Chapter 1

Introduction

In classical optimization and mathematical programming, the usual assumption is that of a single objective or *criterion*, with respect to which the available alternatives are evaluated, in order to achieve an *optimal solution*. In practice, however, decision makers are often confronted with several, usually conflicting criteria, that they want to optimize simultaneously. Think, for example, of the classical shortest path problem. In the application of a route guidance system, a driver probably not only cares about reaching her target as quickly as possible, but might also be concerned about fuel consumption and road charges, minimizing her monetary expenses. Often highways are the fastest connections, but subject to tolls. Municipal roads, on the other hand, are usually free of charge, while causing a longer travel time. Consequently, in most cases there will not exist a route which is both quickest and cheapest—the objectives are in conflict.

1.1 A Motivation of reference point methods and their approximation

Multicriteria optimization and Pareto optimality. This setting of several conflicting criteria, that are to be optimized simultaneously, is the subject of the area of *multicriteria optimization* (also called *multiobjective optimization* or *multiobjective programming*). Due to the non-existence of a clear preference relation, alternative solution concepts have been developed. The central notion in this context is that of *Pareto optimality*. It goes back to the economist and sociologist Vilfredo Pareto (1848–1923). In his *Manuale di Economia Politica* [Par06] from 1906, he writes[1]:

> We will begin by defining a term which is desirable to use in order to avoid prolixity. We will say that the members of a collectivity enjoy

[1] English translation from 1971 [Par71], as cited by Ehrgott in [Ehr12]

> *maximum ophelimity* in a certain position when it is impossible to find a way of moving from that position very slightly in such a manner that the ophelimity enjoyed by each of the individuals of that collectivity increases or decreases. That is to say, any small displacement in departing from that position necessarily has the effect of increasing the ophelimity which certain individuals enjoy, and decreasing that which others enjoy, of being agreeable to some and disagreeable to others.

Pareto refers here to the satisfaction (*ophelimity*) of individuals in a society, rather than to objectives in an optimization problem, but the principle is the same: We say that a solution is *Pareto optimal* (i.e., enjoys maximum ophelimity), if it cannot be improved in one criterion without deteriorating another. A formal definition of this notion is given in Section 1.3.

Informally speaking, a Pareto optimal solution is a solution that is not obviously sub-optimal: There is no other solution that is at least as good in all criteria, and strictly better in at least one criterion. If such a solution existed, it would *dominate* the former solution, and would clearly be preferred. This is why Pareto optimal solutions are often referred to as *non-dominated* solutions. Another term that is used frequently is that of *efficient* solutions.

The observation that Pareto optimal solutions are exactly the non-dominated solutions leads us directly to the major disadvantage of this concept: The dominance relation is not *total*, i.e. there might be two distinct elements that are not comparable—neither dominates the other. In general, there can be a lot of Pareto optimal solutions. Their number can even be exponential (in the size of the instance) for discrete problems (cf. [EP92, Han80]), and infinite in continuous optimization. Since usually it is infeasible to consider (or even to compute or output) all these solutions, a refinement of the concept of Pareto optimality was indispensable.

Weighted sum. For several decades now, researchers have come up with various ways of generating a (total) preference relation on the set of Pareto optimal solutions. Very often this is achieved by the means of *scalarization*, i.e. the aggregation of several objectives (or one *vectorial* value) into *one scalar* value. The search for a solution is thus reduced again to the optimization of a (one-dimensional) *scalarizing function*.

The most straightforward scalarizing function—at least to mathematicians[2]—is the *weighted sum* approach. Assuming that all objectives are quantifiable and of the same sense of optimization (e.g., all minimization), an objective vector is scalarized by multiplying each entry with a criterion-specific weight and summing the products. This weighted sum of the individual criteria is what we then aim to minimize (respectively maximize), inducing a total preference relation on the set of Pareto optimal solutions.

[2] As we will discuss further below, to the "average" human being this concept is not as natural.

