The Power of Compromise

C. Büsing[1], K.S. Goetzmann[2], J. Matuschke[2], and S. Stiller[2]

[1]RWTH Aachen, buesing@or.rwth-aachen.de
[2]TU Berlin, {goetzmann,matuschke,stiller}@math.tu-berlin.de

Abstract

We study a concept in multicriteria optimization called *compromise solutions* (introduced in 1973 by Yu [20]) and a generalized version of this, termed *reference point solutions*. Our main result shows the power of this concept: Approximating reference point solutions is polynomially equivalent to constructing an approximate Pareto set as in [16].

A reference point solution is the solution closest to a given reference point in the objective space. Compromise solutions use the component-wise minimum over all solutions as a reference point. These methods are widely spread in practice. While for a fixed norm it gives a single solution balancing the different criteria, by changing the norm in the objective space each point in the Pareto set can become the reference point solution, thus maintaining the full variability of multicriteria problems. Despite its apparent virtues only few theoretical and even less algorithmic results are known for reference point methods.

We study minimization problems with a constant number of criteria. In addition to the equivalence of approximability of reference point solutions and the Pareto set, our techniques allow us to show that the Pareto set has a constant factor approximation if and only if the single-criterion problem has a constant factor approximation. We further give several general techniques to obtain solutions for reference point methods. The main algorithmic result is an LP-rounding technique that achieves the same approximation factors for reference point solutions as in the single-criterion case for many classical combinatorial problems, including set-cover and several machine scheduling problems. By the established link our algorithmic results also give a short alternative proof for the existence of an FPTAS of the Pareto set in the case of linear optimization over convex sets [16].

1 Introduction

In many applications of combinatorial optimization, trade-offs between conflicting objectives play a crucial role. For example, route guidance systems are a classical application of the shortest path problem. Yet, a good route guidance should allow the driver to make an informed choice to balance travel time and fuel consumption.

It is well-known that even for this basic example, the bicriteria shortest path problem, the number of Pareto optimal (i.e., non-dominated) solutions can grow exponentially with the size of the network. Decision makers may have different preferences how much extra fuel to spend on less travel time. Thus, a central task of multicriteria optimization is to either find a *single* solution based on a priori expressed trade-off preferences of the decision maker, or to identify a set of solutions that is of manageable (in mathematical terms: polynomial) size but still reflects all possible trade-off options at least approximately.

A straightforward way to a single solution is the weighted-sum method: The trade-off preferences are specified by two non-negative weights for time and fuel consumption. The navigation system then chooses a route minimizing the weighted sum of the two objectives. Unfortunately, this method deprives the decision maker of essential solutions: Consider an instance with three possible routes with corresponding objective value vectors $(10, 1)$, $(6, 6)$, and $(1, 10)$, respectively. The route with fuel consumption 6 and travel time 6 will never be the optimum for any choice of weights, despite being a balanced and thus attractive alternative for many drivers.

Formally, this shortcoming of the weighted-sum approach means that it cannot reach every point of the Pareto set. This motivates the concept of compromise solutions and reference point solutions [20], which returns a solution closest in the objective space to a given reference point. (Compromise solutions use the component-wise minimum over all solutions as a reference point.) The trade-off preferences are reflected by the choice of the norm in the objective space. Every point in the Pareto set is a reference point solution for some norm. Reference point methods are widely used in practice, serving as a core concept of MCDM[1] tools (cf. [3, 15] for particular examples and [6] for an overview). Still, they did not attract a lot of theoretical interest so far.

We show that approximating reference point solutions is computationally equivalent to approximating the Pareto set as proposed in Papadimitriou and Yannakakis [16]. Further, we provide general techniques for approximation algorithms, by means of which reference point solutions can often be approximated with the same factor as the single-criterion problem, most notably for the case of LP-rounding. One byproduct are approximation algorithms for the Pareto sets of many hard combinatorial optimization problems.

Related Work. Multicriteria optimization has a long tradition. The central notion of *Pareto optimality* goes back to publications by V. Pareto in the late 19th and early 20th century. Ever since then solution concepts in multicriteria optimization have been studied. The notion of *compromise solutions* was introduced in 1973 by Yu [20] and further studied and extended in the following years by Freimer and Yu [7], Gearhardt [8], Choo and Steuer [5] and many others. Recently, Voorneveld et al. [18] gave an axiomatization of compromise solutions, in particular those w.r.t. the Euclidean norm. The concept was later extended to more general reference points and is incorporated in many MCDM tools [3, 6, 15].

Also the approximation of Pareto sets has been studied for several decades now. It was initiated by Hansen in 1979 [11], followed by several publications on specific problems such as shortest paths [19] and scheduling [4]. More general results on the existence and computability of approximate Pareto sets were presented by Safer in his PhD thesis [17] in 1992, and in 2000 by Papadimitriou and Yannakakis [16]. Some of our results are based on the latter. Recently Mittal and Schulz [13] have used approximate Pareto sets to approximately optimize low-rank functions over polytopes. Their concept is similar to reference point methods as both can be seen as an aggregation of multiple objectives into one.

Multicriteria optimization and in particular compromise solutions are also closely related to robust optimization, in particular to min-max regret robustness. This connection has also been noted and exploited by others, e.g. Aissi et al. [1, 2]. We extend some of their results to reference point methods.

Our Contribution. Due to space limitations, we restrict ourselves to minimization problems throughout this paper. We note that this is *not* without loss of generality and some of our results do not to hold in the context of maximization.

In Section 3, we establish an algorithmic link between reference point solutions and approximation of the Pareto set. As a main result we show that approximating the reference point solution, approximating the compromise solution, and approximating the Pareto set are polynomially equivalent. In this context, we also prove that any point in the Pareto set can be obtained as reference point solution for two classes of popular norms with polynomially sized norm parameters, extending a known result by Gearhardt [8].

Combining these results with an easy constant factor approximation for reference points yields the following interesting corollary: For any discrete minimization problem with a fixed

[1] Multicriteria Decision Making

number of criteria there is a constant factor approximation for the Pareto set if and only if there is a constant factor approximation for the single-criterion version of the problem.

In Section 4, we show how to solve the reference point problem approximately for many combinatorial optimization problems. As a main result in this section, we show that single-objective approximations obtained by LP-rounding directly carry over to approximation algorithms for reference point methods. Along the way, we also prove that reference point solutions for linear objectives on convex sets can be found efficiently. From this we get a short alternative proof for the existence of an FPTAS for the Pareto set of such problems [16].

Finally, we extend a technique by Aissi et al. [1] from robust optimization to multicriteria optimization, allowing us to construct an FPTAS for reference point problems from pseudopolynomial algorithms.

2 Preliminaries

Throughout the paper, we let $\mathcal{P}$ denote a multicriteria discrete minimization problem with k objectives. As usual in multicriteria optimization, we assume the number of objectives to be fixed. For a given instance of $\mathcal{P}$, the set of feasible solutions is denoted by $\mathcal{X}$ and $c : \mathcal{X} \to \mathbb{Z}^k_{\geq 0}$ is the vector of objective functions. The *objective vector set* of the instance is defined by $\mathcal{Y} := c(\mathcal{X}) \subseteq \mathbb{Z}^k_{\geq 0}$. A solution $y \in \mathcal{Y}$ is *Pareto optimal* if there does not exist $y' \in \mathcal{Y} \setminus \{y\}$ with $y' \leq y$. The *Pareto set* $\mathcal{Y}_P$ is the set of all Pareto optimal solutions.

Reference Point Methods. The goal of reference point methods is to find a solution minimizing the distance to a user-specified *reference point* $y^{\text{rp}} \in \mathbb{Z}^k_{\geq 0}$ w.r.t. a given norm $\|\cdot\|$. A particularly interesting reference point is the *ideal point* $y^{\text{id}} \in \mathbb{Z}^k_{\geq 0}$, which is the point in the objective space obtained by optimizing each objective individually, i.e., $y^{\text{id}}_i := \min_{y \in \mathcal{Y}} y_i$. Throughout this paper, we will restrict ourselves to reference points y^{rp} with $y^{\text{rp}} \leq y^{\text{id}}$. We call these points *feasible reference points*.

