Project Number 001907

DELIS

Dynamically Evolving, Large-scale Information Systems

Integrated Project

Member of the FET Proactive Initiative **Complex Systems**

DELIS-TR-539

The Smoothed Number of Pareto Optimal Solutions in Bicriteria Integer Optimization

René Beier and Heiko Röglin and Berthold Vöcking

2007

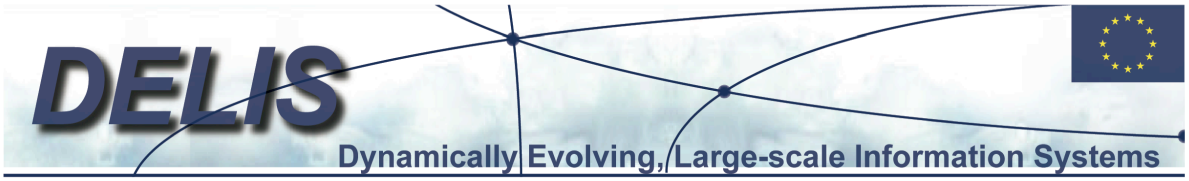

The Smoothed Number of Pareto Optimal Solutions in Bicriteria Integer Optimization*

Rene Beier[1], Heiko Röglin[2], and Berthold Vöcking[2]

[1] Max-Planck-Institut für Informatik
Saarbrücken, Germany
rbeier@mpi-inf.mpg.de
[2] Department of Computer Science
RWTH Aachen, D-52056 Aachen, Germany
{roeglin,voecking}@cs.rwth-aachen.de

Abstract. A well established heuristic approach for solving various bicriteria optimization problems is to enumerate the set of Pareto optimal solutions, typically using some kind of dynamic programming approach. The heuristics following this principle are often successful in practice. Their running time, however, depends on the number of enumerated solutions, which can be exponential in the worst case.
In this paper, we prove an almost tight bound on the expected number of Pareto optimal solutions for general bicriteria integer optimization problems in the framework of smoothed analysis. Our analysis is based on a semi-random input model in which an adversary can specify an input which is subsequently slightly perturbed at random, e. g., using a Gaussian or uniform distribution.
Our results directly imply tight polynomial bounds on the expected running time of the Nemhauser/Ullmann heuristic for the 0/1 knapsack problem. Furthermore, we can significantly improve the known results on the running time of heuristics for the bounded knapsack problem and for the bicriteria shortest path problem. Finally, our results also enable us to improve and simplify the previously known analysis of the smoothed complexity of integer programming.

1 Introduction

We study integer optimization problems having two criteria, say profit and weight, which are to be optimized simultaneously. A common approach for solving such problems is generating the set of *Pareto optimal* solutions, also known as the *Pareto set*. Pareto optimal solutions are optimal compromises of the two criteria in the sense that any improvement of one criterion implies an impairment to the other. In other words, a solution S^* is Pareto optimal if there exists no other solution S that dominates S^*, i. e., has at least the profit and at most the weight of S^* and at least one inequality is strict. Generating the Pareto set

* This work was supported by DFG grant VO 889/2 and by the EU within the 6th Framework Programme under contract 001907 (DELIS).

is of great interest in many scenarios and widely used in practice. Unfortunately, this approach fails to yield reasonable results in the worst case because even integer optimization problems with a simple combinatorial structure can have exponentially many Pareto optimal solutions. In practice, however, generating the Pareto set is often feasible since typically the number of Pareto optimal solutions does not attain its worst-case bound.

The discrepancy between practical experience and worst-case results motivates the study of the number of Pareto optimal solutions in a more realistic scenario. One possible approach is to study the average number of Pareto optimal solutions rather than the worst case number. In order to analyze the average, one has to define a probability distribution on the set of instances, with respect to which the average is taken. In most situations, however, it is not clear how to choose a probability distribution that reflects typical inputs. In order to bypass the limitations of worst-case and average-case analysis, Spielman and Teng defined the notion of *smoothed analysis* [15]. They consider a semi-random input model in which an adversary specifies an input which is then randomly perturbed. One can hope that semi-random input models are more realistic than worst-case and average-case input models since the adversary can specify an arbitrary input with a certain structure, and the subsequent perturbation generates an instance which is still close to the adversarial one.

We consider integer optimization problems in a semi-random setting, in which an adversary can specify an arbitrary set $\mathcal{S} \subseteq \mathcal{D}^n$ of feasible solutions and two objective functions: profit $p\colon \mathcal{S} \to \mathbb{R}$ and weight $w\colon \mathcal{S} \to \mathbb{R}$, where $\mathcal{D} \subset \mathbb{Z}$ denotes a finite set of integers. We assume that the profit is to be maximized and the weight is to be minimized. This assumption is without loss of generality as our results are not affected by changing the optimization direction of any of the objective functions. In our model, the weight function w can be chosen arbitrarily by the adversary, whereas the profit p has to be linear of the form $p(x) = p_1x_1 + \cdots + p_nx_n$. The adversary can choose an arbitrary vector of profits from $[-1, 1]^n$, but in the second step of the semi-random input model, the profits p_i are randomly perturbed by adding an independent Gaussian random variable with mean 0 and standard deviation σ to each profit p_i. The standard deviation σ can be seen as a parameter measuring how close the analysis is to a worst-case analysis: The smaller σ is chosen, the smaller is the influence of the perturbation and, hence, the closer is the analysis to a worst-case analysis. Our probabilistic analysis is not restricted to Gaussian perturbations but is much more general. In fact, it covers arbitrary probability distributions with a bounded density function and a finite absolute mean value. In particular, if one is interested in obtaining a positive domain for the profits, one can restrict the adversary to profits $p_i \in [0, 1]$ and perturb these profits by adding independent random variables that are distributed uniformly over some interval $[0, c]$.

We present a new method for bounding the expected number of Pareto optimal solutions in the aforementioned scenario which yields an upper bound that depends polynomially on the number of variables n, the integer with the largest absolute value in $\mathcal{D}$, and the reciprocal of the standard deviation σ. This imme-

diately implies polynomial upper bounds on the expected running time of several heuristics for generating the Pareto set of problems like, e. g., the Bounded Knapsack problem. Previous results of this kind were restricted to the case of binary optimization problems. For this special case, our method yields an improved upper bound, which matches the known lower bound. Furthermore, we show that our results on the expected number of Pareto optimal solutions yield a significantly simplified and improved analysis of the smoothed complexity of integer programming.

