Discrete representation of the non-dominated set for multi-objective optimization problems using kernels*

Cristina Bazgan† Florian Jamain Daniel Vanderpooten‡

Université Paris-Dauphine, PSL Research University,
CNRS, LAMSADE, 75016 Paris, France
{cristina.bazgan,florian.jamain,daniel.vanderpooten}@lamsade.dauphine.fr

Abstract

In this paper, we are interested in producing discrete and tractable representations of the set of non-dominated points for multi-objective optimization problems, both in the continuous and discrete cases. These representations must satisfy some conditions of *coverage*, i.e. providing a good approximation of the non-dominated set, *spacing*, i.e. without redundancies, and *cardinality*, i.e. with the smallest possible number of points. This leads us to introduce the new concept of $(\varepsilon, \varepsilon')$-*kernels*, or ε-*kernels* when $\varepsilon' = \varepsilon$ is possible, which correspond to ε-Pareto sets satisfying an additional condition of ε'-*stability*. Among these, the kernels of small, or possibly optimal, cardinality are claimed to be good representations of the non-dominated set.

We first establish some general properties on ε-kernels. Then, for the bi-objective case, we propose some generic algorithms computing in polynomial time either an ε-kernel of small size or, for a fixed size k, an ε-kernel with a nearly optimal approximation ratio $1 + \varepsilon$. For more than two objectives, we show that ε-kernels do not necessarily exist but that $(\varepsilon, \varepsilon')$-kernels with $\varepsilon' \leq \sqrt{1+\varepsilon} - 1$ always exist. Nevertheless, we show that the size of a smallest $(\varepsilon, \varepsilon')$-kernel can be very far from the size of a smallest ε-Pareto set.

Keywords: Multiple objective programming, Pareto set, non-dominated points, discrete representation, exact and approximation algorithms, kernel.

1 Introduction

In multi-objective optimization, in opposition to single objective optimization, there is typically no optimal solution i.e. one that is best for all the objectives. The solutions of interest, called *efficient* solutions, are such that any solution which is better on one criterion is necessarily worse on at least one other criterion. In other words, a solution is efficient if its corresponding vector of criterion values is not dominated by any other vector of criterion values corresponding to a feasible solution. These vectors, associated to efficient solutions, are called *non-dominated points*. For many multi-objective optimization problems, one of the main difficulties is the large cardinality of the set of non-dominated points (or *Pareto set*).

*This work was supported by the ANR project ANR-09-BLAN-0361 "GUaranteed Efficiency for PAReto optimal solutions Determination (GUEPARD)"

†Institut Universitaire de France

‡Corresponding author

For problems with continuous variables, the set of non-dominated points is usually infinite. Even in the discrete case, it is well-known that most classic multi-objective combinatorial optimization problems, like shortest path, spanning tree, assignment, knapsack,..., are *intractable*, in the sense that they admit families of instances for which the number of non-dominated points is exponential in the size of the instance [8]. Therefore, in all these cases, it is necessary to determine a good *representation* of the Pareto set so as to provide decision makers with a tractable set of points describing as well as possible the different choices. The notion of representation is understood here in a broad sense, as in [23], as any set of points being representative of the Pareto set. This is more general than the same notion defined, e.g., in [21] where representations are supposed to be subsets of the set to be represented. Therefore we can accept that a representation of the Pareto set might include dominated points. Indeed, provided that a set of points satisfies conditions of coverage, spacing, and cardinality presented hereafter, it fully qualifies to be a good representation. In particular, adopting a broad definition allows us to consider representations whose elements are obtained through approximate optimization (in cases where exact optimization is not available or too costly). Clearly, when dominated points are present in a representation, they must be rather good (so as to satisfy the coverage property). Moreover, when possible, guaranteeing that a representation only contains non-dominated points is a desirable property. In this paper, we propose two algorithms for the bi-objective case. The first one, based on exact optimization, produces a representation consisting of non-dominated points only. The second one, based on approximate optimization, produces a representation that may contain dominated points. In both cases, however, a priori guarantees on the quality of the returned set are provided.

Measures of the quality of a discrete representation of the Pareto set have been discussed in [12, 21]. As outlined in these papers three dimensions are relevant:

- *coverage* which ensures that any non-dominated point is represented or *covered* by at least one point in the representation,
- *spacing*, also called *uniformity*, which ensures that any two points in the representation are sufficiently *spaced*, avoiding redundancies,
- *cardinality* which should be minimal so as to make the representation as tractable as possible.

Coverage is the most important dimension for the representation to be meaningful. However, it must be counterbalanced by the two other dimensions which favor a uniform and small cardinality representation, respectively. While coverage on the one hand and spacing and cardinality on the other hand are clearly conflicting, the relationship between spacing and cardinality is not obvious. At first sight it could seem that improving spacing will lead to a decrease of the number of points in the representation. It must be observed, however, that imposing spacing is an additional constraint that may impact negatively on the cardinality. An interesting result in our paper is that no negative impact is to be expected in the bi-objective case, but it is no longer true when dealing with at least three objectives. This shows the interest of considering all three dimensions.

Coverage and spacing may be implemented in several ways. A distance-oriented perspective is used in [12, 21]. The quality of coverage is then measured by the distance between the points in the Pareto set and the points in the representation (to be minimized) while the quality of spacing is measured by the distance between the points in the representation (to be maximized). Various definitions of distances are possible leading to different types of representations, but the Euclidean norm is often used. Although natural, this geometric vision is not directly related to the decision maker's preferences. Consider a representation containing a point y but not y', based on the fact that y covers y'. In a distance-oriented perspective, this is justified by the fact that y and y' are close enough. In preference-oriented perspective, the justification is that y is preferred to, or at least as good as, y'. We note that in the second case the comparison is oriented, which cannot be represented by a distance. As to spacing, points y and y'' belong to the representation since they are far enough in a distance-oriented perspective, whereas the justification is that they are incomparable in a preference-oriented perspective. In a preference-oriented perspective, the definition of a preference relation is required. When aiming at representing the whole non-dominated set, this relation should generalize the standard Pareto dominance relation, without favoring any type of solution. A natural candidate is the $(1+\varepsilon)$-dominance relation which is an extension of the Pareto dominance relation including a tolerance threshold. Given $\varepsilon > 0$, which represents a tolerance on each objective, this relation is defined as follows between any two points y and y': y $(1+\varepsilon)$-dominates y' if y is at least as good as y' within a factor $1+\varepsilon$ for all the objectives. This leads us to consider that y *covers* y' if y $(1+\varepsilon)$-dominates y', that is if y is at least as good as y' considering the tolerance ε. Moreover, given a tolerance ε', y and y'' are sufficiently *spaced* if neither y $(1+\varepsilon')$-dominates y'' nor y'' $(1+\varepsilon')$-dominates y, that is there is no reason to discard any of the points y and y'' since none of them can be considered at least as good as the other one.

This idea of coverage leads to the concept of an ε-Pareto set, introduced in [17], which is a set P_ε of points such that for any non-dominated point y', there exists a point $y \in P_\varepsilon$ which $(1+\varepsilon)$-dominates y'. Note that there may exist many ε-Pareto sets, some of which can include redundancies and some of which can have a more or less small size. An interesting problem introduced in [24] and further studied in [5] is the efficient construction of ε-Pareto sets of size as small as possible.

In this paper, we focus on the same issue but including also the spacing dimension. Therefore, the ε-Pareto sets studied in this work, called $(\varepsilon, \varepsilon')$-kernels, are required to satisfy an additional property of stability which imposes that the points in an $(\varepsilon, \varepsilon')$-kernel have to be pairwise independent relatively to the $(1+\varepsilon')$-dominance relation, thus controlling spacing.

A variety of methods have been proposed taking coverage, spacing, and/or cardinality into account (see [12, 20] for surveys). Two broad classes of methods can be distinguished: (i) algorithms which generate a set of points satisfying some properties with respect to some of the quality measures, (ii) filtering techniques which start from an initial set of given points - possibly the whole Pareto set - and retain a subset of these so as to ensure properties with respect to some of the quality measures. Among recent references that are not cited in the two previous surveys, we mention [1, 5, 9, 11, 13, 22] and [25] as examples of methods of type (i) and type (ii), respectively.

Methods of type (i) are often based on exact or approximate iterative optimizations which generate the points forming the representation. They are either generic like [5, 13] or specific to a class of problems like [22], [9], [1], and [11] which deal respectively with multi-objective linear programming, multi-objective nonlinear convex problems, multi-objective knapsack problems, and bi-objective cost flow problems. Generic algorithms can also be used as methods of type (ii), where optimizations are simply performed by scanning an explicit list of given points. It should be observed that most methods are specific to some problems and/or restricted to the bi-objective case.

Among the previously mentioned references, [9, 11, 22, 25] are distance-oriented methods. They use a Euclidean norm to define their distance. References [1, 5, 13] are preference-oriented methods. All of them use the $(1+\varepsilon)$-dominance relation. However, they only ensure coverage, and sometimes cardinality, but do not consider spacing.

The algorithms we are proposing are generic preference-oriented methods of type (i). These algorithms can be applied to discrete or continuous, linear or nonlinear, bi-objective optimization problems, depending on the availability of some problem-dependent routines. Besides providing a priori guarantees on the three quality measures, we also guarantee that our generic algorithms are polynomial when the routines are polynomial.

Our paper is organized as follows. In the next section, we define the basic concepts, formalize the notion of $(\varepsilon, \varepsilon')$-kernels, and recall some results of previous related works. In section 3, we study the bi-objective case. We show some general results and present generic polynomial time algorithms to construct small $(\varepsilon, \varepsilon')$-kernels under some conditions. In section 4, we study the case of three or more objectives, pointing out specific difficulties. Section 5 presents some experimental results which demonstrate the practical applicability of our approach. We conclude with some possible extensions to this work.

2 Preliminaries

In this paper, we consider multi-objective optimization problems where we try to minimize $p \geq 2$ criteria, i.e. $\min_{x \in S}\{f_1(x), \ldots, f_p(x)\}$, where $f_1, \ldots, f_p$ are objective functions and S is the set of feasible solutions. In case of some or all objective functions to be maximized, our results are directly extendable.

We distinguish the decision space X which contains the set S of feasible solutions of the instance and the criterion space $Y \subseteq \mathbb{R}^p_+$ which contains the criterion vectors also called, more simply, *points*. We denote by $Z = f(S) \subseteq Y$ the set of the images of feasible solutions called *feasible points*.

We denote by y_i the coordinate on objective f_i of a point $y \in Y$ for $i = 1, \ldots, p$. We say that a point y *dominates* another point y' if it is at least as good in all the objectives, i.e. $y_i \leq y'_i$ for all $i = 1, \ldots, p$. A feasible solution $x \in S$ is called *efficient* if there is no other feasible solution $x' \in S$ such that $f(x) \neq f(x')$ and $f(x')$ dominates $f(x)$. If x is efficient, $z = f(x)$ is called a *non-dominated* point in the criterion space. We denote by P the set of non-dominated points, also called *non-dominated set* or *Pareto set*. A point $z \in Z$ is *weakly non-dominated* if there is no point $z' \in Z$ such that $z'_i < z_i$ for all $i = 1, \ldots, p$.

Given two points $y, y' \in Y$ and any $\varepsilon > 0$, we say that y $(1+\varepsilon)$*-dominates* another point y', denoted by $y \preceq_\varepsilon y'$, if y is at least as good as y' up to a factor $1+\varepsilon$ in all the objectives, i.e. $y_i \leq (1+\varepsilon)y'_i$ for $i = 1, \ldots, p$. The asymmetric part of relation $\preceq_\varepsilon$ is denoted by $\prec_\varepsilon$. Thus, we have $y \prec_\varepsilon y'$ if $y_i \leq (1+\varepsilon)y'_i$ for $i = 1, \ldots, p$ and there exists $k \in \{1, \ldots, p\}$ such that $y_k < y'_k/(1+\varepsilon)$.

For any $\varepsilon > 0$, an ε*-Pareto set* of Z, denoted by P_ε, is a subset of feasible points such that any point in Z, or equivalently in P, is $(1+\varepsilon)$-dominated by at least one point in P_ε. We use this concept to implement the idea of *coverage*.

One way to ensure *spacing* is to impose a condition of stability with respect to a $(1+\varepsilon')$-dominance relation. An ε-Pareto set satisfying this additional condition will be called an $(\varepsilon, \varepsilon')$-kernel and is defined precisely as follows.

Definition 1 *Given a set Z of feasible points and $\varepsilon, \varepsilon' > 0$, an $(\varepsilon, \varepsilon')$-kernel of Z is a set of points $K_{\varepsilon,\varepsilon'} \subset Z$ satisfying the two following conditions:*
(i) for any point $z' \in Z \setminus K_{\varepsilon,\varepsilon'}$, there exists $z \in K_{\varepsilon,\varepsilon'}$ such that $z \preceq_\varepsilon z'$ (ε-coverage).
(ii) for any two distinct points $z, z' \in K_{\varepsilon,\varepsilon'}$, we do not have $z \preceq_{\varepsilon'} z'$ (ε'-stability).

Remark that if $\varepsilon' > \varepsilon$ an $(\varepsilon, \varepsilon')$-kernel does not always exist. This is the case for instance for $Z = \{z^1, z^2\}$ such that neither $z^1 \preceq_\varepsilon z^2$ nor $z^2 \preceq_\varepsilon z^1$ but $z^1 \preceq_{\varepsilon'} z^2$ or $z^2 \preceq_{\varepsilon'} z^1$. Therefore, for a given ε, the goal is to find an $(\varepsilon, \varepsilon')$-kernel with the largest $\varepsilon' \leq \varepsilon$. When $\varepsilon' = \varepsilon$ an $(\varepsilon, \varepsilon')$-kernel is called an ε-kernel.