Besides choosing the reference point, the decision maker can furthermore express her preferences by specifying weights. For any norm $\|\cdot\|$ on $\mathbb{R}^k$ and $\lambda \in \mathbb{Q}^k_{\geq 0}$, let $\|\cdot\|^\lambda$ be the norm defined by $\|y\|^\lambda := \|(\lambda_1 y_1, \ldots, \lambda_k y_k)\|$. We will consider the objective function $r_{y^{\text{rp}},\lambda}(y) = \|y^{\text{rp}}\|^\lambda + \|y - y^{\text{rp}}\|^\lambda$, which we call *relative distance*[2]. Note that, for exact computation, minimizing the relative distance is equivalent to minimizing the absolute distance $\|y - y^{\text{rp}}\|^\lambda$, as the level sets of the two functions are the same. However, in the context of approximation, this change in the objective function plays a crucial role. From a practical point of view, a decision maker will care about the *relative* deviation from the ideal values. From a theoretical point of view, any approximation algorithm for the absolute distance can be turned into an algorithm that solves the single-criterion problem exactly (as the minimal distance to the ideal point is 0, when focusing on a single criterion). As we do not want to restrict our study to exactly solvable problems, we will use the relative distance, which allows us to also consider problems that are only approximately solvable.

We now formally define the problem $\text{RP}(\mathcal{P}, \|\cdot\|)$: Given an instance of $\mathcal{P}$, a feasible reference point $y^{\text{rp}} \in \mathbb{Z}^k_{\geq 0}$, and a weight vector $\lambda \in \mathbb{Q}^k_{\geq 0}$ as input, find a solution $x \in \mathcal{X}$ that minimizes $r_{y^{\text{rp}},\lambda}(c(x))$. Given the particular interest of the ideal point, we will also consider the problem $\text{CP}(\mathcal{P}, \|\cdot\|)$, which is also known as *compromise programming*: Given an instance of $\mathcal{P}$ and $\lambda \in \mathbb{Q}^k_{\geq 0}$, find a solution $x \in \mathcal{X}$ that minimizes $r_{y^{\text{id}},\lambda}(c(x))$.

Note that although the concept of reference point solutions is a generalization of compromise solutions, in terms of complexity CP is *not* a special case of RP. In the former problem, the ideal point is not given and might be hard to compute (in other words: we cannot even evaluate

[2]The relative distance is not a distance in the mathematical sense of a *metric*. In particular the triangle inequality does not hold.

the objective function efficiently), while in the latter case, the reference point is given in the input. However, we will show in Section 3, that this difference only plays a minor role in the context of approximability.

Norms. Throughout this paper, we will restrict ourselves to the norms fulfilling the following two properties. We call a norm $\|\cdot\|$ *monotone* if $y' \leq y''$ implies $\|y'\| \leq \|y''\|$ for any $y', y'' \in \mathbb{R}^k_{\geq 0}$. We call a norm $\|\cdot\|$ *polynomially decidable*, if we can decide whether $\|y'\| \leq \|y''\|$ in time polynomial in the encoding size of y' and y''. We will mainly use the following families of norms: the infinity-norm $\|y\|_\infty := \max_i |y_i|$ (which we will sometimes also denote by $\langle\!\langle y \rangle\!\rangle_\infty$ for convenience), the standard ℓ^p-norm $\|y\|_p := (\sum_i |y_i|^p)^{\frac{1}{p}}$, and the cornered p-norm $\langle\!\langle y \rangle\!\rangle_p := \max_i |y_i| + \frac{1}{p}\sum_i |y_i|$. The motivation for the cornered norm is twofold. Firstly, for general values of p it will be hard to minimize a distance measured in the ℓ^p-norm because of the exponents. The cornered p-norms are simpler, but still have properties similar to the ℓ^p-norms: Their unit spheres are nested within each other, and for increasing values of p they approach the axis parallel square. This allows to control the degree of balancing of the criteria in the reference point solution. Secondly, the infinity-norm (often referred to as Chebyshev-norm in this context) is very popular in MCDM-tools. Often it is *augmented* by a small linear term to avoid weakly Pareto optimal solutions (cf. [5]), similar to the addition of the term $\frac{1}{p}\|y\|_p$.

Note that all ℓ^p- and cornered p-norms are monotone and polynomially decidable.

Approximation of the Pareto Set. We extend the well-known concept of approximation algorithms for the single-objective case to approximability of the Pareto set in a similar way as done in [16], with the slight difference of considering minimization problems and including constant factor approximations. For $\alpha > 1$, an *α-approximate Pareto set* is a set $\mathcal{Y}_\alpha \subseteq \mathcal{Y}$ such that for all $y \in \mathcal{Y}_P$ there is $y' \in \mathcal{Y}_\alpha$ with $y' \leq \alpha y$. An *α-approximation algorithm for the Pareto set* is an algorithm that constructs an α-approximate Pareto set in polynomial time, and an *FPTAS for the Pareto set* is a family of algorithms that contains for all $\varepsilon > 0$ a $(1+\varepsilon)$-approximation algorithm for the Pareto set with running time polynomial in $\frac{1}{\varepsilon}$ and the size of the instance of $\mathcal{P}$.

3 Equivalence of Approximation

In this section, we investigate the relation between approximation of the Pareto set, reference point methods, and compromise programming. Our main theorem states these three notions of approximability are essentially equivalent.

Theorem 3.1. *Let $\mathcal{P}$ be a multicriteria discrete minimization problem. The following statements are equivalent.*

- *There is a constant factor approximation (FPTAS, respectively) for the Pareto set of $\mathcal{P}$.*
- *There is a constant factor approximation (FPTAS, respectively) for $\mathrm{RP}(\mathcal{P}, \|\cdot\|)$ for every monotone and polynomially decidable norm $\|\cdot\|$.*
- *There is a constant factor approximation (FPTAS, respectively) for $\mathrm{RP}(\mathcal{P}, \|\cdot\|_\infty)$.*
- *There is a family of algorithms that, for each $p \geq 1$, contains a constant factor approximation (FPTAS, respectively) for $\mathrm{RP}(\mathcal{P}, \|\cdot\|_p)$ or $\mathrm{RP}(\mathcal{P}, \langle\!\langle \cdot \rangle\!\rangle_p)$, and the running time of all algorithms is bounded by a polynomial in the input size and $\log(p)$.*
- *There is a constant factor approximation (FPTAS, respectively) for $\mathrm{CP}(\mathcal{P}, \|\cdot\|)$ for every monotone and polynomially decidable norm $\|\cdot\|$.*

- *There is a constant factor approximation (FPTAS, respectively) for* $\mathrm{CP}(\mathcal{P}, \|\cdot\|_\infty)$.
- *There is a family of algorithms that, for each* $p \geq 1$*, contains a constant factor approximation (FPTAS, respectively) for* $\mathrm{CP}(\mathcal{P}, \|\cdot\|_p)$ *or* $\mathrm{CP}(\mathcal{P}, \langle\!\langle\cdot\rangle\!\rangle_p)$*, and the running time of all algorithms is bounded by a polynomial in the input size and* $\log(p)$.

Before we discuss the proof of Theorem 3.1 in detail, we turn our attention to a result of independent interest that motivates the algorithmic use of the ℓ^p- and cornered p-norms.

Reference Point Solutions and the Pareto Set. Gearhardt [8] showed that for both the ℓ^p-norm and the cornered p-norm, if p tends to infinity then the distance between the Pareto set and the set of compromise solutions with respect to all non-negative normalized weight vectors tends to zero. This means that for discrete optimization problems with a finite set of feasible solutions there is a finite value p_0 for which the two sets coincide. We show that under the assumption that all objective values have a polynomial encoding (cf. [16]), also p can be chosen in such a way that it is polynomially encodable.

Theorem 3.2. *Let* $y^{\mathrm{rp}} \in \mathbb{Z}^k_{\geq 0}$ *be a feasible reference point. If the objective vector set* $\mathcal{Y}$ *is contained in* $[0, M]^k$*, then the following statements hold true.*

1. *If* $p > \frac{\log k}{\log(1+\frac{1}{M})}$*, then for any Pareto optimal solution* $y \in \mathcal{Y}$ *there is a weight vector* $\lambda \in \mathbb{Q}^k_{\geq 0}$ *such that* y *minimizes* $\|y - y^{\mathrm{rp}}\|^\lambda_p$.
2. *If* $p > kM$*, then for any Pareto optimal solution* $y \in \mathcal{Y}$ *there is a weight vector* $\lambda \in \mathbb{Q}^k_{\geq 0}$ *such that* y *minimizes* $\langle\!\langle y - y^{\mathrm{rp}}\rangle\!\rangle^\lambda_p$.