An optimal solution with respect to the weighted sum (with positive weights) is guaranteed to be Pareto optimal. However, in general not all Pareto optimal solutions can be generated this way. There might be Pareto optimal solutions that, no matter how we choose the weights, are never optimal with respect to the weighted sum scalarization. Consider the example from the beginning, the bicriteria shortest path problem. We want to choose a route, minimizing both travel time and monetary expenses. There are three feasible routes, with objective vectors $(1, 10)$, $(6, 6)$, and $(10, 1)$. No matter which weights we choose, we will always either get the first or the last solution as a scalarization optimum, and never the second one, although this solution is very balanced, and its scalarization value is only slightly worse than the other solutions' values, for identical weights for time and cost.

Another drawback of the weighted sum approach is that human beings tend to have difficulties to handle weights and fully understand their effects. Wierzbicki, for example, stated in 1986 that

> experience in applications of decision support systems shows that weighting coefficients are not easy to be understood well and interpreted by an average user. [Wie86]

This is the reason why alternative ways to model the decision maker's preferences have been developed. One of the most popular among them are reference points.

Reference point methods and compromise programming. Both disadvantages of the weighted sum approach—the disability to generate all Pareto optimal solutions, and the bad interpretability of the parameters—are tackled simultaneously by the concept of *reference point solutions*.

In this concept, the decision maker's preferences are not modeled by weighting the criteria, but through *reference* or *aspiration levels*. In each criterion, an aspired objective value is specified. These values form the coordinates of the so-called *reference point* in the objective space. In the simplest model, one solution is preferred over another if it is closer to this reference point.

A reference point of particular interest is the so-called *ideal point*, which is obtained by optimizing each criterion individually. The reference level in a certain criterion is thus its optimal value, when all other criteria are disregarded. The problem of finding a reference point solution with respect to the ideal point is usually referred to as *compromise programming*, and the corresponding solutions are called compromise solutions.

In some cases, a specification of a weight for each objective might actually be infeasible, while reference points have a clear interpretation. An example are infinite-dimensional objective spaces, e.g. if the object of optimization is a function. Wierzbicki [Wie80b], for example, discusses trajectories of inflation rates and the gross national product over time. Clearly we cannot specify infinitely many weights.

Still, a decision maker might be able to give a target trajectory that she would like to approach as closely as possible.

Already in the 1950s, Herbert A. Simon (1916–2001), who received the Nobel Prize in Economics in 1978 "*for his pioneering research into the decision-making process within economic organizations*"[3], considered a model of decision making based on reference levels. According to his *Behavioral model of rational choice* [Sim55], humans strive to satisfy goals, rather than optimize a utility function. This strategy is termed *satisficing*[4] *behavior* (see the introduction to Part IV in Simon's *Models of Man* [Sim57], and his book *Organizations*, together with James G. March [MS58]). It is not restricted to human individuals, but extends to economic organizations. As Simon noted, *"there is some empirical evidence that business goals are, in fact, stated in satisficing terms*" [Sim57].

The satisficing theory, however, does not differentiate between any two solutions that both attain all aspiration levels, even if one dominates the other. Therefore in the 1980s, Wierzbicki extended it to *quasi-satisficing* decision making:

> A decision maker [. . .] is quasi-satisficing if he optimizes when his reservation or aspiration levels are not yet attained, but he can further optimize or forego the optimization for additional good reasons if his aspiration levels are attained. [Wie86]

According to the quasi-satisficing theory, a decision maker thus first strives to achieve all his reference values, but might also optimize beyond them, in case the reference point is achievable. This allows scalarizing functions that are based on reference points and have optima that are guaranteed to be Pareto optimal.

As we shall see below, scalarizations based on reference points can also have *every* Pareto optimal solution as an optimum. If the measurement of the distance to the reference point, as well as the optimization beyond the reference point, is done appropriately, the set of reference point solutions under varying parameters is thus exactly the set of Pareto optimal solutions.