In Figure 1 we present a small instance to illustrate the interest of this concept. Point z^3, which $(1+\varepsilon)$-dominates all points except z^6, together with point z^4, which $(1+\varepsilon)$-dominates all points except z^1, form an ε-Pareto set of minimal cardinality. In spite of this, due to their proximity, these two points do not *represent* well the whole set of points. Points z^2 and z^5, which also form an ε-Pareto set of minimal cardinality, satisfy the additional stability condition: none of them $(1+\varepsilon)$-dominates the other one. This ε-kernel clearly provides a better *representation* of the whole set of points.

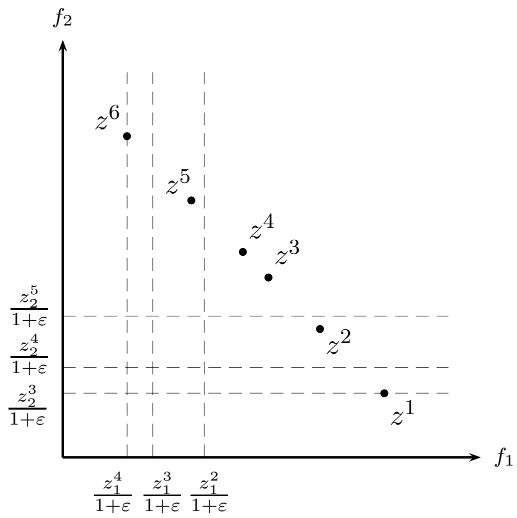

Figure 1: ε-kernels compared to ε-Pareto sets.

In the following, we are interested in the efficient computation of ε-kernels or $(\varepsilon, \varepsilon')$-kernels. For this purpose, we need to represent numbers describing instances of multi-objective optimization problems as well as parameters like ε. Therefore, we shall assume that all these numbers are positive rationals. The representation size of any rational number r will be denoted by $|r|$. Similarly, the representation size of all the rational numbers describing an instance I will be denoted by $|I|$. Moreover, assuming that the objective functions take positive rational values whose numerators and denominators have at most m bits, where $m \leq p(|I|)$ for some polynomial p, any feasible point has a value between 2^{-m} and 2^m. It follows that (the absolute value of) the difference between the values of any two points is at least 2^{-2m} for any criterion.

For a given instance I, there may exist several ε-Pareto sets, and these may have different sizes. It is shown in [17] that, for every standard multi-objective optimization problem, an ε-Pareto set of size polynomial in $|I|$, and $1/\varepsilon$ always exists. Moreover, as also shown in [17], its computation is related to the existence of a routine called GAP_δ where δ is an appropriate function of ε, selected so as to ensure obtaining an ε-Pareto set. This routine is defined as follows.

Routine GAP_δ Given an instance I of a given problem, a point y and a rational $\delta \geq 0$, $GAP_\delta(y)$ either returns a feasible point that dominates y or reports that there does not exist any feasible point z such that $z_i \leq \frac{y_i}{1+\delta}$ for all $i = 1, \ldots, p$.

Observe that, when calling $GAP_\delta(y)$, if there exists a feasible point y' dominating y but such that $y'_k > \frac{y_k}{1+\delta}$ for some $k \in \{1, \ldots, p\}$ and if there does not exist any feasible point z such that $z_i \leq \frac{y_i}{1+\delta}$ for all $i = 1, \ldots, p$, then $GAP_\delta(y)$ may either return y' or report the non-existence of points like z.

We say that routine $GAP_\delta(y)$ runs in polynomial time (resp. fully polynomial time when $\delta > 0$) if its running time is polynomial in $|I|$ and $|y|$ (resp. $|I|$, $|y|$, $|\delta|$ and $1/\delta$). An ε-Pareto set is computable in polynomial time (resp. fully polynomial time) if and only if the routine GAP_δ runs in polynomial time [17].

Since an ε-Pareto set of polynomial size can still be quite large, Vassilvitskii and Yannakakis investigate in [24] the determination of ε-Pareto sets of minimal size. More precisely, they distinguish the two following versions.

Primal version: Given an instance of a multi-objective problem and a rational $\varepsilon > 0$, determine an ε-Pareto set of minimal size.
Dual version: Given an instance of a multi-objective problem and an integer $k > 0$, determine an ε-Pareto set of size at most k with a minimal ε.

These authors also propose *generic* algorithms to deal with these versions. An algorithm is called generic if it does not depend on any particular problem and makes use of general purpose routines for which only the implementation is specific to the problem (GAP_δ is such a general purpose routine). In such algorithms it is only required to have bounds on the minimum and maximum values of the objective functions.

In order to design generic algorithms, Diakonikolas and Yannakakis introduced in [5] two other general purpose routines called $Restrict_\delta$ and $DualRestrict_\delta$ for the bi-objective case.

Routine *Restrict*$_\delta$ Given an instance I, a rational bound $b \geq 0$ and a rational $\delta \geq 0$, $Restrict_\delta(f_1, f_2 \leq b)$ either returns a feasible point z satisfying $z_2 \leq b$ and $z_1 \leq (1+\delta) \cdot \min\{f_1(x) : x \in S \text{ and } f_2(x) \leq b\}$ or correctly reports that there does not exist any feasible point z such that $z_2 \leq b$.

Routine *DualRestrict*$_\delta$ Given an instance I, a rational bound $b \geq 0$ and a rational $\delta \geq 0$, $DualRestrict_\delta(f_1, f_2 \leq b)$ either returns a feasible point z satisfying $z_2 \leq b(1+\delta)$ and $z_1 \leq \min\{f_1(x) : x \in S \text{ and } f_2(x) \leq b\}$ or correctly reports that there does not exist any feasible point z such that $z_2 \leq b$.

We say that routine $Restrict_\delta(f_1, f_2 \leq b)$ or $DualRestrict_\delta(f_1, f_2 \leq b)$ runs in polynomial time (resp. fully polynomial time when $\delta > 0$) if its running time is polynomial in $|I|$ and $|b|$ (resp. $|I|$, $|b|$, $|\delta|$ and $1/\delta$). Routines $Restrict_\delta(f_1, f_2 \leq b)$ and $DualRestrict_\delta(f_2, f_1 \leq b')$ are polynomially equivalent as proved in [5].

Remark that routines $Restrict_\delta(f_1, f_2 < b)$ and $DualRestrict_\delta(f_1, f_2 < b)$ with a strict constraint, can easily be simulated respectively by routines $Restrict_\delta(f_1, f_2 \leq b')$ and $DualRestrict_\delta(f_1, f_2 \leq b')$ using $b' = b'' - 2^{-2m}$ where b'' is the smallest multiple of 2^{-2m} which is larger than or equal to b.

In the routines considered in this paper we assume that the error δ is a rational number, otherwise it is approximated from below by a rational number. We denote by P^*_ε a smallest ε-Pareto set and by opt_ε its cardinality. It follows from [17] that opt_ε is polynomial in the input size and $1/\varepsilon$.

In the bi-objective case, the following results are known for the primal and dual versions.

For the primal version, a generic algorithm that computes an ε-Pareto set of size at most $3opt_\varepsilon$ using routine GAP_δ was established in [24]. Moreover, if GAP_δ runs in polynomial time (resp. fully polynomial time) then the algorithm also runs in polynomial time (resp. fully polynomial time). Then, it is shown in [5] that an ε-Pareto set of size at most $2opt_\varepsilon$ is computable in polynomial time if there exist $Restrict_\delta$ routines computable in polynomial time for both objectives. These approximation results are tight for the class of problems admitting such routines. An algorithm that computes an ε-Pareto set of size at most $k \cdot opt_\varepsilon$ is called a k-approximation algorithm.

For the dual version, Vassilvitskii and Yannakakis [24] state that it is NP-hard even in simple cases but provide a polynomial time approximation scheme (fully polynomial time approximation scheme) when the bi-objective problem admits a GAP_δ routine that runs in polynomial time (fully polynomial time).

In this work, our goal is to establish some general properties on $(\varepsilon, \varepsilon')$-kernels and propose some algorithms for the primal and dual versions in the case of $(\varepsilon, \varepsilon')$-kernels. In the following sections, primal and dual versions refer to $(\varepsilon, \varepsilon')$-kernels instead of ε-Pareto sets.

The proposed concepts, and the resulting algorithms, are independent of the multi-objective problem that is considered. In particular, the applicability of our generic algorithms only depends on the availability of the involved routines ($Restrict_\delta$ and/or $DualRestrict_\delta$) for the considered problem. Therefore, provided that such routines are available, these generic algorithms can be applied to discrete or continuous, linear or nonlinear, multi-objective optimization problems.

3 Two objectives

We first give some general results on ε-kernels in the bi-objective case (section 3.1). Then we consider the computation of ε-kernels when an exact *Restrict* routine, that is $Restrict_\delta$ with $\delta = 0$, is available (section 3.2) and when we only have an approximate *Restrict* routine, that is $Restrict_\delta$ with $\delta > 0$ (section 3.3).

3.1 General results

Relation $\preceq_\varepsilon$, as well as its asymmetric part $\prec_\varepsilon$, are clearly not transitive. Relation $\preceq_\varepsilon$ can even contain cycles. It appears, however, that $\prec_\varepsilon$ cannot contain cycles as shown in the next result.

Lemma 1 *In the bi-objective case, relation $\prec_\varepsilon$ does not contain cycles.*

Proof: Suppose that we have the cycle $z^1 \prec_\varepsilon z^2 \ldots \prec_\varepsilon z^n \prec_\varepsilon z^1$. Thus, for all $i \in \{1, \ldots, n-1\}$ we have (i) $z^i_j \leq (1+\varepsilon)z^{i+1}_j$ for each $j \in \{1,2\}$ and (ii) there exists $j \in \{1,2\}$ such that $z^i_j < z^{i+1}_j/(1+\varepsilon)$. Moreover, we have (i) $z^n_j \leq (1+\varepsilon)z^1_j$ for each $j \in \{1,2\}$ and (ii) there exists $j \in \{1,2\}$ such that $z^n_j < z^1_j/(1+\varepsilon)$.

Considering this cycle, assume that we are t_j times in case (ii) for each $j \in \{1,2\}$. We must have $t_1 + t_2 \geq n$. First, remark that it is not possible that $t_j = 0$ for each $j \in \{1,2\}$. Indeed, assuming without loss of generality that $t_1 = 0$, we get $t_2 = n$ leading to $(1+\varepsilon)^n < 1$. Now, observe that when we are t_j times in case (ii) for criterion j, we are also $n - t_j$ times in case (i). Since $t_j > 0$ for each $j \in \{1,2\}$, we have $z^1_j < (1+\varepsilon)^{n-2t_j} z^1_j$, which implies $t_j < n/2$ for each $j \in \{1,2\}$, contradicting $t_1 + t_2 \geq n$. □

The previous lemma guarantees the existence of ε-kernels in the bi-objective case.

Proposition 1 *In the bi-objective case, an ε-kernel always exists.*

Proof: It is a direct consequence of Lemma 1 since any relation that does not admit cycles in its asymmetric part admits kernels as proved in Duchet [6]. □

In general ε-kernels may contain dominated points. We prove the existence of ε-kernels containing only non-dominated points.

Proposition 2 *In the bi-objective case, an ε-kernel that contains only non-dominated points always exists.*

Proof: Let K_ε be an ε-kernel of the Pareto set P associated to feasible set Z. Proposition 1 implies that such an ε-kernel does exist. K_ε, which contains only non-dominated points by definition, is clearly an ε-Pareto set with respect to Z. □

In the following we give some bounds on the size of any ε-kernel.

Theorem 1 *In the bi-objective case, any ε-kernel has a cardinality less than or equal to $3opt_\varepsilon$.*

Proof: The proof is by contradiction. Let P^*_ε be an ε-Pareto set of minimal size opt_ε. Now assume that there exists an ε-kernel K_ε of size at least $3opt_\varepsilon + 1$. It means that at least one point z^* of P^*_ε $(1+\varepsilon)$-dominates at least 4 points of K_ε.

Let z^i for $i = 1,2,3,4$ be 4 points of K_ε such that $z^* \preceq_\varepsilon z^i$ for each $i = 1,2,3,4$. Assume without loss of generality that $z^{i+1}_1 < z^i_1$ and $z^{i+1}_2 > z^i_2$ for $i = 1,2,3$. Since K_ε is an ε-kernel, the coordinates of the points z^i must satisfy the following inequalities: $z^{i+1}_1 < z^i_1/(1+\varepsilon)$ and $z^{i+1}_2 > z^i_2(1+\varepsilon)$. Using these inequalities and since $z^* \preceq_\varepsilon z^i$ for each $i = 1,2,3,4$, its coordinates satisfy $z^*_1 \leq z^4_1(1+\varepsilon) < z^3_1 < z^2_1/(1+\varepsilon) < z^1_1/(1+\varepsilon)$ and $z^*_2 \leq z^1_2(1+\varepsilon) < z^2_2 < z^3_2/(1+\varepsilon) < z^4_2/(1+\varepsilon)$. Thus no point z^i for $i = 1, \ldots, 4$ $(1+\varepsilon)$-dominates z^*. If another point z of K_ε $(1+\varepsilon)$-dominates z^* the previous inequalities give $z_1 \leq z^*_1(1+\varepsilon) < z^3_1(1+\varepsilon) < z^2_1$

and $z_2 \leq z_2^*(1+\varepsilon) < z_2^2(1+\varepsilon) < z_2^3$, which involves that point z $(1+\varepsilon)$-dominates points z^2 and z^3. This would contradict ε-stability for K_ε. Thus, no point of K_ε $(1+\varepsilon)$-dominates z^*, which contradicts ε-coverage for K_ε. □

If we consider ε-kernels containing non-dominated points only, we obtain a smaller upper bound on their size. The following result is even slightly stronger since it deals with ε-kernels containing *weakly* non-dominated points only.