Proof. We first consider the cornered norm $\langle\!\langle y\rangle\!\rangle^\lambda_p = \max_{i\in[k]}\{\lambda_i y_i\} + \frac{1}{p}\sum_{i\in[k]}\lambda_i y_i$. As demanded in the theorem, let $p > kM$. Further let $y \in \mathcal{Y}$ be a Pareto optimal cost vector, and let $I := \{i \in [k] : y_i = y^{\mathrm{rp}}_i\}$. We set the weight vector λ as follows:

$$\lambda_i = \begin{cases} 1 + k & \text{if } i \in I \\ \frac{1}{y_i - y_i^{\mathrm{rp}}} & \text{otherwise.} \end{cases}$$

The weighted distance of y to the reference point is

$$\langle\!\langle y - y^{\mathrm{rp}}\rangle\!\rangle^\lambda_p = \max_{i\notin I}\{\lambda_i(y_i - y_i^{\mathrm{rp}})\} + \frac{1}{p}\sum_{i\notin I}\lambda_i(y_i - y_i^{\mathrm{rp}}) = 1 + \frac{1}{p}\,(k - |I|) \leq 1 + \frac{k}{p}\,.$$

Consider any $y' \in \mathcal{Y} \setminus \{y\}$. If there is an index $j \in I$ with $y'_j > y_j = y^{\mathrm{rp}}_j$, then since $\mathcal{Y} \subseteq \mathbb{Z}^k$ we know that $y'_j - y^{\mathrm{rp}}_j \geq 1$ and therefore

$$\langle\!\langle y' - y^{\mathrm{rp}}\rangle\!\rangle^\lambda_p \geq \lambda_j(y'_j - y_j^{\mathrm{rp}}) + \frac{1}{p}\,\lambda_j(y'_j - y_j^{\mathrm{rp}}) \geq (1 + \frac{1}{p}) \cdot \lambda_j > 1 + k > 1 + \frac{k}{p}\,,$$

so in this case y is closer to y^{rp} than y'.

Otherwise, since y is Pareto optimal there is some $j \in [k]$ such that $y_j < y'_j$, or with integrality $y'_j - y_j \geq 1$. On the other hand, we know that $y_j - y^{\mathrm{rp}}_j \leq M$. Therefore for this index j,

$$\lambda_j(y'_j - y_j^{\mathrm{rp}}) = \frac{y'_j - y_j^{\mathrm{rp}}}{y_j - y_j^{\mathrm{rp}}} = \frac{y'_j - y_j + y_j - y_j^{\mathrm{rp}}}{y_j - y_j^{\mathrm{rp}}} = 1 + \frac{y'_j - y_j}{y_j - y_j^{\mathrm{rp}}} \geq 1 + \frac{1}{M}\,,$$

and as a consequence

$$\langle\!\langle y' - y^{\mathrm{rp}}\rangle\!\rangle^\lambda_p \geq \max_{i\in[k]}\{\lambda_i(y'_i - y_i^{\mathrm{rp}})\} \geq 1 + \frac{1}{M} > 1 + \frac{k}{p}\,,$$

so again y is closer to y^{rp} than y'.

For the ℓ^p-norm we let $p > \frac{\log k}{\log(1+\frac{1}{M})}$, and set the weight vector λ as before. We get

$$\left(\|y - y^{\mathrm{rp}}\|_p^\lambda\right)^p = \sum_{i \notin I} \left(\frac{y_i - y_i^{\mathrm{rp}}}{y_i - y_i^{\mathrm{rp}}}\right)^p = k - |I| \leq k \,.$$

If there is a $j \in I$ with $y'_j > y_j = y_j^{\mathrm{rp}}$, then

$$\left(\|y' - y^{\mathrm{rp}}\|_p^\lambda\right)^p \geq \left(\lambda_j(y'_j - y_j^{\mathrm{rp}})\right)^p \geq \lambda_j^p > k \,.$$

Otherwise, with the same choice of $j \in [k]$ as above,

$$\left(\|y' - y^{\mathrm{rp}}\|_p^\lambda\right)^p \geq (\lambda_j(y'_j - y_j^{\mathrm{rp}}))^p \geq \left(1 + \frac{1}{M}\right)^p > k \,,$$

where the last inequality holds by the choice of p. Thus again in both cases y is closer to y^{rp} than y', completing the proof. □

From Approximate Pareto Sets to Approximating Reference Point Solutions. We start the proof of Theorem 3.1 by showing that from an α-approximate Pareto set we can always choose an α-approximate solution to RP.

Lemma 3.3. *Let y^{rp} be a feasible reference point, and let $\mathcal{Y}_\alpha$ be an α-approximate Pareto set of $\mathcal{P}$. Then for any monotone norm $\|\cdot\|$, $\min_{y \in \mathcal{Y}_\alpha} r(y) \leq \alpha \cdot \min_{y \in \mathcal{Y}} r(y)$, where $r(y) = \|y^{\mathrm{rp}}\| + \|y - y^{\mathrm{rp}}\|$.*

Proof. Let $y^* \in \mathcal{Y}$ be an optimal solution to $\min_{y \in \mathcal{Y}} r(y)$. By monotonicity, we can w.l.o.g. assume y^* to be Pareto optimal. Thus, there is $y' \in \mathcal{Y}_\alpha$ such that $y' \leq \alpha y^*$. Using monotonicity and triangle inequality, we get

$$\|y' - y^{\mathrm{rp}}\| \leq \|\alpha(y^* - y^{\mathrm{rp}}) + (\alpha - 1)y^{\mathrm{rp}}\| \leq \alpha\|y^* - y^{\mathrm{rp}}\| + (\alpha - 1)\|y^{\mathrm{rp}}\| \,.$$

Reformulation yields

$$\min_{y \in \mathcal{Y}_\alpha} r(y) \leq r(y') = \|y^{\mathrm{rp}}\| + \|y' - y^{\mathrm{rp}}\| \leq \alpha(\|y^{\mathrm{rp}}\| + \|y^* - y^{\mathrm{rp}}\|) = \alpha r(y^*) \,.$$ □

Corollary 3.4. *If there is an α-approximation algorithm for the Pareto set of $\mathcal{P}$, then there is an α-approximation for $\mathrm{RP}(\mathcal{P}, \|\cdot\|)$ for every monotone and polynomially decidable norm $\|\cdot\|$.*

In fact, we can also approximate the compromise solution without knowing the exact ideal point: As $\mathcal{Y}_\alpha$ contains an α-approximate optimal solution for each objective, we can obtain a reference point y^{rp} with $\frac{1}{\alpha} y^{\mathrm{id}} \leq y^{\mathrm{rp}} \leq y^{\mathrm{id}}$. By choosing the point closest to y^{rp} from $\mathcal{Y}_\alpha$ we get an α^2-approximation to the compromise solution:

Corollary 3.5. *If there is an α-approximation algorithm for the Pareto set of $\mathcal{P}$, then there is an α^2-approximation for $\mathrm{CP}(\mathcal{P}, \|\cdot\|)$ for every monotone and polynomially decidable norm $\|\cdot\|$.*

Proof. Let $\mathcal{Y}_\alpha$ be an α-approximation to the Pareto set. Observe that

$$y_i^{\mathrm{rp}} := \left\lceil \frac{1}{\alpha} \min_{y \in \mathcal{Y}_\alpha} y_i \right\rceil$$

yields a feasible reference point with $y^{\mathrm{rp}} \leq y^{\mathrm{id}} \leq \alpha y^{\mathrm{rp}}$.

