Operations Research Letters 43 (2015) 1–6

Contents lists available at ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

Approximate Pareto sets of minimal size for multi-objective optimization problems

Cristina Bazgan [a,b], Florian Jamain [a], Daniel Vanderpooten [a,*]

[a] PSL, Université Paris-Dauphine, LAMSADE UMR 7243, Place du Maréchal de Lattre de Tassigny, 75775 Paris Cedex 16, France

[b] Institut Universitaire de France, France

ARTICLE INFO

Article history:
Received 27 May 2014
Received in revised form
16 October 2014
Accepted 16 October 2014
Available online 27 October 2014

Keywords:
Multi-objective optimization
Pareto set
Non-dominated points
Approximation algorithm
Greedy algorithm

ABSTRACT

We are interested in a problem introduced by Vassilvitskii and Yannakakis (2005), the computation of a minimum set of solutions that approximates within an accuracy ε the Pareto set of a multi-objective optimization problem. We mainly establish a new 3-approximation algorithm for the bi-objective case. We also propose a study of the greedy algorithm performance for the tri-objective case when the points are given explicitly, answering an open question raised by Koltun and Papadimitriou in (2007).

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. Therefore, the standard situation is that any solution can always be improved on at least one objective. The solutions of interest, called *efficient* solutions, are these such that any solution which is better on one objective is necessarily worse on at least one other objective. In other words, a solution is efficient if its corresponding vector of objective values is not dominated by any other vector of objective 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*). Indeed, it is well-known, in particular, that most multi-objective combinatorial optimization problems 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 [4]. Thus, instead of producing the full set of non-dominated points, we may prefer to provide an approximation of this set. This idea is represented by the concept of an ε-Pareto set, which is a set P_ε of solutions that approximately dominates every other solutions, i.e. such that for every solution s, it contains a solution s' that is better within a factor $1+\varepsilon$ than s in all the objectives. The existence of ε-Pareto sets of polynomial size is well-known [10] and polynomial time algorithms that produce ε-Pareto sets have been developed and improved for many multi-objective optimization problems, including Multi-objective Shortest Path [7,13,11], Multi-objective Knapsack [5,1]. However, note that there may exist many ε-Pareto sets, some of which can have very small size and some others very large size. An interesting problem introduced by [12] and continued in [3] is the efficient construction of ε-Pareto sets of size as small as possible. This paper focuses on the same issue.

In the following section, we define the basic concepts, formalize the problem and recall some results of previous related works. Then, in Section 3, we mainly propose a new polynomial time 3-approximation algorithm of the size of a smallest ε-Pareto set for the bi-objective case. In Section 4, we analyze the performance of the greedy algorithm when the points of the objectives space are given explicitly in the input and the number of objectives is three, answering an open question raised in [9]. 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 several objectives, i.e. $\min_{x\in S}\{f_1(x),$

* Corresponding author.
E-mail addresses: bazgan@lamsade.dauphine.fr (C. Bazgan), florian.jamain@lamsade.dauphine.fr (F. Jamain), vdp@lamsade.dauphine.fr (D. Vanderpooten).

http://dx.doi.org/10.1016/j.orl.2014.10.003
0167-6377/

$\ldots, f_p(x)\}$, where $f_1, \ldots, f_p$ are $p \geq 2$ objective functions and S is the set of feasible solutions. In the case where some or all objective functions are 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 $Y \subseteq R^p$ which contains the criterion vectors or 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 criterion f_i of a point $y \in Y$ for $i = 1, \ldots, p$. We say that a point y *dominates* another point y' if y is at least as good as y' 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, called *Pareto set*.

Given a constant $c \geq 1$, a point y *c-dominates* another point y' if y is at least as good as y' up to a factor of c in all the objectives, i.e. $y_i \leq cy'_i$. For any rational $\varepsilon > 0$, an *ε-Pareto set* P_ε is a subset of feasible points such that for all $z \in P$, there exists $z' \in P_\varepsilon$ such that z' $(1+\varepsilon)$-dominates z. In the context of ε-Pareto sets, the central relation is the $(1+\varepsilon)$-dominance relation, denoted by $\preceq_\varepsilon$.