Theorem 2 *In the bi-objective case, any ε-kernel that contains only weakly non-dominated points has a cardinality less than or equal to $2opt_\varepsilon$.*

Proof: The proof is by contradiction. Let P_ε^* be an ε-Pareto set of minimal size opt_ε. Now assume that there exists an ε-kernel K_ε of size at least $2opt_\varepsilon + 1$ containing only weakly non-dominated points. It means that at least one point z^* of P_ε^* $(1+\varepsilon)$-dominates at least 3 points of K_ε.

Let z^i for $i = 1, 2, 3$ be 3 points of K_ε such that $z^* \preceq_\varepsilon z^i$ for each $i = 1, 2, 3$. Assume without loss of generality that $z_1^{i+1} < z_1^i$ and $z_2^{i+1} > z_2^i$. Since K_ε is an ε-kernel, the coordinates of the points z^i must satisfy the following inequalities: $z_1^{i+1} < z_1^i/(1+\varepsilon)$ and $z_2^{i+1} > z_2^i(1+\varepsilon)$ for $i = 1, 2$. Since $z^* \preceq_\varepsilon z^i$ for each $i = 1, 2, 3$, the coordinates of point z^* must satisfy $z_1^* \leq z_1^3(1+\varepsilon) < z_1^2$ and $z_2^* \leq z_2^1(1+\varepsilon) < z_2^2$. This contradicts the fact that z^2 is a weakly non-dominated point. □

Corollary 1 *In the bi-objective case, there exists an ε-kernel with a cardinality less than or equal to $2opt_\varepsilon$.*

Proof: It is a direct consequence of Theorem 2 and Proposition 2. □

We are interested now on ε-kernels of minimal size. An important fact is that an ε-kernel of minimal size is not larger than an ε-Pareto set of minimal size opt_ε.

Theorem 3 *In the bi-objective case, there exists an ε-kernel of size opt_ε.*

A constructive proof of Theorem 3 is given in section 3.2.1, where an algorithm that computes an ε-kernel of size opt_ε is provided (see Theorem 4).

3.2 Algorithms for ε-kernels using exact *Restrict* routines

In this section, we provide algorithms for the primal version (section 3.2.1) and the dual version (section 3.2.2) considering that a $Restrict_0$ routine is available for both objectives. In particular, such a *polynomial* routine is available for (continuous) multi-objective linear programming. Even if no polynomial $Restrict_0$ routine is available for most discrete and/or nonlinear problems, optimal (non polynomial) routines will guarantee obtaining an ε-kernel of minimal size.

3.2.1 Primal version

We propose a generic algorithm that produces an ε-kernel of minimal size that contains only non-dominated points. This improves significantly over the two generic algorithms proposed in [13]. The first algorithm requires a more demanding exact *Restrict* routine, where restrictions are imposed on both objectives, while the second one only requires a $Restrict_0$ routine

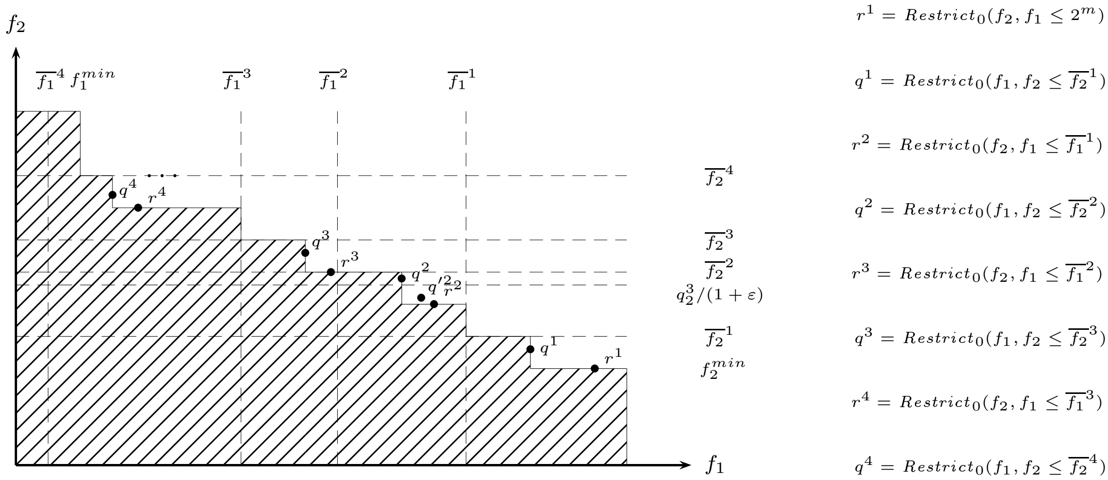

Figure 2: Illustration of Algorithm 1

for one objective. Each of these algorithms produces an ε-Pareto set whose size is only guaranteed to be at most three times the minimal size. In comparison, our algorithm guarantees to produce an ε-Pareto set of minimal size which in addition satisfies the ε-stability condition.

Algorithm description The algorithm proceeds in two phases. The first phase (greedy phase) corresponds to a slightly modified version of the algorithm presented in [5] which returns a set $\{q^1, \ldots, q^s\}$ of non-dominated points as an ε-Pareto set of minimal size. The second phase (verification phase) ensures ε-stability by checking, and possibly modifying, the returned set. We denote by f_1^{min} and f_2^{min} the minimum values on the first and second objectives respectively. In the first phase, the algorithm iteratively generates points $r^1, q^1, \ldots, r^s, q^s$ in decreasing order according to f_1 and increasing order according to f_2. Point r^1 corresponds to an optimal solution on objective f_2. Point q^1 is the non-dominated point with the best possible value on f_1 which $(1+\varepsilon)$-dominates r^1. Point r^i is a point with the smallest value on f_2 that is not $(1+\varepsilon)$-dominated by the point q^{i-1}. Point q^i is the non-dominated point with the smallest value on f_1 that $(1+\varepsilon)$-dominates point r^i. The first phase of the algorithm stops when it determines a point q^s that $(1+\varepsilon)$-dominates the feasible points that have a first coordinate equal to f_1^{min}. At the end of the first phase, ε-stability is ensured on the first objective, but not on the second one. In the second phase, points q^i are checked, starting from q^s, in decreasing order according to f_2. If point q^i $(1+\varepsilon)$-dominates point q^{i-1}, we replace point q^{i-1} by the non-dominated point with the smallest f_1 value which is not $(1+\varepsilon)$-dominated by q^i while having a strictly larger value on f_1 than q^{i-1}. This ensures ε-stability on the second objective, while preserving ε-stability on the first one.

A formal description of this algorithm is given in Algorithm 1.

Before analyzing this algorithm, we illustrate its behavior in Figure 2 where 4 points q^1, q^2, q^3, q^4 are selected during the first phase. During the second phase, the algorithm detects that point q^3 $(1+\varepsilon)$-dominates point q^2, showing that ε-stability is not satisfied. Therefore, it replaces q^2 by q'^2 which is not $(1+\varepsilon)$-dominated by q^3 but $(1+\varepsilon)$-dominates all the points that were $(1+\varepsilon)$-dominated by q^2 only. This way, ε-stability is restored, while preserving ε-coverage. The resulting ε-kernel consists of points q^1, q'^2, q^3, q^4.

Algorithm 1: Algorithm Greedy and Verification

input : An instance of a bi-objective problem for which routines $Restrict_0(f_1, f_2 \leq b)$ and $Restrict_0(f_2, f_1 \leq b)$ are available

output : An ε-kernel of size opt_ε

1 $f_1^{min} \leftarrow f_1(Restrict_0(f_1, f_2 \leq 2^m))$; $f_2^{min} \leftarrow f_2(Restrict_0(f_2, f_1 \leq 2^m))$;
2 $r^1 \leftarrow Restrict_0(f_2, f_1 \leq 2^m)$;
3 $\overline{f_2}^1 \leftarrow (1+\varepsilon) r_2^1$;
4 $q^1 \leftarrow Restrict_0(f_1, f_2 \leq \overline{f_2}^1)$;
5 $q^1 \leftarrow Restrict_0(f_2, f_1 \leq q_1^1)$;
6 $\overline{f_1}^1 \leftarrow q_1^1/(1+\varepsilon)$;
7 $Q \leftarrow \{q^1\}$;
8 $i \leftarrow 1$;
/* greedy phase */
9 **while** $\overline{f_1}^i > f_1^{min}$ **do**
10     $i \leftarrow i+1$;
11     $r^i \leftarrow Restrict_0(f_2, f_1 < \overline{f_1}^{i-1})$;
12     $\overline{f_2}^i \leftarrow (1+\varepsilon) r_2^i$;
13     $q^i \leftarrow Restrict_0(f_1, f_2 \leq \overline{f_2}^i)$;
14     $q^i \leftarrow Restrict_0(f_2, f_1 \leq q_1^i)$;
15     $\overline{f_1}^i \leftarrow q_1^i/(1+\varepsilon)$;
16     $Q \leftarrow Q \cup \{q^i\}$;
/* verification phase */
17 $i \leftarrow i-1$;
18 **while** $q_2^{i+1}/(1+\varepsilon) > f_2^{min}$ **do**
19   **if** $q_2^{i+1}/(1+\varepsilon) \leq q_2^i$ **then**
20         $Q \leftarrow Q - \{q^i\}$;
21         $q^i \leftarrow Restrict_0(f_1, f_2 < q_2^{i+1}/(1+\varepsilon))$;
22         $q^i \leftarrow Restrict_0(f_2, f_1 \leq q_1^i)$;
23         $Q \leftarrow Q \cup \{q^i\}$;
24     $i \leftarrow i-1$;
25 **return** Q;

Algorithm analysis We show now that Algorithm 1 produces an ε-kernel of minimal size. Let $R = \{r^1, \ldots, r^s\}$ and $Q = \{q^1, \ldots, q^s\}$ be the set of feasible points produced by the algorithm. We first show some preliminary results regarding points in Q and R.

Proposition 3 *Set Q contains only non-dominated points.*

Proof: Points $q^i \in Q$ are computed in two steps, both in the greedy phase (steps 13-14) and in the verification phase (steps 21-22). The first step returns a point q^i such that there exists no point $z \in Z$ such that $z_1 < q_1^i$ and $z_2 \leq q_i^2$. Thus, at this step, q^i is only guaranteed to be weakly non-dominated since there may exist a point z such that $z_1 = q_1^i$ and $z_2 < q_2^i$. The second step rules out this possibility, ensuring that q^i is non-dominated. □

Observe that the algorithm proposed in [5], which corresponds to the greedy phase, does not include this second step optimization. Therefore, the returned ε-Pareto set in [5] consists of weakly non-dominated points.

Lemma 2 *During the verification phase, if a point q'^i replaces a point q^i in Q, we have (i) $q'^i_2 < q^i_2$ and (ii) $q'^i_1 > q^i_1$.*

Proof: (i) Point q'^i computed at steps 21-22 satisfies $q'^i_2 < q^{i+1}_2/(1+\varepsilon) \leq q^i_2$.

(ii) Since points in Q are non-dominated, including q^i and q'^i, (i) implies that $q'^i_1 > q^i_1$. □

Lemma 3 *Any feasible point $z \in Z$ $(1+\varepsilon)$-dominates at most one point from R.*

Proof: Suppose, by contradiction, that z $(1+\varepsilon)$-dominates two points from R. Clearly, the most favorable situation is when these points are consecutive. Thus, let r^i and r^{i-1} be two consecutive points in R such that z $(1+\varepsilon)$-dominates them. Assuming that $z \preceq_\varepsilon r^{i-1}$, we have $z_2 \leq (1+\varepsilon)r^{i-1}_2$. By steps 13-14, this inequality implies that $q^{i-1}_1 \leq z_1$, which implies $q^{i-1}_1/(1+\varepsilon) \leq z_1/(1+\varepsilon)$. From step 11, we have $r^i_1 < q^{i-1}_1/(1+\varepsilon)$ and thus $r^i_1 < z_1/(1+\varepsilon)$, contradicting $z \preceq_\varepsilon r^i$. □

Lemma 4 *The only point in R which is $(1+\varepsilon)$-dominated by q^i is r^i, for $i = 1, \ldots, s$.*

Proof: By Lemma 3, we just need to show that $q^i \preceq_\varepsilon r^i$, for $i = 1, \ldots, s$. We proceed by induction. By steps 13-14, the assertion is clear if q^i has not be modified. In particular, for q^s which is not modified, the assertion is true. Assuming now that $q^{i+1} \preceq_\varepsilon r^{i+1}$, we prove that $q^i \preceq_\varepsilon r^i$. The only case that could me problematic is when q^i has been modified during the second phase. By Lemma 3, we have not $(q^{i+1} \preceq_\varepsilon r^i)$, which means that $q^{i+1}_2 > (1+\varepsilon)r^i_2$. Hence, by steps 21-22, we get $q^i_1 \leq r^i_1$. Moreover, regarding the second criterion, since q^i computed during the first phase $(1+\varepsilon)$-dominates r^i, we have $q^i_2 \leq (1+\varepsilon)r^i_2$. Considering that q^i has been modified, using Lemma 2-(i) we get $q^i_2 < (1+\varepsilon)r^i_2$. Therefore, we get finally $q^i \preceq_\varepsilon r^i$. □

We can now prove that Q satisfies the two conditions required to be an ε-kernel.

Proposition 4 *Set Q satisfies the ε-coverage condition.*

Proof: We show that the points in Q cover all the feasible points by partitioning the range of feasible values on f_1. More precisely, we show that:

(i) Point q^1 $(1+\varepsilon)$-dominates all the feasible points with an f_1 value greater than or equal to $q^1_1/(1+\varepsilon)$.

(ii) For each $i = 2, \ldots, s$, point q^i $(1+\varepsilon)$-dominates all the feasible points that have their f_1 value in the interval $\big[q^i_1/(1+\varepsilon), q^{i-1}_1/(1+\varepsilon)\big)$.

(iii) There is no feasible point with a f_1 value smaller than $q^s_1/(1+\varepsilon)$.

