations of reference points.

from being selected (crowding distance penalty).

We see that while MinMaxRanking, AvgRanking and AvgDistance succeed in focusing on parts of the Pareto front which satisfy the preferences of both DMs (both reference points), MinMaxDistance focuses on the right part of the front favoring the reference point $[0.5, 0.1]$. The same tendency can be observed with AvgDistance. However, taking the average instead of the maximum of the normalized Euclidean distances produces less extreme values.

Figures 3(a) and 3(b) explain the behavior of the *MinMax* approach (the *average* approach behaves equivalently). We plot all denominators $f_i^{max} - f_i^{min}$ of the normalized Euclidean distance over all evaluations. For both, MinMaxDistance and MinMaxRanking the range of f_2 is higher in the first 10000 evaluations. Then, in case of MinMaxDistance f_2 constantly falls below f_1 , while in MinMaxRanking f_1 and f_2 are more or less equal. In other words, in both cases in the first 10000 evaluations, the denominator for the normalization of f_2 is much higher than for f_1 . Consequently changes in the f_2 value of a solution have less effect on the crowding distance than changes in f_1 . The EA focuses on minimizing f_1 and the population is pushed to the left side of the Pareto front. From the 10000th evaluation on only for MinMaxDistance the behavior is the other way around. From now on, the EA focuses on minimizing f_2 . This behavior is displayed in figure 3(c) for the whole search process. After 2000 evaluations the EA focuses on the left upper part and moves to the right part from the 10000th evaluation on. In the ranking approaches there is no such bias after the 10000th evaluation.

Due to this high influence of the different ranges of the objective functions for the distance based approach, the following experiments focus on the ranking based approaches.

ZDT2.

ZDT2 has a non-convex Pareto front. The results for a combination of one feasible $([0.2, 1.5])$ and one infeasible reference point $([0.6, 0.2])$ for the two distance based approaches are presented in figure 3(d) with an epsilon of 0.003. Here again, the Pareto front is drawn with a parallel translation of plus 0.1 in each dimension. Both the non-convexity of the front as well as the feasibility/infeasibility of reference points do not cause any problems for the methods. They easily focus on the part between the reference points on the true Pareto front. More tests were done with different numbers and locations of reference points. All runs show the desired behavior. For an example with three reference points see figure 3(e).

ZDT3.

The ZDT3 problem has a non-convex and discontinuous Pareto front. Both variations behave very similarly on a case with four reference points (see figure 3(f)). The concentration on the upper left parts of the front result from the concentration of three out of the four reference points in this same corner. Here again, the proposed approach has no difficulties with the form of the Pareto front or different reference points.

Conclusion.

The experiments revealed that the rank based consensus approaches work well with different combinations of reference points and for different shapes of the Pareto front. The grouping parameter ϵ controls nicely the extension of the focused front but must be adopted separately for each test problem and each method. It is interesting to note that with the grouping parameter the DMs can decide a priori

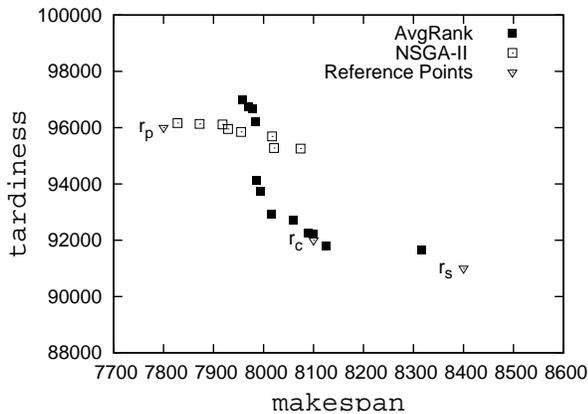


Figure 4: Comparison of AvgRanking approach and NSGA-II for a PFSP problem with 50 jobs and 20 machines.

whether they would like to have a small concentrated set of solutions for the final choice or whether they want a more varied set which also includes more extreme solutions which are more in favor of only one of the DMs. That is, the lower the grouping parameter is chosen, the more equal the balance of power between the DM should be distributed. If, on the other hand, the preferences of one or several DM are more important than others, then the grouping parameter should be assigned a higher value.

