

Smoothed Analysis of Pareto Curves in Multiobjective Optimization

Heiko Röglin*

Abstract

In a multiobjective optimization problem a solution is called Pareto-optimal if no criterion can be improved without deteriorating at least one of the other criteria. Computing the set of all Pareto-optimal solutions is a common task in multiobjective optimization to filter out unreasonable trade-offs.

For most problems the number of Pareto-optimal solutions increases only moderately with the input size in applications. However, for virtually every multiobjective optimization problem there exist worst-case instances with an exponential number of Pareto-optimal solutions. In order to explain this discrepancy, we analyze a large class of multiobjective optimization problems in the model of smoothed analysis and prove a polynomial bound on the expected number of Pareto-optimal solutions.

We also present algorithms for computing the set of Pareto-optimal solutions for different optimization problems and discuss related results on the smoothed complexity of optimization problems.

1 Algorithms for Computing Pareto Curves

Suppose you would like to book a flight to your favorite conference. Your decision is then probably guided by different factors, like the price, the number of stops, and the arrival time. Usually you won't find a flight that is optimal in every respect and you have to choose the best trade-off. This is characteristic for many decisions faced every day by people, companies, and other economic entities.

The notion of “best trade-off” is hard to formalize and often there is no consensus on how different criteria should be traded off against each other. However, there is little disagreement that in a reasonable outcome no criterion can be improved without deteriorating at least one of the other criteria. Outcomes with this property are called *Pareto-optimal* and they play a crucial role in multi-criteria decision making as they help to filter out unreasonable solutions. In this section we discuss algorithms for computing the set of Pareto-optimal solutions for different problems.

1.1 Knapsack Problem

The *knapsack problem* is a well-known NP-hard optimization problem. An instance of this problem consists of a set of items, each with a profit and a weight, and a capacity. The goal is to find a subset of the items that maximizes the total profit among all subsets whose total weight does not exceed the capacity. Let $p = (p_1, \dots, p_n)^T \in \mathbb{R}_{\geq 0}^n$ and $w = (w_1, \dots, w_n)^T \in \mathbb{R}_{\geq 0}^n$ denote the profits

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

and weights, respectively, and let $W \in \mathbb{R}_{\geq 0}$ denote the capacity. Formally the knapsack problem can be stated as follows:

$$\begin{aligned} & \text{maximize} && p^T x = p_1 x_1 + \cdots + p_n x_n \\ & \text{subject to} && w^T x = w_1 x_1 + \cdots + w_n x_n \leq W, \\ & && \text{and } x = (x_1, \dots, x_n)^T \in \{0, 1\}^n. \end{aligned}$$

The knapsack problem has attracted a great deal of attention, both in theory and in practice. Theoreticians are interested in the knapsack problem because of its simple structure; it can be expressed as a binary program with one linear objective function and one linear constraint. On the other hand, knapsack-like problems often occur in applications, and practitioners have developed numerous heuristics for solving them. These heuristics work very well on random and real-world instances and they usually find optimal solutions quickly even for very large instances.

In the following, we assume that an arbitrary instance \mathcal{I} of the knapsack problem is given. We use the term *solution* to refer to a vector $x \in \{0, 1\}^n$, and we say that a solution is *feasible* if $w^T x \leq W$. We say that a solution x *contains item* i if $x_i = 1$ and that it *does not contain item* i if $x_i = 0$.

One naive approach for solving the knapsack problem is to enumerate all feasible solutions and to select the one with maximum profit. This approach is not efficient as there are typically exponentially many feasible solutions. In order to decrease the number of solutions that have to be considered, we view the knapsack problem as a bicriteria optimization problem and restrict the enumeration to only the Pareto-optimal solutions.

Definition 1. A solution y *dominates* a solution x if $p^T y \geq p^T x$ and $w^T y \leq w^T x$, with at least one of these inequalities being strict. A solution x is called *Pareto-optimal* if it is not dominated by any other solution. The *Pareto set* or *Pareto curve* is the set of all Pareto-optimal solutions.

Once the Pareto set is known, the given instance of the knapsack problem can be solved optimally in time linear in the size of this set due to the following observation.

Lemma 2. *There always exists an optimal solution that is also Pareto-optimal.*

Proof. Take an arbitrary optimal solution x and assume that it is not Pareto-optimal. There cannot be a solution y with $p^T y > p^T x$ and $w^T y \leq w^T x$ because then y would be a better solution than x . Hence, if x is not Pareto-optimal then it is dominated by a solution y with $p^T y = p^T x$ and $w^T y < w^T x$. Then either y is Pareto-optimal or we repeat the argument to find a solution z with $p^T z = p^T y$ and $w^T z < w^T y$. This construction terminates after a finite number of iterations with an optimal solution that is also Pareto-optimal. \square

We denote the Pareto set by $\mathcal{P} \subseteq \{0, 1\}^n$. It can happen that there are two or more Pareto-optimal solutions with the same profit and the same weight. Then \mathcal{P} is assumed to contain only one of these solutions, which can be chosen arbitrarily. Due to the previous lemma the solution

$$x^* = \arg \max_{x \in \mathcal{P}} \{p^T x \mid w^T x \leq W\},$$

is an optimal solution of the given instance of the knapsack problem.

In the following we present an algorithm invented by Nemhauser and Ullmann (1969) to compute the Pareto set of a given instance of the knapsack problem. We will refer to this algorithm, which is based on dynamic programming, as the *Nemhauser-Ullmann algorithm*. For

each $i \in \{0, 1, \dots, n\}$ it computes the Pareto set \mathcal{P}_i of the restricted instance \mathcal{I}_i that contains only the first i items of the given instance \mathcal{I} . Then $\mathcal{P}_n = \mathcal{P}$ is the set we are looking for. Let

$$\mathcal{S}_i = \{x \in \{0, 1\}^n \mid x_{i+1} = \dots = x_n = 0\}$$

denote the set of solutions that do not contain the items $i + 1, \dots, n$. Formally, solutions of the instance \mathcal{I}_i are binary vectors of length i . We will, however, represent them as binary vectors of length n from \mathcal{S}_i . For a solution $x \in \{0, 1\}^n$ and an item $i \in \{1, \dots, n\}$ we denote by x^{+i} the solution that is obtained by adding item i to solution x :

$$x_j^{+i} = \begin{cases} x_j & \text{if } j \neq i, \\ 1 & \text{if } j = i. \end{cases}$$

Furthermore, for a set $\mathcal{S} \subseteq \{0, 1\}^n$ of solutions let

$$\mathcal{S}^{+i} = \{y \in \{0, 1\}^n \mid \exists x \in \mathcal{S} : y = x^{+i}\}.$$

If for some $i \in \{1, \dots, n\}$, the set \mathcal{P}_{i-1} is known then the set \mathcal{P}_i can be computed with the help of the following lemma. For the lemma we assume a consistent tie-breaking between solutions that have the same profit and the same weight. In particular, if $p^T x = p^T y$ and $w^T x = w^T y$ for two solutions x and y and the tie-breaking favors x over y then it should also favor x^{+i} over y^{+i} for any i .