For a given instance I, there may exist several ε-Pareto sets, and these may have different sizes. It is shown in [10] that, for every classical multi-objective optimization problem, an ε-Pareto set of size polynomial in the input size and $1/\varepsilon$ always exists. Moreover its computation is related to the computation of the following routine GAP_δ.

Given an instance I of a given problem, a point y and a rational $\delta \geq 0$, the routine $GAP_\delta(y)$ 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 $i = 1, \ldots, p$.

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 (resp. fully polynomial time) [10].

Since an ε-Pareto set of polynomial size can still be quite large, Vassilvitskii and Yannakakis investigate in [12] the determination of ε-Pareto sets of minimal size. These authors also propose *generic* algorithms to deal with this problem. 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. Assuming in the following that the objective functions take positive rational values whose numerators and denominators have at most m bits, any feasible point has a value between 2^{-m} and 2^m and moreover the difference between the values of any two solutions is at least 2^{-2m} for any criterion. From [10], opt_ε is polynomial in the input size and $1/\varepsilon$.

In order to use generic algorithms, Diakonikolas and Yannakakis introduced in [3] two other general purpose routines called $Restrict_\delta$ and $DualRestrict_\delta$ for the bi-objective case.

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

Given an instance I, a bound b and a rational $\delta \geq 0$, the routine $DualRestrict_\delta(f_1, f_2 \leq b)$ 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 [3].

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 [10] that opt_ε is polynomial in the input size and $1/\varepsilon$.

We are interested in generic algorithms that compute in polynomial time an ε-Pareto set of minimal size. For the bi-objective case, a generic algorithm that computes an ε-Pareto set of size at most $3opt_\varepsilon$ was established in [12] using routines GAP_δ. Moreover, if the routine 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 [3] that an ε-Pareto set of size at most $2opt_\varepsilon$ is computable in polynomial time if there exists routines $Restrict_\delta$ 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 $k.opt_\varepsilon$ is called a k-approximation algorithm.

3. Two objectives

We first present a hardness result for the Bi-objective Knapsack problem then we propose a new generic algorithm that approximates the size of a smallest ε-Pareto set to a factor 3, which is much simpler and, in some cases, more efficient than the one presented in [12].

3.1. Approximation hardness for Bi-objective Knapsack

Diakonikolas and Yannakakis [3] showed that the size of a smallest ε-Pareto set of Bi-objective Shortest Path and Bi-objective Spanning Tree cannot be approximated within a factor better than 2 in polynomial time, unless P = NP. These results are tight since these two problems admit a routine $Restrict_\delta$ that runs in polynomial time, and thus an ε-Pareto set of size at most $2opt_\varepsilon$ is computable in polynomial time as shown in [3]. Vassilvitski and Yannakakis [12] showed that the size of a smallest ε-Pareto set of an artificial variant of Knapsack, called Bi-objective 2-Type-Knapsack, cannot be approximated within a factor better than 3 in polynomial time, unless P = NP. This result is also tight since this problem has a routine GAP_δ that runs in polynomial time, and thus an ε-Pareto set of size at most $3opt_\varepsilon$ is computable in polynomial time as shown in [12].

In this part, we investigate the status of the classical version, called Bi-objective Knapsack, with as input a set Q of items, a capacity c and for each item i two values $v_1(i), v_2(i)$ and a weight $w(i)$. Values and weights are positive rationals. A solution is a nonempty subset Q' of items with total values $v_1(Q') = \sum_{i \in Q'} v_1(i)$, $v_2(Q') = \sum_{i \in Q'} v_2(i)$ and a total weight $w(Q') = \sum_{i \in Q'} w(i) \leq c$. The goal is to maximize the values. First, note that the size of a smallest ε-Pareto set of Bi-objective Knapsack is approximable in polynomial time to a factor 3 since this problem admits an FPTAS, which is equivalent to the existence of a polynomial time routine GAP_δ [5]. We prove that the size of a smallest ε-Pareto set of Bi-objective Knapsack is not approximable in polynomial time within a factor better than 3, if P ≠ NP.

Theorem 1. *For* Bi-objective Knapsack *the size of a smallest ε-Pareto set cannot be approximated within a factor better than* 3 *in polynomial time, unless* P = NP.