Now let $y' := \arg\min_{y \in \mathcal{Y}_\alpha} \|y^{\mathrm{rp}}\| + \|y - y^{\mathrm{rp}}\|$, which by Corollary 3.4 is an α-approximation to the reference point solution for y^{rp}. Thus, for the compromise solution y^*, we get

$$\|y' - y^{\mathrm{rp}}\| \leq \alpha\|y^* - y^{\mathrm{rp}}\| + (\alpha - 1)\|y^{\mathrm{rp}}\| \,.$$

By this fact and monotonicity, we get

$$\|y' - y^{\mathrm{id}}\| \le \|y' - y^{\mathrm{rp}}\| \le \alpha \|y^* - y^{\mathrm{rp}}\| + (\alpha - 1)\|y^{\mathrm{rp}}\| \ .$$

Observe that $y^{\mathrm{id}} - y^{\mathrm{rp}} \le (\alpha - 1)y^{\mathrm{rp}}$ and thus, again by monotonicity

$$\|y^* - y^{\mathrm{rp}}\| \le \|y^* - y^{\mathrm{rp}}\| + (\alpha - 1)\|y^{\mathrm{rp}}\| \ .$$

This finally yields,

$$\begin{aligned} \|y^{\mathrm{id}}\| + \|y' - y^{\mathrm{id}}\| &\le \|y^{\mathrm{id}}\| + \alpha(\|y^* - y^{\mathrm{rp}}\| + (\alpha - 1)\|y^{\mathrm{rp}}\|) + (\alpha - 1)\|y^{\mathrm{rp}}\| \\ &\le \alpha^2 \|y^{\mathrm{id}}\| + \alpha \|y^* - y^{\mathrm{id}}\| \ , \end{aligned}$$

which concludes the proof. □

From Approximating a Reference Point to an Approximate Pareto Set. In order to show the converse of the result proven above, we use a characterization from [16], stating that approximability of the Pareto set is equivalent to tractability of the so-called GAP problem.

Definition 3.6 (GAP Problem)**.** Given an instance of $\mathcal{P}$ and a vector $y \in \mathbb{Q}^k_{\ge 0}$ as input, the GAP problem for approximation factor $\alpha > 1$, denoted by GAP$(\mathcal{P}, \alpha)$, is to find a solution $y' \in \mathcal{Y}$ with $y' \le y$ or to guarantee that there is no solution $y'' \in \mathcal{Y}$ with $y'' \le \frac{1}{\alpha} y$.

Theorem 3.7 (Papadimitriou & Yannakakis, 2000)**.** *Let $\mathcal{P}$ be a multicriteria discrete minimization problem and let $\alpha > 1$. If there is an α-approximation algorithm for the Pareto set, then* GAP$(\mathcal{P}, \alpha)$ *is solvable in polynomial time. If* GAP$(\mathcal{P}, \alpha)$ *is solvable in polynomial time, then there is an α^2-approximation algorithm for the Pareto set.*

We now show how to use an approximation algorithm for RP to solve the GAP problem with a slight increase in the approximation factor. In fact, our result does not even require the algorithm to solve RP for every given reference point. It suffices to find a particular reference point on an instance-by-instance basis that can be approximated. We formalize this by introducing two algorithms, the first acting as an oracle computing a suitable reference point, which then can be approximated by the second algorithm.[3]

Lemma 3.8. *Let $\alpha > 1$ and set $\beta := \frac{\alpha^2}{2\alpha - 1}$. There is a polynomial time algorithm for* GAP$(\mathcal{P}, \alpha)$*, if there are two polynomial time algorithms A_1, A_2 such that,*

- *given an instance of $\mathcal{P}$, algorithm A_1 computes a feasible reference point $y^{\mathrm{rp}} \in \mathbb{Q}^k_{\ge 0}$ for that instance, and,*
- *additionally given y^{rp} and $\lambda \in \mathbb{Q}^k_{\ge 0}$, algorithm A_2 computes in polynomial time a solution $y' \in \mathcal{Y}$ with $r(y') \le \beta \min_{z \in \mathcal{Y}} r(z)$ for $r(z) = \|y^{\mathrm{rp}}\|^{\lambda}_{\infty} + \|y^{\mathrm{rp}} - z\|^{\lambda}_{\infty}$.*

Proof. Let $y \in \mathbb{Q}^k_{\ge 0}$ be the input to the GAP problem. W.l.o.g., we can assume that $y \ge \alpha y^{\mathrm{rp}}$ for the reference point y^{rp} computed by A_1, as otherwise there is no $y' \le \frac{1}{\alpha} y$ and GAP can be answered negatively.

We will solve the GAP problem with a single call of the β-approximation algorithm for RP$(\mathcal{P}, \|\cdot\|_\infty)$. For $i \in [k]$, let $\lambda_i := \frac{1}{y_i - y_i^{\mathrm{rp}}}$ if $y_i > y_i^{\mathrm{rp}}$ and $\lambda_i := 2$ if $y_i = y_i^{\mathrm{rp}} = 0$, and let y' be a β-approximation to $\min_{z \in \mathcal{Y}} r(z)$.

[3] Note that it is not sufficient for the first algorithm to simply return a trivial feasible reference point such as 0, as it has to ensure that the second algorithm can provide an approximation for this point. E.g., in the proof of Corollary 3.9, it needs to return a point close to the ideal point.

If $r(y') \le r(y)$, we return y' as answer to the Gap problem:

$$\lambda_i(y'_i - y^{\text{rp}}_i) \le \|y' - y^{\text{rp}}\|^{\lambda}_{\infty} \le \|y - y^{\text{rp}}\|^{\lambda}_{\infty} \le 1$$

for all $i \in [k]$ by choice of the weights. Dividing by λ_i yields $y'_i \le y_i$ if $y_i > 0$ or $y'_i \le \frac{1}{2}$ if $y_i = 0$. In the latter case, integrality of y'_i implies $y'_i = 0$.

If $r(y') > r(y)$, we answer Gap negatively: Let $y'' \in \mathcal{Y}$. We show that there is an $i \in [k]$ with $y''_i > \frac{1}{\alpha} y_i$. First observe that $\beta r(y'') \ge r(y') > r(y)$, which implies

$$\beta \|y'' - y^{\text{rp}}\|^{\lambda}_{\infty} > \|y - y^{\text{rp}}\|^{\lambda}_{\infty} - (\beta - 1)\|y^{\text{rp}}\|^{\lambda}_{\infty} .$$

Plugging in the definition of the weights and using $y \ge \alpha y^{\text{rp}}$ yields

$$\beta \tfrac{y''_i - y^{\text{rp}}_i}{y_i - y^{\text{rp}}_i} > \tfrac{y_j - y^{\text{rp}}_j}{y_j - y^{\text{rp}}_j} - (\beta - 1)\tfrac{y^{\text{rp}}_{j'}}{y_{j'} - y^{\text{rp}}_{j'}} \ge 1 - \tfrac{\beta - 1}{\alpha - 1} ,$$

with i, j, j' being the indices of those components attaining the maxima in the norms. (If either of the denominators is 0 then $y''_i > \frac{1}{\alpha} y_i$ follows directly.) Using the fact that $1 < \beta \le \alpha$, we get

$$\beta y''_i > (1 - \tfrac{\beta-1}{\alpha-1})(y_i - y^{\text{rp}}_i) + \beta y^{\text{rp}}_i \ge (1 - \tfrac{\beta-1}{\alpha-1}) y_i .$$

It is easy to verify that $\beta = \frac{\alpha^2}{2\alpha - 1}$ now implies $y''_i > \frac{1}{\alpha} y_i$ and the negative answer to the Gap problem is correct. □

As a particular application of Lemma 3.8 we can show now that also an approximation to CP suffices to approximate the Pareto set:

Corollary 3.9. *Let $\alpha > 1$ and set $\beta := \sqrt{\frac{\alpha^2}{2\alpha-1}}$. There is a polynomial time algorithm for* Gap$(\mathcal{P}, \alpha)$*, if there is a β-approximation algorithm for* CP$(\mathcal{P}, \|\cdot\|_\infty)$.

Proof. We show that algorithm A_1 and A_2 exist, as required by Lemma 3.8.

Algorithm A_1: For every $i \in [k]$, let $\bar{y}^{(i)}$ be a β-approximation to $\min_{z \in \mathcal{Y}} r_{y^{\text{id}}, \bar{\lambda}}(z)$ for weights $\bar{\lambda}_i = 1$ and $\bar{\lambda}_j = 0$ for $j \in [k] \setminus \{i\}$. Then $y^{\text{rp}}_i := \frac{1}{\beta} \bar{y}^{(i)}_i$ defines a feasible reference point with $y^{\text{rp}} \le y^{\text{id}} \le \beta y^{\text{rp}}$.