5.2 Permutation Flow Shop Scheduling Problem (PFSP)

The PFSP considers the task of scheduling n jobs, denoted by J_1, \dots, J_n , on a set of m machines, where one machine can only process one job at a time. All jobs are available at time zero and have the same processing route (M_1, M_2, \dots, M_m) . For each job J_i and each machine M_j a processing time is given by p_{ij} describing the time J_i needs on machine M_j . For a feasible schedule σ , the completion time of J_i is denoted by $C_i(\sigma)$. In the single-objective case of the PFSP, the objective is usually to minimize the makespan $C_{max}(\sigma) = \max_{i=1 \dots n} [C_i(\sigma)]$. In this paper we consider not only the minimization of the makespan, but also the minimization of the total tardiness of the jobs. A job is tardy if it is completed after its due date d_i . The total tardiness is therefore defined as $T(\sigma) = \sum_{i=1 \dots n} \max[C_i(\sigma) - d_i, 0]$. Both problems, minimizing the makespan as well as minimizing the tardiness, are NP-hard [24].

In a real-world flow shop problem, usually several divisions are involved in the scheduling decision [2]. The production division, for instance, wants to achieve output maximization and would therefore be interested mostly in minimizing the makespan. In contrast to that, the sales division mostly wants to achieve customer satisfaction and therefore aims at minimizing the tardiness. The controlling has to ensure profit maximization and therefore should be interested in optimizing both objectives concurrently.

We test the AvgRanking consensus finding approach on two flow shop benchmark problems of different complexity. The first problem is one with 50 jobs and 20 machines taken from the literature [13]. The second problem from Taillard is more complex with 100 jobs and 10 machines [18]. It was previously extended to the bi-objective case [19]. Three dif-

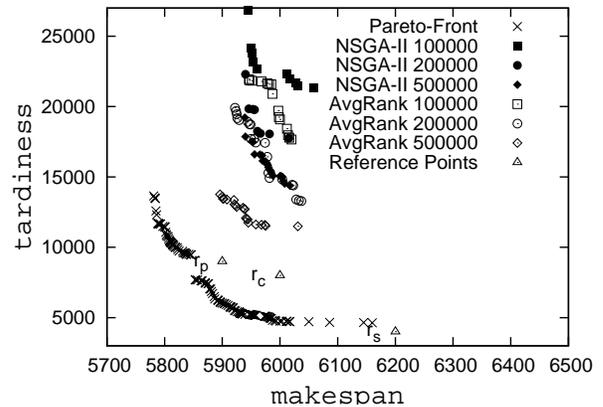


Figure 5: Comparison of AvgRanking approach and NSGA-II for a PFSP problem with 100 jobs and 10 machines.

ferent reference points for each problem are set according to the preferences of every division described above (production r_p , sales r_s , and controlling r_c). In the experiments we use OX crossover and shift mutation. The population size is 100, the crossover probability 0.9 and the mutation probability 0.05.

Figure 4 displays the non-dominated solutions for the first problem after 10 runs with different seeds, each with 50,000 evaluations. We can see that the AvgRanking approach finds more and better solutions than the NSGA-II in the region which represents a consensus between all reference points. The NSGA-II produced only good results for r_p . The convergence of NSGA-II is slower since it aims at approximating the whole Pareto front.

In figure 5 one can see the development of the non-dominated individuals for the AvgRanking approach and the standard NSGA-II considering different numbers of evaluations for the second problem. Again the AvgRanking approach shows considerably better results than the standard NSGA-II. Due to the concentration on the region of interest for the three DMs, it achieves better solutions while performing less evaluation steps than the standard NSGA-II.