1.1 Previous Results

Multi-objective optimization is a well studied research area. Various algorithms for generating the Pareto set of various optimization problems like, e. g., the (bounded) knapsack problem [11, 8], the bicriteria shortest path problem [4, 14] and the bicriteria network flow problem [5, 10], have been proposed. The running time of these algorithms depends crucially on the number of Pareto optimal solutions and, hence, none of them runs in polynomial time in the worst case. In practice, however, generating the Pareto set is tractable in many situations. For instance, Müller-Hannemann and Weihe [9] study the number of Pareto optimal solutions in multi-criteria shortest path problems experimentally. They consider examples that arise from computing the set of best train connections (in view of travel time, fare, and number of train changes) and conclude that in this application scenario generating the complete Pareto set is tractable even for large instances. For more examples, we refer the reader to [6].

One way of coping with the bad worst-case behavior is to relax the requirement of finding the complete Pareto set. Papadimitriou and Yannakakis present a general framework for finding *approximate Pareto sets*. A solution S is ε-dominated by another solution S' if $p(S)/p(S') \leq 1+\varepsilon$ and $w(S')/w(S) \leq 1+\varepsilon$. We say that $\mathcal{P}_\varepsilon$ is an ε-approximation of a Pareto set $\mathcal{P}$ if for any solution $S \in \mathcal{P}$ there is a solution $S' \in \mathcal{P}_\varepsilon$ that ε-dominates it. Papadimitriou and Yannakakis show that for any Pareto set $\mathcal{P}$, there is an ε-approximation of $\mathcal{P}$ with polynomially many points (w. r. t. the input size and $1/\varepsilon$) [12]. Furthermore they give necessary and sufficient conditions under which there is an FPTAS to generate $\mathcal{P}_\varepsilon$. Vassilvitskii and Yannakakis [16] show how to compute ε-approximate Pareto curves of almost minimal size.

Beier and Vöcking analyze the expected number of Pareto optimal solutions for binary optimization problems [2]. They consider the aforementioned scenario with $\mathcal{D} = \{0, 1\}$ and show that the expected number of Pareto optimal solutions is bounded from above by $O(n^4/\sigma)$. This result implies that the Nemhauser/Ullmann algorithm [11] has polynomial expected running time. Furthermore, they also present a lower bound of $\Omega(n^2)$ on the expected number of Pareto optimal solutions for profits that are chosen uniformly from the interval $[0, 1]$.

In [3] Beier and Vöcking analyze the smoothed complexity of binary optimization problems. They consider optimization problems with one objective function in which the set of feasible solutions is given as $\mathcal{S} \cap \mathcal{B}_1 \cap \ldots \cap \mathcal{B}_m$, where

$\mathcal{S} \subseteq \{0,1\}^n$ denotes a fixed ground set and $\mathcal{B}_i$ denotes a halfspace induced by a linear constraint of the form $w_{i,1}x_1 + \cdots + w_{i,n}x_n \leq t_i$. Similar to the aforementioned model it is assumed that the coefficients $w_{i,j}$ are perturbed by adding independent random variables to them. Based on the probabilistic analysis of certain structural properties, Beier and Vöcking show that a binary optimization problem in this form has polynomial smoothed complexity if and only if there exists a pseudo-polynomial (w. r. t. the $w_{i,j}$) time algorithm for solving the problem. The term *polynomial smoothed complexity* is defined analogously to the way polynomial complexity is defined in average-case complexity theory, adding the requirement that the running time should be polynomially bounded not only in the input size but also in $1/\sigma$. This characterization is extended to the case of integer optimization problems where $\mathcal{D} \subset \mathbb{Z}$ is a finite set of integers by Röglin and Vöcking [13].

1.2 Our Results

In this paper, we present a new approach for bounding the expected number of Pareto optimal solutions for bicriteria integer optimization problems. This approach yields the first bounds for integer optimization problems and improves the known bound for the binary case significantly. We show that the expected number of Pareto optimal solutions is bounded from above by $O(n^2k^2\log(k)/\sigma)$ if $\mathcal{D} = \{0,\ldots,k-1\}$. We also present a lower bound of $\Omega(n^2k^2)$, assuming that the profits are chosen uniformly at random from the interval $[-1,1]$. For the case in which the adversary is restricted to linear weight functions, we present a lower bound of $\Omega(n^2k\log k)$. Furthermore, for the binary case $\mathcal{D} = \{0,1\}$, the upper bound simplifies to $O(n^2/\sigma)$, which improves the previously known bound by a factor of $\Theta(n^2)$ and matches the lower bound in [2] in terms of n. Hence, our method yields tight bounds in terms of n and almost tight bounds in terms of k for the expected number of Pareto optimal solutions and, thereby, even simplifies the proof in [2]. In the following, we list some applications of these results.

Knapsack Problem. The Nemhauser/Ullmann algorithm solves the knapsack problem by enumerating all Pareto optimal solutions [11]. Its running time on an instance with n items is $\Theta(\sum_{i=1}^{n} q_i)$, where q_i denotes the number of Pareto optimal solutions of the knapsack instance that consists only of the first i items. Beier and Vöcking analyze the expected number of Pareto optimal solutions and show that the expected running time of the Nemhauser/Ullmann algorithm is bounded by $O(n^5/\sigma)$ if all profits are perturbed by adding Gaussian or uniformly distributed random variables with standard deviation σ [2]. Based on our improved bounds on the expected number of Pareto optimal solutions, we conclude the following corollary.

Corollary 1. *For semi-random knapsack instances in which the profits are perturbed by adding independent Gaussian or uniformly distributed random variables with standard deviation σ, the expected running time of the Nemhauser/Ullmann algorithm is $O(n^3/\sigma)$*

For uniformly distributed profits Beier and Vöcking present a lower bound on the expected running time of $\Omega(n^3)$. Hence, we obtain tight bounds on the running time of the Nemhauser/Ullmann algorithm in terms of the number of items n. This lower bound can easily be extended to the case of Gaussian perturbations.

Bounded Knapsack Problem. In the bounded knapsack problem, a number $k \in \mathbb{N}$ and a set of n items with weights and profits are given. It is assumed that k instances of each of the n items are given. In [7] it is described how an instance with n items of the bounded knapsack problem can be transformed into an instance of the (binary) knapsack problem with $\Theta(n \log(k+1))$ items. Using this transformation, the bounded knapsack problem can be solved by the Nemhauser/Ullmann algorithm with running time $\Theta(\sum_{i=1}^{n \log(k+1)} q_i)$, where q_i denotes the number of Pareto optimal solutions of the binary knapsack instance that consists only of the first i items. Based on our results on the expected number of Pareto optimal solutions, we obtain the following corollary.