Algorithm A_2: Let $\lambda \in \mathbb{Q}^k$. Let y' be a β-approximation to $\min_{z \in \mathcal{Y}} r_{y^{\text{id}}, \lambda}(z)$, and let $y^* = \text{argmin}_{z \in \mathcal{Y}} r_{y^{\text{rp}}, \lambda}(z)$ be an optimal solution to RP. We show that $r_{y^{\text{rp}}, \lambda}(y') \le \beta^2 r_{y^{\text{rp}}, \lambda}(y^*)$, which concludes the proof.

$$\begin{aligned}
\|y^{\text{rp}}\|_\infty + \|y' - y^{\text{rp}}\|_\infty &\le \|y^{\text{rp}}\|_\infty + \|y' - y^{\text{id}}\|_\infty + \|y^{\text{id}} - y^{\text{rp}}\|_\infty \le \beta \|y^{\text{rp}}\|_\infty + \|y' - y^{\text{id}}\|_\infty \\
&\le \beta \|y^{\text{rp}}\|_\infty + \beta \|y^* - y^{\text{id}}\|_\infty + (\beta - 1)\|y^{\text{id}}\|_\infty \\
&\le \beta^2 \|y^{\text{rp}}\|_\infty + \beta \|y^* - y^{\text{id}}\|_\infty .
\end{aligned}$$ □

Corresponding versions of Lemma 3.8 and Corollary 3.9 with the same approximation factors can be shown for the $\|\cdot\|_p$- and $\langle\!\langle \cdot \rangle\!\rangle_p$-norms. These results have been moved to the appendix.

4 Approximating Reference Point Solutions

Approximation by Weighted Sum. Although not all Pareto optimal solutions can be reached by minimizing a weighted sum, this method still provides an easy way to transfer approximability results from the single-criterion world to reference point methods.

Theorem 4.1. *If there is an α-approximation for $\min_{y \in \mathcal{Y}} \lambda^T y$, then there is a $k\alpha$-approximation for* RP$(\mathcal{P}, \|\cdot\|_\infty)$.

Proof. Let y^{rp} be a feasible reference point and $\lambda \in \mathbb{Q}^k_{\geq 0}$. Let $y^* = \operatorname{argmin} r_{y^{\mathrm{rp}},\lambda}(y)$ and let $y' \in \mathcal{Y}$ be an α-approximation to $\min_{y\in\mathcal{Y}} \lambda^T y$. Then

$$\begin{aligned}\|y^{\mathrm{rp}}\|^{\lambda}_{\infty} + \|y' - y^{\mathrm{rp}}\|^{\lambda}_{\infty} &\leq \|y^{\mathrm{rp}}\|^{\lambda}_{\infty} + k\lambda^T(y' - y^{\mathrm{rp}}) \leq \|y^{\mathrm{rp}}\|^{\lambda}_{\infty} + \alpha k \lambda^T y^* - k\lambda^T y^{\mathrm{rp}} \\ &\leq \|y^{\mathrm{rp}}\|^{\lambda}_{\infty} + \alpha k \lambda^T (y^* - y^{\mathrm{rp}}) + (\alpha - 1) k \lambda^T y^{\mathrm{rp}} \\ &\leq k\alpha(\|y^{\mathrm{rp}}\|^{\lambda}_{\infty} + \|y^* - y^{\mathrm{rp}}\|^{\lambda}_{\infty}) \,.\end{aligned}$$

□

In combination with Theorem 3.1, this implies the following result.

Corollary 4.2. *For any multicriteria combinatorial minimization problem $\mathcal{P}$ with a constant number of linear objectives, there is a constant factor approximation for the Pareto set of $\mathcal{P}$ if and only if there is a constant factor approximation for the single-criterion version of $\mathcal{P}$.*

Convex Optimization with Linear Objectives. For optimization problems where the solution space is convex and the objectives are linear (e.g. linear programming), we can compute reference point solutions w.r.t. the cornered norm exactly:

Theorem 4.3 (Reference Point Solutions for Convex Optimization)**.** *For a multicriteria minimization problem $\min_{x\in\mathcal{X}} Cx$ with a convex solution set $\mathcal{X} \subseteq \mathbb{Q}^n$ for which a polynomial separation algorithm exists, and a cost matrix $C \in \mathbb{Q}^{k\times n}$, the problem $\min_{x\in\mathcal{X}} r(Cx)$ with $r(y) = \langle\langle y^{\mathrm{rp}}\rangle\rangle_p + \langle\langle y - y^{\mathrm{rp}}\rangle\rangle_p$, for any feasible reference point y^{rp} and any $p \in [1, \infty]$, is again a convex optimization problem with linear objectives and thus solvable in polynomial time.*

Proof. The problem can be formulated as follows:

$$\min_{x\in\mathcal{X}} r(Cx) = \|y^{\mathrm{rp}}\|_{\infty} + \begin{cases} \min & \Delta + \frac{1}{p}\cdot \mathbb{1}^{\mathrm{T}} Cx \\ \text{s.t.} & Cx - y^{\mathrm{rp}} \leq \Delta \cdot \mathbb{1} \\ & x \in \mathcal{X} \\ & \Delta \in \mathbb{R}\,. \end{cases}$$

Here $\mathbb{1}$ denotes the vector of ones of corresponding dimension. In the optimum $\Delta = \max_i\{c_{i\cdot}x - y^{\mathrm{rp}}_i\}$, and therefore the two programs are equivalent. The objective is clearly linear, and the solution space is $\mathcal{X}\times\mathbb{R}$ intersected with the halfspaces defined by the inequalities $c_{i\cdot}x - y^{\mathrm{rp}}_i \leq \Delta$, $i \in [k]$, and thus convex.

Since we can solve the separation problem for the original set $\mathcal{X}$, we can also solve it for the set with the added inequalities. By the equivalence of separation and optimization (cf. [9]) we can solve $\min_{x\in\mathcal{X}} r(Cx)$ in polynomial time. □

Remark. A special case of convex optimization problems are linear programs (LPs). From our result it follows that we can exactly compute reference point solutions for multicriteria LPs. It also yields a nice alternative proof of the existence of an FPTAS for the Pareto set, which has first been proven in [16] using an involved geometric argument.

A different argument for the approximability of Pareto sets of linear programs has independently been noted by Mittal and Schulz [13]. They use it to approximately optimize low-rank functions over polytopes.

Corollary 4.4. *Let $\mathcal{P}$ be a multicriteria minimization problem with convex feasible set and linear objective functions. Assume there is a positive polynomial π such that $\mathcal{Y} \subseteq \{y \in \mathbb{Q}^k : y_i \geq \frac{1}{\pi(|I|)}\ \forall i \in [k]\}$, where $|I|$ is the encoding length of the instance. If there is a polynomial time algorithm for the separation problem of $\mathcal{P}$, then there is an FPTAS for the Pareto set.*

Remark (Convex sets and the integrality assumption)**.** Note that our general integrality assumption $\mathcal{Y} \subseteq \mathbb{Z}^k_{\geq 0}$ for discrete optimization problems introduced in Section 2 does not hold for the case of convex optimization problems in Theorem 4.3 and Corollary 4.4. However, by assuming $y_i \geq \frac{1}{\pi(|I|)}$ for all occurring objective values in Corollary 4.4, we ensure that all prerequisites stated in [16] for Theorem 3.7 are still fulfilled. Furthermore observe that, while our proof of Lemma 3.8 also assumed integral objectives, we used this integrality assumption only for showing that if the solution y' computed by algorithm A_2 fulfills $r(y') \leq r(y)$ then $y_i = 0$ implies $y'_i = 0$. However, we can ignore this case, as by our assumption all objectives are strictly positive and thus $y_i = 0$ already implies that the answer to Gap is negative. Thus, both Lemma 3.8 and Theorem 3.7 are still valid for convex optimization problems fulfilling the condition of Corollary 4.4.

Proof of Corollary 4.4. By Theorem 4.3, we can compute an optimal solution to $\mathrm{RP}(\mathcal{P}, \|\cdot\|_\infty)$ for any reference point in polynomial time. Thus, by Lemma 3.8 we can solve $\textsc{Gap}(\mathcal{P}, 1+\varepsilon)$ in polynomial time for any $\varepsilon > 0$ (with running time independent of ε), which by Theorem 3.7 gives an FPTAS for the Pareto set. □

Approximation through LP Rounding. One of the most successful techniques for the design of approximation algorithms for integer problems is *LP rounding*: The problem is formulated as a linear integer program (IP), then the integrality constraints are relaxed and the resulting LP is solved, and finally the optimal fractional solution is rounded to a feasible integral solution, losing only a certain factor in the objective.

We show that these algorithms can be adapted such that they also approximate the reference point version of the problem with the same approximation factor.