6. CONCLUSION

Some of the latest approaches in multi-objective evolutionary algorithms (MOEA) do not attempt to approximate the whole Pareto optimal set of solutions but instead focus on those solutions which are preferred by the decision maker (DM). Therefore, it is assumed that the DM has at least some vague notion about his/her preferences which can help to focus the search process on only the relevant parts of the search space. Although this new research on integrating user preferences into a MOEA makes MOEAs more attractive for real-world scenarios, a crucial aspect relevant in many real-world scenarios misses: usually, it is a heterogenous group of DMs which decides for one final solution. In this paper, we have addressed this topic and propose a reference point based MOEA which finds a set of consensus decisions among the Pareto optimal set. We have shown that when normalized distances from solutions to the reference points are minimized in the algorithm, difference in the ranges of the objectives causes strong biases. These biases prevent the search from focusing on the consensus solutions. However,

when an additional ranking scheme is assigned according to the normalized distances, this bias vanishes after some evaluations and the approach succeeds in finding solutions in the parts which are preferred by all decision makers. Experiments with three artificial test problems with different Pareto fronts and with two realistic flow shop problems have validated the proposed approach. In future work, more studies should be done supporting multi-criteria group decision problems with MOEAs as we think that such approaches are very relevant for real-world scenarios.

7. REFERENCES

- [1] K. J. Arrow. *Social Choice and Individual Values*. John Wiley, New York, 1951.
- [2] T. Bagchi. Pareto-optimal solutions for multi-objective production scheduling problems. In *Proceedings of the First International Conference on Evolutionary Multi-Criterion Optimization*, volume 1993 of *Lecture Notes in Computer Science*, pages 458–471. Springer-Verlag, 2001.
- [3] J. Branke and K. Deb. Integrating user preferences into evolutionary multi-objective optimization. In Y. Jin, editor, *Knowledge Incorporation in Evolutionary Computation under the book series on Studies in Fuzziness and Soft Computing*, volume 167, pages 461–477. Springer-Verlag, 2004.
- [4] J. Branke, T. Kaufler, and H. Schmeck. Guidance in evolutionary multi-objective optimization. *Advances in Engineering Software*, 32:499–507, 2001.
- [5] C. Coello. Handling preferences in evolutionary multiobjective optimization: A survey. In *Proceedings of the Congress on Evolutionary Computation (CEC-00)*, volume 1, pages 30–37, 2000.
- [6] C. Coello, G. Lamont, and D. Van Veldhuizen. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Springer, New York, 2007.
- [7] D. Cvetkovic and I. C. Parmee. Preferences and their application in evolutionary multiobjective optimization. *IEEE Transactions on Evolutionary Computation*, 6(1):42–57, 2002.
- [8] K. Deb. Solving goal programming problems using multi-objective genetic algorithms. In *Proceedings of the Congress on Evolutionary Computation (CEC 99)*, pages 77–84, 1999.
- [9] K. Deb. Multi-objective evolutionary algorithms: Introducing bias among pareto-optimal solutions. In A. Ghosh and S. Tsutsui, editors, *Advances in Evolutionary Computing: Theory and Applications, Chapter 10*, pages 263–292. Springer-Verlag, London, 2003.
- [10] K. Deb and A. Kumar. Interactive evolutionary multi-objective optimization and decision-making using reference direction method. In *Proceedings of the Conference on Genetic and Evolutionary Computation (GECCO-07)*, pages 781–788. ACM, 2007.
- [11] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.