This, together with the improved interpretability of the parameters, shows that reference point methods are a more powerful concept than the weighted sum technique. The conceptual advantage, however, comes at the cost of a higher complexity. Whereas optimization of the weighted sum scalarization for linear objectives boils down to the optimization of the classical single-criterion version of the problem, computing an optimal reference point solution is NP-hard already for simple combinatorial optimization problems such as the shortest path problem or the minimum spanning tree problem.[5]

[3] see http://www.nobelprize.org/nobel_prizes/economics/laureates/1978/. Retrieved on May 2, 2013.

[4] a portmanteau of *satisfy* and *suffice*

[5] See the part on robust optimization in Section 1.5, as well as Section 3.6, for further remarks on the hardness.

Approximation. After more than four decades of research on NP-hardness, it is widely believed, although without proof, that the complexity classes P and NP are not equal. If this is the case, then NP-hard problems cannot be solved efficiently, i.e. in a running time that is polynomial in the encoding length of the input. One can, of course, still solve these problems exactly, investing a lot of time. An alternative approach is to maintain the polynomial running time, and instead give up on optimality, still striving to get as close as possible to an optimal solution. This leads to the notion of *approximation algorithms*, which compute solutions whose objective values are provably within a certain factor of the optimum.

In the light of the complexity of finding optimal reference point solutions, in this work we analyze their approximability. In particular, we relate the approximability of reference point solutions to another concept in multicriteria optimization, *approximate Pareto sets*. These sets are an alternative way to deal with the huge number of Pareto optimal solutions. Instead of computing all of them, we content ourselves with a smaller set which contains, for every Pareto optimal solution, a solution that approximates the former, i.e. whose objective values are all within a certain factor of those of the Pareto optimal solution (see Section 1.3 for a formal definition).

In 2000, Papadimitriou and Yannakakis [PY00] showed that, under mild conditions, such an approximate set of reasonable (i.e. polynomial) size exists, and they gave a sufficient and necessary condition on the tractability of computing such a set. We use their results in order to show that approximability of the Pareto set is essentially equivalent to the approximability of reference point solutions.

1.2 Our contribution: The power of compromise

In this thesis, we discuss theory and application of reference point methods, including compromise programming. We mainly restrict to minimization problems and *utopian* reference points, i.e. reference points that are not achievable in any criterion.[6]

The concept of reference point solutions is widely spread in practice. There are numerous publications about their application, and they form an integral part of many state-of-the-art software tools in multicriteria decision making (MCDM). Their theory, on the other hand, is not so well studied. In particular, there are only few publications on the relation between reference point methods and the approximation of the Pareto set, although the latter is arguably the most prominent concept for multicriteria optimization in theoretical computer science.

Our main theoretical result (Theorem 2.1) shows the power of reference point methods—the *power of compromise*: Being able to approximate reference point solutions or compromise solutions is as good as being able to approximate the whole Pareto set. The problems are equivalent in terms of approximability, i.e., the existence of a polynomial time algorithm with a provable approximation guarantee. This

[6] An extension to arbitrary reference points is discussed in Chapter 5.

is true for both constant factor approximation and approximation schemes. Our result thus establishes a strong link between the two lines of research mentioned above: The work on approximate Pareto sets in the domain of theoretical computer science, and the applications of reference point methods in practical MCDM.

An immediate consequence of our result is the existence of approximation algorithms for reference point solutions in cases where the Pareto set can be approximated, e.g. the shortest path problem and the minimum spanning tree problem. On the other hand, we also show approximation results for reference point solutions directly, implying approximability of the Pareto set.

A notable example of such an approximation result is our extension of LP rounding, a technique used to design approximation algorithms for single-criterion optimization problems. We show that this algorithmic paradigm, under fairly general conditions, can be extended to reference point solutions. This implies in particular the approximability of the Pareto sets for the multicriteria version of several classical optimization problems, e.g. Set Cover and certain scheduling problems.