Theorem 4.5. *Consider a multicriteria minimization problem $\min_{x \in \mathcal{X}} Cx$ with a solution set $\mathcal{X} \subseteq \mathbb{Z}^n_{\geq 0}$ and a cost matrix $C \in \mathbb{Q}^{k \times n}$. If there exist*

- *a convex relaxation $\mathcal{X}'$ for which the separation problem can be solved in polynomial time, and*
- *a polynomial time rounding procedure $\mathcal{R} : \mathcal{X}' \to \mathcal{X}$ such that for all $c \in \mathbb{Q}^n_{\geq 0}$ and all $x' \in \mathcal{X}'$ it holds that $c^{\mathrm{T}}\mathcal{R}(x') \leq \alpha c^{\mathrm{T}} x'$,*

then for any feasible reference point y^{rp} and any $p \in [1, \infty]$ there is a α-approximation algorithm for $\min_{x \in \mathcal{X}} r(Cx)$ with $r(y) = \langle\langle y^{\mathrm{rp}} \rangle\rangle_p + \langle\langle y - y^{\mathrm{rp}} \rangle\rangle_p$.

Proof. From Theorem 4.3 it follows that we can compute in polynomial time a fractional solution $x' \in \mathcal{X}'$ minimizing $r(Cx)$. Let $x = \mathcal{R}(x')$. Then

$$\begin{aligned} r(Cx) &= \max_{i \in [k]}\{y_i^{\mathrm{rp}}\} + \max_{i \in [k]}\{(Cx)_i - y_i^{\mathrm{rp}}\} + \frac{1}{p} \cdot \sum_{i \in [k]} (Cx)_i \\ &\leq \max_{i \in [k]}\{y_i^{\mathrm{rp}}\} + \max_{i \in [k]}\{\alpha (Cx')_i - y_i^{\mathrm{rp}}\} + \alpha \cdot \frac{1}{p} \cdot \sum_{i \in [k]} (Cx')_i \\ &= \max_{i \in [k]}\{y_i^{\mathrm{rp}}\} + \max_{i \in [k]}\left\{\alpha\big((Cx')_i - y_i^{\mathrm{rp}}\big) + (\alpha - 1) y_i^{\mathrm{rp}}\right\} + \alpha \cdot \frac{1}{p} \cdot \sum_{i \in [k]} (Cx')_i \\ &\leq \alpha \cdot \max_{i \in [k]}\{y_i^{\mathrm{rp}}\} + \alpha \cdot \max_{i \in [k]}\left\{(Cx')_i - y_i^{\mathrm{rp}}\right\} + \alpha \cdot \frac{1}{p} \cdot \sum_{i \in [k]} (Cx')_i \\ &= \alpha \cdot r(Cx') \,. \end{aligned}$$

□

This immediately results in the approximability, (with a factor independent of k,) of reference point solutions and the Pareto set for several classical combinatorial optimization problems. We give two examples here.

For Set Cover, in 1982 Hochbaum [12] presented an LP-based κ-approximation algorithm, where κ is the maximum cardinality of a set, thus there is a κ-approximation algorithm for the corresponding reference point version and a $\mathcal{O}(\kappa^2)$-approximation algorithm for the Pareto set. A notable special case is Vertex Cover, where $\kappa = 2$.

For the scheduling problem of minimizing the weighted sum of completion times on a single machine with release dates $(1|r_j|\sum w_j C_j)$ Hall et al. [10] gave a 3-approximation algorithm based on an LP-relaxation, resulting in a 3-approximation for compromise solutions, which gives a constant factor approximation for the Pareto set as well. Möhring et al. [14] extended this to stochastic scheduling with random processing times $(P|p_j \sim \text{stoch}, r_j|E[\sum w_j C_j])$, for which we consequently also get constant factor approximation for the multicriteria problems.

Remark. While we usually restrict ourselves to the case of a constant number of criteria, the results on convex optimization and LP-rounding also hold for a polynomial number of criteria. This is due to the fact that we can still solve the linear program if we add a polynomial number of constraints.

From Pseudopolynomial Algorithms to Approximation Schemes. Multicriteria optimization and in particular the concept of compromise solutions is closely related to robust optimization. If each criterion is considered as one *scenario* in the robust setting, then a compromise solution w.r.t. $\|\cdot\|_\infty$ is exactly the same as a *min-max regret robust* solution.

Aissi et al. [1] consider this robust setting and show that if upper and lower bounds on the optimum can be computed that only differ by a polynomial factor, and if there is a pseudopolynomial algorithm whose running time depends on the size of the instance and the upper bound, then there is an FPTAS for the min-max regret robust problem. We show that this result can be extended to reference point solutions.

Theorem 4.6. *Consider a multicriteria minimization problem with a set of feasible solutions* $\mathcal{X} \subseteq \{0,1\}^n$ *and cost matrix* $C \in \mathbb{Z}^{k\times n}_{\geq 0}$. *For any* $p \in [1, \infty]$, *if*

1. *for any instance* $I = (\mathcal{X}, C)$, *and any feasible reference point* y^{rp}, *a lower and an upper bound* L *and* U *on* $\min_{x\in\mathcal{X}} r(Cx)$ *can be computed in time* $\pi_1(|I|)$, *such that* $U \leq \pi_2(|I|)L$, *where* π_1 *and* π_2 *are non-decreasing polynomials,*

2. *and there exists an algorithm that solves* $\min_{x\in\mathcal{X}} r(Cx)$ *for any instance* $I = (\mathcal{X}, C)$ *in time* $\pi_3(|I|, U)$, *where* π_3 *is a non-decreasing polynomial,*

then there is an FPTAS for $\min_{x\in\mathcal{X}} r(Cx)$, *where* $r(y) = \langle\!\langle y^{\text{rp}} \rangle\!\rangle_p + \langle\!\langle y - y^{\text{rp}} \rangle\!\rangle_p$.

By $|I|$ we denote the encoding length of the instance I.

Proof. To compute a $(1+\varepsilon)$-approximation to the reference point solution, we set $\varepsilon' = \varepsilon(1+\frac{k}{p})^{-1}$ and apply the pseudopolynomial algorithm to a modified instance $\overline{I}$ with cost coefficients $\overline{c}_{ij} := \lfloor \frac{3n}{\varepsilon' L} c_{ij} \rfloor$. Then

$$\frac{\varepsilon' L}{3n} \cdot \overline{c}_{ij} \leq c_{ij} < \frac{\varepsilon' L}{3n}(\overline{c}_{ij} + 1) \ .$$

The reference point for the modified instance is defined by $\overline{y}_i^{\text{rp}} := \lfloor \frac{3n}{\varepsilon' L} y_i^{\text{rp}} \rfloor$. This reference point is feasible for the modified instance, and it holds that

$$\frac{\varepsilon' L}{3n} \overline{y}_i^{\text{rp}} \leq y_i^{\text{rp}} < \frac{\varepsilon' L}{3n} \overline{y}_i^{\text{rp}} + \frac{\varepsilon' L}{3n} < \frac{\varepsilon' L}{3n} \overline{y}_i^{\text{rp}} + \frac{\varepsilon' L}{3} \ . \tag{1}$$

Let x^* and $\overline{x}^*$ be reference point solutions for I and $\overline{I}$, respectively. We now bound the value of $\overline{x}^*$ w.r.t. the original costs c. Let r and $\overline{r}$ denote the relative distance for the original and the modified costs, respectively. We get