Lemma 3. *For every $i \in \{1, \dots, n\}$, the set \mathcal{P}_i is a subset of $\mathcal{P}_{i-1} \cup \mathcal{P}_{i-1}^{+i}$.*

Proof. Let $x \in \mathcal{P}_i$. Based on the value of x_i we distinguish two cases.

First we consider the case $x_i = 0$. We claim that in this case $x \in \mathcal{P}_{i-1}$. Assume for contradiction that $x \notin \mathcal{P}_{i-1}$. Then there exists a solution $y \in \mathcal{P}_{i-1} \subseteq \mathcal{S}_{i-1} \subseteq \mathcal{S}_i$ that dominates x . Since $y \in \mathcal{S}_i$, solution x cannot be Pareto-optimal among the solutions in \mathcal{S}_i . Hence, $x \notin \mathcal{P}_i$, contradicting the choice of x .

Now we consider the case $x_i = 1$. We claim that in this case $x \in \mathcal{P}_{i-1}^{+i}$. Since $x \in \mathcal{S}_i$ and $x_i = 1$, there exists a solution $y \in \mathcal{S}_{i-1}$ such that $x = y^{+i}$. We need to show that $y \in \mathcal{P}_{i-1}$. Assume for contradiction that there exists a solution $z \in \mathcal{P}_{i-1}$ that dominates y . Then $p^T z \geq p^T y$ and $w^T z \leq w^T y$ and one of these inequalities is strict. By adding item i to the solutions y and z , we obtain $p^T z^{+i} \geq p^T y^{+i}$ and $w^T z^{+i} \leq w^T y^{+i}$, with one of these inequalities being strict. Hence, the solution z^{+i} dominates the solution $x = y^{+i}$. Since $z^{+i} \in \mathcal{S}_i$, this implies $x \notin \mathcal{P}_i$, contradicting the choice of x . \square

Due to the previous lemma, the Pareto set \mathcal{P}_i can be computed easily if the Pareto set \mathcal{P}_{i-1} is already known. For this one only needs to compute the set $\mathcal{P}_{i-1} \cup \mathcal{P}_{i-1}^{+i}$ and remove solutions from this set that are dominated by other solutions from this set. Using additionally that $\mathcal{P}_0 = \mathcal{S}_0 = \{0^n\}$, we obtain the following algorithm to solve the knapsack problem optimally (see Figure 1 for an illustration).

Algorithm 1 Nemhauser-Ullmann algorithm

- 1: $\mathcal{P}_0 := \{0^n\}$;
 - 2: **for** $i = 1, \dots, n$ **do**
 - 3: $\mathcal{Q}_i := \mathcal{P}_{i-1} \cup \mathcal{P}_{i-1}^{+i}$;
 - 4: $\mathcal{P}_i := \{x \in \mathcal{Q}_i \mid \nexists y \in \mathcal{Q}_i : y \text{ dominates } x\}$;
 - 5: **return** $x^* := \arg \max_{x \in \mathcal{P}_n} \{p^T x \mid w^T x \leq W\}$;
-

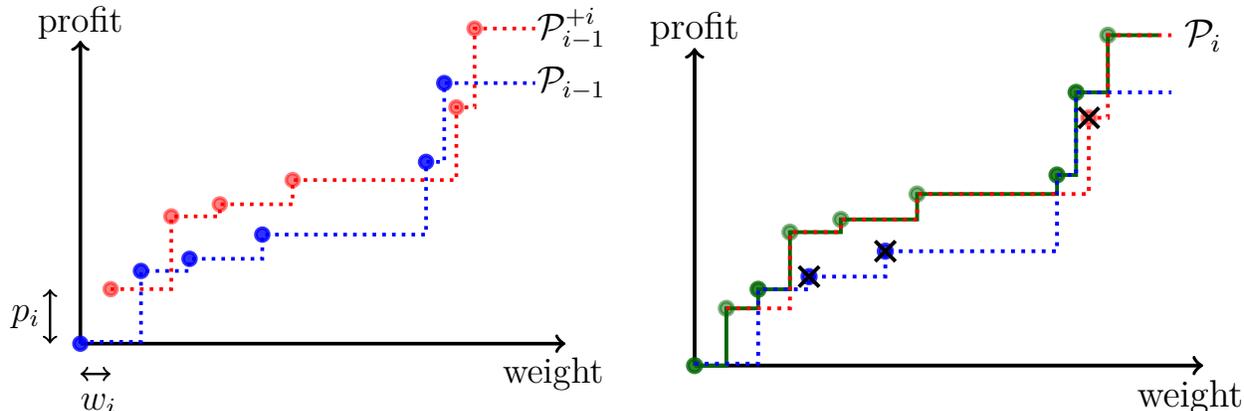


Figure 1: Illustration of one iteration of the for-loop of the Nemhauser-Ullmann algorithm: The set \mathcal{P}_{i-1}^{+i} is a copy of the set \mathcal{P}_{i-1} that is shifted by (w_i, p_i) . The set \mathcal{P}_i is obtained by removing dominated solutions.

In Line 4 a tie-breaking is assumed so that \mathcal{P}_i does never contain two solutions with identical profits and weights.

Observe that all steps of the Nemhauser-Ullmann algorithm except for Line 5 are independent of the capacity W . In order to speed up the algorithm, one could remove solutions with weights larger than W already from \mathcal{Q}_i in Line 3.

We analyze the running time of the Nemhauser-Ullmann algorithm using the model of a unit-cost RAM. In this model, arithmetic operations like adding and comparing two numbers can be performed in constant time regardless of their bit-lengths. We use this model for the sake of simplicity and in order to keep the focus on the important details of the running time analysis.