(i) Let z be a feasible point with $z_1 \geq q^1_1/(1+\varepsilon)$ and, by definition, $z_2 \geq f^{min}_2$. Point q^1 computed in steps 4-5 satisfies $q^1_2 \leq (1+\varepsilon)f^{min}_2 \leq (1+\varepsilon)z_2$, which shows that q^1 $(1+\varepsilon)$-dominates z. If point q^1 is modified during the verification phase, using Lemma 2-(i) we also have $q^1_2 \leq (1+\varepsilon)z_2$.

(ii) Let z be a feasible point satisfying $q^i_1/(1+\varepsilon) \leq z_1 < q^{i-1}_1/(1+\varepsilon)$. In order to prove that z is $(1+\varepsilon)$-dominated by q^i, we have to show that $q^i_2 \leq z_2(1+\varepsilon)$. We consider three cases.

- If points q^i and q^{i-1} have not been modified during the verification phase, then q^i, which is defined in steps 13-14, verifies $q_2^i \leq (1+\varepsilon)r_2^i$. From step 11, we have $r_2^i \leq z_2$, which leads to $q_2^i \leq (1+\varepsilon)z_2$.
- If point q^i has been modified but not point q^{i-1}, then by Lemma 2-(i), the inequality is preserved.
- Finally if point q^{i-1} has been modified during the verification phase, step 21 ensures that there is no feasible point z' such that $z'_2 < q_2^i/(1+\varepsilon)$ and $z'_1 < q_1^{i-1}$. Since $z_1 < q_1^{i-1}/(1+\varepsilon)$, it follows that $z_1 < q_1^{i-1}$ and thus $z_2 \geq q_2^i/(1+\varepsilon)$.

(iii) Point q^s, which is not modified in the verification phase, is the last point obtained in the while loop 9-16. By step 15 and condition in step 9, we have $q_1^s/(1+\varepsilon) \leq f_1^{min}$. □

Proposition 5 *Set Q satisfies the ε-stability condition.*

Proof: We just need to show that ε-stability holds for consecutive points in Q, that is for all $i = 2, \ldots, s$ we have (i) not ($q^{i-1} \preceq_\varepsilon q^i$) and ($ii$) not ($q^i \preceq_\varepsilon q^{i-1}$).

(i) From Lemma 4, we have not($q^{i-1} \preceq_\varepsilon r^i$). This occurs because we have on the first criterion $q_1^{i-1} > (1+\varepsilon)r_1^i$. Since we have $r_1^i \geq q_1^i$, we get $q_1^{i-1} > (1+\varepsilon)q_1^i$, that is not($q^{i-1} \preceq_\varepsilon q^i$).

(ii) Test 19-23 ensures that $q_2^{i-1} < q_2^i/(1+\varepsilon)$. □

Combining the previous results, we obtain the main result of this section.

Theorem 4 *For any $\varepsilon > 0$, Algorithm 1 computes an ε-kernel of minimal size opt_ε that contains only non-dominated points using $O(opt_\varepsilon)$ routine calls to $Restrict_0$.*

Proof: Q is an ε-kernel containing only non-dominated points from Propositions 3, 4, and 5. Moreover, set Q has minimal size opt_ε since, from Lemma 3, at least $|R|$ points are required for any ε-Pareto set, whereas Algorithm 1 returns a set Q with $|Q|{=}|R|$.

Since the algorithm uses at most $3|Q| + 2|Q| = 5|Q|$ times the *Restrict*$_0$ routine, the number of routine calls is bounded by $5opt_\varepsilon$. □

Since opt_ε is polynomially bounded in the input size and $1/\varepsilon$ [17], we have the following corollary.

Corollary 2 *For any $\varepsilon > 0$, if Restrict$_0$ routines are computable in polynomial time for both objectives, then we can determine an ε-kernel of minimal size that contains only non-dominated points in polynomial time in the size of the input and $1/\varepsilon$.*

3.2.2 Dual version

We show that the minimal ratio $1+\varepsilon^*$ is approximable within any factor $1+\theta$ in polynomial time in the input size and $1/\theta$.

Theorem 5 *Let k be a nonnegative integer and let $1+\varepsilon^*$ be the minimal ratio for which an ε^*-kernel of size at most k exists. For any rational $\theta > 0$, we can determine an ε-kernel of size at most k with $1+\varepsilon \leq (1+\varepsilon^*)(1+\theta)$. This can be done using $O(k\log(m/\theta))$ routine calls to Restrict$_0$.*

Proof: We first apply Algorithm 1 with $\varepsilon = \theta$. If the returned ε-kernel has size at most k, then the required condition is satisfied. Otherwise, the minimal ratio $1+\varepsilon^*$ belongs to the range $[1+\theta, 2^{2m}]$, where the upper bound corresponds to the extreme situation with $k=1$ and $Z = \{z^1 = (2^m, 1/2^m), z^2 = (1/2^m, 2^m)\}$. Let $1+\varepsilon_i = (1+\theta)^i$ be the candidate ratios for $i = 1, \ldots, \lceil 2m/\log(1+\theta)\rceil$. We perform a binary search on i values. At each step we call Algorithm 1 in order to obtain an ε_i-kernel of minimal size. If this size is greater than k then we continue the search in the right part, otherwise in the left part. Observe that, at each step, the search is between the indices i_ℓ and i_r such that the size of the smallest i_ℓ-kernel is more than k and the size of the smallest i_r-kernel is at most k. Thus, $1+\varepsilon_\ell < 1+\varepsilon^* \leq 1+\varepsilon_r$. The search is stopped when $i_r = i_\ell + 1$, i.e. when $1+\varepsilon_r = (1+\varepsilon_\ell)(1+\theta)$. Then the current ε_{i_r}-kernel is of size at most k and such that $1+\varepsilon_{i_r} = (1+\varepsilon_{i_\ell})(1+\theta) \leq (1+\varepsilon^*)(1+\theta)$.

The number of calls to Algorithm 1 is $O(\log(2m/\log(1+\theta))) \approx O(\log(m/\theta))$. Since we can stop each call to Algorithm 1 when it tries to compute a $(k+1)^{th}$ point, each such call uses $O(k)$ calls to *Restrict*$_0$. Thus, the total running time is $O(k\log(m/\theta))$ *Restrict*$_0$ calls. □

Corollary 3 *Let k be a nonnegative integer and let $1+\varepsilon^*$ be the minimal ratio for which an ε^*-kernel of size at most k exists. If Restrict$_0$ routines are computable in polynomial time for both objectives, then we can determine an ε-kernel of size at most k with $1+\varepsilon \leq (1+\varepsilon^*)(1+\theta)$ in polynomial time in the size of the input and $1/\theta$.*

3.3 Algorithms for $(\varepsilon, \varepsilon')$-kernels using approximate *Restrict* routines

In this section, we provide algorithms for the primal version (section 3.3.1) and the dual version (section 3.3.2) considering that a *Restrict*$_\delta$ routine is available for both objectives. Such polynomial routines are available for various problems: fully polynomial time routines for shortest path [16, 10] and polynomial time routines for spanning tree [19], matching and matroid intersection [3].

Assuming that fully polynomial time *Restrict*$_\delta$ routines, with $\delta > 0$, are available for both objectives, Diakonikolas and Yannakakis [5] showed that (i) there is no polynomial time generic algorithm based on these routines able to compute an ε-Pareto set of size better than $2opt_\varepsilon$, but (ii) it is possible to compute an ε-Pareto set of size $2opt_\varepsilon$ in polynomial time. Then, from Theorem 1, it follows that, using such routines, we can only hope to compute an ε-kernel of size between $2opt_\varepsilon$ and $3opt_\varepsilon$ in polynomial time. In fact, using the same routines, we even show that finding an ε-kernel in polynomial time cannot be guaranteed.

Proposition 6 *Consider the class of bi-objective problems that possess a fully polynomial time Restrict$_\delta$ routine, with $\delta > 0$, for both objectives. Then, for any $\varepsilon > 0$, there is no polynomial time generic algorithm using Restrict$_\delta$ that computes an ε-kernel.*

Proof: Consider the following set of feasible points $Z = \{z, z^1, z^2, z^3, z^4\}$ (see Figure 3) where: $z = (z_1, z_2)$, with $z_1, z_2 \geq 1/\varepsilon$, $z^1 = ((z_1+1)(1+\varepsilon), z_2/(1+\varepsilon)^2)$, $z^2 = (z_1+1, z_2)$, $z^3 = (z_1, z_2+1)$ and $z^4 = (z_1/(1+\varepsilon)^2, (z_2+1)(1+\varepsilon))$. Then, note that each point of $\{z, z^2, z^3\}$ $(1+\varepsilon)$-dominates only these three points, and that z^1 $(1+\varepsilon)$-dominates z^2 and z^4 $(1+\varepsilon)$-dominates z^3. Then, there are exactly three minimal ε-Pareto sets: $P_\varepsilon = \{z, z^1, z^4\}$, $P'_\varepsilon = \{z^2, z^1, z^4\}$, $P''_\varepsilon = \{z^3, z^1, z^4\}$ and only P_ε is an ε-kernel.

We show that a generic algorithm using *Restrict*$_\delta$ is guaranteed to return the ε-kernel only if $1/\delta$ is exponential in the size of the input. Let $z_1 = z_2 = M$, where M is an integer value

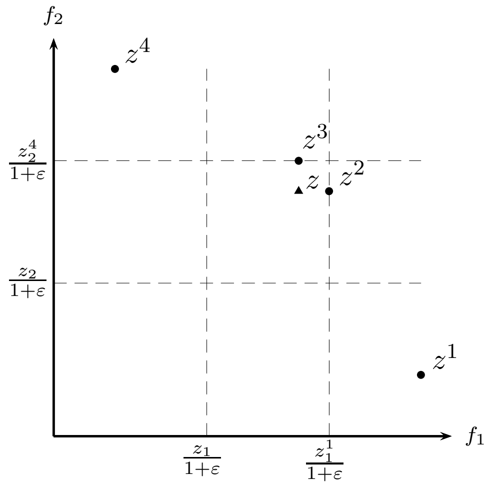

Figure 3: No polynomial time generic algorithm can compute an ε-kernel (Proposition 6).

exponential in the size of the input and $1/\varepsilon$. Assume that we call $Restrict_\delta(f_1, f_2 \leq C)$ with $C \in [M, M+1)$. Then, it can return point z^2 instead of z as long as $\delta \geq 1/M$. Symmetrically, if we call $Restrict_\delta(f_2, f_1 \leq C)$ with $C \in [M, M+1)$ we can obtain z^3 instead of z. But, since we want a polynomial time algorithm, $1/\delta$ has to be polynomial in $\log M$. Therefore, a polynomial time generic algorithm cannot guarantee to compute the unique ε-kernel which contains point z. □

In spite of this negative result, if we relax the stability condition using $\varepsilon' < \varepsilon$, we show that $(\varepsilon, \varepsilon')$-kernel can be computed in polynomial time. Therefore, in the following, we assume that $\varepsilon' < \varepsilon$.

3.3.1 Primal version

We propose an algorithm that produces an $(\varepsilon, \varepsilon')$-kernel of size at most twice the size of a minimal ε-Pareto set.

Algorithm description The algorithm proceeds in two phases. The first phase (greedy phase) corresponds to the algorithm presented in [5] which returns a 2-approximation algorithm for finding an ε-Pareto set of minimal size. The second phase (verification phase) is basically the same as Algorithm 1 but using ε' instead of ε.

The algorithm is shown to produce an $(\varepsilon, \varepsilon')$-kernel when $\delta < (1+\varepsilon)/(1+\varepsilon') - 1$ (Propositions 7 and 8) and the size of this $(\varepsilon, \varepsilon')$-kernel is proved to be at most $2opt_\varepsilon$ if $\delta \leq \sqrt[3]{1+\varepsilon} - 1$ (Theorem 6). Therefore, we assume that $\delta < \min\{(1+\varepsilon)/(1+\varepsilon') - 1, \sqrt[3]{1+\varepsilon} - 1\}$.

A formal description of this algorithm is given in Algorithm 2.

Note that, when a $Restrict_\delta$ routine is available only for one objective, we have another version of this algorithm that requires $\delta < \min\{\sqrt{(1+\varepsilon)/(1+\varepsilon')} - 1, \sqrt[3]{1+\varepsilon} - 1\}$ by replacing step 19 by $q^i \leftarrow DualRestrict_\delta(f_1, f_2 < q_2^{i+1}/(1+\varepsilon')(1+\delta)^2)$.

Algorithm analysis We show now that Algorithm 2 produces an $(\varepsilon, \varepsilon')$-kernel whose size is at most $2opt_\varepsilon$. Let $Q = \{q^1, \ldots, q^s\}$ be the set of feasible points produced by the algorithm.

First, observe that in steps 23-24 Algorithm 2 discards points that are proved unnecessary in the next result. The returned set may thus be of smaller cardinality than the ε-Pareto set obtained at the end of the greedy phase.

Lemma 5 *During the verification step, if a point q'^i, replacing a point q^i, is such that $q'^i_1 \geq q^{i-1}_1/(1+\varepsilon')$, then point q'^i is unnecessary.*

Proof: Point q'^i, with $q'^i_1 \geq q^{i-1}_1/(1+\varepsilon')$, is computed in step 19 using $Restrict_\delta(f_1, f_2 < q^{i+1}_2/(1+\varepsilon))$ where $\delta < (1+\varepsilon)/(1+\varepsilon') - 1$. This implies that any feasible point z satisfying $z_2 < q^{i+1}_2/(1+\varepsilon)$ is such that $z_1 \geq q'^i_1/(1+\delta) > q'^i_1(1+\varepsilon')/(1+\varepsilon) \geq q^{i-1}_1/(1+\varepsilon)$. Therefore, there is no feasible point z such that $z_1 < q^{i-1}_1/(1+\varepsilon)$ and $z_2 < q^{i+1}_2/(1+\varepsilon)$. Thus, a point that is $(1+\varepsilon)$-dominated by point q'^i is $(1+\varepsilon)$-dominated by point q^{i-1} or q^{i+1}. $\square$

In the following, for proving the correctness of our algorithm, the case of points which are not included (steps 23-24) can be ignored. Indeed, when this happens, the consequence of reindexing at step 23 is that points q^{i+1} and q^{i-1} become respectively points q^{i+1} and q^i at the next iteration, without any impact on the ε-coverage condition as shown by Lemma 5.