Corollary 2. *The expected running time of the Nemhauser/Ullmann algorithm on semi-random bounded knapsack instances in which the profits are perturbed by adding independent Gaussian or uniformly distributed random variables with standard deviation σ is bounded from above by $O(n^3k^2(\log^2(k+1))/\sigma)$ and bounded from below by $\Omega(n^3k\log^2(k+1))$.*

Hence, our results yield tight bounds in terms of n for the expected running time of the Nemhauser/Ullmann algorithm.

Bicriteria Shortest Path Problem. Different algorithms have been proposed for enumerating the Pareto set in bicriteria shortest path problems [4, 14]. In [4] a modified version of the Bellman/Ford algorithm is suggested. Beier shows that the expected running time of this algorithm is $O(nm^5/\sigma)$ for graphs with n nodes and m edges [1]. We obtain the following improved bound.

Corollary 3. *For semi-random bicriteria shortest path problems in which one objective function is linear and its coefficients are perturbed by adding independent Gaussian or uniformly distributed random variables with standard deviation σ, the expected running time of the modified Bellman/Ford algorithm is $O(nm^3/\sigma)$*

Smoothed Complexity of Integer Programming. We were not able to bound the expected number of Pareto optimal solutions for optimization problems with more than two objective functions. One approach for tackling multicriteria problems is to solve a constrained problem in which all objective functions except for one are made constraints. Our results for the bicriteria case can be used to improve the smoothed analysis of integer optimization problems with multiple constraints. In [13] we show that an integer optimization problem has polynomial smoothed complexity if and only if there exists a pseudo-polynomial time algorithm for solving the problem. To be more precise, we consider integer optimization problems in which an objective function is to be maximized over a

feasible region that is defined as the intersection of a fixed ground set $\mathcal{S} \subseteq \mathcal{D}^n$ with halfspaces $\mathcal{B}_1, \ldots, \mathcal{B}_m$ that are induced by m linear constraints of the form $w_{i,1}x_1 + \cdots + w_{i,n}x_n \leq t_i$, where the $w_{i,j}$ are independently perturbed by adding Gaussian or uniformly distributed random variables with standard deviation σ to them.

The term *polynomial smoothed complexity* is defined such that it is robust under different machine models analogously to the way polynomial average-case complexity is defined. One disadvantage of this definition is that polynomial smoothed/average-case complexity does not imply expected polynomial running time. For the binary case it is shown in [3] that problems that admit a pseudo-linear algorithm, i. e., an algorithm whose running time is bounded by $O(\text{poly}(N)W)$, where N denotes the input size and W the largest coefficient $|w_{i,j}|$ in the input, can be solved in expected polynomial time in the smoothed model. Based on our analysis of the expected number of Pareto optimal solutions, we generalize this result to the integer case.

Theorem 4. *Every integer optimization problem that can be solved in time $O(\text{poly}(N)W)$, where N denotes the input size and $W = \max_{i,j} |w_{i,j}|$, allows an algorithm with expected polynomial (in N and $1/\sigma$) running time for perturbed instances, in which an independent Gaussian or uniformly distributed random variables with standard deviation σ is added to each coefficient.*

In the following section, we introduce the probabilistic model we analyze, which is more general than the Gaussian and uniform perturbations described above. After that, in Sections 3 and 4, we present the upper and lower bounds on the expected number of Pareto optimal solutions. Finally, in Section 5, we present the applications of our results to the smoothed analysis of integer programming.

2 Model and Notations

For the sake of a simple presentation, using the framework of smoothed analysis, we described our results in the introduction not in their full generality. Our probabilistic analysis assumes that the adversary can choose, for each p_i, a probability distribution according to which p_i is chosen independently of the other profits. We prove an upper bound that depends linearly on the maximal density of the distributions and on the expected distance to zero. The maximal density of a continuous probability distribution, i. e., the supremum of the density function, is a parameter of the distribution, which we denote by ϕ. Similar to the standard deviation σ for Gaussian random variables, ϕ can be seen as a measure specifying how close the analysis is to a worst-case analysis. The larger ϕ, the more concentrated the probability mass can be. For Gaussian and uniformly distributed random variables, we have $\phi \sim 1/\sigma$.

In the following, we assume that p_i is a random variable with density f_i and that $f_i(x) \leq \phi_i$ for all $x \in \mathbb{R}$. Furthermore, we denote by μ_i the expected absolute value of p_i, i. e., $\mu_i = \mathbf{E}[|p_i|] = \int_{x \in \mathbb{R}} |x| f_i(x)\, dx$. Let $\phi = \max_{i \in [n]} \phi_i$ and $\mu = \max_{i \in [n]} \mu_i$. We denote by $[n]$ the set $\{1, \ldots, n\}$, and we use the notations $d = |\mathcal{D}|$ and $D = \max\{a - b \mid a, b \in \mathcal{D}\}$.

3 Upper Bound on the Expected Number of Pareto Optimal Solutions

While the profit function is assumed to be linear with stochastic coefficients, the weight function $w : \mathcal{S} \to \mathbb{R}$ can be chosen arbitrarily. We model this by assuming an explicit ranking of the solutions in $\mathcal{S}$, which can be chosen by the adversary. This way, we obtain a bicriteria optimization problem that aims at maximizing the rank as well as the profit. Observe that the weight function can map several solutions to the same value whereas the rank of a solution is always unique. This strict ordering, however, can only increase the number of Pareto optimal solutions.

Theorem 5. *Let $\mathcal{S} \subseteq \mathcal{D}^n$ be a set of arbitrarily ranked solutions with a finite domain $\mathcal{D} \subset \mathbb{Z}$. Define $d = |\mathcal{D}|$ and $D = \max\{a - b \mid a, b \in \mathcal{D}\}$. Assume that each profit p_i is a random variable with density function $f_i : \mathbb{R} \to \mathbb{R}_{\geq 0}$. Suppose $\mu_i = \mathbf{E}[|p_i|]$ and $\phi_i = \sup_{x \in \mathbb{R}} f_i(x)$. Let q denote the number of Pareto optimal solutions. Then*

$$\mathbf{E}[q] \leq 2DdH_d \left(\sum_{i=1}^{n} \phi_i\right) \left(\sum_{i=1}^{n} \mu_i\right) + O(dn) \ ,$$

where H_d is the d-th harmonic number. For $\mathcal{D} = \{0, \ldots, k-1\}$ and $\mu = \max_{i \in [n]} \mu_i$ and $\phi = \max_{i \in [n]} \phi_i$ the bound simplifies to

$$\mathbf{E}[q] = O(\mu\phi n^2 k^2 \log k) \ .$$

Note that the number of Pareto optimal solutions is not affected when all profits are scaled by some constant $c \neq 0$. This property is also reflected by the above bound. The random variable cp_i has maximal density ϕ_i/c and the expected absolute value is $c\mu_i$. Hence, the product $\phi\mu$ is invariant under scaling too.