Theorem 4. *The Nemhauser-Ullmann algorithm solves the knapsack problem optimally. There exists an implementation with running time $\Theta(\sum_{i=0}^{n-1} |\mathcal{P}_i|)$.*

Proof. The correctness of the algorithm follows immediately from the previous discussion. In order to achieve the claimed running time, we do not compute the sets \mathcal{P}_i explicitly, but only the values of the solutions in these sets. That is, instead of \mathcal{P}_i only the set $\text{val}(\mathcal{P}_i) := \{(p^T x, w^T x) \mid x \in \mathcal{P}_i\}$ is computed. Analogously to the computation of \mathcal{P}_i , one can compute $\text{val}(\mathcal{P}_i)$ easily if $\text{val}(\mathcal{P}_{i-1})$ is known. If we store for each element of $\text{val}(\mathcal{P}_i)$ a pointer to the element of $\text{val}(\mathcal{P}_{i-1})$ from which it originates, then in Step 5 the solution x^* can be efficiently reconstructed from the sets $\text{val}(\mathcal{P}_i)$ and these pointers.

The running times of Steps 1 and 5 are $O(1)$ and $O(n + |\mathcal{P}|)$, respectively, where the term n accounts for the running time of reconstructing the solution x^* once its value $(p^T x^*, w^T x^*)$ is determined. In every iteration i of the for-loop, the running time of Step 3 to compute $\text{val}(\mathcal{Q}_i)$ is $\Theta(|\mathcal{P}_{i-1}|)$ because on a unit-cost RAM the set $\text{val}(\mathcal{P}_{i-1}^{+i})$ can be computed in time $\Theta(|\mathcal{P}_{i-1}|)$ from the set $\text{val}(\mathcal{P}_{i-1})$.

In a straightforward implementation, the running time of Step 4 is $\Theta(|\mathcal{Q}_i|^2) = \Theta(|\mathcal{P}_{i-1}|^2)$ because we need to compare every pair of values from $\text{val}(\mathcal{Q}_i)$ and each comparison takes time $O(1)$. Step 4 can be implemented more efficiently. For this, we store the values in each set $\text{val}(\mathcal{P}_i)$ sorted in non-decreasing order of weights. If $\text{val}(\mathcal{P}_{i-1})$ is sorted in this way, then, without any additional computational effort, the computation of the set $\text{val}(\mathcal{Q}_i)$ in Step 3 can be implemented such that $\text{val}(\mathcal{Q}_i)$ is also sorted: The sorted set $\text{val}(\mathcal{P}_{i-1}^{+i})$ can be computed in time $\Theta(|\mathcal{P}_{i-1}|)$. Then, in order to compute the set $\text{val}(\mathcal{Q}_i)$, only the two sorted sets $\text{val}(\mathcal{P}_{i-1})$ and $\text{val}(\mathcal{P}_{i-1}^{+i})$ need

to be merged in time $\Theta(|\mathcal{P}_{i-1}|)$. If the set $\text{val}(\mathcal{Q}_i)$ is sorted, Step 4 can be implemented to run in time $\Theta(|\mathcal{Q}_i|)$ as a sweep algorithm going once through $\text{val}(\mathcal{Q}_i)$ in non-decreasing order of weights (see Exercise 1). \square

The above theorem ensures that the Nemhauser-Ullmann algorithm solves the knapsack problem efficiently if all Pareto sets \mathcal{P}_i have polynomial size.¹ As the knapsack problem is NP-hard, it is not surprising that there are instances with exponentially many Pareto-optimal solutions. If one sets $p_i = w_i = 2^i$ for each item $i \in \{1, \dots, n\}$ then even every solution from $\{0, 1\}^n$ is Pareto-optimal.

1.2 Shortest Path Problem

Shortest path problems often come naturally with multiple objectives. Think for example of automotive navigation systems in which one can usually choose between the shortest, cheapest, and fastest route. Let us consider the *bicriteria single-source shortest path problem*. An instance of this problem is described by a directed graph $G = (V, E)$ with costs $c : E \rightarrow \mathbb{R}_{>0}$, weights $w : E \rightarrow \mathbb{R}_{>0}$, and a source vertex $s \in V$. The goal is to compute for each $v \in V$ the set \mathcal{P}^v of Pareto-optimal s - v -paths according to the following definition.

Definition 5. For an s - v -path P we denote by $w(P) = \sum_{e \in P} w(e)$ its weight and by $c(P) = \sum_{e \in P} c(e)$ its costs. An s - v -path P_1 *dominates* an s - v -path P_2 if $w(P_1) \leq w(P_2)$ and $c(P_1) \leq c(P_2)$, with at least one of these inequalities being strict. An s - v -path P is called *Pareto-optimal* if it is not dominated by any other s - v -path.

A well-known algorithm for the single-criterion single-source shortest path problem (with only weights but no costs on the edges) is the Bellman-Ford algorithm. It stores a distance label for each vertex which is initially set to infinity for each vertex except the source s for which it is set to zero. Then it performs a sequence of relax operations on the edges as shown in the following pseudocode.

Algorithm 2 Bellman-Ford algorithm

```

1: dist( $s$ ) = 0;
2: for  $v \in V \setminus \{s\}$  do dist( $v$ ) =  $\infty$ ;
3: for  $i = 1, \dots, |V| - 1$  do
4:   for each  $(u, v) \in E$  do
5:     RELAX( $u, v$ );
6: procedure RELAX( $u, v$ )
7:   if dist( $v$ ) > dist( $u$ ) +  $w(u, v)$  then
8:     dist( $v$ ) := dist( $u$ ) +  $w(u, v$ );

```

It can be shown that after termination the distance label $\text{dist}(v)$ of each vertex v equals the length of the shortest s - v -path in G . By standard methods one can adapt the algorithm so that for each vertex the actual shortest s - v -path is computed. One can also easily adapt this algorithm to the bicriteria shortest path problem if one replaces each distance label $\text{dist}(v)$ by a list L_v of s - v -paths. Initially L_s contains only the trivial path of length 0 from s to s and all other lists L_v

¹Let us remark that the sizes of the Pareto sets are in general not monotone and there are instances where $|\mathcal{P}_{i+1}| < |\mathcal{P}_i|$ for some i . Hence, it does not suffice if only the Pareto set $\mathcal{P} = \mathcal{P}_n$ has polynomial size. However, we are not aware of any class of instances where $|\mathcal{P}_n|$ is polynomially bounded while $|\mathcal{P}_i|$ is superpolynomial for some i .

are empty. In every relax operation for an edge (u, v) a new set $L_u^{+(u,v)}$ is obtained from L_u by appending the edge (u, v) to each path from L_u . Then the paths from $L_u^{+(u,v)}$ are added to L_v . Finally L_v is cleaned up by removing all paths from L_v that are dominated by other paths from this list. This is shown in the following pseudocode.