Lemma 6 *During the verification step, if a point q'^i replaces a point q^i in Q, we have (i) $q'^i_2 < q^i_2$ and (ii) $q'^i_1 \geq q^i_1$.*

Algorithm 2: Algorithm Greedy and Verification Extended

input : An instance of a bi-objective problem for which routines $Restrict_\delta(f_1, f_2 \leq b)$ and $Restrict_\delta(f_2, f_1 \leq b)$ are available

output : An $(\varepsilon, \varepsilon')$-kernel of size at most $2opt_\varepsilon$

1 $f_1^{min} \leftarrow f_1(DualRestrict_\delta(f_1, f_2 \leq 2^m))$; $f_2^{min} \leftarrow f_2(DualRestrict_\delta(f_2, f_1 \leq 2^m))$;
2 $r^1 \leftarrow Restrict_\delta(f_2, f_1 \leq 2^m)$;
3 $\overline{f_2}^1 \leftarrow \frac{1+\varepsilon}{(1+\delta)^2} r_2^1$;
4 $q^1 \leftarrow DualRestrict_\delta(f_1, f_2 \leq \overline{f_2}^1)$;
5 $\overline{f_1}^1 \leftarrow q_1^1/(1+\varepsilon)$;
6 $Q \leftarrow \{q^1\}$;
7 $i \leftarrow 1$;
`/* greedy phase */`
8 **while** $\overline{f_1}^i > f_1^{min}$ **do**
9     $i \leftarrow i + 1$;
10     $r^i \leftarrow Restrict_\delta(f_2, f_1 < \overline{f_1}^{i-1})$;
11     $\overline{f_2}^i \leftarrow \frac{1+\varepsilon}{1+\delta} max\{\overline{f_2}^{i-1}, r_2^i/(1+\delta)\}$;
12     $q^i \leftarrow DualRestrict_\delta(f_1, f_2 \leq \overline{f_2}^i)$;
13     $\overline{f_1}^i \leftarrow q_1^i/(1+\varepsilon)$;
14     $Q \leftarrow Q \cup \{q^i\}$;
`/* verification phase */`
15 $s \leftarrow i$, $i \leftarrow i - 1$;
16 **while** $q_2^{i+1}/(1+\varepsilon) > f_2^{min}$ **do**
17   **if** $q_2^{i+1}/(1+\varepsilon') \leq q_2^i$ **then**
18         $Q \leftarrow Q - \{q^i\}$;
19         $q^i \leftarrow Restrict_\delta(f_1, f_2 < q_2^{i+1}/(1+\varepsilon))$;
20   **if** $q_1^i < q_1^{i-1}/(1+\varepsilon')$ **then**
21             $Q \leftarrow Q \cup \{q^i\}$;
22   **else**
23   reindex $\{q^{i+1}, \ldots, q^s\}$ by $\{q^i, \ldots, q^{s-1}\}$;
24             $s \leftarrow s - 1$;
25     $i \leftarrow i - 1$;
26 **return** Q;

Proof: (i) Point q'^i computed at step 19 satisfies $q_2'^i < q_2^{i+1}/(1+\varepsilon) < q_2^{i+1}/(1+\varepsilon') \leq q_2^i$.

(ii) Remark that point q^i was computed in step 12 using routine $DualRestrict_\delta$ during the greedy phase. It follows that there is no feasible point z such that $z_1 < q_1^i$ and $z_2 < q_2^i/(1+\delta)$. Since $\varepsilon' < (1+\varepsilon)/(1+\delta) - 1$, point q'^i is computed in step 19 such that $q_2'^i < q_2^{i+1}/(1+\varepsilon) < q_2^{i+1}/(1+\varepsilon')(1+\delta) \leq q_2^i/(1+\delta)$. It follows that $q_1'^i \geq q_1^i$. □

We can now prove that Q satisfies the two conditions required to be an $(\varepsilon, \varepsilon')$-kernel.

Proposition 7 *Set Q satisfies the ε-coverage condition.*

Proof: We show that the points in Q cover all the feasible points by partitioning the range of feasible values on f_1. More precisely, we show that:

(i) Point q^1 $(1+\varepsilon)$-dominates all the feasible points with an f_1 value greater than or equal to $q_1^1/(1+\varepsilon)$.

(ii) For each $i = 2, \ldots, s$, the point q^i $(1+\varepsilon)$-dominates all the feasible points that have their f_1 value in the interval $\left[q_1^i/(1+\varepsilon), q_1^{i-1}/(1+\varepsilon)\right)$.

(iii) There is no feasible point with a f_1 value smaller than $q_1^s/(1+\varepsilon)$.

(i) Let z be a feasible point with $z_1 \geq q_1^1/(1+\varepsilon)$ and, by definition, $z_2 \geq f_2^{min}$. Point q^1 computed in step 4 satisfies $q_2^1 \leq (1+\varepsilon)f_2^{min} \leq (1+\varepsilon)z_2$, which shows that q^1 $(1+\varepsilon)$-dominates z. If point q^1 is modified during the verification phase, using Lemma 6-(i) we also have $z_2 \geq q_2^1/(1+\varepsilon)$.

(ii) Let z be a feasible point satisfying $q_1^i/(1+\varepsilon) \leq z_1 < q_1^{i-1}/(1+\varepsilon)$. In order to prove that z is $(1+\varepsilon)$-dominated by q^i, we have to show that $q_2^i \leq (1+\varepsilon)z_2$. We consider three cases.

- If points q^i and q^{i-1} have not been modified during the verification phase, then q^i, which is defined in step 12, verifies $q_2^i \leq (1+\varepsilon) \cdot \max\{\overline{f_2}^{i-1}, r_2^i/(1+\delta)\}$. From step 10 we have $z_2 \geq r_2^i/(1+\delta)$ and from step 12 for $i-1$ we have $z_2 \geq \overline{f_2}^{i-1}$. Thus we have $\max\{\overline{f_2}^{i-1}, r_2^i/(1+\delta)\} \leq z_2$ which leads to $q_2^i \leq (1+\varepsilon)z_2$.
- If point q^i has been modified but not point q^{i-1}, then by Lemma 6-(i), the inequality is preserved.
- Finally if point q^{i-1} has been modified during the verification phase, step 19 ensures that there is no feasible point z' such that $z'_2 < q_2^i/(1+\varepsilon)$ and $z'_1 < q_1^{i-1}/(1+\delta)$. Since $z_1 < q_1^{i-1}/(1+\varepsilon)$ it follows that $z_1 < q_1^{i-1}/(1+\delta)$ and thus $z_2 \geq q_2^i/(1+\varepsilon)$.

(iii) Point q^s, which is not modified during the verification phase, is the last point obtained in the while loop 8-14. By step 13 and condition in step 8, we have $q_1^s/(1+\varepsilon) \leq f_1^{min}$. □

Proposition 8 *Set Q satisfies the ε'-stability condition.*

Proof: We just need to show that ε'-stability holds for consecutive points in Q, that is for all $i = 2, \ldots, s$ we have (i) not $(q^{i-1} \preceq_{\varepsilon'} q^i)$ and (ii) not $(q^i \preceq_{\varepsilon'} q^{i-1})$.

(i) We consider three cases.

- If points q^i and q^{i-1} have not been modified during the verification phase, then point r^i, computed in step 10, is such that $r_1^i < q_1^{i-1}/(1+\varepsilon)$. Moreover since point q^i, computed in step 12, is such that $q_1^i \leq r_1^i$, we get $q_1^i < q_1^{i-1}/(1+\varepsilon) < q_1^{i-1}/(1+\varepsilon')$, that is not $(q^{i-1} \preceq_{\varepsilon'} q^i)$.
- If point q^i is modified and point q^{i-1} is not modified, then since q^i is added to Q in step 21, it satisfies $q_1^i < q_1^{i-1}/(1+\varepsilon')$, that is not $(q^{i-1} \preceq_{\varepsilon'} q^i)$.
- The final case is when point q^{i-1} changes during the verification phase and is replaced by a point q'^{i-1}. Then, according to Lemma 6-(ii) the inequality is preserved.

(ii) Test 17-24 and the definition of point q^{i-1} at step 19 ensures that $q_2^{i-1} < q_2^i/(1+\varepsilon) < q_2^i/(1+\varepsilon')$. □

Lemma 7 *Any point $z \in Z$ $(1+\varepsilon)$-dominates at most two points from R.*

Proof: Suppose, by contradiction, that z $(1+\varepsilon)$-dominates three points from R. Clearly, the most favorable situation is when these points are consecutive. Thus, let r^i, r^{i-1}, and r^{i-2} be three consecutive points in R such that z $(1+\varepsilon)$-dominates them. Assuming that $z \preceq_\varepsilon r^{i-2}$, we have $z_2 \leq (1+\varepsilon)r_2^{i-2}$. By step 11, for $i-2$ and $i-1$, we get $\overline{f_2}^{i-2} \geq \frac{1+\varepsilon}{(1+\delta)^2} r_2^{i-2}$ and $\overline{f_2}^{i-1} \geq \frac{1+\varepsilon}{1+\delta}\overline{f_2}^{i-2}$ and thus $\overline{f_2}^{i-2} \geq \frac{(1+\varepsilon)^2}{(1+\delta)^3} r_2^{i-2}$. Since $(1+\delta)^3 < 1+\varepsilon$, we have $z_2 \leq \overline{f_2}^{i-1}$. From this last inequality, by step 12, for $i-1$, we have $q_1^{i-1} \leq z_1$, which implies $q_1^{i-1}/(1+\varepsilon) \leq z_1/(1+\varepsilon)$. From step 10, we have $r_1^i < q_1^{i-1}/(1+\varepsilon)$ and thus $r_1^i < z_1/(1+\varepsilon)$, contradicting $z \preceq_\varepsilon r^i$. □

Combining the previous results, we obtain the following result.

Theorem 6 *For any $\varepsilon, \varepsilon'$ such that $\varepsilon > \varepsilon' > 0$, Algorithm 2 computes an $(\varepsilon, \varepsilon')$-kernel of size less than or equal to $2opt_\varepsilon$ using $O(opt_\varepsilon)$ routine calls to Restrict$_\delta$ or DualRestrict$_\delta$, where $\delta < \min\{(1+\varepsilon)/(1+\varepsilon') - 1, \sqrt[3]{1+\varepsilon} - 1\}$.*

Proof: Q is an $(\varepsilon, \varepsilon')$-kernel from Propositions 7 and 8. Moreover, set Q has a size less than or equal to $2opt_\varepsilon$ since, from Lemma 7, at least $\lceil |R|/2 \rceil$ points are required for any ε-Pareto set, whereas Algorithm 2 returns a set Q with $|Q| \leq |R|$.

Since the algorithm uses at most $2|Q| + |Q| = 3|Q|$ times the routines *Restrict$_\delta$* or *DualRestrict$_\delta$*, the number of routine calls is bounded by $3opt_\varepsilon$. □

Since opt_ε is polynomially bounded in the input size and $1/\varepsilon$ [17], we have the following corollary.

Corollary 4 *For any $\varepsilon, \varepsilon'$ such that $\varepsilon > \varepsilon' > 0$, if routines Restrict$_\delta$ and DualRestrict$_\delta$ with $\delta > 0$ are computable in (fully) polynomial time for both objectives, then we can determine an $(\varepsilon, \varepsilon')$-kernel of size less than or equal to $2opt_\varepsilon$ in (fully) polynomial time.*

We recall that it is not possible to produce an ε-Pareto set of size opt_ε in polynomial time using *Restrict$_\delta$* routines [5]. Nevertheless, Vassilvitskii and Yannakakis showed in [24] that it is possible to produce in polynomial time an ε-Pareto set of size bounded by $opt_{\hat{\varepsilon}}$ for any $\hat{\varepsilon} < \varepsilon$. In the following we present a similar result for $(\varepsilon, \varepsilon')$-kernels. More precisely, we show that Algorithm 2 used with $\delta < \min\{\sqrt{(1+\varepsilon)/(1+\hat{\varepsilon})} - 1, (1+\varepsilon)/(1+\varepsilon') - 1\}$ computes an $(\varepsilon, \varepsilon')$-kernel of size bounded by $opt_{\hat{\varepsilon}}$, for any $\hat{\varepsilon} < \varepsilon$ and $\varepsilon' < \varepsilon$. Let Q be the set of feasible points produced by the algorithm.

Since $\delta < (1+\varepsilon)/(1+\varepsilon') - 1$, Q is an $(\varepsilon, \varepsilon')$-kernel from Propositions 7 and 8. Therefore, we only need to show that set Q has a size less than or equal to $opt_{\hat{\varepsilon}}$.

Proposition 9 *When $\delta \leq \sqrt{(1+\varepsilon)/(1+\hat{\varepsilon})} - 1$, Algorithm 2 returns a set Q with $|Q| \leq opt_{\hat{\varepsilon}}$.*

Proof: Let $P^*_{\hat{\varepsilon}} = \{p^{*1}, \ldots, p^{*k}\}$ be an $\hat{\varepsilon}$-Pareto set of minimal size, where points p^{*i} for $i = 1, \ldots, k$ are in increasing order of their coordinates on f_2 and decreasing order of their coordinates on f_1. Let $\tilde{Q} = \{\tilde{q}^1, \ldots, \tilde{q}^r\}$ be the set of points returned by the greedy phase of Algorithm 2. We have $|\tilde{Q}| \geq |Q|$ due to the possible omission of points in steps 23-24 of the verification step. We show now that $|\tilde{Q}| \leq |P^*_{\hat{\varepsilon}}|$. For this purpose, we show by induction on i that for each point $\tilde{q}^i$ in $\tilde{Q}$ there exists a point p^{*i} in $P^*_{\hat{\varepsilon}}$ such that $\tilde{q}^i_1 \leq p^{*i}_1$.