Proof (Theorem 5). We use the following classification of Pareto optimal solutions. We say that a Pareto optimal solution x is of class $c \in \mathcal{D}$ if there exists an index $i \in [n]$ with $x_i \neq c$ such that the succeeding Pareto optimal solution y satisfies $y_i = c$, where succeeding Pareto optimal solution refers to the highest ranked Pareto optimal solution that is lower ranked than x. The lowest ranked Pareto optimal solution, which does not have a succeeding Pareto optimal solution, is not contained in any of the classes. A Pareto optimal solution can be in several classes but it is at least in one class. Let q_c denote the number of Pareto optimal solutions of class c. Since $q \leq 1 + \sum_{c \in \mathcal{D}} q_c$ it holds $\mathbf{E}[q] \leq 1 + \sum_{c \in \mathcal{D}} \mathbf{E}[q_c]$.

Lemma 6 enables us to bound the expected number of class-0 Pareto optimal solutions. In order to bound $\mathbf{E}[q_c]$ for values $c \neq 0$ we analyze a modified sequence of solutions. Starting from the original sequence $\mathcal{S} = x^1, x^2, \ldots, x^l$ ($x^j \in \mathcal{D}^n$), we obtain a modified sequence $\mathcal{S}^c$ by subtracting $(c, \ldots, c)$ from each solution vector x^j. This way, the profit of each solution is reduced by $c \sum p_i$. Observe that this operation does not affect the set of Pareto optimal solutions. A

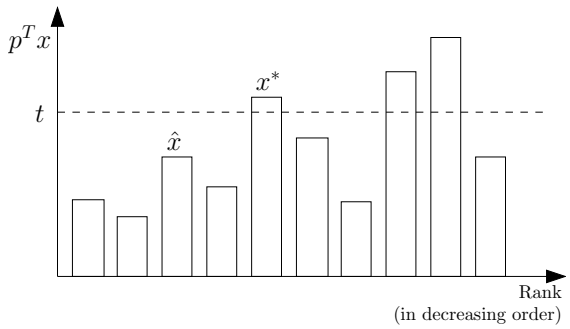

Fig. 1. If $\hat{x}$ is an ordinary class-0 Pareto optimal solution, then there must be an index i with $x_i^* = 0$ and $\hat{x}_i \neq 0$.

solution x is class-c Pareto optimal in $\mathcal{S}$ if and only if the corresponding solution $x-(c,\ldots,c)$ is class-0 Pareto optimal in $\mathcal{S}^c$. Hence, the number of class-c Pareto optimal solutions in $\mathcal{S}$ corresponds to the number of class-0 Pareto optimal solutions in $\mathcal{S}^c$. We apply Lemma 6 for the solution set $\mathcal{S}^c$ with a corresponding domain $\mathcal{D}^c = \{z - c : z \in \mathcal{D}\}$. Since the difference between the largest and the smallest element of the domain does not change, applying Lemma 6 yields that $\mathbf{E}[q]$ is bounded from above by

$$1 + \sum_{c\in\mathcal{D}} \mathbf{E}[q_0(S^c)] \leq 1 + \sum_{c\in\mathcal{D}} \left(D \left(\sum_{v\in\mathcal{D}^c\setminus\{0\}} |v|^{-1} \right) \left(\sum_{i=1}^{n} \phi_i \right) \left(\sum_{i=1}^{n} \mu_i \right) + n \right) \ ,$$

and the theorem follows. □

Lemma 6. *Let $\mathcal{S} \subseteq \mathcal{D}^n$ be a set of arbitrarily ranked solutions with a finite domain $\mathcal{D} \subset \mathbb{Z}$ with $0 \in \mathcal{D}$. Let D denote the difference between the largest and the smallest element in $\mathcal{D}$. Let q_0 denote the number of class-0 Pareto optimal solutions. Then*

$$\mathbf{E}[q_0] \leq D \left(\sum_{v\in\mathcal{D}\setminus\{0\}} |v|^{-1} \right) \left(\sum_{i=1}^{n} \phi_i \right) \left(\sum_{i=1}^{n} \mu_i \right) + n \ .$$

Proof. The key idea is to prove an upper bound on the probability that there exists a class-0 Pareto optimal solution whose profit falls into a small interval $(t-\varepsilon, t)$, for arbitrary t and ε. We will classify class-0 Pareto optimal solutions to be *ordinary* or *extraordinary*. Considering only ordinary solutions allows us to prove a bound that depends not only on the length ε of the interval but also on $|t|$, the distance to zero. This captures the intuition that it becomes increasingly unlikely to observe solutions whose profits are much larger than the expected profit of the most profitable solution. The final bound is obtained by observing that there can be at most n extraordinary class-0 Pareto optimal solutions.

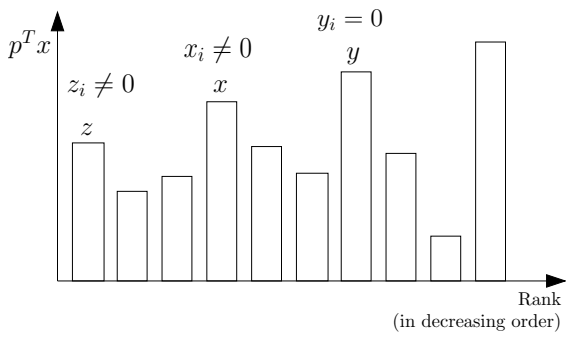

Fig. 2. In this case x is an extraordinary class-0 Pareto optimal solution.

We want to bound the probability that there exists an ordinary class-0 Pareto optimal solution whose profit lies in the interval $(t-\varepsilon, t)$. Define x^* to be the highest ranked solution from $\mathcal{S}$ satisfying $p^{\mathrm{T}}x \geq t$. If x^* exists then it is Pareto optimal. Let $\hat{x}$ denote the Pareto optimal solution that precedes x^*, i. e., $\hat{x}$ has the largest profit among all solutions that are higher ranked than x^* (see Fig. 1). We aim at bounding the probability that $\hat{x}$ is an ordinary class-0 Pareto optimal solution and falls into the interval $(t-\varepsilon, t)$.