Algorithm 3 Bicriteria Bellman-Ford algorithm

```

1:  $L_s = \{\text{path of length 0 from } s \text{ to } s\}$ ;
2: for  $v \in V \setminus \{s\}$  do  $L_v = \emptyset$ ;
3: for  $i = 1, \dots, |V| - 1$  do
4:   for each  $(u, v) \in E$  do
5:     RELAX( $u, v$ );
6: procedure RELAX( $u, v$ )
7:   Obtain  $L_u^{+(u,v)}$  from  $L_u$  by appending the edge  $(u, v)$  to each path from  $L_u$ .
8:    $L_v := L_v \cup L_u^{+(u,v)}$ ;
9:   Remove dominated paths from  $L_v$ .
```

Analogously to the Nemhauser-Ullmann algorithm, the running time of the Bicriteria Bellman-Ford algorithm depends crucially on the sizes of the lists L_v that appear throughout the algorithm. We have to look at the algorithm in slightly more detail to give an upper bound on its running time. The algorithm performs $M := (|V| - 1) \cdot |E|$ relax operations, which we denote by R_1, \dots, R_M . For a relax operation R_k that relaxes the edge (u, v) , we define $u(R_k) = u$ and $v(R_k) = v$. Let $k \in [M]$ and consider the first k relax operations. These define for every vertex $v \in V$ a set S_v^k of s - v -paths that can be discovered by the first k relax operations. To be more precise, S_v^k contains exactly those s - v -paths that appear as a subsequence in $(u(R_1), v(R_1)), \dots, (u(R_k), v(R_k))$. In the single-criterion version, after k relax operations the distance label $\text{dist}(v)$ contains the length of the shortest path in S_v^k . In the bicriteria version, the list L_v contains after k relax operations all paths from S_v^k that are Pareto-optimal within this set (i.e. that are not dominated by other paths from this set). We will denote the list L_v after k relax operations by L_v^k in the following.

Theorem 6. *After termination of the Bicriteria Bellman-Ford algorithm the list L_v equals for every vertex $v \in V$ the set of Pareto-optimal s - v -paths. There exists an implementation with running time $\Theta\left(\sum_{k=1}^M \left(|L_{u(R_k)}^{k-1}| + |L_{v(R_k)}^{k-1}|\right)\right)$.*

Proof. The correctness of the algorithm follows by an inductive argument along the lines of the analysis of the single-criterion version (see Exercise 4). The analysis of the running time is similar to the proof of Theorem 4. The dominating factor is the time to remove dominated paths from L_v in Line 9 of the pseudocode. A naive implementation has running time $\Theta(|L_{u(R_k)}^{k-1}| \cdot |L_{v(R_k)}^{k-1}|)$ for the k th relax operation, while the running time $\Theta(|L_{u(R_k)}^{k-1}| + |L_{v(R_k)}^{k-1}|)$ can be achieved by sweeping through the lists when they are sorted in non-decreasing order of weight. \square

While in applications where the bicriteria shortest path problem occurs, it has been observed that the number of Pareto-optimal solutions is usually not very large, one can easily construct instances of the bicriteria shortest path problem in which the number of Pareto-optimal paths is exponential in the size of the graph (see Exercise 3).

The reader might wonder why we adapted the Bellman-Ford algorithm and not Dijkstra's algorithm to the bicriteria single-source shortest path problem. Indeed there is a generalization of Dijkstra's algorithm to the bicriteria shortest path problem due to Hansen (1979), which also

performs a sequence of operations similar to the relax operations of the Bellman-Ford algorithm. However, in contrast to the Bellman-Ford algorithm the sequence of relax operations is not fixed beforehand but it depends on the actual costs and weights of the edges. For this reason, it is not clear how to analyze the expected running time and in particular the analysis that we present in Section 2 does not apply to the generalization of Dijkstra’s algorithm.

1.3 Multiple Objectives and Other Optimization Problems

For the sake of simplicity, we have discussed only problems with two objectives above. However, one can easily adapt the definition of Pareto-optimal solutions and both presented algorithms to more than two objectives. Consider the multidimensional knapsack problem, a version of the knapsack problem in which every item still has a single profit but instead of a single weight, it has a weight vector from $\mathbb{R}_{\geq 0}^{d-1}$ for some $d \geq 2$, and also the capacity is a vector from $\mathbb{R}_{\geq 0}^{d-1}$. This problem gives rise to a multiobjective optimization problem with d objectives: maximize the profit $p^T x$ and minimize for each $i \in [d - 1]$ the i th weight $(w^{(i)})^T x$. Similarly it is often natural to consider multiobjective shortest path problems with more than two objectives.

In order to compute the Pareto set of an instance of the multidimensional knapsack problem or the multiobjective shortest path problem, no modification to the pseudocode of the Nemhauser-Ullmann algorithm (Algorithm 1) and the Bicriteria Bellman-Ford algorithm (Algorithm 3) are necessary. However, the implementation and the analysis of the running time have to be adapted. The crucial difference is that the removal of dominated solutions from \mathcal{Q}_i and L_v cannot be implemented in time linear in the sizes of these sets anymore because the sweeping approach, which assumes that the solutions are sorted with respect to one of the objectives, fails for more than two objectives. If one uses the naive implementation, which pairwise compares the solutions, then the running times of the algorithms become $\Theta(\sum_{i=0}^{n-1} |\mathcal{P}_i|^2)$ and $O(\sum_{i=1}^M |L_{u(R_i)}^{i-1}| \cdot |L_{v(R_i)}^{i-1}|)$, respectively.