Initialization ($i = 1$). The fact that P^*_{ε} contains at least one point is trivially true. We need to show that $\tilde{q}^1_1 \leq p^{*1}_1$. Since point $\tilde{q}^1$ is computed in step 4 using $DualRestrict_\delta(f_1, f_2 \leq \overline{f_2}^1)$, to show the statement it suffices to prove that $\overline{f_2}^1 \geq p^{*1}_2$. Since $P^*_{\hat{\varepsilon}}$ is an $\hat{\varepsilon}$-Pareto set where its points p^{*j} for $j = 1, \ldots, k$ are in increasing order of their coordinates on f_2, it follows that point p^{*1} must $(1+\hat{\varepsilon})$-dominates f_2^{min} and so $p^{*1}_2 \leq (1+\hat{\varepsilon}) f_2^{min}$. Since $\delta \leq \sqrt{(1+\varepsilon)/(1+\hat{\varepsilon})} - 1$, it follows that $p^{*1}_2 \leq \frac{1+\varepsilon}{(1+\delta)^2} f_2^{min}$. From step 2 we have $r^1_2 \geq f_2^{min}$ and from step 3 we have $\overline{f_2}^1 = \frac{1+\varepsilon}{(1+\delta)^2} r^1_2$, thus it follows that $\overline{f_2}^1 \geq p^{*1}_2$.

Induction step. Assume the result is true until index $i-1$, we prove it for index i. By the termination condition of the greedy phase of Algorithm 2 (step 8), we have $\tilde{q}^{i-1}_1 > (1+\varepsilon) f_1^{min}$ and by the induction hypothesis that $p^{*i-1}_1 \geq \tilde{q}^{i-1}_1$, it follows that $p^{*i-1}_1 > (1+\varepsilon) f_1^{min}$. Thus, point p^{*i-1} does not $(1+\varepsilon)$-dominate the feasible points that have a first coordinate equal to f_1^{min}, and so $P^*_{\hat{\varepsilon}}$ must contain another point p^{*i}. Since point $\tilde{q}^i$ is computed in step 12 using $DualRestrict_\delta(f_1, f_2 \leq \overline{f_2}^i)$, to show the statement it suffices to prove that $\overline{f_2}^i \geq p^{*i}_2$. Since P^*_{ε} is an $\hat{\varepsilon}$-Pareto set where its points p^{*j} for $j = 1, \ldots, k$ are in increasing order of their coordinates on f_2, it follows that point p^{*i} must $(1+\hat{\varepsilon})$-dominates point r^i and so $p^{*i}_2 \leq (1+\hat{\varepsilon}) r^i_2$. Since $\delta \leq \sqrt{(1+\varepsilon)/(1+\hat{\varepsilon})} - 1$, it follows that $p^{*i}_2 \leq \frac{1+\varepsilon}{(1+\delta)^2} r^i_2$. From step 11 we have $\overline{f_2}^i \geq \frac{1+\varepsilon}{(1+\delta)^2} r^i_2$, thus it follows that $\overline{f_2}^i \geq p^{*i}_2$. □

The second main result of this section follows.

Theorem 7 *For any $\hat{\varepsilon}, \varepsilon, \varepsilon'$ such that $\varepsilon > \hat{\varepsilon} > 0$ and $\varepsilon > \varepsilon' > 0$, Algorithm 2 computes an $(\varepsilon, \varepsilon')$-kernel of size less than or equal to $opt_{\hat{\varepsilon}}$ using $O(opt_{\hat{\varepsilon}})$ routine calls to $Restrict_\delta$ or $DualRestrict_\delta$, with $\delta < \min\{\sqrt{(1+\varepsilon)/(1+\hat{\varepsilon})} - 1, (1+\varepsilon)/(1+\varepsilon') - 1\}$.*

Proof: Set Q returned by Algorithm 2 is an $(\varepsilon, \varepsilon')$-kernel since Propositions 7 and 8 hold. Moreover, the size of Q is less than or equal to $opt_{\hat{\varepsilon}}$ by Proposition 9. Since the algorithm uses $3|Q|$ times the $Restrict_\delta$ or $DualRestrict_\delta$ routines, the number of routine calls is bounded by $3opt_{\hat{\varepsilon}}$. □

Corollary 5 *For any $\hat{\varepsilon}, \varepsilon, \varepsilon'$ such that $\varepsilon > \hat{\varepsilon} > 0$ and $\varepsilon > \varepsilon' > 0$, if routines $Restrict_\delta$ and $DualRestrict_\delta$ with $\delta > 0$ are computable in (fully) polynomial time for both objectives, then we can determine an $(\varepsilon, \varepsilon')$-kernel of size less than or equal to $opt_{\hat{\varepsilon}}$ in (fully) polynomial time.*

3.3.2 Dual version

We show that the minimal ratio $1+\varepsilon^*$ is approximable within any factor $1+\theta$ in polynomial time in the input size and $1/\theta$.

Theorem 8 *Let k be a nonnegative integer and let $1+\varepsilon^*$ be the minimal ratio for which an ε^*-kernel of size at most k exists. For any rational $\theta > 0$, we can determine an $(\varepsilon, \varepsilon')$-kernel* *with $1+\varepsilon \leq (1+\varepsilon^*)(1+\theta)$, for all $\varepsilon' < \varepsilon$, of size at most k using $O(k \log(m/\theta))$ routine calls to Restrict$_\delta$ or DualRestrict$_\delta$.*

Proof: We first apply Algorithm 2 with $\varepsilon = \theta$, $\varepsilon' < \varepsilon$, and $\delta < \min\{\sqrt[4]{1+\theta}-1, (1+\varepsilon)/(1+\varepsilon')-1\}$, where $\delta < \sqrt[4]{1+\theta}-1$ results from considering $1+\hat{\varepsilon} = \sqrt{1+\theta}$ in Theorem 7. If the returned $(\varepsilon, \varepsilon')$-kernel has size at most k, then the required condition is satisfied. Otherwise, from Theorem 7, the minimal ratio $1+\varepsilon^*$ belongs to the range $[\sqrt{1+\theta}, 2^{2m}]$. Let $1+\varepsilon_i = (\sqrt{1+\theta})^i$ be the candidate ratios for $i = 2, \ldots, \lceil 4m/\log(1+\theta) \rceil$ and let $1+\hat{\varepsilon}_i = (1+\varepsilon_i)/\sqrt{1+\theta}$. We perform a binary search on i values. At each step we call Algorithm 2 with $\delta < \min\{\sqrt[4]{1+\theta}-1, (1+\varepsilon_i)/(1+\varepsilon'_i)-1\}$, where ε'_i is an arbitrary number such that $\varepsilon'_i < \varepsilon_i$, in order to obtain an $(\varepsilon_i, \varepsilon'_i)$-kernel of size at most $opt_{\hat{\varepsilon}_i}$ (see Theorem 7). If this size is greater than k then we continue the search in the right part, otherwise in the left part. Observe that, at each step, the search is between the indices i_ℓ and i_r such that the size of the $(\varepsilon_{i_\ell}, \varepsilon'_{i_\ell})$-kernel is more than k and the size of the $(\varepsilon_{i_r}, \varepsilon'_{i_r})$-kernel is at most k. Thus, $(1+\varepsilon_{i_\ell})/\sqrt{1+\theta} < 1+\varepsilon^* \leq 1+\varepsilon_{i_r}$. The search is stopped when $i_r = i_\ell + 1$, i.e. when $1+\varepsilon_{i_r} = (1+\varepsilon_{i_\ell})\sqrt{1+\theta}$. Then, the current $(\varepsilon_{i_r}, \varepsilon'_{i_r})$-kernel is of size at most $opt_{\hat{\varepsilon}_{i_r}} \leq k$ and such that $1+\varepsilon_{i_r} = (1+\varepsilon_{i_\ell})\sqrt{1+\theta} \leq (1+\varepsilon^*)(1+\theta)$.

The number of calls to Algorithm 2 is $O(\log(4m/\log(1+\theta))) \approx O(\log(m/\theta))$. Since we can stop each call to Algorithm 2 when it tries to compute a $(k+1)^{th}$ point, each such call uses $O(k)$ calls to *Restrict$_\delta$* or *DualRestrict$_\delta$*. Thus, the total running time is $O(k \log(m/\theta))$ *Restrict$_\delta$* or *DualRestrict$_\delta$* calls. □

Corollary 6 *Let k be a nonnegative integer and let $1+\varepsilon^*$ be the minimal ratio for which an ε^*-kernel of size at most k exists. If routines Restrict$_\delta$ and DualRestrict$_\delta$ with $\delta > 0$ are computable in (fully) polynomial time for both objectives, for any rational $\theta > 0$, we can determine an $(\varepsilon, \varepsilon')$-kernel with $1+\varepsilon \leq (1+\varepsilon^*)(1+\theta)$, for all $\varepsilon' < \varepsilon$, of size at most k in (fully) polynomial time.*

4 More than two objectives

For more than two objectives, the concept of ε-kernel is not really operational since an ε-kernel does not always exist.

Proposition 10 *For $p \geq 3$ objectives, an ε-kernel may not exist.*

Proof: Let $p = 3$ and z^1, z^2, and z^3 be three points with the following coordinates: $z^1 = (a(1+\varepsilon), b/(1+\varepsilon), c)$, $z^2 = (a, b(1+\varepsilon), c/(1+\varepsilon))$, $z^3 = (a/(1+\varepsilon), b, c(1+\varepsilon))$ where a, b, and c are three nonnegative rational numbers.

Clearly z^1 $(1+\varepsilon)$-dominates z^2, z^2 $(1+\varepsilon)$-dominates z^3 and z^3 $(1+\varepsilon)$-dominates z^1. Since any ε-kernel must satisfy the ε-stability condition, it follows that an ε-kernel must contain at most one point. Moreover, no point $(1+\varepsilon)$-dominates the two others. Since any ε-kernel must satisfy the ε-coverage condition, it follows that an ε-kernel must contain at least two points. This is clearly impossible. □

Moreover, even if an ε-kernel exists, we have no guarantee on its size like Theorems 1, 2, and 3 for the bi-objective case. On the opposite, we can show that a smallest ε-kernel may have a very large size compared with opt_ε.

Proposition 11 *For $p \geq 3$ objectives, the size of a smallest ε-kernel, when it exists, can be greater than $k \cdot opt_\varepsilon$ for any integer k.*

Proof: Let $p = 3$ and z^1, z^2, and z^3 be defined as in the proof of Proposition 10. Let $z = (z_1^2, z_2^3, z_3^1) = (a, b, c)$. Fix any rational $\hat{\varepsilon} > \varepsilon$ and consider the $3k$ points $z^{1j} = (z_1^1(1+\hat{\varepsilon})^j, z_2^1/(1+\varepsilon), z_3^1(1+\hat{\varepsilon})^{k-j})$, $z^{2j} = (z_1^2(1+\hat{\varepsilon})^{k-j}, z_2^2(1+\hat{\varepsilon})^j, z_3^2/(1+\varepsilon))$ and $z^{3j} = (z_1^3/(1+\varepsilon), z_2^3(1+\hat{\varepsilon})^{k-j}, z_3^3(1+\hat{\varepsilon})^j)$ for $j = 1, \ldots, k$.

For this instance, the only cases of $(1+\varepsilon)$-dominance are: $z^1 \preceq_\varepsilon z^2$, $z^2 \preceq_\varepsilon z^3$, $z^3 \preceq_\varepsilon z^1$, $z \preceq_\varepsilon z^i$ and $z^i \preceq_\varepsilon z$ for $i = 1, 2, 3$, and $z^i \preceq_\varepsilon z^{ij}$ for $i = 1, 2, 3$ and $j = 1, \ldots, k$.

The set constituted by points z^1, z^2, and z^3 is clearly an ε-Pareto set of minimal size. Moreover, any ε-kernel must contain point z and thus points z^{ij} for $i = 1, 2, 3$ and $j = 1, \ldots, k$. This is the only ε-kernel and it contains $3k + 1$ points. □

However, if we consider $\varepsilon' \leq \sqrt{1+\varepsilon} - 1$, we can show that an $(\varepsilon, \varepsilon')$-kernel always exists. For this purpose, we recall the notion of quasi-kernel (also called semi-kernel).

Definition 2 *Given a directed graph $G = (V, A)$, a quasi-kernel is a set $S \subseteq V$ such that (i) for any $v \in V - S$, there exists $v' \in S$ such that $(v', v) \in A$ or there exist $v' \in S$ and $v'' \in V - S$ such that $(v', v'') \in A$ and $(v'', v) \in A$ (ii) for any $u, v \in S$, $u \neq v$, $(u, v) \notin A$.*

The following result is well-known.

Theorem 9 *(Chvátal and Lovász [4]) Any finite directed graph G admits a quasi-kernel.*

Applied in our context, this gives rise to the following result.

Proposition 12 *For any number of objectives $p \geq 3$ and any finite set Z of points an $(\varepsilon, \varepsilon')$-kernel exists if and only if $\varepsilon' \leq \sqrt{1+\varepsilon} - 1$.*

Proof: $\Leftarrow$ Consider the graph $G = (Z, \preceq_{\varepsilon'})$ and apply Theorem 9.

$\Rightarrow$ Assuming that $\varepsilon' > \sqrt{1+\varepsilon} - 1$, we show the existence of an instance which does not admit an $(\varepsilon, \varepsilon')$-kernel.

