Lower Bounds for the Smoothed Number of Pareto optimal Solutions*

Tobias Brunsch and Heiko Röglin

Department of Computer Science, University of Bonn, Germany
brunsch@cs.uni-bonn.de, heiko@roeglin.org

Abstract. In 2009, Röglin and Teng showed that the smoothed number of Pareto optimal solutions of linear multi-criteria optimization problems is polynomially bounded in the number n of variables and the maximum density ϕ of the semi-random input model for any fixed number of objective functions. Their bound is, however, not very practical because the exponents grow exponentially in the number $d+1$ of objective functions. In a recent breakthrough, Moitra and O'Donnell improved this bound significantly to $O\big(n^{2d}\phi^{d(d+1)/2}\big)$.
An "intriguing problem", which Moitra and O'Donnell formulate in their paper, is how much further this bound can be improved. The previous lower bounds do not exclude the possibility of a polynomial upper bound whose degree does not depend on d. In this paper we resolve this question by constructing a class of instances with $\Omega((n\phi)^{(d-\log(d))\cdot(1-\Theta(1/\phi))})$ Pareto optimal solutions in expectation. For the bi-criteria case we present a higher lower bound of $\Omega(n^2\phi^{1-\Theta(1/\phi)})$, which almost matches the known upper bound of $O(n^2\phi)$.

1 Introduction

In multi-criteria optimization problems we are given several objectives and aim at finding a solution that is simultaneously optimal in all of them. In most cases the objectives are conflicting and no such solution exists. The most popular way to deal with this problem is to just concentrate on the relevant solutions. If a solution is *dominated* by another solution, i.e., it is worse than the other solution in at least one objective and not better in the others, then this solution does not have to be considered for our optimization problem. All solutions that are not dominated by any other solution are called *Pareto optimal* and form the so-called *Pareto set*. For a general introduction to multi-criteria optimization problems, we refer the reader to the book of Matthias Ehrgott [Ehr05].

Smoothed Analysis For many multi-criteria optimization problems the worst-case size of the Pareto set is exponential. However, worst-case analysis is often too pessimistic, whereas average-case analysis assumes a certain distribution

* A part of this work was done at Maastricht University and was supported by a Veni grant from the Netherlands Organisation for Scientific Research.

on the input universe. Usually it is hard if not impossible to find a distribution resembling practical instances. *Smoothed analysis*, introduced by Spielman and Teng [ST04] to explain the efficiency of the simplex algorithm in practice despite its exponential worst-case running time, is a combination of both approaches and has been successfully applied to a variety of fields like machine learning, numerical analysis, discrete mathematics, and combinatorial optimization in the past decade (see [ST09] for a survey). Like in a worst-case analysis the model of smoothed analysis still considers adverserial instances. In contrast to the worst-case model, however, these instances are subsequently slightly perturbed at random, for example by Gaussian noise. This assumption is made to model that often the input an algorithm gets is subject to imprecise measurements, rounding errors, or numerical imprecision. In a more general model of smoothed analysis, introduced by Beier and Vöcking [BV04], the adversary is even allowed to specify the probability distribution of the random noise. The influence he can exert is described by a parameter ϕ denoting the maximum density of the noise.

Optimization Problems and Smoothed Input Model Beier and Vöcking [BV04] have initiated the study of binary bi-criteria optimization problems. In their model, which has been extended to multi-criteria problems by Röglin and Teng [RT09], one considers optimization problems that can be specified in the following form. There are an arbitrary set $\mathcal{S} \subseteq \{0,1\}^n$ of *solutions* and $d+1$ objective functions $w^j : \mathcal{S} \to \mathbb{R}$, $j = 0, \dots, d$, given. While w^0 can be an arbitrary function, which is to be minimized, the functions $w^1, \dots, w^d$, which are to be maximized, are linear of the form $w^j(s) = w^j_1 s_1 + \ldots + w^j_n s_n$ for $s = (s_1, \dots, s_n) \in \mathcal{S}$. Formally, the problem can be described as follows:

minimize $w^0(s)$, and **maximize** $w^j(s)$ for all $j = 1, \dots, d$
subject to s in the feasible region $\mathcal{S}$.

As there are no restrictions on the set $\mathcal{S}$ of solutions, this model is quite general and can encode many well-studied problems like, e.g., the multi-criteria knapsack, shortest path, or spanning tree problem. Let us remark that the choice which objective functions are to be maximized and minimized is arbitrary and just chosen for ease of presentation. All results also hold for other combinations of objective functions.

In the framework of smoothed analysis the coefficients $w^j_1, \dots, w^j_n$ of the linear functions w^j are drawn according to (adversarial) probability density functions $f_{i,j} : [-1,1] \to \mathbb{R}$ that are bounded by the maximum density parameter ϕ, i.e., $f_{i,j} \leq \phi$ for $i = 1, \dots, n$ and $j = 1, \dots, d$. The adversary could, for example, choose for each coefficient an interval of length $1/\phi$ from which it is chosen uniformly at random. Hence, the parameter ϕ determines how powerful the adversary is. For large ϕ he can specify the coefficients very precisely, and for $\phi \to \infty$ the smoothed analysis becomes a worst-case analysis. The coefficients are restricted to the interval $[-1,1]$ because otherwise, the adversary could diminish the effect of the perturbation by choosing large coefficients.