Asymptotically one can do better by using known algorithms for the maximum vector problem to filter out the dominated solutions. In this problem a set of m vectors in \mathbb{R}^k is given and one wants to compute the set of Pareto-optimal vectors among them. The fastest known algorithm for this problem is due to Kung et al. (1975). It relies on divide and conquer and its running time is $O(m \log^{k-2} m)$. For d objectives this yields running times of $\Theta\left(\sum_{i=0}^{n-1} |\mathcal{P}_i| \log^{d-2}(|\mathcal{P}_i|)\right)$ and

$$O\left(\sum_{i=1}^M (|L_{u(R_i)}^{i-1}| + |L_{v(R_i)}^{i-1}|) \cdot \log^{d-2}(|L_{u(R_i)}^{i-1}| + |L_{v(R_i)}^{i-1}|)\right)$$

for the Nemhauser-Ullmann algorithm and the Bellman-Ford algorithm, respectively.

The Nemhauser-Ullmann algorithm and the Bicriteria Bellman-Ford algorithm are only two examples of many algorithms in the literature for computing Pareto sets of various multiobjective optimization problems. Similar algorithms exist, for example, for the multiobjective network flow problem. As a rule of thumb, algorithms that solve the single-criterion version of an optimization problem by dynamic programming can usually be adapted to compute the Pareto set of the multiobjective version.

On the other hand, there are also problems for which it is unknown if there exist algorithms that compute the Pareto set in time polynomial in its size and the sizes of the Pareto sets of appropriate subproblems. The multiobjective spanning tree problem is one such example, where the best known way to compute the Pareto set is essentially to first compute the set of all spanning trees and then to remove the dominated ones. An even stronger requirement is that of an efficient output-sensitive algorithm, which computes the Pareto set in time polynomial in its size and the input size. Bökler et al. (2017) show that such an algorithm exists for the multiobjective version

of the minimum-cut problem and that no such algorithm exists for the bicriteria shortest path problem, unless $P=NP$. For many other multiobjective problems, including the knapsack problem and the multiobjective spanning tree problem, it is an open question whether efficient output-sensitive algorithms exist.

1.4 Approximate Pareto Curves

For virtually every multiobjective optimization problem the number of Pareto-optimal solutions can be exponential in the worst case. One way of coping with this problem is to relax the requirement of finding the complete Pareto set. A solution x is ε -dominated by a solution y if y is worse than x by at most a factor of $1 + \varepsilon$ in each objective (i.e., $w(y)/w(x) \leq 1 + \varepsilon$ for each criterion w that is to be minimized and $p(x)/p(y) \leq 1 + \varepsilon$ for each criterion p that is to be maximized). We say that \mathcal{P}_ε is an ε -approximation of a Pareto set \mathcal{P} if for any solution in \mathcal{P} , there is a solution in \mathcal{P}_ε that ε -dominates it.

In his pioneering work, Hansen (1980) presents an approximation scheme for computing ε -approximate Pareto sets of the bicriteria shortest path problem. Papadimitriou and Yannakakis (2000) show that for any instance of a multiobjective optimization problem, there is an ε -approximation of the Pareto set whose size is polynomial in the input size and $1/\varepsilon$ but exponential in the number of objectives. Furthermore, they define the *gap version* of a multiobjective optimization problem with d objectives as follows: given an instance of the problem and a vector $b \in \mathbb{R}^d$, either return a solution whose objective vector dominates b or report (correctly) that there does not exist any solution whose objective vector is better than b by more than a $(1 + \varepsilon)$ factor in all objectives. They show that an FPTAS for approximating the Pareto set of a multiobjective optimization problem exists if and only if the gap version of the problem can be solved in polynomial time. In particular, this implies that if the exact single-criterion version of a problem (i.e., the question “Is there a solution with weight exactly x ?”) can be solved in pseudopolynomial time, then its multiobjective version admits an FPTAS for approximating the Pareto set. This is the case, for example, for the spanning tree problem, the all-pair shortest path problem, and the perfect matching problem.

Vassilvitskii and Yannakakis (2005) show how to compute ε -approximate Pareto sets whose size is at most three times as large as the smallest such set for bicriteria problems whose gap versions can be solved in polynomial time. Diakonikolas and Yannakakis (2007) improve this factor to two and show that this is the best possible that can be achieved in polynomial time, unless $P=NP$.

2 Number of Pareto-optimal solutions

Both for the knapsack problem and the bicriteria shortest path problem, the number of Pareto-optimal solutions increases only moderately with the input size in applications. This is in contrast to the exponential worst-case behavior (see Exercise 3). To explain this discrepancy, we will analyze the number of Pareto-optimal solutions in the framework of smoothed analysis. First we will focus on the knapsack problem but we will see afterwards that the proven bound also holds for a much larger class of problems including the bicriteria shortest path problem and many other natural bicriteria optimization problems. We will then also briefly discuss known results for problems with more than two objectives.

2.1 Knapsack Problem

Let us consider the knapsack problem. In a worst-case analysis the adversary is allowed to choose the profits p_1, \dots, p_n and the weights w_1, \dots, w_n exactly (he can also choose the capacity but the

number of Pareto-optimal solutions is independent of this). This makes him very powerful and makes it possible to choose an instance in which every solution is Pareto-optimal. In order to limit the power of the adversary to construct such artificial instances that do not resemble typical inputs, we add some randomness to his decisions.

Let $\phi \geq 1$ be a parameter. In the following analysis, we assume that the adversary can still determine the profits exactly while for each weight he can only choose an interval of length $1/\phi$ from which it is chosen uniformly at random independently of the other weights. This means that the adversary can specify each weight only with a precision of $1/\phi$. We normalize the weights and restrict the adversary to intervals that are subsets of $[0, 1]$. This normalization is necessary to ensure that the effect of the noise cannot be ruled out by scaling all weights in the input by some large number.

Observe that the parameter ϕ measures the strength of the adversary. If $\phi = 1$ then all weights are chosen uniformly at random from $[0, 1]$, which resembles an average-case analysis. On the other hand, in the limit for $\phi \rightarrow \infty$ the adversary can determine the weights (almost) exactly and the model approaches a classical worst-case analysis. Hence, it is not surprising that the bound that we will prove for the expected number of Pareto-optimal solutions grows with ϕ . However, we will see that it grows only polynomially with n and ϕ , which implies that already a small amount of random noise suffices to rule out the worst case and to obtain a benign instance in expectation.