We classify solutions to be ordinary or extraordinary as follows. Let x be a class-0 Pareto optimal solution and let y be the succeeding Pareto optimal solution, which must exist as the lowest ranked Pareto optimal solution is not class-0 Pareto optimal. We say that x is extraordinary if for all indices $i \in [n]$ with $x_i \neq 0$ and $y_i = 0$, $z_i \neq 0$ holds for all Pareto optimal solutions z that preceed x. In other words, for those indices i that make x class-0 Pareto optimal, y is the highest ranked Pareto optimal solution that is independent of p_i (see Fig. 2). For every index $i \in [n]$ there can be at most one extraordinary class-0 Pareto optimal solution. In the following we will restrict ourselves to solutions $\hat{x}$ that are ordinary. Define

$$\Lambda(t) \;=\; \begin{cases} t - p^{\mathrm{T}}\hat{x} & \text{if } x^* \text{ and } \hat{x} \text{ exist and } \hat{x} \text{ is ordinary class-0 Pareto optimal} \\ \perp & \text{otherwise.} \end{cases}$$

Let $\mathcal{P}^0$ denote the set of ordinary class-0 Pareto optimal solutions. Whenever $\Lambda(t) < \varepsilon$, then there exists a solution $x \in \mathcal{P}^0$ with $p^{\mathrm{T}}x \in (t-\varepsilon, t)$, namely $\hat{x}$. The reverse is not true because it might be the case that $\hat{x} \notin \mathcal{P}^0$ but that there exists another solution $x \in \mathcal{P}^0$ with $p^{\mathrm{T}}x \in (t-\varepsilon, t)$. If, however, ε is smaller than the minimum distance between two Pareto optimal solutions, then the existence of a solution $x \in \mathcal{P}^0$ with $p^{\mathrm{T}}x \in (t-\varepsilon, t)$ implies $\hat{x} = x$ and hence $\Lambda(t) < \varepsilon$. Let $\mathcal{A}(t, \varepsilon)$ denote the event that there is at most one Pareto optimal solution with

a profit in the interval $(t-\varepsilon, t)$. Then

$$\begin{aligned}\mathbf{Pr}\left[\Lambda(t) < \varepsilon\right] &\geq \mathbf{Pr}\left[(\Lambda(t) < \varepsilon) \wedge \mathcal{A}(t,\varepsilon)\right] \\ &= \mathbf{Pr}\left[(\exists x \in \mathcal{P}^0 : p^{\mathrm{T}} x \in (t-\varepsilon, t)) \wedge \mathcal{A}(t,\varepsilon)\right] \\ &\geq \mathbf{Pr}\left[\exists x \in \mathcal{P}^0 : p^{\mathrm{T}} x \in (t-\varepsilon, t)\right] - \mathbf{Pr}\left[\neg\mathcal{A}(t,\varepsilon)\right],\end{aligned}$$

and therefore

$$\lim_{\varepsilon\to 0} \frac{\mathbf{Pr}\left[\Lambda(t) < \varepsilon\right]}{\varepsilon} \geq \lim_{\varepsilon\to 0} \frac{\mathbf{Pr}\left[\exists x \in \mathcal{P}^0 : p^{\mathrm{T}} x \in (t-\varepsilon, t)\right]}{\varepsilon} - \lim_{\varepsilon\to 0} \frac{\mathbf{Pr}\left[\neg\mathcal{A}(t,\varepsilon)\right]}{\varepsilon}.$$

In the full version we show that for every $t \neq 0$ the probability of that two solutions lie in the interval $(t-\varepsilon, t)$ decreases like ε^2 for $\varepsilon \to 0$. Hence, for every $t \neq 0$, $\lim_{\varepsilon\to 0} \frac{\mathbf{Pr}[\neg\mathcal{A}(t,\varepsilon)]}{\varepsilon} = 0$.

Since the expected number of ordinary class-0 Pareto optimal solutions can be written as

$$\int_{-\infty}^{\infty} \lim_{\varepsilon\to 0} \frac{\mathbf{Pr}\left[\exists x \in \mathcal{P}^0 : p^{\mathrm{T}} x \in (t-\varepsilon, t)\right]}{\varepsilon}\, dt \leq \int_{-\infty}^{\infty} \lim_{\varepsilon\to 0} \frac{\mathbf{Pr}\left[\Lambda(t) < \varepsilon\right]}{\varepsilon}\, dt \ ,$$

it remains to analyze the term $\mathbf{Pr}\left[\Lambda(t) < \varepsilon\right]$. In order to analyze this probability we define a set of auxiliary random variables such that $\Lambda(t)$ is guaranteed to always take a value also taken by one of the auxiliary random variables. Then we analyze the auxiliary random variables and use a union bound to conclude the desired bound for $\Lambda(t)$.

Define $\mathcal{D}' = \mathcal{D}\backslash\{0\}$ and $\mathcal{S}^{x_i=v} = \{x \in \mathcal{S} \mid x_i = v\}$ for all $i \in [n]$ and $v \in \mathcal{D}$. Let $x^{*(i)}$ denote the highest ranked solution from $\mathcal{S}^{x_i=0}$ with profit at least t. For each $i \in [n]$ and $v \in \mathcal{D}'$ we define the set $\mathcal{L}^{(i,v)}$ as follows. If $x^{*(i)}$ does not exist or $x^{*(i)}$ is the highest ranked solution in $\mathcal{S}^{x_i=0}$ then we define $\mathcal{L}^{(i,v)} = \emptyset$. Otherwise $\mathcal{L}^{(i,v)}$ consists of all solutions from $\mathcal{S}^{x_i=v}$ that have a higher rank than $x^{*(i)}$. Let $\hat{x}^{(i,v)}$ denote the lowest ranked Pareto optimal solution from the set $\mathcal{L}^{(i,v)}$, i. e., $\hat{x}^{(i,v)}$ has the largest profit among all solutions in $\mathcal{L}^{(i,v)}$. Finally we define for each $i \in [n]$ and $v \in \mathcal{D}'$ the auxiliary random variable

$$\Lambda_i^v(t) \ = \ \begin{cases} t - p^{\mathrm{T}}\hat{x}^{(i,v)} & \text{if } \hat{x}^{(i,v)} \text{ exists,} \\ \perp & \text{otherwise.} \end{cases}$$

If $\Lambda_i^v(t) \in (0, \varepsilon)$ (which excludes $\Lambda_i^v(t) = \perp$) then the following three events must co-occur:

1. $\mathcal{E}_1$: There exists an $x \in \mathcal{S}^{x_i=0}$ with $p^{\mathrm{T}} x \geq t$.
2. $\mathcal{E}_2$: There exists an $x \in \mathcal{S}^{x_i=0}$ with $p^{\mathrm{T}} x < t$.
3. $\mathcal{E}_3$: $\hat{x}^{(i,v)}$ exists and its profit falls into the interval $(t-\varepsilon, t)$.