Previous Work Beier and Vöcking [BV04] showed that for $d = 1$ the expected size of the Pareto set of the optimization problem above is $O(n^4\phi)$ regardless of how the set $\mathcal{S}$, the objective function w^0 and the densities $f_{i,j}$ are chosen. Later, Beier, Röglin, and Vöcking [BRV07] improved this bound to $O(n^2\phi)$ by analyzing the so-called *loser gap*. Röglin and Teng [RT09] generalized the notion of this gap to higher dimensions, i.e., $d \geq 2$, and gave the first polynomial bound in n and ϕ for the smoothed number of Pareto optimal solutions. Furthermore, they were able to bound higher moments. The degree of the polynomial, however, was $d^{\Theta(d)}$. Recently, Moitra and O'Donnell [MO10] showed a bound of $O(n^{2d}\phi^{d(d+1)/2})$, which is the first polynomial bound for the expected size of the Pareto set with degree polynomial in d. An "intriguing problem" with which Moitra and O'Donnell conclude their paper is whether their upper bound could be significantly improved, for example to $f(d, \phi)n^2$. Moitra and O'Donnell suspect that for constant ϕ there should be a lower bound of $\Omega\left(n^d\right)$. In this paper we resolve this question almost completely.

Our Contribution For the bi-criteria case, i.e., $d = 1$, we prove a lower bound of $\Omega\left(\min\left\{n^2\phi^{1-\Theta(1/\phi)}, 2^{\Theta(n)}\right\}\right)$. This is the first bound with dependence on n and ϕ and it nearly matches the upper bound $O(\min\left\{n^2\phi, 2^n\right\})$. For $d \geq 2$ we prove a lower bound of $\Omega\left(\min\left\{(n\phi)^{(d-\log(d))\cdot(1-\Theta(1/\phi))}, 2^{\Theta(n)}\right\}\right)$. Note that throughout the paper "log" denotes the binary logarithm. This is the first bound for the general multi-criteria case. Still, there is a significant gap between this lower bound and the upper bound of $O(\min\left\{n^{2d}\phi^{d(d+1)/2}, 2^n\right\})$, but the exponent of n is nearly $d - \log(d)$. Hence our lower bound is close to the lower bound of $\Omega\left(n^d\right)$ conjectured by Moitra and O'Donnell.

Restricted Knapsack Problem To prove the lower bounds stated above we consider a variant of the knapsack problem where we have n objects $a_1, \ldots, a_n$, each with a weight w_i and a profit vector $p_i \in [0, 1]^d$ for a positive integer d. By a vector $s \in \{0, 1\}^n$ we describe which objects to put into the knapsack. In contrast to the unrestricted variant not all combinations of objects are allowed. Instead, all valid combinations are described by a set $\mathcal{S} \subseteq \{0, 1\}$. We want to simultaneously minimize the total weight and maximize all total profits of a solution s. Thus, the *restricted knapsack problem*, denoted by $K_{\mathcal{S}}(\{a_1, \ldots, a_n\})$, can be written as

$$\textbf{minimize} \sum_{i=1}^{n} w_i \cdot s_i, \quad \text{and} \quad \textbf{maximize} \sum_{i=1}^{n} (p_i)_j \cdot s_i \text{ for all } j = 1, \ldots, d$$
$$\textbf{subject to } s \text{ in the feasible region } \mathcal{S}.$$

For $\mathcal{S} = \{0, 1\}^n$ we just write $K(\{a_1, \ldots, a_n\})$ instead of $K_{\mathcal{S}}(\{a_1, \ldots, a_n\})$.

Note that the instances of the restricted knapsack problem that we use to prove the lower bounds are not necessarily interesting on its own because they have a somewhat artificial structure. However, they are interesting as they show that the known upper bounds in the general model cannot be significantly improved.

2 The Bi-criteria Case

In this section we present a lower bound for the expected number of Pareto optimal solutions in bi-criteria optimization problems that shows that the upper bound of Beier, Röglin, and Vöcking [BRV07] cannot be significantly improved.

Theorem 1. *There is a class of instances for the restricted bi-criteria knapsack problem for which the expected number of Pareto-optimal solutions is lower bounded by*

$$\Omega\left(\min\left\{n^2\phi^{1-\Theta(1/\phi)}, 2^{\Theta(n)}\right\}\right),$$

where n is the number of objects and ϕ is the maximum density of the profits' probability distributions.

Note that the exponents of n and ϕ in this bound are asymptotically the same as the exponents in the upper bound $O(\min\{n^2\phi, 2^n\})$ proved by Beier, Röglin, and Vöcking [BRV07].

For our construction we use the following bound from Beier and Vöcking.

Theorem 2 ([BV04]). *Let $a_1, \dots, a_n$ be objects with weights $2^1, \dots, 2^n$ and profits $p_1, \dots, p_n$ that are independently and uniformly distributed in $[0,1]$. Then, the expected number of Pareto optimal solutions of $K(\{a_1, \dots, a_n\})$ is $\Omega\left(n^2\right)$.*

Note that scaling all profits does not change the Pareto set and hence Theorem 2 remains true if the profits are chosen uniformly from $[0, a]$ for an arbitrary $a > 0$. We will exploit this observation later in our construction.

The idea how to create a large Pareto set is what we call the copy step. Let us consider an additional object b with weight 2^{n+1} and fixed profit q. In Figure 1 all solutions are represented by a weight-profit pair in the weight-profit space. The set of solutions using object b can be considered as the set of solutions that do not use object b, but shifted by $(2^{n+1}, q)$. If the profit q is chosen sufficiently large, i.e., larger than the sum of the profits of the objects $a_1, \dots, a_n$, then there is no domination between solutions from different copies and hence the Pareto optimal solutions of $K(\{a_1, \dots, a_n, b\})$ are just the copies of the Pareto optimal solutions of $K(\{a_1, \dots, a_n\})$. Lemma 3 formalizes this observation.