Theorem 7. *Consider an instance \mathcal{I} of the knapsack problem with arbitrary profits $p_1, \dots, p_n \in \mathbb{R}_{\geq 0}$ in which every weight w_i is chosen uniformly at random from an arbitrary interval $A_i \subseteq [0, 1]$ of length $1/\phi$ independently of the other weights. Then the expected number of Pareto-optimal solutions in \mathcal{I} is bounded from above by $n^2\phi + 1$.*

The proof of Theorem 7, which we present in detail below, can be summarized as follows. Since all weights take values between 0 and 1, all solutions have weights between 0 and n . We divide the interval $[0, n]$ uniformly into a large number k of subintervals of length n/k each. For large enough k it is unlikely that there exist two Pareto-optimal solutions whose weights lie in the same subinterval because the weights are continuous random variables. Assuming that this does not happen, the number of Pareto-optimal solutions equals the number of subintervals that contain a Pareto-optimal solution. The most important and non-trivial step is then to bound, for each subinterval, the probability that it contains a Pareto-optimal solution. Once we have proven an upper bound for this, the theorem follows by summing up this upper bound over all subintervals due to linearity of expectation.

Before we prove the theorem, we state one simple but crucial property of the random variables that we consider.

Lemma 8. *Let X be a random variable that is chosen uniformly at random from some interval A of length $1/\phi$. Furthermore let I be an interval of length ε . Then $\Pr[X \in I] \leq \phi\varepsilon$.*

Proof. Since X is chosen uniformly at random from A , we obtain

$$\Pr[X \in I] = \frac{|A \cap I|}{|A|} \leq \frac{|I|}{|A|} \leq \frac{\varepsilon}{1/\phi} = \phi\varepsilon. \quad \square$$

Proof of Theorem 7. Every solution $x \in \{0, 1\}^n$ has a weight $w^T x$ in the interval $[0, n]$ because each weight w_i lies in $[0, 1]$. We partition the interval $(0, n]$ uniformly into $k \in \mathbb{N}$ intervals I_0^k, \dots, I_{k-1}^k for some large number k to be chosen later. Formally, let $I_i^k = (ni/k, n(i+1)/k]$. We say that the interval I_i^k is *non-empty* if there exists a Pareto-optimal solution $x \in \mathcal{P}$ with $w^T x \in I_i^k$.

We denote by X^k the number of non-empty intervals I_i^k plus one. The term $+1$ accounts for the solution 0^n , which is always Pareto-optimal and does not belong to any interval I_i^k . Nevertheless, the variable X^k can be much smaller than $|\mathcal{P}|$ because many Pareto-optimal solutions could lie in the same interval I_i^k . We will ensure that every interval I_i^k contains at most one Pareto-optimal solution with high probability by choosing k sufficiently large. Then, with high probability, $|\mathcal{P}| = X^k$.

In the following, we make this argument more formal. For $k \in \mathbb{N}$, let \mathcal{F}_k denote the event that there exist two different solutions $x, y \in \{0, 1\}^n$ with $|w^T x - w^T y| \leq n/k$. Since each interval I_i^k has length n/k , every interval I_i^k contains at most one Pareto-optimal solution if \mathcal{F}_k does not occur.

Lemma 9. *For every $k \in \mathbb{N}$, $\Pr[\mathcal{F}_k] \leq \frac{2^{2n+1}n\phi}{k}$.*

Proof. There are 2^n choices for x and y each. We prove the lemma by a union bound over all these choices. Let $x, y \in \{0, 1\}^n$ with $x \neq y$ be fixed. Then there exists an index i with $x_i \neq y_i$. Assume without loss of generality that $x_i = 0$ and $y_i = 1$. We use the principle of deferred decisions and assume that all weights w_j except for w_i are already fixed. Then $w^T x - w^T y = \alpha - w_i$ for some constant α that depends on x and y and the fixed profits w_j . It holds that

$$\begin{aligned} \Pr\left[|w^T x - w^T y| \leq \frac{n}{k}\right] &\leq \sup_{\alpha \in \mathbb{R}} \Pr_{w_i}\left[|\alpha - w_i| \leq \frac{n}{k}\right] \\ &= \sup_{\alpha \in \mathbb{R}} \Pr_{w_i}\left[w_i \in \left[\alpha - \frac{n}{k}, \alpha + \frac{n}{k}\right]\right] \leq \frac{2n\phi}{k}, \end{aligned}$$

where the last inequality follows from Lemma 8². Now a union bound over all choices for x and y concludes the proof. \square

The most non-trivial part in the analysis is the following lemma, which states for an arbitrary interval an upper bound for the probability that it contains a Pareto-optimal solution. We defer the proof of this lemma to the end of this section.

Lemma 10. *For every $t \geq 0$ and every $\varepsilon > 0$,*

$$\Pr[\exists x \in \mathcal{P} \mid w^T x \in (t, t + \varepsilon)] \leq n\phi\varepsilon.$$

The following lemma is the main building block in the proof of the theorem.

Lemma 11. *For every $k \in \mathbb{N}$, $\mathbf{E}[X^k] \leq n^2\phi + 1$.*

Proof. Let X_i^k denote a random variable that is 1 if the interval I_i^k is non-empty and 0 otherwise. Then

$$X^k = 1 + \sum_{i=0}^{k-1} X_i^k$$

and by linearity of expectation

$$\mathbf{E}[X^k] = \mathbf{E}\left[1 + \sum_{i=0}^{k-1} X_i^k\right] = 1 + \sum_{i=0}^{k-1} \mathbf{E}[X_i^k]. \quad (1)$$

²Formally, we condition on the outcome of the w_j with $j \neq i$. This outcome determines the value of α . Then we apply the law of total probability, but instead of integrating over all possible outcomes of the w_j with $j \neq i$, we derive an upper bound by looking only at the worst choice for α .