We also prove that minimizing a weighted sum of the objectives gives a constant factor approximation of the reference point solution. This has an interesting and surprising consequence: Whenever the weighted sum can be approximated within a constant factor, there also is a constant factor approximation algorithm for the Pareto set. This is in particular the case if all objectives are linear and the single-criterion problem has an approximation algorithm.

For maximization objectives, the picture looks more diverse: If the reference point can be chosen freely, the equivalence of approximability continues to hold. For compromise programming, i.e., when the reference point is restricted to the ideal point, an approximate Pareto set still gives approximate compromise solutions, but not vice versa.

All the results above hold in this generality only for utopian (i.e., unachievable) reference points. In Chapter 5, we present an extension to arbitrary, possibly achievable reference points. Obviously, the implication from the approximability of reference point solutions to the approximability of the Pareto set carries over, as we can simply choose the same utopian reference point as before. In the reverse direction, the resulting approximation factors depend on the norm. As a consequence, a constant factor approximation of the Pareto set always gives a constant factor approximation of the reference point solution, but for approximation schemes the implication only holds for a certain family of norms, the so-called cornered norms or augmented Chebychev norms.

All results in this first part are based on joint work with Christina Büsing (RWTH Aachen), Jannik Matuschke and Sebastian Stiller (both TU Berlin). Most results are also contained in a prospective journal paper, a preliminary version of which has been published online [GBMS13].

The theoretical results in Part I of the thesis are complemented in Part II by two examples from the practice of reference point methods.

The first example is an application of a reference point method to the area of sustainability. As sustainability deals with compromising between economic, environmental and social objectives, it can be seen as inherently multicriteria. Methods from multicriteria decision analysis are thus an appropriate tool to tackle these kind of problems.

In the presented project, the task was to evaluate the municipalities of Andalusia, Spain, in terms of sustainability. To this end, a row of economic, environmental, social and financial indicators were aggregated into a pair of synthetic sustainability indicators, using a double reference point method (with both aspiration and reservation levels). We present the methodology, discuss some particularities of the application, and analyze strengths and weaknesses of the applied method.

The second example is concerned with a computational issue. In practice, multicriteria optimization problems are often solved by heuristics, mainly using evolutionary algorithms. These algorithms generate a heuristic approximation of the Pareto set, or of a certain area of interest of this set. However, they usually do not have a provable approximation guarantee. It is therefore even more important to compare different algorithms empirically, which is not an easy task, due to the high dimension of the objective space.

We compare three genetic algorithms for multicriteria optimization problems, all of them based on reference points. Our comparative study makes use of performance indicators for multicriteria heuristics, two of which are closely related to approximation factors. The study gives valuable insights about the performance of the algorithms, that could not be revealed by previous studies based on plots of the objective vector set. We therefore believe that our study can help to improve the algorithms.

Both applications are based on joint work with Francisco Ruiz, Mariano Luque and Rubén Saborido (all Universidad de Málaga).

1.3 Preliminaries

Throughout this thesis, unless explicitly stated otherwise, we let $\mathcal{P}$ denote a multicriteria discrete optimization problem with k objectives. We usually assume k to be a fixed number. With some exceptions in Chapters 3, 5 and 6, we consider only minimization objectives. As we want to study approximation, we also restrict to non-negative objective values. In the context of discrete optimization, we can assume without loss of generality that all given numbers are integers.

An instance I of $\mathcal{P}$ is thus given by the set of feasible solutions $\mathcal{X} \subseteq \mathbb{Z}^n$, and the vector of objective functions $c : \mathcal{X} \to \mathbb{Z}^k_{\geq 0}$, that are to be minimized simultaneously. We write $I = (\mathcal{X}, c)$. The encoding length of the instance is denoted by $|I|$. Since the set $\mathcal{X}$ is usually given implicitly, in general $|I|$ is much less than $|\mathcal{X}|$. The *cost* *vector set* of the instance $I = (\mathcal{X}, c)$ is defined by $\mathcal{Y} := c(\mathcal{X}) \subseteq \mathbb{Z}^k_{\geq 0}$. Most of the time we will only be interested in the cost vectors, and not in the actual solutions. We thus usually omit the feasibilty set $\mathcal{X}$ and only deal with the set $\mathcal{Y}$.