Lemma 3. *Let $a_1, \dots, a_n$ be objects with weights $2^1, \dots, 2^n$ and non-negative profits $p_1, \dots, p_n$ and let b be an object with weight 2^{n+1} and profit $q > \sum_{i=1}^{n} p_i$. Furthermore, let $\mathcal{P}$ denote the Pareto set of $K(\{a_1, \dots, a_n\})$ and let $\mathcal{P}'$ denote the Pareto set of $K(\{a_1, \dots, a_n, b\})$. Then, $\mathcal{P}'$ is the disjoint union of $\mathcal{P}'_0 := \{(s, 0) :\ s \in \mathcal{P}\}$ and $\mathcal{P}'_1 := \{(s, 1) :\ s \in \mathcal{P}\}$ and thus $|\mathcal{P}'| = 2 \cdot |\mathcal{P}|$.*

Now we use the copy idea to construct a large Pareto set. Let $a_1, \dots, a_{n_p}$ be objects with weights $2^1, \dots, 2^{n_p}$ and with profits $p_1, \dots, p_{n_p} \in P := [0, 1/\phi]$ where $\phi > 1$, and let $b_1, \dots, b_{n_q}$ be objects with weights $2^{n_p+1}, \dots, 2^{n_p+n_q}$ and with profits $q_i \in Q_i := (m_i - \lceil m_i \rceil / \phi, m_i]$, where $m_i = (n_p + 1)/(\phi - 1) \cdot ((2\phi - 1)/(\phi - 1))^{i-1}$. The choice of the intervals Q_i is due to the fact that we have to ensure $q_i > \sum_{j=1}^{n_p} p_j + \sum_{j=1}^{i-1} q_j$ to apply Lemma 3 successively for the objects

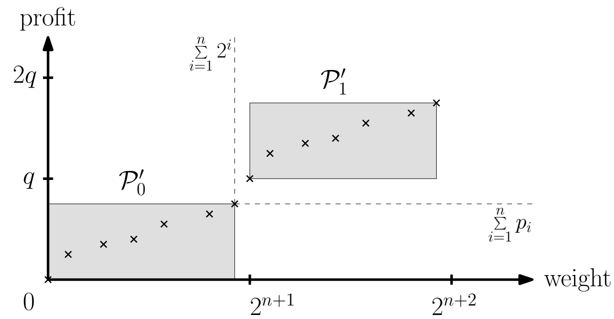

Fig. 1. The copy step. The Pareto set $\mathcal{P}'$ consist of two copies of the Pareto set $\mathcal{P}$.

$b_1, \ldots, b_{n_q}$. We will prove this inequality in Lemma 4. More interesting is the fact that the size of an interval Q_i is $\lceil m_i \rceil / \phi$ which might be larger than $1/\phi$. To explain this consider the case $m_i > 1$ for some index i. For this index the interval Q_i is not a subset of $[-1, 1]$ as required for our model. Instead of avoiding such large values m_i by choosing n_q small enough, we will split Q_i into $\lceil m_i \rceil$ intervals of equal size which must be at least $1/\phi$. This so-called split step will be explained later.

Lemma 4. *Let* $p_1, \ldots, p_{n_p} \in P$ *and let* $q_i \in Q_i$. *Then,* $q_i > \sum_{j=1}^{n_p} p_j + \sum_{j=1}^{i-1} q_j$ *for all* $i = 1, \ldots, n_q$.

Note that with Lemma 4 we implicitely show that the lower boundaries of the intervals Q_i are non-negative.

Proof. Using the definition of m_i, we get

$$\begin{aligned} q_i &> m_i - \frac{\lceil m_i \rceil}{\phi} \geq m_i - \frac{m_i + 1}{\phi} = \frac{\phi - 1}{\phi} \cdot m_i - \frac{1}{\phi} \\ &= \frac{n_p + 1}{\phi} \cdot \left(\frac{2\phi - 1}{\phi - 1} \right)^{i-1} - \frac{1}{\phi} \, . \end{aligned}$$

On the other hand we have

$$\begin{aligned} \sum_{j=1}^{n_p} p_j + \sum_{j=1}^{i-1} q_j \leq \sum_{j=1}^{n_p} \frac{1}{\phi} + \sum_{j=1}^{i-1} m_j &= \frac{n_p}{\phi} + \sum_{j=1}^{i-1} \frac{n_p + 1}{\phi - 1} \cdot \left(\frac{2\phi - 1}{\phi - 1} \right)^{j-1} \\ &= \frac{n_p}{\phi} + \frac{n_p + 1}{\phi - 1} \cdot \frac{\left(\frac{2\phi - 1}{\phi - 1} \right)^{i-1} - 1}{\frac{2\phi - 1}{\phi - 1} - 1} \\ &= \frac{n_p}{\phi} + \frac{n_p + 1}{\phi} \cdot \left(\left(\frac{2\phi - 1}{\phi - 1} \right)^{i-1} - 1 \right) \\ &= \frac{n_p + 1}{\phi} \cdot \left(\frac{2\phi - 1}{\phi - 1} \right)^{i-1} - \frac{1}{\phi} \, . \end{aligned}$$

□

Combining Theorem 2, Lemma 3 and Lemma 4, we immediately get a lower bound for the knapsack problem using the objects $a_1, \dots, a_{n_p}$ and $b_1, \dots, b_{n_q}$ with profits chosen from P and Q_i, respectively.