Since X_i^k is a 0-1-random variable, its expected value can be written as

$$\mathbf{E} \left[X_i^k \right] = \Pr[X_i^k = 1] = \Pr[\exists x \in \mathcal{P} \mid w^T x \in I_i^k]. \quad (2)$$

Using that each interval I_i^k has length n/k , Lemma 10 and (2) imply

$$\mathbf{E} \left[X_i^k \right] \leq \frac{n^2 \phi}{k}.$$

Together with (1) this implies

$$\mathbf{E} \left[X^k \right] = 1 + \sum_{i=0}^{k-1} \mathbf{E} \left[X_i^k \right] \leq 1 + k \cdot \frac{n^2 \phi}{k} = n^2 \phi + 1. \quad \square$$

With the help of Lemmas 9 and 11, we can finish the proof of the theorem as follows:

$$\begin{aligned} \mathbf{E} [|\mathcal{P}|] &= \sum_{i=1}^{2^n} (i \cdot \Pr[|\mathcal{P}| = i]) \\ &= \sum_{i=1}^{2^n} (i \cdot \Pr[|\mathcal{P}| = i \wedge \mathcal{F}_k] + i \cdot \Pr[|\mathcal{P}| = i \wedge \neg \mathcal{F}_k]) \\ &\stackrel{(1)}{=} \sum_{i=1}^{2^n} (i \cdot \Pr[\mathcal{F}_k] \cdot \Pr[|\mathcal{P}| = i \mid \mathcal{F}_k]) + \sum_{i=1}^{2^n} (i \cdot \Pr[X^k = i \wedge \neg \mathcal{F}_k]) \\ &\leq \Pr[\mathcal{F}_k] \cdot \sum_{i=1}^{2^n} (i \cdot \Pr[|\mathcal{P}| = i \mid \mathcal{F}_k]) + \sum_{i=1}^{2^n} (i \cdot \Pr[X^k = i]) \\ &\stackrel{(2)}{\leq} \frac{2^{2n+1} n \phi}{k} \cdot \sum_{i=1}^{2^n} (2^n \cdot \Pr[|\mathcal{P}| = i \mid \mathcal{F}_k]) + \mathbf{E} [X^k] \\ &\stackrel{(3)}{\leq} \frac{2^{3n+1} n \phi}{k} + n^2 \phi + 1. \end{aligned} \quad (3)$$

Let us comment on some of the steps in the previous calculation.

- The upper bound 2^n on the indices of the sums follows because $|\mathcal{P}|$ can never exceed the total number of solutions, which is 2^n .
- The rewriting of the first term in (1) follows from the definition of the conditional probability and the rewriting of the second term follows because $X^k = |\mathcal{P}|$ when the event $\neg \mathcal{F}_k$ occurs.
- (2) follows from Lemma 9 and the definition of the expected value.
- (3) follows from the identity $\sum_{i=1}^{2^n} \Pr[|\mathcal{P}| = i \mid \mathcal{F}_k] = 1$ and Lemma 11.

Since (3) holds for every $k \in \mathbb{N}$, it must be $\mathbf{E} [|\mathcal{P}|] \leq n^2 \phi + 1$. \square

It only remains to prove Lemma 10. An easy way to derive an upper bound for the probability that there exists a Pareto-optimal solution in the interval $(t, t + \varepsilon]$ is to apply a union bound over all solutions. Since there is an exponential number of solutions, this does not lead to a useful bound. The key improvement in the proof of Lemma 10 is to apply the union bound only over the n dimensions.

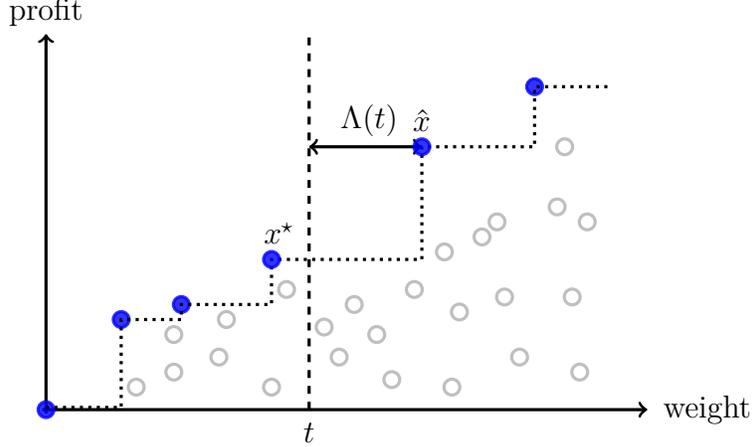


Figure 2: Definitions of the winner x^* , the loser \hat{x} , and the random variable $\Lambda(t)$.

Proof of Lemma 10. Fix $t \geq 0$ and $\varepsilon > 0$. First of all we define a random variable $\Lambda(t)$. In order to define $\Lambda(t)$, we define the *winner* x^* to be the most valuable solution satisfying $w^T x \leq t$, i.e.,

$$x^* = \arg \max \{p^T x \mid x \in \{0, 1\}^n \text{ and } w^T x \leq t\}.$$

For $t \geq 0$, such a solution x^* must always exist. We say that a solution x is a *loser* if it has a higher profit than x^* . By the choice of x^* , losers do not satisfy the constraint $w^T x \leq t$ (hence their name). We denote by \hat{x} the loser with the smallest weight (see Figure 2), i.e.,

$$\hat{x} = \arg \min \{w^T x \mid x \in \{0, 1\}^n \text{ and } p^T x > p^T x^*\}.$$

If there does not exist a solution x with $p^T x > p^T x^*$, then \hat{x} is undefined, which we denote by $\hat{x} = \perp$. Based on \hat{x} , we define the random variable $\Lambda(t)$ as

$$\Lambda(t) = \begin{cases} w^T \hat{x} - t & \text{if } \hat{x} \neq \perp, \\ \infty & \text{if } \hat{x} = \perp. \end{cases}$$

The random variable $\Lambda(t)$ satisfies the following equivalence:

$$\Lambda(t) \leq \varepsilon \iff \exists x \in \mathcal{P}: w^T x \in (t, t + \varepsilon]. \quad (4)$$

To see this, assume that there exists a Pareto-optimal solution whose weight lies in $(t, t + \varepsilon]$, and let y denote the Pareto-optimal solution with the smallest weight in $(t, t + \varepsilon]$. Then $y = \hat{x}$ and hence $\Lambda(t) = w^T \hat{x} - t \in (0, \varepsilon]$. Conversely, if $\Lambda(t) \leq \varepsilon$, then \hat{x} must be a Pareto-optimal solution whose weight lies in the interval $(t, t + \varepsilon]$. Together this yields Equivalence (4). Hence,

$$\Pr[\exists x \in \mathcal{P} \mid w^T x \in (t, t + \varepsilon]] = \Pr[\Lambda(t) \leq \varepsilon]. \quad (5)$$

It only remains to bound the probability that $\Lambda(t)$ does not exceed ε . In order to analyze this probability, we define a set of auxiliary random variables $\Lambda^1(t), \dots, \Lambda^n(t)$ such that $\Lambda(t)$ is guaranteed to always take a value also taken by at least one of the auxiliary random variables. Then we analyze the auxiliary random variables and use a union bound to conclude the desired bound for $\Lambda(t)$.