We use the notation $[k]$ to denote the set $\{1, 2, \ldots, k\}$. For $y, y' \in \mathbb{R}^k$, we write $y \leq y'$ if and only if the inequality holds component-wise, i.e. $y_i \leq y'_i$ for all $i \in [k]$. With this, the notion of Pareto optimality can be defined as follows:

Definition 1.1 (Domination, Pareto optimality, Pareto set)**.** Let $\mathcal{P}$ be a multicriteria minimization problem with cost vector set $\mathcal{Y}$. For $y, y' \in \mathcal{Y}$, we say that y' *dominates* y, if $y' \leq y$ and $y'_i < y_i$ for some $i \in [k]$. A solution $y \in \mathcal{Y}$ is called *Pareto optimal* if there does not exist $y' \in \mathcal{Y}$ that dominates y. The *Pareto set* $\mathcal{Y}_P$ is the set of all Pareto optimal solutions.

Similar to Papadimitriou and Yannakakis [PY00], we will assume throughout this work that for any instance I, we can compute an exponential bound on the objective values of all solutions, i.e., a number $M > 0$ such that $\mathcal{Y} \subseteq [0, M]^k$ and such that there is a polynomial π with $M \leq 2^{\pi(|I|)}$. This is not a major restriction in usual discrete optimization problems. For example, if we can compute a vector u of upper bounds on all feasible solutions (i.e. $x \leq u$ for all $x \in \mathcal{X}$) and the cost functions are linear with non-negative coefficients c_{ij}, then $M = \max_{i \in [k]} \left\{ \sum_{j \in [n]} c_{ij} u_j \right\}$ is such an exponential upper bound.

Reference point methods. The decision maker's preferences are modeled by reference points, i.e. aspiration levels in all criteria. A special reference point of particular interest is the *ideal point*, which is defined as the point in the objective space obtained by optimizing each objective individually. In Chapters 2–4, we will restrict to the ideal point and those reference points that are beyond it (and thus unachievable), called *utopian* reference points.

Definition 1.2 ((Utopian) reference point, ideal point)**.** Let $\mathcal{Y} \subseteq \mathbb{Z}^k_{\geq 0}$ be the objective vector set of a multicriteria minimization problem. A *reference point* is a point $y^{\mathrm{rp}} \in \mathbb{Z}^k_{\geq 0}$ in the objective space. The *ideal point* $y^{\mathrm{id}} \in \mathbb{Z}^k_{\geq 0}$ is defined by $y^{\mathrm{id}}_i = \min_{y \in \mathcal{Y}} y_i$ for $i \in [k]$. A reference point y^{rp} is called *utopian* if $y^{\mathrm{rp}} \leq y^{\mathrm{id}}$.

In addition to the reference point, we will also allow as input a vector $\lambda \in \mathbb{Q}^k_{\geq 0}$ of weights, to adjust a fixed norm $\|\cdot\|$ on $\mathbb{R}^k$ by letting $\|\cdot\|^\lambda$ be the norm defined by $\|y\|^\lambda = \|(\lambda_1 y_1, \ldots, \lambda_k y_k)\|$.

Given a utopian reference point, the goal is to find a solution that is closest to this point w.r.t. $\|\cdot\|^\lambda$. Conceive of this distance as the price to pay in order to attain a compromise among the criteria. The objective value of a reference point solution is *the value of the reference point*, $\|y^{\mathrm{rp}}\|^\lambda$, *degraded by the price of compromise.* For minimization, the reference point objective function thus reads:

$$z_{y^{\mathrm{rp}},\lambda}(y) = \|y^{\mathrm{rp}}\|^\lambda + \|y - y^{\mathrm{rp}}\|^\lambda \; . \tag{1.1}$$

If at least some of the reference levels are achievable, this definition of course no longer makes sense, as in that case we do not want to minimize the distance to the reference point, but optimize beyond it. We will discuss in Chapter 5 how to extend the objective function to this case.