Corollary 5. *Let $a_1, \dots, a_{n_p}$ and $b_1, \dots, b_{n_q}$ be as above, but the profits p_i are chosen uniformly from P and the profits q_i are arbitrarily chosen from Q_i. Then, the expected number of Pareto optimal solutions of $K(\{a_1, \dots, a_{n_p}, b_1, \dots, b_{n_q}\})$ is $\Omega\left(n_p^2 \cdot 2^{n_q}\right)$.*

Proof. Because of Lemma 4, we can apply Lemma 3 for each realization of the profits $p_1, \dots, p_{n_p}$ and $q_1, \dots, q_{n_q}$. This implies that the expected number of Pareto optimal solutions is 2^{n_q} times the expected size of the Pareto set of $K(\{a_1, \dots, a_{n_p}\})$ which is $\Omega\left(n_p^2\right)$ according to Theorem 2. $\square$

The profits of the objects b_i grow exponentially and leave the interval $[0, 1]$. As mentioned earlier, we resolve this problem by splitting each object b_i into $k_i := \lceil m_i \rceil$ objects $b_i^{(1)}, \dots, b_i^{(k_i)}$ with the same total weight and the same total profit, i.e., each with weight $2^{n_p+i}/k_i$ and profit $q_i^{(l)} \in Q_i/k_i := (m_i/k_i - 1/\phi, m_i/k_i]$. As the intervals Q_i are subsets of $\mathbb{R}_+$, the intervals Q_i/k_i are subsets of $[0, 1]$. It remains to ensure that for any fixed index i all objects $b_i^{(l)}$ are treated as a group. This can be done by restricting the set $\mathcal{S}$ of solutions. Let $\mathcal{S}_i = \{(0, \dots, 0), (1, \dots, 1)\} \subseteq \{0, 1\}^{k_i}$. Then, the set $\mathcal{S}$ of solutions is defined as $\mathcal{S} := \{0, 1\}^{n_p} \times \prod_{i=1}^{n_q} \mathcal{S}_i$. By choosing the set of solutions that way, the objects $b_i^{(1)}, \dots, b_i^{(k_i)}$ can be viewed as substitute for object b_i. Thus, a direct consequence of Corollary 5 is the following.

Corollary 6. *Let $\mathcal{S}$, $a_1, \dots, a_{n_p}$ and $b_i^{(l)}$ be as above, let the profits $p_1, \dots, p_{n_p}$ be chosen uniformly from P and let the profits $q_i^{(1)}, \dots, q_i^{(k_i)}$ be chosen uniformly from Q_i/k_i. Then, the expected number of Pareto optimal solutions of $K_{\mathcal{S}}(\{a_1, \dots, a_{n_p}\} \cup \{b_i^{(l)} : i = 1, \dots, n_q,\ l = 1, \dots, k_i\})$ is $\Omega\left(n_p^2 \cdot 2^{n_q}\right)$.*

The remainder contains just some technical details. First, we give an upper bound for the number of objects $b_i^{(l)}$.

Lemma 7. *The number of objects $b_i^{(l)}$ is upper bounded by $n_q + \frac{n_p+1}{\phi} \cdot \left(\frac{2\phi-1}{\phi-1}\right)^{n_q}$.*

Proof. The number of objects $b_i^{(l)}$ is $\sum_{i=1}^{n_q} k_i = \sum_{i=1}^{n_q} \lceil m_i \rceil \le n_q + \sum_{i=1}^{n_q} m_i$, and

$$\begin{aligned}\sum_{i=1}^{n_q} m_i &= \frac{n_p+1}{\phi-1} \cdot \sum_{i=1}^{n_q} \left(\frac{2\phi-1}{\phi-1}\right)^{i-1} \le \frac{n_p+1}{\phi-1} \cdot \frac{\left(\frac{2\phi-1}{\phi-1}\right)^{n_q}}{\frac{2\phi-1}{\phi-1} - 1} \\ &= \frac{n_p+1}{\phi} \cdot \left(\frac{2\phi-1}{\phi-1}\right)^{n_q} .\end{aligned}$$

$\square$

Now we are able to prove Theorem 1.

Proof (Theorem 1). Without loss of generality let $n \geq 4$ and $\phi \geq \frac{3+\sqrt{5}}{2} \approx 2.62$. For the moment let us assume $\phi \leq (\frac{2\phi-1}{\phi-1})^{\frac{n-1}{3}}$. This is the interesting case leading to the first term in the minimum in Theorem 1. We set $\hat{n}_q := \frac{\log(\phi)}{\log((2\phi-1)/(\phi-1))} \in [1, \frac{n-1}{3}]$ and $\hat{n}_p := \frac{n-1-\hat{n}_q}{2} \geq \frac{n-1}{3} \geq 1$. All inequalities hold because of the bounds on n and ϕ. We obtain the numbers n_p and n_q by rounding, i.e., $n_p := \lfloor \hat{n}_p \rfloor \geq 1$ and $n_q := \lfloor \hat{n}_q \rfloor \geq 1$. Now we consider objects $a_1, \ldots, a_{n_p}$ with weights 2^i and profits chosen uniformly from P, and objects $b_i^{(l)}$, $i = 1, \ldots, n_q$, $l = 1, \ldots, k_i$, with weights $2^{n_p+i}/k_i$ and profits chosen uniformly from Q_i/k_i. Observe that P and all Q_i/k_i have length $1/\phi$ and thus the densities of all profits are bounded by ϕ. Let N be the number of all these objects. By Lemma 7, this number is bounded by

$$\begin{aligned} N &\leq n_p + n_q + \frac{n_p+1}{\phi} \cdot \left(\frac{2\phi-1}{\phi-1}\right)^{n_q} \leq \hat{n}_p + \hat{n}_q + \frac{\hat{n}_p+1}{\phi} \cdot \left(\frac{2\phi-1}{\phi-1}\right)^{\hat{n}_q} \\ &= \hat{n}_p + \hat{n}_q + \frac{\hat{n}_p+1}{\phi} \cdot \phi = 2\hat{n}_p + \hat{n}_q + 1 = n\,. \end{aligned}$$