$$\begin{aligned} r(C\overline{x}^*) &= \langle\!\langle y^{\mathrm{rp}} \rangle\!\rangle_p + \langle\!\langle C\overline{x}^* - y^{\mathrm{rp}} \rangle\!\rangle_p \\ &\leq \frac{\varepsilon' L}{3n} \langle\!\langle \overline{y}^{\mathrm{rp}} \rangle\!\rangle_p + \frac{\varepsilon' L}{3}\left(1 + \frac{k}{p}\right) \\ &\quad + \max_{i\in[k]} \left\{ \frac{\varepsilon' L}{3n}(\overline{c}_i \overline{x}^* - \overline{y}_i^{\mathrm{rp}}) \right\} + \frac{\varepsilon' L}{3} + \frac{1}{p} \sum_{i\in[k]} \left(\frac{\varepsilon' L}{3n}(\overline{c}_i \overline{x}^* - \overline{y}_i^{\mathrm{rp}}) + \frac{\varepsilon' L}{3} \right) \\ &= \frac{\varepsilon' L}{3n} \langle\!\langle \overline{y}^{\mathrm{rp}} \rangle\!\rangle_p + \frac{\varepsilon' L}{3n} \langle\!\langle \overline{C}\overline{x}^* - \overline{y}_i^{\mathrm{rp}} \rangle\!\rangle_p + \frac{2\varepsilon' L}{3}\left(1 + \frac{k}{p}\right) \\ &\leq \frac{\varepsilon' L}{3n} \langle\!\langle \overline{y}^{\mathrm{rp}} \rangle\!\rangle_p + \frac{\varepsilon' L}{3n} \langle\!\langle \overline{C}x^* - \overline{y}_i^{\mathrm{rp}} \rangle\!\rangle_p + \frac{2\varepsilon' L}{3}\left(1 + \frac{k}{p}\right) \\ &\leq \langle\!\langle y^{\mathrm{rp}} \rangle\!\rangle_p + \frac{\varepsilon' L}{3n}\left(\frac{3n}{\varepsilon' L} \langle\!\langle Cx^* - y^{\mathrm{rp}} \rangle\!\rangle_p + n\Big(1 + \frac{k}{p}\Big) \right) + \frac{2\varepsilon' L}{3}\left(1 + \frac{k}{p}\right) \\ &= \langle\!\langle y^{\mathrm{rp}} \rangle\!\rangle_p + \langle\!\langle Cx^* - y^{\mathrm{rp}} \rangle\!\rangle_p + \varepsilon' L \left(1 + \frac{k}{p}\right) \\ &= r(Cx^*) + \varepsilon L \\ &\leq (1+\varepsilon) r(Cx^*) \ . \end{aligned}$$

It remains to be shown that $\overline{x}^*$ can be computed in time polynomial in $|I|$ and $\frac{1}{\varepsilon}$. For this, denote by $\overline{L}$ and $\overline{U}$ the lower and upper bounds on the optimal value $\overline{\mathrm{OPT}}$ of the modified instance $\overline{I}$. According to the requirements of the theorem we can compute L and then $\overline{x}^*$ in time

$$\begin{aligned} \pi_1(|I|) + \pi_3(|\overline{I}|, \overline{U}) &\leq \pi_1(|I|) + \pi_3(|\overline{I}|, \pi_2(|\overline{I}|)\overline{L}) \\ &\leq \pi_1(|I|) + \pi_3(|\overline{I}|, \pi_2(|\overline{I}|)\overline{\mathrm{OPT}}) \\ &\leq \pi_1(|I|) + \pi_3\left(|\overline{I}|, \pi_2(|\overline{I}|)\big(\tfrac{3n}{\varepsilon'}\pi_2(|I|) + n(1 + \tfrac{k}{p})\big)\right) \ , \end{aligned}$$

where the last inequality holds because

$$\begin{aligned} \overline{\mathrm{OPT}} &\leq \frac{3n}{\varepsilon' L} \langle\!\langle y^{\mathrm{rp}} \rangle\!\rangle_p + \frac{3n}{\varepsilon' L} \langle\!\langle Cx^* - y^{\mathrm{rp}} \rangle\!\rangle_p + n\Big(1 + \frac{k}{p}\Big) \\ &\leq \frac{3n}{\varepsilon' L} \cdot U + n\Big(1 + \frac{k}{p}\Big) \\ &\leq \frac{3n}{\varepsilon'} \cdot \pi_2(|I|) + n\left(1 + \frac{k}{p}\right) \ . \end{aligned}$$

Finally note that $|\overline{I}| \leq \pi_4(|I|, \log\frac{1}{\varepsilon}, \log\frac{k}{p})$ for some polynomial π_4, thus the above calculations prove that the running time is indeed polynomial.

For compromise solutions, we can not choose the reference point of the modified instance as we see fit. However, also for the ideal point eq. (1) still holds. To see this, denote the respective ideal points by y^{id} and $\overline{y}^{\mathrm{id}}$ and let $x^{(i)}, \overline{x}^{(i)}$ for $i \in [k]$ be optimal solutions of $\min_{x\in\mathcal{X}} c_i x$ and $\min_{x\in\mathcal{X}} \overline{c}_i x$.

It holds that

$$\begin{aligned} y_i^{\mathrm{id}} &= c_i x^{(j)} \geq c_i \overline{x}^{(j)} \geq \frac{\varepsilon' L}{3n} \overline{c}_i \overline{x}^{(j)} = \frac{\varepsilon' L}{3n} \overline{y}_i^{\mathrm{id}} \ , \\ y_i^{\mathrm{id}} &= c_i x^{(j)} \leq \frac{\varepsilon' L}{3n} (\overline{c}_i + \mathbb{1}^{\mathrm{T}}) x^{(j)} \leq \frac{\varepsilon' L}{3n} \overline{c}_i x^{(j)} + \frac{\varepsilon' L}{3} \leq \frac{\varepsilon' L}{3n} \overline{y}_i^{\mathrm{id}} + \frac{\varepsilon' L}{3} \ , \end{aligned}$$

so eq. (1) also holds for the ideal points and the rest of the proof works as before. □

Remark. For the running time it is essential that p is fixed or at least bounded from below by a positive constant (e.g. $p \geq 1$), as the running time is only polynomial in $\frac{1}{p}$. Since for $p \to 0$ compromise solutions at some point is equivalent to the weighted sum technique this is only a minor restriction.

Remark. Theorem 4.6 also holds for $\mathrm{CP}(\mathcal{P}, \langle\!\langle\cdot\rangle\!\rangle_p)$.

Similarly to Proposition 1 in [1], we can show that the necessary bounds U and L can be computed if the single-objective problem is tractable. This is a direct implication of the weighted sum approximation described in Theorem 4.1.

Corollary 4.7. *If there is an α-approximation for the single-criterion version of $\mathcal{P}$, then for all instances of $\mathrm{RP}(\mathcal{P}, \|\cdot\|)$ we can compute L such that $L \leq \min_{y\in\mathcal{Y}} r(y) \leq \alpha k L$.*

The pseudopolynomial algorithms for the shortest path problem (SP) and the minimum spanning tree problem (MST) presented in [1] can be used to compute reference point solutions as well, as they both compute all (non-dominated) regret vectors (that obey the upper bound U), and the reference point solution always has a non-dominated regret vector.

Corollary 4.8. *There is an FPTAS for $\mathrm{RP}(\mathrm{SP}, \langle\!\langle\cdot\rangle\!\rangle_p)$ and $\mathrm{RP}(\mathrm{MST}, \langle\!\langle\cdot\rangle\!\rangle_p)$ for any $p \in [1, \infty]$.*

5 Conclusion

Reference point methods are a popular tool of practitioners in multicriteria optimization. To the best of our knowledge, this paper provides the first extensive theoretical study of these methods in the context of approximability.

Our main result gives a new twist to approximation in multicriteria optimization: We show that the approximability of the Pareto set of a *multi*criteria problem can in fact be reduced to the approximation of a *single* objective, namely the relative distance to a feasible reference point. We believe that this insight can spark off new and interesting research in this field. The applicability of the presented methods is supported by several examples, such as the approximation of reference point solutions through LP rounding or dynamic programming.

Acknowledgment. The authors would like to thank Günter Ziegler for his help in simplifying the proof of Lemma 3.3.

References

[1] H Aissi, C Bazgan, and D Vanderpooten. Approximating min-max (regret) versions of some polynomial problems. In *Computing and Combinatorics*, volume 4112 of *LNCS*, pages 428–438. Springer Berlin / Heidelberg, 2006.

[2] H Aissi, C Bazgan, and D Vanderpooten. Approximation of min-max and min-max regret versions of some combinatorial optimization problems. *European J. Oper. Res.*, 179(2):281–290, 2007.

[3] R Caballero, M Luque, J Molina, and F Ruiz. Promoin: an interactive system for multiobjective programming. *Journal of Information Technology & Decision Making*, 1(4):635–656, 2002.

[4] TCE Cheng, A Janiak, and MY Kovalyov. Bicriterion single machine scheduling with resource dependent processing times. *SIAM J. Optim.*, 8(2):617–630, 1998.

[5] E U Choo and R E Steuer. An interactive weighted Tchebycheff procedure for multiple objective programming. *Math. Program.*, 26(3):326–344, 1983.

[6] M Ehrgott, J Figuera, and S Greco, editors. *Multiple Criteria Decision Analysis: State of the Art Surveys.* Springer New York, 2005.

[7] M Freimer and P L Yu. Some new results on compromise solutions for group decision problems. *Management Sci.*, 22(6):688–693, 1976.