The events $\mathcal{E}_1$ and $\mathcal{E}_2$ only depend on the profits p_j, $j \neq i$. The existence and identity of $\hat{x}^{(i,v)}$ is completely determined by those profits as well. Hence, if we fix all profits except for p_i, then $\hat{x}^{(i,v)}$ is fixed and its profit is $c + vp_i$ for some constant c that depends on the profits already fixed. Observe

that the random variable $c + vp_i$ has density at most $\phi_i/|v|$. Hence we obtain $\mathbf{Pr}\left[p^{\mathrm{T}}\hat{x}^{(i,v)} \in (t-\varepsilon, t) \middle| \hat{x}^{(i,v)} \text{ exists }\right] \leq \varepsilon\frac{\phi_i}{|v|}$. Define

$$P^+ = \sum_{j:p_j>0} p_j \quad \text{and} \quad P^- = \sum_{j:p_j<0} p_j \ .$$

Moreover let d^+ and d^- denote the largest and the smallest element in $\mathcal{D}$. For $t \geq 0$, the event $\mathcal{E}_1$ implies $t \leq d^+P^+ + d^-P^-$, and hence $\mathbf{Pr}[\mathcal{E}_1] \leq \mathbf{Pr}[d^+P^+ + d^-P^- \geq t]$. For $t \leq 0$, the event $\mathcal{E}_2$ implies $t > d^+P^- + d^-P^+$ and hence $\mathbf{Pr}[\mathcal{E}_2] \leq \mathbf{Pr}[d^+P^- + d^-P^+ \leq t]$. By combining these results we get

$$\mathbf{Pr}\left[\Lambda_i^v(t) \in (0,\varepsilon)\right] \leq \begin{cases} \mathbf{Pr}\left[d^+P^+ + d^-P^- \geq t\right]\varepsilon\frac{\phi_i}{|v|}, \text{ for } t \geq 0, \text{ and} \\ \mathbf{Pr}\left[d^+P^- + d^-P^+ \leq t\right]\varepsilon\frac{\phi_i}{|v|}, \text{ for } t \leq 0. \end{cases}$$

Next we argue that $\Lambda(t) < \varepsilon$ implies $\Lambda_i^v(t) \in (0,\varepsilon)$ for at least one pair $(i,v) \in [n] \times \mathcal{D}'$. So assume that $\Lambda(t) < \varepsilon$. By definition, x^* and $\hat{x}$ exist and $\hat{x}$ is an ordinary class-0 Pareto optimal solution. Since $\hat{x}$ is class-0 Pareto optimal and x^* is the succeeding Pareto optimal solution, there exists an index $i \in [n]$ such that

(a) $x_i^* = 0$ and $\hat{x}_i = v \neq 0$ for some $v \in \mathcal{D}'$, and
(b) x^* is not the highest ranked solution in $\mathcal{S}^{x_i=0}$.

The second condition is a consequence of the assumption, that $\hat{x}$ is not extraordinary, i. e., there exists a Pareto optimal solution z with $z_i = 0$ that has higher rank than $\hat{x}$. Recall that $x^{*(i)}$ is defined to be the highest ranked solution in $\mathcal{S}^{x_i=0}$ with $p^{\mathrm{T}}x \geq t$. As $x^* \in \mathcal{S}^{x_i=0}$, $x^* = x^{*(i)}$. Moreover, $\mathcal{L}^{(i,v)}$ consists of all solutions from $\mathcal{S}^{x_i=v}$ that have a higher rank than x^*. Thus, $\hat{x} \in \mathcal{L}^{(i,v)}$. By construction, $\hat{x}$ has the largest profit among the solutions in $\mathcal{L}^{(i,v)}$ and, therefore $\hat{x}^{(i,v)} = \hat{x}$ and $\Lambda_i^v(t) = \Lambda(t)$. Applying a union bound yields, for all $t \geq 0$,

$$\begin{aligned} \mathbf{Pr}\left[\Lambda(t) < \varepsilon\right] &\leq \sum_{i=1}^{n}\sum_{v\in\mathcal{D}'} \mathbf{Pr}\left[\Lambda_i^v(t) < \varepsilon\right] \\ &\leq \sum_{i=1}^{n}\sum_{v\in\mathcal{D}'} \mathbf{Pr}\left[d^+P^+ + d^-P^- \geq t\right]\varepsilon\frac{\phi_i}{|v|} \\ &\leq \mathbf{Pr}\left[d^+P^+ + d^-P^- \geq t\right]\varepsilon\sum_{i=1}^{n}\sum_{v\in\mathcal{D}'}\frac{\phi_i}{|v|} \ . \end{aligned}$$

For $t \leq 0$ we get analogously

$$\mathbf{Pr}\left[\Lambda(t) < \varepsilon\right] \leq \mathbf{Pr}\left[d^+P^- + d^-P^+ \leq t\right]\varepsilon\sum_{i=1}^{n}\sum_{v\in\mathcal{D}'}\frac{\phi_i}{|v|} \ .$$

Now we can bound the expected number of class-0 Pareto optimal solutions, taking into account that at most n of them can be extraordinary.

$$\begin{aligned}
\mathbf{E}[q_0] &\le n + \int_{-\infty}^{\infty} \lim_{\varepsilon\to 0} \frac{\mathbf{Pr}\left[\Lambda(t) \le \varepsilon\right]}{\varepsilon}\, dt \\
&\le n + \int_{0}^{\infty} \lim_{\varepsilon\to 0} \frac{\mathbf{Pr}\left[d^+P^+ + d^-P^- \ge t\right] \varepsilon \sum_{i=1}^{n} \sum_{v} \frac{\phi_i}{|v|}}{\varepsilon} dt \\
&\quad + \int_{-\infty}^{0} \lim_{\varepsilon\to 0} \frac{\mathbf{Pr}\left[d^+P^- + d^-P^+ \le t\right] \varepsilon \sum_{i=1}^{n} \sum_{v} \frac{\phi_i}{|v|}}{\varepsilon} dt \\
&\le n + \left(\sum_{v} \frac{1}{|v|}\right)\left(\sum_{i=1}^{n} \phi_i\right)\left(\int_{0}^{\infty} \mathbf{Pr}\left[d^+P^+ + d^-P^- \ge t\right] dt \right. \\
&\qquad\qquad \left. + \int_{0}^{\infty} \mathbf{Pr}\left[-d^+P^- - d^-P^+ \ge t\right]\, dt\right)
\end{aligned}$$