Hence, the number N of objects we actually use is at most n, as required. As set of solutions we consider $\mathcal{S} := \{0,1\}^{n_p} \times \prod_{i=1}^{n_q} \mathcal{S}_i$. Due to Corollary 6, the expected size of the Pareto set of $K_{\mathcal{S}}(\{a_1, \ldots, a_{n_p}\} \cup \{b_i^{(l)} : i = 1, \ldots, n_q,\ l = 1, \ldots, k_i\})$ is

$$\begin{aligned} \Omega\left(n_p^2 \cdot 2^{n_q}\right) &= \Omega\left(\hat{n}_p^2 \cdot 2^{\hat{n}_q}\right) = \Omega\left(\hat{n}_p^2 \cdot 2^{\frac{\log(\phi)}{\log\left(\frac{2\phi-1}{\phi-1}\right)}}\right) = \Omega\left(n^2 \cdot \phi^{\frac{1}{\log\left(\frac{2\phi-1}{\phi-1}\right)}}\right) \\ &= \Omega\left(n^2 \cdot \phi^{1-\Theta(1/\phi)}\right)\,, \end{aligned}$$

where the last step holds because

$$\frac{1}{\log\left(2 + \frac{c_1}{\phi-c_2}\right)} = 1 - \frac{\log\left(1 + \frac{c_1}{2\phi-2c_2}\right)}{\log\left(2 + \frac{c_1}{\phi-c_2}\right)} = 1 - \frac{\Theta\left(\frac{c_1}{2\phi-2c_2}\right)}{\Theta(1)} = 1 - \Theta\left(\frac{1}{\phi}\right)$$

for any constants $c_1, c_2 > 0$. We formulated this calculation slightly more general than necessary as we will use it again in the multi-criteria case.

For $\phi > (\frac{2\phi-1}{\phi-1})^{\frac{n-1}{3}}$ we construct the same instance as above, but for maximum density $\phi' > 1$ where $\phi' = (\frac{2\phi'-1}{\phi'-1})^{\frac{n-1}{3}}$. Since $n \geq 4$, ϕ' exists, is unique and $\phi' \in \left[\frac{3+\sqrt{5}}{2}, \phi\right)$. This yields $\hat{n}'_p = \hat{n}'_q = \frac{n-1}{3}$ and, as above, the expected size of the Pareto set is $\Omega\left((\hat{n}'_p)^2 \cdot 2^{\hat{n}'_q}\right) = \Omega\left(n^2 \cdot 2^{\Theta(n)}\right) = \Omega\left(2^{\Theta(n)}\right)$. □

3 The Multi-criteria Case

In this section we present a lower bound for the expected number of Pareto optimal solutions in multi-criteria optimization problems. We concentrate our attention to $d \geq 2$ as we discussed the case $d = 1$ in the previous section.

Theorem 8. *For any fixed integer $d \geq 2$ there is a class of instances for the restricted $(d+1)$-dimensional knapsack problem for which the expected number of Pareto-optimal solutions is lower bounded by*

$$\Omega\left(\min\left\{(n\phi)^{(d-\log(d))\cdot(1-\Theta(1/\phi))}, 2^{\Theta(n)}\right\}\right),$$

where n is the number of objects and ϕ is the maximum density of the profit's probability distributions.

Unfortunately, Theorem 8 does not generalize Theorem 1. This is due to the fact that, though we know an explicit formula for the expected number of Pareto optimal solutions if all profits are uniformly chosen from $[0,1]$, we were not able to find a simple non-trivial lower bound for it. Hence, in the general multi-criteria case, we concentrate on analyzing the copy and split steps.

In the bi-criteria case we used an additional object b to copy the Pareto set (see Figure 1). For that we had to ensure that every solution using this object has higher weight than all solutions without b. The same had to hold for the profit. Since all profits are in $[0,1]$, the profit of every solution must be in $[0,n]$. As the Pareto set of the first $n_p \leq n/2$ objects has profits in $[0, n/(2\phi)]$, we could fit $n_q = \Theta(\log(\phi))$ copies of this initial Pareto set into the interval $[0,n]$.

In the multi-criteria case, every solution has a profit in $[0,n]^d$. In our construction, the initial Pareto set consists only of a single solution, but we benefit from the fact that the number of mutually non-dominating copies of the initial Pareto set that we can fit into the hypercube $[0,n]^d$ grows quickly with d.

Let us consider the case that we have some Pareto set $\mathcal{P}$ whose profits lie in some hypercube $[0,a]^d$. We will create $\binom{d}{d_h}$ copies of this Pareto set; one for every vector $x \in \{0,1\}^d$ with exactly $d_h = \lceil d/2 \rceil$ ones. Let $x \in \{0,1\}^d$ be such a vector. Then we generate the corresponding copy C_x of the Pareto set $\mathcal{P}$ by shifting it by $a + \varepsilon$ in every dimension i with $x_i = 1$. If all solutions in these copies have higher weights than the solutions in the initial Pareto set $\mathcal{P}$, then the initial Pareto set stays Pareto optimal. Furthermore, for each pair of copies C_x and C_y, there is one index i with $x_i = 1$ and $y_i = 0$. Hence, solutions from C_y cannot dominate solutions from C_x. Similarly, one can argue that no solution in the initial copy can dominate any solution from C_x. This shows that all solutions in copy C_x are Pareto optimal. All the copies (including the initial one) have profits in $[0, 2a+\varepsilon]^d$ and together $|\mathcal{P}| \cdot \left(1 + \binom{d}{d_h}\right) \geq |\mathcal{P}| \cdot 2^d/d$ solutions.