If we choose the ideal point as a reference point, this problem is referred to as *compromise programming*. Formally, the problems we consider in this thesis are thus defined as follows:

Definition 1.3 (Reference point solutions, compromise programming)**.** Let $\mathcal{P}$ be a multicriteria minimization problem, and $\|\cdot\|$ a norm on $\mathbb{R}^k$.

The problem of *reference point solutions*, $\mathrm{RP}(\mathcal{P}, \|\cdot\|)$ for short, is defined as follows: Given an instance $I = (\mathcal{X}, c)$ of $\mathcal{P}$, a utopian reference point $y^{\mathrm{rp}} \in \mathbb{Z}^k_{\geq 0}$, and a weight vector $\lambda \in \mathbb{Q}^k_{\geq 0}$ as input, find a solution $x \in \mathcal{X}$ that minimizes $z_{y^{\mathrm{rp}},\lambda}(c(x))$, where z is defined as in (1.1).

The problem of *compromise programming*, $\mathrm{CP}(\mathcal{P}, \|\cdot\|)$ for short, is defined as follows: Given an instance $I = (\mathcal{X}, c)$ of $\mathcal{P}$ and $\lambda \in \mathbb{Q}^k_{\geq 0}$, find a solution $x \in \mathcal{X}$ that minimizes $z_{y^{\mathrm{id}},\lambda}(c(x))$.

Approximation. In theoretical computer science, one way to tackle NP-hard problems is the design of algorithms that compute, in polynomial time, solutions that are guaranteed to be within a certain factor of the optimum. These algorithms are called *approximation algorithms*. In single-criterion optimization they are defined as follows:

Definition 1.4 (α-approximation algorithm)**.** Let $\mathcal{P}$ be a single-criterion minimization problem, and let $\alpha > 1$. An α-approximation algorithm is an algorithm that for any instance $I = (\mathcal{X}, c)$ of $\mathcal{P}$ computes a solution $x' \in \mathcal{X}$ in time polynomial in $|I|$ with $c(x') \leq \alpha \cdot \min_{x \in \mathcal{X}} c(x)$.

Sometimes we can even choose the approximation factor α arbitrarily close to 1. These algorithms are called‘ *approximation schemes*.

Definition 1.5 ((Fully) polynomial time approximation scheme)**.** Let $\mathcal{P}$ be a single-criterion minimization problem. A *polynomial time approximation scheme* (PTAS) is a family of algorithms that for every $\varepsilon > 0$ contains a $(1 + \varepsilon)$-approximation algorithm. A PTAS is called *fully polynomial time approximation scheme* (FPTAS), if the running time of the algorithms is polynomial in $\frac{1}{\varepsilon}$ (and $|I|$).

This concept has been extended to multicriteria optimization, e.g. by Papadimitriou and Yannakakis [PY00]. They only consider approximation schemes, however. To include constant factor approximations, we use a slightly different notation:

Definition 1.6 (α-approximate Pareto set, (F)PTAS for the Pareto set)**.** Let $\mathcal{P}$ be a multicriteria minimization problem with a set of cost vectors $\mathcal{Y} \subseteq \mathbb{Z}^k_{\geq 0}$, and let $\alpha > 1$.

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 of $\mathcal{P}$* is an algorithm that, for any instance I of $\mathcal{P}$, constructs an α-approximate Pareto set in time polynomial in $|I|$.

An *(F)PTAS for the Pareto set of $\mathcal{P}$* is a family of algorithms that, for all $\varepsilon > 0$, contains a $(1+\varepsilon)$-approximation algorithm for the Pareto set of $\mathcal{P}$ (with running time polynomial in $\frac{1}{\varepsilon}$ for FPTAS).

We conclude this preliminary section with some remarks on the reference point objective function, on a complexity issue, and on the norms we want to consider.