[8] W B Gearhart. Compromise solutions and estimation of the noninferior set. *J. Optim. Theory Appl.*, 28:29–47, 1979.

[9] M Grötschel, L Lovász, and A Schrijver. The Ellipsoid Method and its Consequences in Combinatorial Optimization. *Combinatorica*, 1(2):169–197, 1981.

[10] L A Hall, A S Schulz, D B Shmoys, and J Wein. Scheduling to minimize average completion time: Off-line and on-line approximation algorithms. *Math. Oper. Res.*, 22:513–544, 1997.

[11] P Hansen. Bicriterion path problems. In *Proc. 3rd Conf. Multiple Criteria Decision Making Theory and Application*, volume 177 of *LNEMS*, pages 109–127. Springer Verlag, 1979.

[12] D S Hochbaum. Approximation algorithms for the set covering and vertex cover problems. *SIAM J. Comput.*, 11(3):555–556, 1982.

[13] S Mittal and A S Schulz. An FPTAS for optimizing a class of low-rank functions over a polytope. *Math. Program.*, published online 2012.

[14] R H Möhring, A S Schulz, and M Uetz. Approximation in Stochastic Scheduling: The Power of LP-Based Priority Policies. *J. ACM*, 46(6):924–942, 1999.

[15] S Opricovic and G-H Tzeng. Compromise solution by MCDM methods: A comparative analysis of VIKOR and TOPSIS. *European J. Oper. Res.*, 156:445–455, 2004.

[16] C H Papadimitriou and M Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *Proc. of FOCS*, pages 86–92, 2000.

[17] H M Safer. *Fast Approximation Schemes for Multi-Criteria Combinatorial Optimization.* PhD thesis, MIT, 1992.

[18] M Voorneveld, A van den Nouweland, and R McLean. Axiomatizations of the euclidean compromise solution. *Internat. J. Game Theory*, 40(3):427–448, 2011.

[19] A Warburton. Approximation of pareto optima in multiple-objective, shortest-path problems. *Oper. Res.*, 35(1):70–79, 1987.

[20] P L Yu. A class of solutions for group decision problems. *Management Sci.*, 19(8):936–946, 1973.

Appendix

Lemma (3.8 revisited, for $\langle\!\langle\cdot\rangle\!\rangle_p$ and $\|\cdot\|_p$)**.** *Let $\alpha > 1$ and set $\beta := \frac{\alpha^2}{2\alpha-1}$. There is a polynomial time algorithm for* GAP$(\mathcal{P}, \alpha)$*, if there are two polynomial time algorithms A_1, A_2 such that,*

- *given an instance of $\mathcal{P}$, algorithm A_1 computes in polynomial time a feasible reference point $y^{\mathrm{rp}} \in \mathbb{Q}^k_{\geq 0}$ for that instance, and,*
- *additionally given y^{rp} and $\lambda \in \mathbb{Q}^k_{\geq 0}$ and $p \geq 1$, algorithm A_2 computes in polynomial time a solution $y' \in \mathcal{Y}$ with $r(y') \leq \beta \min_{z \in \mathcal{Y}} r(z)$ for $r(z) = \langle\!\langle y^{\mathrm{rp}} \rangle\!\rangle_p^\lambda + \langle\!\langle y^{\mathrm{rp}} - z \rangle\!\rangle_p^\lambda$ or $r(z) = \|y^{\mathrm{rp}}\|_p^\lambda + \|y^{\mathrm{rp}} - z\|_p^\lambda$, respectively.*

Proof. Let $y \in \mathbb{Q}^k_{\geq 0}$ be the input to the GAP problem. W.l.o.g., we can assume that $y \geq \alpha y^{\mathrm{rp}}$ for the reference point y^{rp} computed by A_1, as otherwise there is no $y' \leq \frac{1}{\alpha} y$ and GAP can answered negatively.

We will solve the GAP problem with a single call of the β-approximation algorithm for RP$(\mathcal{P}, \langle\!\langle\cdot\rangle\!\rangle_p)$ (or RP$(\mathcal{P}, \|\cdot\|_p)$, respectively) with

$$p := \max\left\{ \frac{\log k}{\log(1 + \frac{1}{2M})}, 2kMq \right\} ,$$

where q is the largest denominator of all the components in y and M is an upper bound on the objectives in $\mathcal{Y}$.

Let $I := \{i \in [k] : y_i = y_i^{\mathrm{rp}} = 0\}$. For $i \in [k]$, $\lambda_i = \begin{cases} 2 & \text{if } i \in I \\ \frac{1}{y_i - y_i^{\mathrm{rp}}} & \text{otherwise.} \end{cases}$

Let y' be a β-approximation to $\min_{z \in \mathcal{Y}} r(z)$.

If $r(y') \leq r(y)$, we return y' as answer to the GAP problem: We observe

$$\lambda_i (y'_i - y_i^{\mathrm{rp}}) \leq \langle\!\langle y' - y^{\mathrm{rp}} \rangle\!\rangle_p^\lambda \leq \langle\!\langle y - y^{\mathrm{rp}} \rangle\!\rangle_p^\lambda \leq 1 + \tfrac{k}{p} .$$

If $i \in I$, we have $y'_i \leq \frac{1}{2}(1 + \frac{1}{2Mq}) < 1$. If $i \notin I$, then $y'_i \leq (1 + \frac{k}{p}) y_i \leq y_i + \frac{1}{2q} < y_i + 1$. In both cases, integrality of y' implies $y'_i \leq y_i$. The same holds for the $\|\cdot\|_p$-norm with $1 + \frac{k}{p}$ replaced by $\sqrt[p]{k}$—in this case, the choice of p guarantees $\sqrt[p]{k} \cdot y'_i < y'_i + 1$.

If $r(y') > r(y)$, we answer GAP negatively: Let $y'' \in \mathcal{Y}$. We show that there is an $i \in [k]$ with $y''_i > \frac{1}{\alpha} y_i$. This is true if $y''_i > 0 = y_i$ for any $i \in I$. Thus, we can restrict to the projection of $\mathbb{Q}^k$ to the components in $[k] \setminus I$, and w.l.o.g. assume $I = \emptyset$. First observe that $\beta r(y'') \geq r(y') > r(y)$, which implies

$$\beta \langle\!\langle y'' - y^{\mathrm{rp}} \rangle\!\rangle_p^\lambda > \langle\!\langle y - y^{\mathrm{rp}} \rangle\!\rangle_p^\lambda - (\beta - 1) \langle\!\langle y^{\mathrm{rp}} \rangle\!\rangle_p^\lambda .$$

It is easy to verify that $\langle\!\langle z \rangle\!\rangle_p^\lambda \leq (1 + \frac{k}{p}) \|z\|_\infty^\lambda$ for all $z \in \mathbb{Q}^k$, and furthermore $\langle\!\langle y - y^{\mathrm{rp}} \rangle\!\rangle_p^\lambda = 1 + \frac{k}{p}$ as $I = \emptyset$. This yields

$$(1 + \tfrac{k}{p}) \beta \|y'' - y^{\mathrm{rp}}\|_\infty^\lambda > (1 + \tfrac{k}{p}) \|y - y^{\mathrm{rp}}\|_\infty^\lambda - (1 + \tfrac{k}{p})(\beta - 1) \|y^{\mathrm{rp}}\|_\infty^\lambda ,$$

which brings us back to the case of the $\|\cdot\|_\infty$-norm. The same holds true for the $\|\cdot\|_p$-norm with the factor $1 + \frac{k}{p}$ replaced by $\sqrt[p]{k}$. □

Corollary (3.9 revisited, for $\langle\!\langle\cdot\rangle\!\rangle_p$ and $\|\cdot\|_p$)**.** *Let $\alpha > 1$ and set $\beta := \sqrt{\frac{\alpha^2}{2\alpha-1}}$. There is a polynomial time algorithm for* GAP$(\mathcal{P}, \alpha)$*, if there is a β-approximation algorithm for* CP$(\mathcal{P}, \langle\!\langle\cdot\rangle\!\rangle_p)$ *(*CP$(\mathcal{P}, \|\cdot\|_p)$*, respectively) for every $p \geq 1$ and the running time of all algorithms is bounded by a polynomial in the instance size and* $\log(p)$*.*

Proof. The proof is identical to that of Corollary 3.9 given in the paper. In fact, the second part of this proof only uses properties of monotone norms. □