We start with an initial Pareto set of a single solution with profit in $[0, 1/\phi]^d$, and hence we can make $\Theta(\log(n\phi))$ copy steps before the hypercube $[0,n]^d$ is filled. In each of these steps the number of Pareto optimal solutions increases by a factor of at least $2^d/d$, yielding a total number of at least $(2^d/d)^{\Theta(\log(n\phi))} = (n\phi)^{\Theta(d-\log(d))}$ Pareto optimal solutions.

In the following, we describe how these copy steps can be realized in the restricted knapsack problem. Again, we have to make a split step because the profit of every object must be in $[0,1]^d$. Due to such technicalities, the actual bound we prove looks slightly different than the one above. It turns out that we need (before splitting) d new objects $b_1, \ldots, b_d$ for each copy step in contrast to the bi-criteria case, where (before splitting) a single object b was enough.

Let $n_q \geq 1$ be an arbitrary positive integer and let $\phi \geq 2d$ be a real. We consider objects $b_{i,j}$ with weights $2^i/d_h$ and profit vectors

$$q_{i,j} \in Q_{i,j} := \prod_{k=1}^{j-1} \left[0, \frac{\lceil m_i \rceil}{\phi}\right] \times \left(m_i - \frac{\lceil m_i \rceil}{\phi}, m_i\right] \times \prod_{k=j+1}^{d} \left[0, \frac{\lceil m_i \rceil}{\phi}\right] ,$$

where m_i is recursively defined as

$$m_0 := 0 \quad \text{and} \quad m_i := \frac{1}{\phi - d} \cdot \left(\sum_{l=0}^{i-1} (m_l \cdot (\phi + d) + d) \right), \; i = 1, \ldots, n_q \,. \tag{1}$$

The explicit formula for this recurrence is

$$m_i = \frac{d}{\phi + d} \cdot \left(\left(\frac{2\phi}{\phi - d} \right)^i - 1 \right), \; i = 1, \ldots, n_q \,.$$

The d-dimensional interval $Q_{i,j}$ is of the form that the j^{th} profit of object $b_{i,j}$ is large and all the other profits are small. By using object $b_{i,j}$ the copy of the Pareto set is shifted in direction of the j^{th} unit vector. As mentioned in the motivation we will choose exactly d_h such objects to create additional copies. To give a better intuition for the form of the single intervals the d-dimensional interval $Q_{i,j}$ is constructed of we refer the reader to the explanation in the bi-criteria case.

Let $H(x)$ be the *Hamming weight* of a 0-1-vector x, i.e., the number of ones in x, and let $\hat{\mathcal{S}} := \{x \in \{0,1\}^d \,:\, H(x) \in \{0, d_h\}\}$ denote the set of all 0-1-vectors of length d with 0 or d_h ones. As set $\mathcal{S}$ of solutions we consider $\mathcal{S} := \hat{\mathcal{S}}^{n_q}$.

Lemma 9. *Let the set $\mathcal{S}$ of solutions and the objects $b_{i,j}$ be as above. Then, each solution $s \in \mathcal{S}$ is Pareto optimal for $K_{\mathcal{S}}(\{b_{i,j} \,:\, i = 1, \ldots, n_q, \; j = 1, \ldots, d\})$.*

Proof. We show the statement by induction over n_q and discuss the base case and the inductive step simultaneously because of similar arguments. Let $\mathcal{S}' := \hat{\mathcal{S}}^{n_q - 1}$ and let $(s, s_{n_q}) \in \mathcal{S}' \times \hat{\mathcal{S}}$ be an arbitrary solution from $\mathcal{S}$. Note that for $n_q = 1$ we get $s = \lambda$, the 0-1-vector of length 0. First we show that there is no domination within one copy, i.e., there is no solution of type $(s', s_{n_q}) \in \mathcal{S}$ that dominates (s, s_{n_q}). For $n_q = 1$ this is obviously true. For $n_q \geq 2$ the existence of such a solution would imply that s' dominates s in the knapsack problem $K_{\mathcal{S}'}(\{b_{i,j} \,:\, i = 1, \ldots, n_q - 1, \; j = 1, \ldots, d\})$. This contradicts the inductive hypothesis.

Now we prove that there is no domination between solutions from different copies, i.e., there is no solution of type $(s', s'_{n_q}) \in \mathcal{S}$ with $s'_{n_q} \neq s_{n_q}$ that dominates (s, s_{n_q}). If $s_{n_q} = \mathbf{0}$, then the total weight of the solution (s, s_{n_q}) is at most $\sum_{i=1}^{n_q - 1} 2^i < 2^{n_q}$. The right side of this inequality is a lower bound for the weight of solution (s', s'_{n_q}) because $s'_{n_q} \neq s_{n_q}$. Hence, (s', s'_{n_q}) does not dominate (s, s_{n_q}). Finally, let us consider the case $s_{n_q} \neq \mathbf{0}$. There must be an index $j \in [d]$

where $(s_{n_q})_j = 1$ and $(s'_{n_q})_j = 0$. We show that the j^{th} total profit of (s, s_{n_q}) is higher than the j^{th} profit of (s', s'_{n_q}). The former one is strictly bounded from below by $m_{n_q} - \lceil m_{n_q} \rceil / \phi$, whereas the latter one is bounded from above by

$$\sum_{i=1}^{n_q-1} \left((d_h - 1) \cdot \frac{\lceil m_i \rceil}{\phi} + \max\left\{ \frac{\lceil m_i \rceil}{\phi}, m_i \right\} \right) + d_h \cdot \frac{\lceil m_{n_q} \rceil}{\phi} .$$