Let $Z = \{z^1, z^2, z^3\}$ where z^1, z^2, and z^3 are three points in the criterion space and assume that their coordinates are the following: $z^1 = (a(1+\varepsilon'), b/(1+\varepsilon'), c)$, $z^2 = (a, b(1+\varepsilon'), c/(1+\varepsilon'))$, $z^3 = (a/(1+\varepsilon'), b, c(1+\varepsilon'))$ with a, b, and c three nonnegative rational numbers.

Remark that z^1 $(1+\varepsilon')$-dominates z^2, z^2 $(1+\varepsilon')$-dominates z^3 and z^3 $(1+\varepsilon')$-dominates z^1. In order to satisfy the ε'-stability condition an $(\varepsilon, \varepsilon')$-kernel contains at most one point among z^1, z^2, and z^3. Moreover, since $\varepsilon' > \sqrt{1+\varepsilon} - 1$, no point $(1+\varepsilon)$-dominates the two others and thus in order to satisfy the ε-coverage condition, an $(\varepsilon, \varepsilon')$-kernel must contain at least two points. This is clearly impossible. □

Moreover, when the points Z are given explicitly and $\varepsilon' \leq \sqrt{1+\varepsilon} - 1$ it is possible to compute an $(\varepsilon, \varepsilon')$-kernel in polynomial time. Indeed, the problem can be reduced to finding a kernel in a directed acyclic graph [7]. We briefly describe the method of Duchet et al. from [7]. Consider the directed graph $G = (Z, \preceq_{\varepsilon'})$ and any arbitrary order $<$ on the vertices. We first partition the set of arcs into two disjoint subsets $A_1 = \{(i, j) \in \preceq_{\varepsilon'} : i < j\}$, $A_2 = \{(i, j) \in \preceq_{\varepsilon'} : i > j\}$. The two directed graphs (Z, A_1) and (Z, A_2) contain no cycle. Since a (unique) kernel can be easily computed in polynomial time in directed acyclic graphs, first construct the kernel

K of (Z, A_1) and then the kernel K' of (K, A_2). The resulting subset K' is a quasi-kernel of G, i.e. an $(\varepsilon, \varepsilon')$-kernel.

In the general case, when the points of the criterion space are not given explicitly, we have the following result.

Proposition 13 *For $p \geq 3$ objectives and any $0 < \varepsilon' \leq \sqrt[3]{1+\varepsilon} - 1$, an $(\varepsilon, \varepsilon')$-kernel is computable in polynomial time when the associated GAP_δ routine runs in polynomial time.*

Proof: First we construct a grid in the criterion space as in the proof of the efficient constructability of an ε-Pareto set presented in [17]. Consider a subdivision of the criterion space into hyperrectangles such that, in each dimension, the ratio of the largest to the smallest coordinate of each hyperrectangle is $\sqrt[6]{1+\varepsilon}$. In each corner point, call the GAP_δ routine with $\delta = \sqrt[6]{1+\varepsilon} - 1$ and denote by S the resulting set of points. Set S (after removing the dominated points) is clearly an $(\sqrt[3]{1+\varepsilon} - 1)$-Pareto set.

On set S, we use the method of Duchet et al. [7] to construct a quasi-kernel in a directed graph. Thus, we obtain a subset $K \subseteq S$ which is an $((\sqrt[3]{1+\varepsilon})^2 - 1, \sqrt[3]{1+\varepsilon} - 1)$-kernel for the points in S. Since S is an $(\sqrt[3]{1+\varepsilon} - 1)$-Pareto set, it implies that K is an $((\sqrt[3]{1+\varepsilon})^2 \cdot \sqrt[3]{1+\varepsilon} - 1, \sqrt[3]{1+\varepsilon} - 1)$-kernel i.e. an $(\varepsilon, \varepsilon')$-kernel. □

Nevertheless, we can show a result similar to Proposition 11 for $(\varepsilon, \varepsilon')$-kernels.

Proposition 14 *For $p \geq 3$ objectives and any $0 < \varepsilon' \leq \sqrt{1+\varepsilon} - 1$, the size of a smallest $(\varepsilon, \varepsilon')$-kernel can be greater than $k \cdot opt_\varepsilon$ for any integer k.*

Proof: Let $p = 3$ and z^1, z^2, and z^3 be three points with the following coordinates: $z = (a, b, c)$, $z^1 = (a(1+\varepsilon'), b/(1+\varepsilon'), c)$, $z^2 = (a, b(1+\varepsilon'), c/(1+\varepsilon'))$, $z^3 = (a/(1+\varepsilon'), b, c(1+\varepsilon'))$ where a, b, and c are three nonnegative rational numbers. Fix any rational $\hat{\varepsilon}$ and consider $6k$ points $z^{1j} = (z^1_1(1+\hat{\varepsilon})^j, z^1_2/(1+\varepsilon), z^1_3(1+\hat{\varepsilon})^{2k-j})$, $z^{2j} = (z^2_1(1+\hat{\varepsilon})^{2k-j}, z^2_2(1+\hat{\varepsilon})^j, z^2_3/(1+\varepsilon))$ and $z^{3j} = (z^3_1/(1+\varepsilon), z^3_2(1+\hat{\varepsilon})^{2k-j}, z^3_3(1+\hat{\varepsilon})^j)$ for $j = 1, \ldots, 2k$.

Remark that points z, z^1, z^2, and z^3 $(1+\varepsilon)$-dominates each other and $z^i \preceq_\varepsilon z^{ij}$ for $i = 1, 2, 3$ and $j = 1, \ldots, k$. For this instance, the only cases of $(1+\varepsilon')$-dominance are: $z^1 \preceq_{\varepsilon'} z^2$, $z^2 \preceq_{\varepsilon'} z^3$, $z^3 \preceq_{\varepsilon'} z^1$, $z \preceq_{\varepsilon'} z^i$ and $z^i \preceq_{\varepsilon'} z$ for $i = 1, 2, 3$.

The set constituted by points z^1, z^2, and z^3 is clearly an ε-Pareto set of minimal size. Moreover, a smallest $(\varepsilon, \varepsilon')$-kernel contains a point z^i with $i = 1, 2, 3$ and all the points $z^{i'j}$ for $i' = 1, 2, 3$ with $i \neq i'$ and $j = 1, \ldots, 2k$, and it contains $4k + 1$ points. □

5 Experiments

We show now the implementation of our exact algorithm in the bi-objective case in order to generate an ε-kernel (Algorithm 1). We first illustrate this algorithm in the context of (continuous) Multi-Objective Linear Programming (Section 5.1). Then we report experiments on two standard multi-objective combinatorial optimization problems, the bi-objective shortest path problem (Section 5.2) and the bi-objective assignment problem (Section 5.3). These experiments are performed on a PC (i7-2600, 3.4GHz, 8GB) using CPLEX 12.6.3 with one thread. Reported computation times are CPU times expressed in seconds. We recall from Theorem 4 that the discrete representations provided by Algorithm 1 (i) guarantee both ε-coverage and ε-stability,(ii) are of minimum size among all the representations guaranteeing these properties, and (iii) contain only non-dominated points.

5.1 Bi-objective linear programming

The use of Algorithm 1 is particularly effective for bi-objective linear programs since polynomial routines $Restrict_0$ are available. In practice, we just need to solve a sequence of linear programs.

To illustrate our algorithm, we apply it on a manpower planning problem stated in the reference textbook by Williams [26]. In a context where new machinery is installed, a company must decide, over a 3 years horizon, whether to recruit, retrain, or make redundant some employees of different categories. The first objective is to minimize the number of employees made redundant (f_1) while the second objective is to minimize the total cost of retraining, redundancy, hiring additional employees (f_2). This gives rise to a linear program with 60 (continuous) variables and 24 constraints precisely described in [26]. Williams also provides the two extreme optimal solutions: the one minimizing redundancy, corresponding to point y^{*1}, leads to a (rounded) number of 842 employees made redundant for a total cost of £1 438 383, and the one minimizing the total cost, corresponding to point y^{*2}, leads to a (rounded) number of 1 424 employees made redundant for a total cost of £498 677. The author observes that this second solution saves £939 706 but results in 582 extra redundancies and concludes that the cost of saving each job could, therefore, be regarded as £1 615.

ε-kernels are represented in Figure 4 for different values of ε. The corresponding number of points and CPU times are reported for each case.

Considering the ε-kernel for $\varepsilon = 0.001$, which provides a precise representation of the non-dominated set, the previous conclusion can be refined. Indeed, for each of the optimal points, and particularly for y^{*1}, a very small decay on the optimal value of the optimized criterion leads to a substantial improvement on the other criterion. Quite interestingly, the ε-kernels for larger values of ε tend to focus on the central points and ignore the extreme optimal points. In particular for $\varepsilon = 0.05$, the representation which contains only 6 points proposes a point $y^1 = (877, 967\,055)$ with the best evaluation on objective f_1 and a point $y^2 = (1\,299, 523\,611)$ with the best evaluation on objective f_2. We check that y^i does indeed $(1+\varepsilon)$-dominate y^{*i}, $i = 1, 2$, which means that the loss on optimality remains within the tolerance margin. Above all the gain on the other criterion is much larger than 5%: for y^1, the gain on criterion f_2 is 1 438 383 - 967 055 = 471 328, representing a 32.77% gain and for y^2, the gain on criterion f_1 is 1 424 - 1 299 = 125, representing a 8.37% gain.

This explains why points between y^{*i} and y^i, $i = 1, 2$, are not part of the representation.

5.2 Bi-objective shortest path problem

We consider here the well known pathological family of instances introduced in Hansen [14], and depicted in Figure 5. Each of the 2^n feasible paths from vertex v_0 to vertex v_n corresponds to a non-dominated point, illustrating the intractability of the bi-objective shortest path problem.

We tested our algorithm when $n = 25$, corresponding to $2^{25} = 33\,554\,432$ non-dominated points. Due to the size of the non-dominated set, it is practically impossible to compute this set. However, an appropriate representation can be computed extremely quickly using Algorithm 1. Information on the size and time required to compute ε-kernels of minimum size is reported in Table 1.

The graphical representation of this instance is given in Figure 6. Representing about 33.5 millions points with only 21 points, while guaranteeing that any other point can be at most

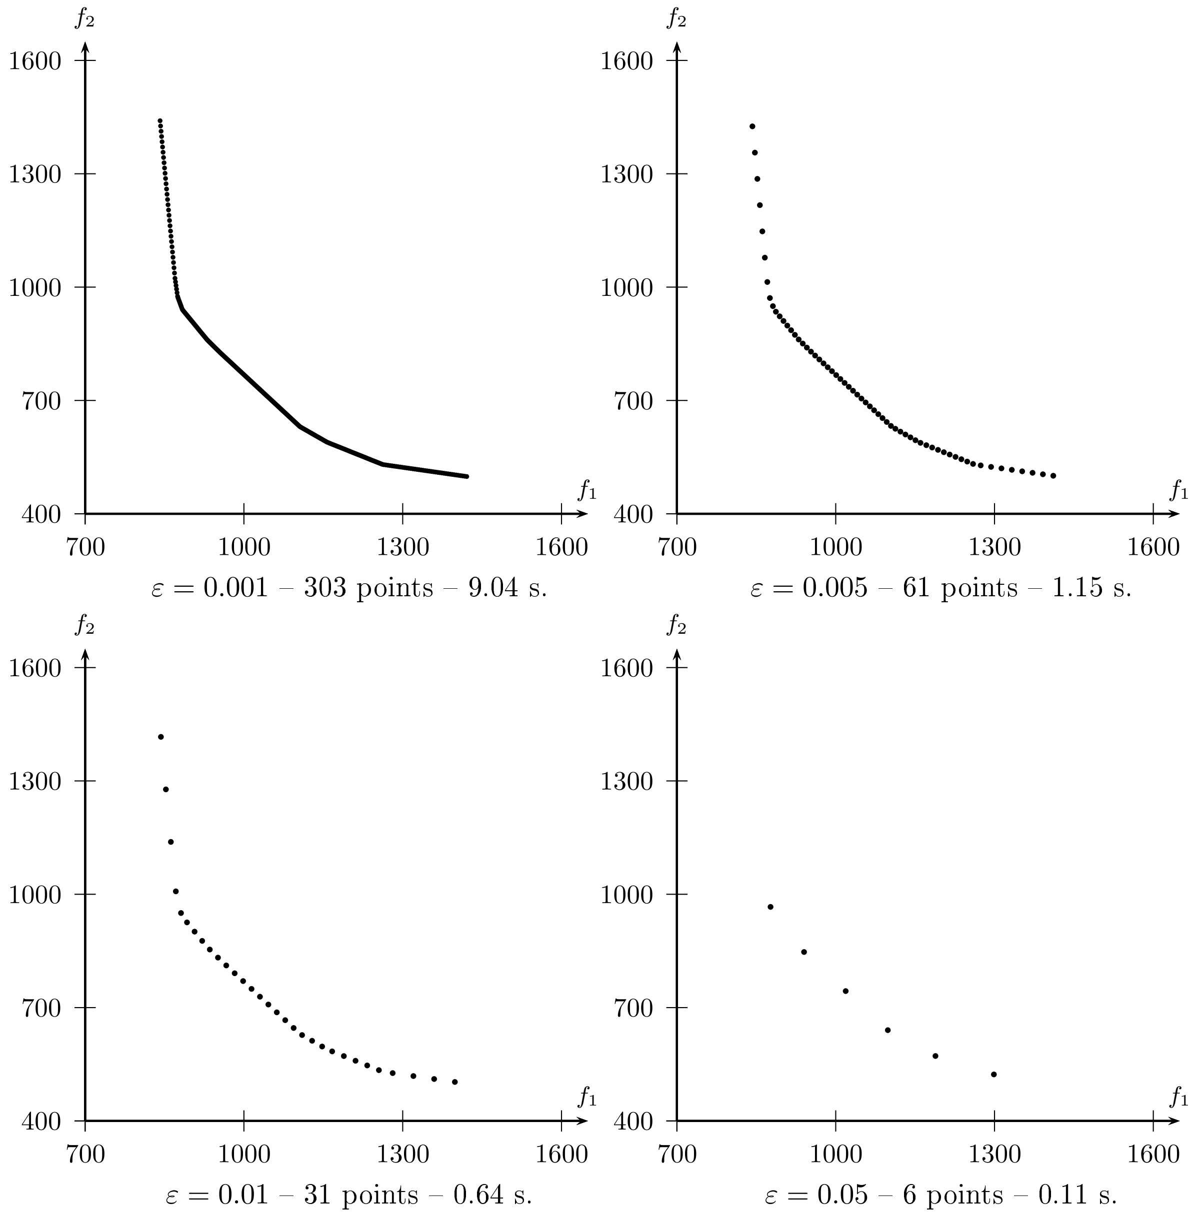

Figure 4: The manpower planning problem: different ε-kernels

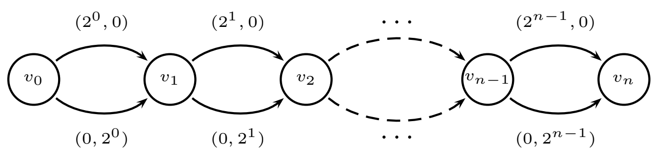

Figure 5: Intractable instances for the bi-objective shortest path problem

5% better than one of these points, is quite remarkable. We also notice that, unlike in the previous linear programming example, the points are extremely well-dispersed. This is due to the fact that in these very specific instances there is a constant tradeoff of one unit between consecutive non-dominated points.