- [12] K. Deb and J. Sundar. Reference point based multi-objective optimization using evolutionary algorithms. In *Proceedings of the Conference on Genetic and Evolutionary Computation (GECCO-06)*, pages 635–642. ACM, 2006.
- [13] E. Demirkol, S. Mehta, and R. Uzsoy. Benchmarks for shop scheduling problems. *European Journal of Operational Research*, 109(1):137–141, August 1998.
- [14] C. M. Fonseca and P. J. Fleming. Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In *Proceedings of the Fifth International Conference on Genetic Algorithms*, pages 416–423. Morgan Kaufmann, 1993.
- [15] M. Köksalan, M. Karwan, and S. Zionts. An improved method for solving multiple criteria problems involving discrete alternatives. *IEEE Transactions on Systems, Man, and Cybernetics*, 14:24–34, 1984.
- [16] S. P. Phelps and M. Köksalan. An interactive evolutionary metaheuristic for multiobjective combinatorial optimization. *Management Science*, 49(12):1726–1738, 2003.
- [17] A. Sen. The impossibility of a paretian liberal. *Journal of Political Economy*, 78:152–157, 1970.
- [18] E. Taillard. Benchmarks for basic scheduling problems. *European Journal of Operational Research*, 64(2):278–285, 1993.
- [19] E. G. Talbi, M. Rahoual, M. H. Mabed, and C. Dhaenens. A hybrid evolutionary approach for multicriteria optimization problems: Application to the flow shop. In E. Zitzler, K. Deb, L. Thiele, C. A. C. Coello, and D. Corne, editors, *Lecture Notes in Computer Science*, pages 416–428. Springer, Berlin, 2001.
- [20] K. C. Tan, T. H. Lee, and E. F. Khor. Evolutionary algorithms with goal and priority information for multi-objective optimization. In *Proceedings of the Congress on Evolutionary Computation (CEC-99)*, pages 106–113, 1999.
- [21] M. Tanaka and T. Tanino. Global optimization by the genetic algorithm in a multiobjective decision support system. In *Proceedings of the 10th International conference on Multiple Criteria Decision Making*, volume 2, pages 261–270, 1992.
- [22] T. Tanino, M. Tanaka, and C. Hojo. An interactive multicriteria decision making method by using a genetic algorithm. In *Proceedings of the 2nd International Conference on Systems Science and Systems Engineering*, pages 381–386, 1993.
- [23] L. Thiele, K. Miettinen, P. Korhonen, and J. Molina. A preference-based interactive evolutionary algorithm for multiobjective optimization. Technical report, 2007.
- [24] V. T'kindt and J. C. Billaut. *Multicriteria Scheduling: Theory, Models and Algorithms*. Springer, Berlin, 2002.
- [25] D. S. Todd and P. Sen. Directed multiple objective search of design spaces using genetic algorithms and neural networks. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-99)*, volume 2, pages 1738–1743, 1999.
- [26] E. Turban. *Decision Support Systems and Expert Systems: Managerial Perspectives*. MacMillan, 1988.
- [27] D. A. V. Veldhuizen and G. B. Lamont. Multiobjective evolutionary algorithms: Analysing the state-of-the-art. *IEEE Transactions on Evolutionary Computation*, 8(2):125–147, 2000.
- [28] A. Wierzbicki. The use of reference objectives in multiobjective optimization. In G. Fandel and T. Gal, editors, *Lectures Notes in Economics and Mathematical Systems*, volume 177, pages 468–486. Springer, New York, 1980.
- [29] J. M. Yeh, C. Lin, B. Kreng, and J. Y. Gee. A modified procedure for synthesising ratio judgements in the analytic hierarchy process. *Journal of the Operational Research Society*, 50(8):867–873, 1999.
- [30] S. Zahir. Clusters in a group: Decision making in the vector space formulation of the analytic hierarchy process. *European Journal of Operations Research*, 112(3):620–634, 1999.
- [31] E. Zitzler, K. Deb, and L. Thiele. Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. *Evolutionary Computation*, 8(2):173–195, 2000.
- [32] E. Zitzler and S. Künzli. Indicator-Based Selection in Multiobjective Search. In *Conference on Parallel Problem Solving from Nature (PPSN VIII)*, pages 832–842. Springer, 2004.