Solution (s', s'_{n_q}) can use at most d_h objects of each group $b_{i,1}, \ldots, b_{i,d}$. Each of them, except one, can contribute at most $\frac{\lceil m_i \rceil}{\phi}$ to the j^{th} total profit. One can contribute either at most $\frac{\lceil m_i \rceil}{\phi}$ or at most m_i. This argument also holds for the n_q^{th} group, but by the choice of index j we know that each object chosen by s'_{n_q} contributes at most $\frac{\lceil m_i \rceil}{\phi}$ to the j^{th} total profit. It is easy to see that $\lceil m_i \rceil / \phi \leq m_i$ because of $\phi > d \geq 1$. Hence, our bound simplifies to

$$\begin{aligned}
&\sum_{i=1}^{n_q-1} \left((d_h - 1) \cdot \frac{\lceil m_i \rceil}{\phi} + m_i \right) + d_h \cdot \frac{\lceil m_{n_q} \rceil}{\phi} \\
&\quad \leq \sum_{i=1}^{n_q-1} \left(d \cdot \frac{m_i + 1}{\phi} + m_i \right) + (d-1) \cdot \frac{m_{n_q} + 1}{\phi} && (d \geq 2) \\
&\quad = \frac{1}{\phi} \cdot \left(\sum_{i=1}^{n_q-1} (m_i \cdot (\phi + d) + d) + d \cdot (m_{n_q} + 1) \right) - \frac{m_{n_q} + 1}{\phi} \\
&\quad = \frac{1}{\phi} \cdot \left(\sum_{i=0}^{n_q-1} (m_i \cdot (\phi + d) + d) + d \cdot m_{n_q} \right) - \frac{m_{n_q} + 1}{\phi} && (m_0 = 0) \\
&\quad = \frac{1}{\phi} \cdot ((\phi - d) \cdot m_{n_q} + d \cdot m_{n_q}) - \frac{m_{n_q} + 1}{\phi} && \text{(Equ. (1))} \\
&\quad \leq m_{n_q} - \frac{\lceil m_{n_q} \rceil}{\phi} .
\end{aligned}$$

This implies that (s', s'_{n_q}) does not dominate (s, s_{n_q}). □

Immediately, we get a statement about the expected number of Pareto optimal solutions if we randomize.

Corollary 10. *Let $\mathcal{S}$ and $b_{i,j}$ be as above, but the profit vectors $q_{i,j}$ are arbitrarily drawn from $Q_{i,j}$. Then, the expected number of Pareto optimal solutions for $K_{\mathcal{S}}(\{b_{i,j} : i = 1, \ldots, n_q,\ j = 1, \ldots, d\})$ is at least $(2^d/d)^{n_q}$.*

Proof. This result follows from Lemma 9 and $|\hat{S}| = 1 + \binom{d}{d_h} = 1 + \max_{i=1,\ldots,d} \binom{d}{i} \geq 1 + (\sum_{i=1}^{d} \binom{d}{i})/d = 1 + (2^d - 1)/d \geq 2^d/d$. □

As in the bi-criteria case we now split each object $b_{i,j}$ into $k_i := \lceil m_i \rceil$ objects $b_{i,j}^{(1)}, \ldots, b_{i,j}^{(k_i)}$ with weights $2^i/(k_i \cdot d_h)$ and with profit vectors

$$q_{i,j}^{(l)} \in Q_{i,j}/k_i := \prod_{k=1}^{j-1} \left[0, \frac{1}{\phi}\right] \times \left(\frac{m_i}{k_i} - \frac{1}{\phi}, \frac{m_i}{k_i}\right] \times \prod_{k=j+1}^{d} \left[0, \frac{1}{\phi}\right].$$

Then, we adapt our set $\mathcal{S}$ of solutions such that for any fixed indices i and j either all objects $b_{i,j}^{(1)}, \ldots, b_{i,j}^{(k_i)}$ are put into the knapsack or none of them. Corollary 10 yields the following result.

Corollary 11. *Let $\mathcal{S}$ and $b_{i,j}^{(l)}$ be as described above, but let the profit vectors $p_{i,j}^{(1)}, \ldots, p_{i,j}^{(k_i)}$ be chosen uniformly from $Q_{i,j}/k_i$. Then, the expected number of Pareto optimal solutions of $K_{\mathcal{S}}(\{b_{i,j}^{(l)} :\ i = 1, \ldots, n_q,\ j = 1, \ldots, d,\ l = 1, \ldots, k_i\})$ is at least $(2^d/d)^{n_q}$.*

Still, the lower bound is expressed in n_q and not in the number of objects used. So the next step is to analyze the number of objects.

Lemma 12. *The number of objects $b_{i,j}^{(l)}$ is upper bounded by $d \cdot n_q + \frac{2d^2}{\phi-d} \cdot \left(\frac{2\phi}{\phi-d}\right)^{n_q}$.*

Proof. The number of objects $b_{i,j}^{(l)}$ is $\sum_{i=1}^{n_q}(d \cdot k_i) = d \cdot \sum_{i=1}^{n_q} \lceil m_i \rceil \leq d \cdot n_q + d \cdot \sum_{i=1}^{n_q} m_i$, and

$$\sum_{i=1}^{n_q} m_i \leq \frac{d}{\phi+d} \cdot \sum_{i=1}^{n_q} \left(\frac{2\phi}{\phi-d}\right)^i \leq \frac{d}{\phi+d} \cdot \frac{\left(\frac{2\phi}{\phi-d}\right)^{n_q+1}}{\left(\frac{2\phi}{\phi-d}\right) - 1}$$
$$\leq \frac{d}{\phi} \cdot \left(\frac{2\phi}{\phi-d}\right) \cdot \left(\frac{2\phi}{\phi-d}\right)^{n_q} = \frac{2d}{\phi-d} \cdot \left(\frac{2\phi}{\phi-d}\right)^{n_q}. \qquad \square$$

Now we can prove Theorem 8.