As $0 \in \mathcal{D}$, it holds $d^+ \ge 0$ and $d^- \le 0$. Hence we have $d^+P^+ + d^-P^- \ge 0$, $-d^+P^- - d^-P^+ \ge 0$, and

$$\begin{aligned}
&\int_{0}^{\infty} \mathbf{Pr}\left[d^+P^+ + d^-P^- \ge t\right]\, dt + \int_{0}^{\infty} \mathbf{Pr}\left[-d^+P^- - d^-P^+ \ge t\right]\, dt \\
&\quad = \mathbf{E}\left[d^+P^+ + d^-P^-\right] + \mathbf{E}\left[-d^+P^- - d^-P^+\right] \\
&\quad = (d^+ - d^-)\,\mathbf{E}\left[P^+ - P^-\right] = (d^+ - d^-)\,\mathbf{E}\left[\sum_{i=1}^{n} |p_i|\right] = D\sum_{i=1}^{n} \mu_i \ . \qquad \square
\end{aligned}$$

4 Lower Bounds on the Expected Number of Pareto optimal Solutions

In this section we present a lower bound of $\Omega(n^2 k \log(1+k))$ on the number of Pareto optimal solutions for $\mathcal{D} = \{0, \ldots, k\}$, generalizing a bound for the binary domain presented in [2]. In Theorem 8 we prove the stronger bound $\Omega(n^2k^2)$ under slightly stronger assumptions. The weaker bound provides a vector of weights $w_1, \ldots, w_n$, such that the bound holds for a linear weight function $w^{\mathrm{T}}x$. For the stronger bound we can only prove that there is some weight function $w\colon \mathcal{S} \to \mathbb{R}$ for which the bound holds but this function might not be linear. In combinatorial optimization, however, many problems have linear objective functions. The proofs of the theorems in this section will be contained in the full version of this paper.

Theorem 7. *Let $\mathcal{D} = \{0, \ldots, k\}$. Suppose profits are drawn independently at random according to a continuous probability distribution with non-increasing density function $f : \mathbb{R}_{\ge 0} \to \mathbb{R}_{\ge 0}$. Let q denote the number of Pareto optimal solutions over $\mathcal{S} = \mathcal{D}^n$. Then there is a vector of weights $w_1, \ldots, w_n \in \mathbb{R}_{>0}$ for which*

$$\mathbf{E}[q] \ \ge \ \frac{H_k}{4} k(n^2 - n) + kn + 1 \ ,$$

where H_k is the k-th harmonic number. If the profits are drawn according to the uniform distribution over some interval $[0, c]$ with $c > 0$ then the above inequality holds with equality.

Similarly, a lower bound of $\Omega(n^2 k \log k)$ can be obtained for the case that f is the density of a Gaussian random variable with mean 0. Since all weights w_i are larger than 0, a solution with a negative profit cannot be contained in a Pareto optimal solution. Hence, we can ignore those items. Restricted to the interval $[0, \infty)$ the density of a Gaussian random variable with mean 0 is non-increasing and, hence, we can apply Theorem 7.

Now we consider general weight functions and show a lower bound of $\Omega(n^2 k^2)$ on the expected number of Pareto optimal solutions for $\mathcal{D} = \{0, \ldots, k\}$ and $\mathcal{S} = \mathcal{D}^n$. We assume that k is a function of n with $(5(c+1)+1) \log n \leq k \leq n^c$ for some constant c. We use the probabilistic method to show that, for each sufficiently large $n \in \mathbb{N}$, a ranking exists for which the expected number of Pareto optimal solutions is lower bounded by $n^2 k^2 / \kappa$ for some constant κ depending only on c, that is, we create a ranking at random (but independently of the profits) and show that the expected number of Pareto optimal solutions (where the expectation is taken over both the random ranking and the random profits) satisfies the desired lower bound. This implies that, for each sufficiently large $n \in \mathbb{N}$, there must exist a deterministic ranking on $\{0, \ldots, k\}^n$ for which the expected number of Pareto optimal solutions (where the expectation is now taken only over the random profits) is at least $n^2 k^2 / \kappa$.

Theorem 8. *Let $(5(c+1)+1) \log n \leq k \leq n^c$ for some $c \geq 2$ and assume that n is a multiple of $c+2$. There exists a constant κ depending only on c and a ranking on $\{0, \ldots, k\}^n$ such that the expected number of Pareto optimal solutions is lower bounded by $n^2 k^2 / \kappa$ if each profit p_i is chosen independently, uniformly at random from the interval $[-1, 1]$.*

5 Smoothed Complexity of Integer Programming

In [13], we analyze the smoothed complexity of integer programming. We consider integer programs in which an objective function is to be maximized over a feasible region that is defined as the intersection of a fixed ground set $\mathcal{S} \subseteq \mathcal{D}^n$ with a halfspace $\mathcal{B}$ that is induced by a linear constraint $w_1 x_1 + \cdots + w_n x_n \leq t$, where the w_i are independent random variables which can be represented by densities that are bounded by ϕ. We show that an integer optimization problem in this form has polynomial smoothed complexity if and only if there exists a pseudo-polynomial algorithm (w. r. t. the w_i) for solving it.

The main technical contribution in [13] is the analysis of the random variables *loser gap* and *feasibility gap*. The feasibility gap Γ is defined as the slack of the optimal solution from the threshold t. To be more precise, let x^* denote the optimal solution, that is, x^* denotes the solution from $\mathcal{S} \cap \mathcal{B}$ that maximizes the objective function. Then the feasibility gap can be defined as $\Gamma = t - w^{\mathrm{T}} x^*$. A solution $x \in \mathcal{S}$ is called a loser if it has a higher objective value than x^* but is

infeasible due to the linear constraint, that is, $w^{\mathrm{T}}x > t$. We denote the set of all losers by $\mathcal{L}$. Furthermore, we define the minimal loser $\overline{x} \in \mathcal{L}$ to be the solution from $\mathcal{L}$ with minimal weight, that is, $\overline{x} = \operatorname{argmin}\{w^{\mathrm{T}}x \mid x \in \mathcal{L}\}$. The loser gap $\mathit{\Lambda}$ denotes the slack of the minimal loser from the threshold t, i. e., $\mathit{\Lambda} = w^{\mathrm{T}}\overline{x} - t$.

If both the loser and the feasibility gap are not too small, then rounding all weights w_i with sufficient accuracy does not change the optimal solution. Rounding the weights can only affect the optimal solution if either x^* becomes infeasible or a loser x becomes feasible. The former event can only occur if the feasibility gap is small; the latter event can only occur if the loser gap is small. In a rather technical and lengthy analysis we show the following lemma on the probability that the loser or the feasibility gap is small.