ε	size	CPU time (s.)
0.01	101	5.27
0.05	21	1.53
0.1	11	0.61

Table 1: Different ε-kernels for the bi-objective shortest path problem (Hansen instance, $n = 25$)

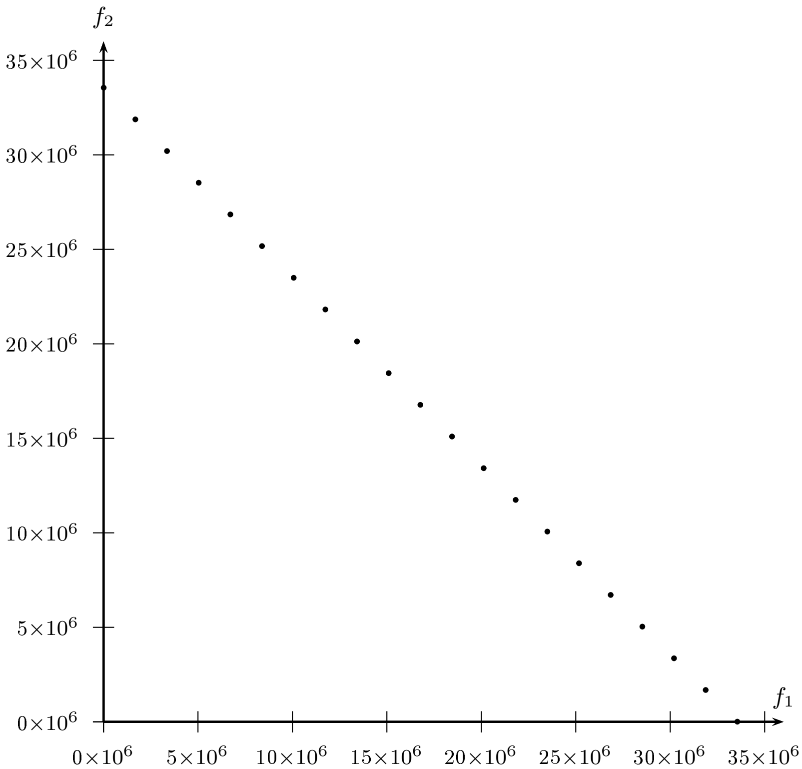

Figure 6: ε-kernel for the bi-objective shortest path problem (Hansen instance, n=25), $\varepsilon = 0.05$

5.3 Bi-objective assignment problem

We consider now the bi-objective assignment problem, which consists in assigning n resources to n tasks taking into account two total cost functions to be minimized. A resource is assigned to one and only one task and a task is assigned to one and only one resource. Each resource-task assignment involves two costs. Each total cost of an assignment is computed by adding up the costs of every chosen resource-task assignment.

We test our algorithm using the largest and most difficult instance used in [18], where a specific two-phase algorithm is proposed for the bi-objective assignment problem. This instance, called 2AP100-1A100, is available on GuepardLib, a library of multi-objective combinatorial optimization instances (`http://guepard.lip6.fr/Main/GuepardLib`). For this instance, $n = 100$ and each cost is generated randomly, independently, uniformly in $\{0, \ldots, 99\}$.

We first used a standard e-constraint approach to generate the whole non-dominated set. Information on the size and time required to compute ε-kernels of minimum size is reported in Table 2.

ε	size	CPU time (s.)
0	947	684.31
(e-constraint)		
0.01	197	263.95
0.05	40	50.45
0.1	21	23.20

Table 2: Different ε-kernels for the bi-objective assignment problem (2AP100-1A100 instance)

The graphical representation of this instance is given in Figure 7. We observe, here again, the modulation of the dispersion of points depending on the shape of the non-dominated set.

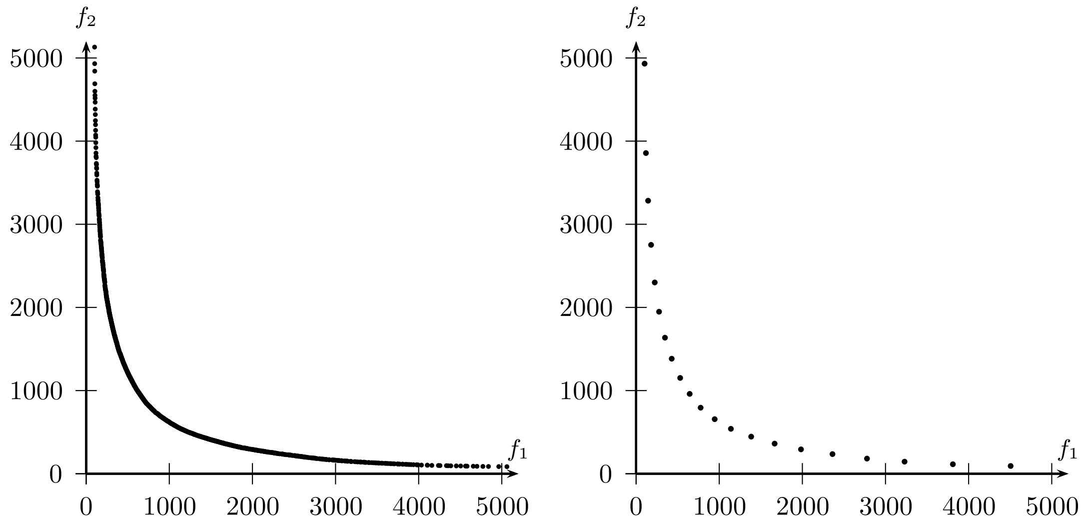

Non-dominated set – 947 points

ε-kernel – $\varepsilon = 0.1$ - 21 points

Figure 7: Instance 2AP100-1A100: non-dominated set and ε-kernel

6 Conclusions

The purpose of this work was to produce discrete and tractable representations of the set of non-dominated points for multi-objective optimization problems. We considered that representations should satisfy some conditions of *coverage*, *spacing*, and *cardinality*. For this purpose, we introduced the concept of $(\varepsilon, \varepsilon')$-kernel which is a particular ε-Pareto set that satisfies an additional condition of stability implementing *spacing*. We proposed some generic methods to produce $(\varepsilon, \varepsilon')$-kernels. Our algorithms run in polynomial time if and only if the routines called in the algorithms run in polynomial time.

The situation for the bi-objective case is quite clear and the concept of $(\varepsilon, \varepsilon')$-kernel, or even ε-kernel, seems quite relevant to provide a good discrete representation of the non-dominated set. Our experiments demonstrate the practical applicability of our algorithm. For more than two objectives, we showed that imposing a condition of spacing may impact negatively on the cardinality. Since a coverage condition must necessarily be imposed, the choice is between emphasizing spacing or cardinality. If the condition on spacing prevails, we showed that it is possible to construct an $(\varepsilon, \varepsilon')$-kernel, with $\varepsilon' \leq \sqrt[3]{1+\varepsilon} - 1$, provided that the GAP_δ routine is available, but without any guarantee on its cardinality. If the condition on cardinality prevails, known guarantees are very weak, even without any condition on spacing. The only known results deal with the case where the points of the objective space are explicitly given. In this case, finding an ε-Pareto set of minimal size can be formulated as a minimum set cover problem. Thus, it is $\log n$-approximable by the greedy algorithm [15] and it is proved in [2] that the greedy algorithm cannot perform better on these specific set cover instances. Moreover, for three objectives, Koltun and Papadimitriou [15] show the existence of a polynomial algorithm which returns an ε-Pareto set of size at most $c \cdot opt_\varepsilon$ where c is a large constant. Obtaining a better approximation of the size of a smallest ε-Pareto set for this specific case, as well as for more general cases, are challenging open questions. From a practical point of view, designing algorithms for more than two objectives, that would focus either on spacing or cardinality, is also a very interesting question.

References

[1] C. Bazgan, H. Hugot, and D. Vanderpooten. Implementing an efficient fptas for the 0–1 multi-objective knapsack problem. *European Journal of Operational Research*, 198(1):47–56, 2009.

[2] C. Bazgan, F. Jamain, and D. Vanderpooten. Approximate pareto sets of minimal size for multi-objective optimization problems. *Operations Research Letters*, 43(1):1–6, 2015.

[3] A. Berger, V. Bonifaci, F. Grandoni, and G. Schäfer. Budgeted matching and budgeted matroid intersection via the gasoline puzzle. *Mathematical Programming, Ser. A*, 128(1-2):355–372, 2011.

[4] V. Chvátal and L. Lovász. Every directed graph has a semi-kernel. *Lecture Notes in Mathematics*, 411:175, 1974.

[5] I. Diakonikolas and M. Yannakakis. Small approximate Pareto sets for bi-objective shortest paths and other problems. *SIAM Journal on Computing*, 39(4):1340–1371, 2009.

[6] P. Duchet. Graphes noyau-parfaits. *Annals of Discrete Mathematics*, 9:93–101, 1980.

[7] P. Duchet, Y.O. Hamidoune, and H. Meyniel. Sur les quasi-noyaux d'un graphe. *Discrete Mathematics*, 65(3):231–235, 1987.

[8] M. Ehrgott. *Multicriteria Optimization.* Springer, 2000.

[9] M. Ehrgott, L. Shao, and A. Schöbel. An approximation algorithm for convex multi-objective programming problems. *Journal of Global Optimization*, 50(3):397–416, 2011.

[10] F. Ergün, R.K. Sinha, and L. Zhang. An improved FPTAS for Restricted Shortest Path. *Information Processing Letters*, 83(5):287–291, 2002.

[11] A. Eusébio, J.R. Figueira, and M. Ehrgott. On finding representative non-dominated points for bi-objective integer network flow problems. *Computers & Operations Research*, 48:1–10, 2014.

[12] S.L. Faulkenberg and M.M. Wiecek. On the quality of discrete representations in multiple objective programming. *Optimization and Engineering*, 11(3):423–440, 2010.

[13] C. Filippi and E. Stevanato. Approximation schemes for bi-objective combinatorial optimization and their application to the TSP with profits. *Computers & Operations Research*, 40(10):2418–2428, 2013.

[14] P. Hansen. Bicriterion path problems. In *Multiple Criteria Decision Making: Theory and Application (G. Fandel and T. Gal, Eds.), LNEMS 177, Springer*, pages 109–127, 1979.

[15] V. Koltun and C. Papadimitriou. Approximately dominating representatives. *Theoretical Computer Science*, 371(3):148–154, 2007.

[16] D.H. Lorenz and D. Raz. A simple efficient approximation scheme for the restricted shortest path problem. *Operations Research Letters*, 28(5):213–219, 2001.

[17] C. Papadimitriou and M. Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS'00)*, pages 86–92, 2000.

[18] A. Przybylski, X. Gandibleux, and M. Ehrgott. Two phase algorithms for the bi-objective assignment problem. *European Journal of Operational Research*, 185(2):509–533, 2008.

[19] R. Ravi and M.X. Goemans. The constrained minimum spanning tree problem (extended abstract). In *Proceedings of the 5th Scandinavian Workshop on Algorithms and Theory (SWAT'96), LNCS 1097*, pages 66–75, 1996.

[20] S. Ruzika and M.M. Wiecek. Survey paper: approximation methods in multiobjective programming. *Journal of Optimization Theory and Applications* 126(3):473–571, 2005.

[21] S. Sayin. Measuring the quality of discrete representations of efficient sets in multiple objective mathematical programming. *Mathematical Programming*, 87(3):543–560, 2000.

[22] L. Shao and M. Ehrgott. Discrete representation of non-dominated sets in multi-objective linear programming *European Journal of Operational Research*, 255(3): 687—698, 2016.

[23] D. Vanderpooten, L. Weerasena, and M.M. Wiecek. Covers and approximations in multiobjective optimization. *Journal of Global Optimization*, to appear, 2016, DOI 10.1007/s10898-016-0426-4.

[24] S. Vassilvitskii and M. Yannakakis. Efficiently computing succinct trade-off curves. *Theoretical Computer Science*, 348(2-3):334–356, 2005.

[25] D. Vaz, L. Paquete, C.M. Fonseca, K. Klamroth, M. Stiglmayr. Representation of the nondominated set in biobjective discrete optimization. *Computers & Operations Research*, 63:172–186, 2015.

[26] H.P. Williams. *Model building in Mathematical Programming.* Wiley, 5th edition, 2013.