Proof. Without loss of generality let $n \geq 16d$ and $\phi \geq 2d$. For the moment let us assume $\phi - d \leq \frac{4d^2}{n} \cdot \left(\frac{2\phi}{\phi-d}\right)^{\frac{n}{2d}}$. This is the interesting case leading to the first term in the minimum in Theorem 8. We set $\hat{n}_q := \frac{\log\left((\phi-d) \cdot \frac{n}{4d^2}\right)}{\log\left(\frac{2\phi}{\phi-d}\right)} \in \left[1, \frac{n}{2d}\right]$ and obtain $n_q := \lfloor \hat{n}_q \rfloor \geq 1$ by rounding. All inequalities hold because of the bounds on n and ϕ. Now we consider objects $b_{i,j}^{(l)}$, $i = 1, \ldots, n_q$, $j = 1, \ldots, d$, $l = 1, \ldots, k_i$, with weights $2^i/(k_i \cdot d)$ and profit vectors $q_{i,j}$ chosen uniformly from $Q_{i,j}/k_i$. All these intervals have length $1/\phi$ and hence all densities are bounded by ϕ. Let N be the number of objects. By Lemma 12, this number is bounded by

$$N \leq d \cdot n_q + \frac{2d^2}{\phi-d} \cdot \left(\frac{2\phi}{\phi-d}\right)^{n_q} \leq d \cdot \hat{n}_q + \frac{2d^2}{\phi-d} \cdot \left(\frac{2\phi}{\phi-d}\right)^{\hat{n}_q}$$
$$\leq d \cdot \hat{n}_q + \frac{2d^2}{\phi-d} \cdot (\phi-d) \cdot \frac{n}{4d^2} \leq n.$$

Hence, the number N of objects we actually use is at most n, as required. As set $\mathcal{S}$ of solutions we use the set described above, encoding the copy step and the split step. Due to Corollary 11, for fixed $d \geq 2$ the expected number of Pareto optimal solutions of $K_{\mathcal{S}}(\{b_{i,j}^{(l)} : i = 1, \ldots, n_q,\ j = 1, \ldots, d,\ l = 1, \ldots, k_i\})$ is

$$\begin{aligned}
\Omega\left(\left(\frac{2^d}{d}\right)^{n_q}\right) &= \Omega\left(\left(\frac{2^d}{d}\right)^{\hat{n}_q}\right) = \Omega\left(\left(\frac{2^d}{d}\right)^{\frac{\log\left((\phi-d)\cdot\frac{n}{4d^2}\right)}{\log\left(\frac{2\phi}{\phi-d}\right)}}\right) \\
&= \Omega\left(\left((\phi-d)\cdot\frac{n}{4d^2}\right)^{\frac{\log\left(\frac{2^d}{d}\right)}{\log\left(\frac{2\phi}{\phi-d}\right)}}\right) = \Omega\left((\phi\cdot n)^{\frac{d-\log(d)}{\log\left(\frac{2\phi}{\phi-d}\right)}}\right) \\
&= \Omega\left((\phi\cdot n)^{(d-\log(d))\cdot(1-\Theta(1/\phi))}\right) ,
\end{aligned}$$

where the last step holds because of the same reason as in the proof of Theorem 1.

In the case $\phi - d > \frac{4d^2}{n} \cdot \left(\frac{2\phi}{\phi-d}\right)^{\frac{n}{2d}}$ we construct the same instance above, but for a maximum density $\phi' > d$ where $\phi' - d = \frac{4d^2}{n} \cdot \left(\frac{2\phi'}{\phi'-d}\right)^{\frac{n}{2d}}$. Since $n \geq 16d$, the value ϕ' exists, is unique and $\phi' \in [65d, \phi)$. Futhermore, we get $\hat{n}_q = \frac{n}{2d}$. As above, the expected size of the Pareto set is $\Omega\left((2^d/d)^{\hat{n}_q}\right) = \Omega\left((2^d/d)^{n/(2d)}\right) = \Omega\left(2^{\Theta(n)}\right)$. □

References

[BRV07] René Beier, Heiko Röglin, and Berthold Vöcking. The smoothed number of Pareto optimal solutions in bicriteria integer optimization. In *Proc. of the 12th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 53–67, 2007.

[BV04] René Beier and Berthold Vöcking. Random knapsack in expected polynomial time. *Journal of Computer and System Sciences*, 69(3):306–329, 2004.

[Ehr05] Matthias Ehrgott. *Multicriteria Optimization*. Springer-Verlag, second edition, 2005.

[MO10] Ankur Moitra and Ryan O'Donnell. Pareto optimal solutions for smoothed analysts. Technical report, CoRR (abs/1011.2249), 2010. http://arxiv.org/abs/1011.2249. To appear in Proc. of the 43rd Annual ACM Symposium on Theory of Computing (STOC), 2011.

[RT09] Heiko Röglin and Shang-Hua Teng. Smoothed analysis of multiobjective optimization. In *Proc. of the 50th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 681–690, 2009.

[ST04] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *Journal of the ACM*, 51(3):385–463, 2004.

[ST09] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis: an attempt to explain the behavior of algorithms in practice. *Communications of the ACM*, 52(10):76–84, 2009.