Lemma 9. *(Separating Lemma [13]) Let $\mathcal{S} \subseteq \mathcal{D}^n$ with $0^n \notin \mathcal{S}$ be chosen arbitrarily, let $\mu = \max_{i\in[n]} \mathbf{E}[|w_i|]$, $d = |\mathcal{D}|$, and $d_{\max} = \max\{|a| \mid a \in \mathcal{D}\}$. Then, for all $\varepsilon \in [0, (32\mu n^5 d^7 d_{\max}\phi^2)^{-1}]$,*

$$\mathbf{Pr}[\mathit{\Gamma} \leq \varepsilon] \leq 2(\varepsilon \cdot 32\mu n^5 d^7 d_{\max}\phi^2)^{1/3} \text{ and } \mathbf{Pr}[\mathit{\Lambda} \leq \varepsilon] \leq 2(\varepsilon \cdot 32\mu n^5 d^7 d_{\max}\phi^2)^{1/3}.$$

In the full version of this paper we present a much simpler proof for the following improved version of the previous lemma.

Theorem 10. *Let $\mathcal{S} \subseteq \mathcal{D}^n$ with $0^n \notin \mathcal{S}$ be chosen arbitrarily, and let $D = \max\{a - b \mid a, b \in \mathcal{D}\} \leq 2d_{\max}$. There exists a constant κ such that, for all $\varepsilon \geq 0$,*

$$\mathbf{Pr}[\mathit{\Gamma} \leq \varepsilon] \leq \varepsilon\kappa\phi^2\mu n^3 Dd \log^2 d \text{ and } \mathbf{Pr}[\mathit{\Lambda} \leq \varepsilon] \leq \varepsilon\kappa\phi^2\mu n^3 Dd \log^2 d \ .$$

In [13] we show that Lemma 9 can also be used to analyze integer optimization problems with more than one linear constraint. We consider integer optimization problems in which an objective function is to be maximized over a feasible region that is defined as the intersection of a fixed ground set $\mathcal{S} \subseteq \mathcal{D}^n$ with halfspaces $\mathcal{B}_1, \ldots, \mathcal{B}_m$ that are induced by m linear constraints of the form $w_{i,1}x_1 + \cdots + w_{i,n}x_n \leq t_i$, where the $w_{i,j}$ are independent random variables which can be represented by densities that are bounded by ϕ.

The feasibility gap $\mathit{\Gamma}$ for multiple constraints is defined to be the minimal slack of the optimal solution x^* from one of the thresholds, i. e., $\mathit{\Gamma} = \min_{i\in[m]}(t_i - (w_{i,1}x_1 + \cdots + w_{i,n}x_n))$. The loser gap $\mathit{\Lambda}$ for multiple constraints is defined as $\mathit{\Lambda} = \min_{x\in\mathcal{L}} \max_{i\in[m]}(w_{i,1}x_1 + \cdots + w_{i,n}x_n - t_i)$. In [13] we show how Lemma 9 gives rise to bounds on the sizes of loser and feasibility gap for multiple constraints. Based on this observation we show that an integer optimization problem with multiple constraints has polynomial smoothed complexity if and only if there exists a pseudo-polynomial algorithm (w. r. t. the $w_{i,j}$) for solving it. By applying the same arguments, our bounds in Theorem 10 yield the following corollary.

Corollary 11. *Let $\mathcal{S} \subseteq \mathcal{D}^n$ with $0^n \notin \mathcal{S}$ be chosen arbitrarily, let $D = \max\{a - b \mid a, b \in \mathcal{D}\} \leq 2d_{\max}$, and let the set of feasible solutions be given as $\mathcal{S} \cap \mathcal{B}_1 \cap \ldots \cap \mathcal{B}_m$. There exists a constant κ such that, for all $\varepsilon \geq 0$,*

$$\mathbf{Pr}[\mathit{\Gamma} \leq \varepsilon] \leq \varepsilon\kappa\phi^2\mu mn^3 Dd \log^2 d \text{ and } \mathbf{Pr}[\mathit{\Lambda} \leq \varepsilon] \leq \varepsilon\kappa\phi^2\mu mn^3 Dd \log^2 d \ .$$

The main improvement upon our previous analysis is that the bounds in Corollary 11 depend only linearly on ε instead of $\varepsilon^{1/3}$. Due to this improvement we can prove Theorem 4 in the same way as its binary version in [3], which is not possible with the bounds derived in [13].

References

1. René Beier. *Probabilistic Analysis of Discrete Optimization Problems*. PhD thesis, Universität des Saarlandes, 2004.
2. René Beier and Berthold Vöcking. Random knapsack in expected polynomial time. *Journal of Computer and System Sciences*, 69(3):306–329, 2004.
3. René Beier and Berthold Vöcking. Typical properties of winners and losers in discrete optimization. *SIAM Journal on Computing*, 35(4):855–881, 2006.
4. H.W. Corley and I.D Moon. Shortest paths in networks with vector weights. *Journal of Optimization Theory and Application*, 46(1):79–86, 1985.
5. Matthias Ehrgott. Integer solutions of multicriteria network flow problems. *Investigacao Operacional*, 19:61–73, 1999.
6. Matthias Ehrgott and Xavier Gandibleux. *Multiple Criteria Optimization*, volume 491 of *Lecture Notes in Economics and Mathematical Systems*, chapter Multiobjective Combinatorial Optimization. Springer-Verlag, 2000.
7. H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack Problems*. Springer, Berlin, Germany, 2004.
8. Kathrin Klamroth and Margaret M. Wiecek. Dynamic programming approaches to the multiple criteria knapsack problem. *Naval Research Logistics*, 47(1):57–76, 2000.
9. Matthias Müller-Hannemann and Karsten Weihe. Pareto shortest paths is often feasible in practice. In *Proceedings of the 5th International Workshop on Algorithm Engineering (WAE)*, pages 185–198, 2001.
10. Adli Mustafa and Mark Goh. Finding integer efficient solutions for bicriteria and tricriteria network flow problems using dinas. *Computers & OR*, 25(2):139–157, 1998.
11. George L. Nemhauser and Zev Ullmann. Discrete dynamic programming and capital allocation. *Management Science*, 15:494–505, 1969.
12. Christos H. Papadimitriou and Mihalis 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)*, pages 86–92. IEEE Computer Society, 2000.
13. Heiko Röglin and Berthold Vöcking. Smoothed analysis of integer programming. In *Proceedings of the 11th International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, volume 3509 of *Lecture Notes in Computer Science*, pages 276–290. Springer, 2005.
14. Anders J. V. Skriver and Kim Allan Andersen. A label correcting approach for solving bicriterion shortest-path problems. *Computers & OR*, 27(6):507–524, 2000.
15. 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.
16. Sergei Vassilvitskii and Mihalis Yannakakis. Efficiently computing succinct trade-off curves. *Theoretical Computer Science*, 348(2-3):334–356, 2005.