The constant in the objective. Since $\|y^{\text{rp}}\|^\lambda$ is a constant, exact minimization of (1.1) boils down to minimizing the distance $\|y - y^{\text{rp}}\|^\lambda$, as the level sets of this function are identical to that of the reference point objective function. Still, for judging the quality of an approximation, this short-cut is not permissible, as the following trivial example shows.

Consider a multicriteria problem defined by k unrelated copies of a single-criterion optimization problem, for which we have a tight approximation algorithm with factor α. Let the distance be measured in any norm, and choose the ideal point as a reference point. As the single criteria problems are unrelated, one expects that solving each problem separately by the approximation algorithm gives an $O(\alpha)$-approximation for the reference point solution. This is indeed true for the reference point objective function. However, for minimizing the distance, the ratio to the optimum is infinite, because the optimum attains the ideal point for the unrelated problems, and thus the optimal distance is zero.

Conversely, any approximation algorithm for the distance $\|y - y^{\text{rp}}\|$ could 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. Thus, we cannot hope for approximating the distance $\|y - y^{\text{rp}}\|$ for any problem that is NP-hard in the single-criterion version. In contrast to that, for the objective $z(y)$ we do get positive approximation results also for NP-hard problems.

Also note that an α-approximate solution for the distance to the reference point also is an α-approximate solution for the reference point objective function. Therefore, any result in this thesis that depends on the existence of an approximation for reference point methods in particular holds if the $\|y - y^{\text{rp}}\|$ can be approximately minimized.

Caveat on complexity. We want to remark 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, while in the latter case the reference point is part of the input. This leads to different consequences if the underlying single-criterion problem cannot be solved in polynomial time. In this case, the objective function of CP is hard to evaluate. However, in the context of approximability this is only a minor issue, as Corollary 2.5 shows.

For RP, on the other hand, it becomes hard to verify feasibility of the input. In our setting, a specified reference point is only feasible if it is utopian, i.e. $y^{\mathrm{rp}} \leq y^{\mathrm{id}}$. The best we can expect from an algorithm is to *approximately distinguish* between feasible and infeasible instances, i.e., an α-approximation algorithm needs to accept all feasible inputs and reject all instances where $y_i^{\mathrm{rp}} > \alpha \cdot y_i^{\mathrm{id}}$ for some $i \in [k]$, but it might also accept instances with slightly infeasible reference points, as long as $y^{\mathrm{rp}} \leq \alpha y^{\mathrm{id}}$.

We will get rid of this artifact with the extension discussed in Chapter 5.

Norms. Throughout this thesis, we will restrict to norms fulfilling the following two properties.

Definition 1.7 (Monotone and polynomially decidable norms)**.** A norm $\|\cdot\|$ on $\mathbb{R}^k$ is called *monotone*, if for any $y', y'' \in \mathbb{R}^k_{\geq 0}$, $y' \leq y''$ implies $\|y'\| \leq \|y''\|$. It is called *polynomially decidable*, if we can decide whether $\|y'\| \leq \|y''\|$ in time polynomial in the encoding length of y' and y''.

Remark. We need monotonicity in order to get Pareto optimality of reference point solutions. For a non-monotone norm, it might be the case that y dominates y', but still $z(y') < z(y)$. If the norm is monotone, however, there always exists at least one Pareto optimal reference point solution, since in this case from $y \leq y'$ it follows that $z(y) \leq z(y')$.

Remark. Note that not all norms are monotone. Consider for example the norm in $\mathbb{R}^2$ defined by

$$\|y\| = \|Ay\|_1 = \mathbb{1}^{\mathrm{T}} Ay, \qquad \text{with } \mathbb{1} = \begin{pmatrix} 1 \\ 1 \end{pmatrix},\ A = \begin{pmatrix} 1 & 1 \\ -2 & 2 \end{pmatrix}.$$

Consider now the vectors $y' = (0, 3)$ and $y'' = (4, 4)$. Then we have $y' \leq y''$, but $\|y'\| = 9 > 8 = \|y''\|$.