Let $i \in [n]$ be fixed. The random variable $\Lambda^i(t)$ is defined similarly to $\Lambda(t)$, but only solutions that do not contain item i are eligible as winners and only solutions that contain item i are eligible as losers. We make this more formal in the following. For $j \in \{0, 1\}$, we define

$$\mathcal{S}^{x_i=j} = \{x \in \{0, 1\}^n \mid x_i = j\},$$

and we define $x^{*,i}$ to be

$$x^{*,i} = \arg \max\{p^T x \mid x \in \mathcal{S}^{x_i=0} \text{ and } w^T x \leq t\}.$$

That is, $x^{*,i}$ is the winner among the solutions that do not contain item i . We restrict our attention to losers that contain item i and define

$$\hat{x}^i = \arg \min\{w^T x \mid x \in \mathcal{S}^{x_i=1} \text{ and } p^T x > p^T x^{*,i}\}.$$

If there does not exist a solution $x \in \mathcal{S}^{x_i=1}$ with $p^T x > p^T x^{*,i}$, then \hat{x}^i is undefined, i.e., $\hat{x}^i = \perp$. Based on \hat{x}^i , we define the random variable $\Lambda^i(t)$ as

$$\Lambda^i(t) = \begin{cases} w^T \hat{x}^i - t & \text{if } \hat{x}^i \neq \perp, \\ \infty & \text{if } \hat{x}^i = \perp. \end{cases}$$

Lemma 12. *For every choice of profits and weights, either $\Lambda(t) = \infty$ or there exists an index $i \in [n]$ such that $\Lambda(t) = \Lambda^i(t)$.*

Proof. Assume that $\Lambda(t) \neq \infty$. Then there exist a winner x^* and a loser \hat{x} . Since $x^* \neq \hat{x}$, there must be an index $i \in [n]$ with $x_i^* \neq \hat{x}_i$. Since all weights are non-negative and $w^T x^* < w^T \hat{x}$, there must even be an index $i \in [n]$ with $x_i^* = 0$ and $\hat{x}_i = 1$. We claim that for this index i , $\Lambda(t) = \Lambda^i(t)$. In order to see this, we first observe that $x^* = x^{*,i}$. This follows because x^* is the solution with the highest profit among all solutions with weight at most t . Since it belongs to $\mathcal{S}^{x_i=0}$, it is in particular the solution with the highest profit among all solutions that do not contain item i and have weight at most t . Since $x^* = x^{*,i}$, by similar arguments it follows that $\hat{x} = \hat{x}^i$. This directly implies that $\Lambda(t) = \Lambda^i(t)$. \square

Lemma 13. *For every $i \in [n]$ and every $\varepsilon \geq 0$,*

$$\Pr[\Lambda^i(t) \in (0, \varepsilon]] \leq \phi\varepsilon.$$

Proof. In order to prove the lemma, it suffices to exploit the randomness of the weight w_i . We apply the principle of deferred decisions and assume that all other weights are fixed arbitrarily. Then the weights of all solutions from $\mathcal{S}^{x_i=0}$ and hence also the solution $x^{*,i}$ are fixed because w_i does not influence the solutions in $\mathcal{S}^{x_i=0}$ and the profits p_1, \dots, p_n are fixed. If the solution $x^{*,i}$ is fixed, then also the set of losers $\mathcal{L} = \{x \in \mathcal{S}^{x_i=1} \mid p^T x > p^T x^{*,i}\}$ is fixed. Since, by definition, all solutions from \mathcal{L} contain item i the identity of the solution \hat{x}^i does not depend on w_i . (Of course, the weight $w^T \hat{x}^i$ depends on w_i . However, which solution will become \hat{x}^i is independent of w_i .) This implies that, given the fixed values of the weights w_j with $j \neq i$, we can rewrite the event $\Lambda^i(t) \in (0, \varepsilon]$ as $w^T \hat{x}^i - t \in (0, \varepsilon]$ for a fixed solution \hat{x}^i . For a constant $\alpha \in \mathbb{R}$ depending on the fixed values of the weights w_j with $j \neq i$, we can rewrite this event as $w_i \in (\alpha, \alpha + \varepsilon]$. By Lemma 8, the probability of this event is bounded from above by $\phi\varepsilon$. \square

Combining Lemmas 12 and 13 yields

$$\Pr[\Lambda(t) \leq \varepsilon] \leq \Pr[\exists i \in [n]: \Lambda^i(t) \in (0, \varepsilon)] \leq \sum_{i=1}^n \Pr[\Lambda^i(t) \in (0, \varepsilon)] \leq n\phi\varepsilon.$$

Together with (5) this proves the lemma. \square

Theorem 7 implies the following result on the running time of the Nemhauser-Ullmann algorithm.

Corollary 14. *Consider an instance \mathcal{I} of the knapsack problem with arbitrary profits $p_1, \dots, p_n \in \mathbb{R}_{\geq 0}$ in which every weight w_i is chosen uniformly at random from an arbitrary interval $A_i \subseteq [0, 1]$ of length $1/\phi$ independently of the other weights. Then the expected running time of the Nemhauser-Ullmann algorithm is $O(n^3\phi)$.*

Proof. It follows from Theorem 4 that the expected running time of the Nemhauser-Ullmann algorithm is bounded from above by

$$O\left(\mathbf{E}\left[\sum_{i=0}^{n-1} |\mathcal{P}_i|\right]\right),$$

where \mathcal{P}_i denotes the Pareto set of the restricted instance that consists only of the first i items. Using linearity of expectation and Theorem 7, we obtain that this term is bounded from above by

$$O\left(\sum_{i=0}^{n-1} \mathbf{E}[|\mathcal{P}_i|]\right) = O\left(\sum_{i=0}^{n-1} (i^2\phi + 1)\right) = O(n^3\phi). \quad \square$$

The decision to add randomness only to the weights is arbitrary. Of course if both the profits and the weights are chosen independently uniformly at random from intervals of length $1/\phi$ then the upper bound still applies. With minor modifications, the analysis can also be adapted to the case that only the profits are random while the weights are adversarial.

2.2 General Model

Theorem 7 can be extended in several ways. First of all, the noise model can be generalized to a much wider class of distributions. In fact the only property that we used about the random weights is Lemma 8, which says that the probability to fall into any interval of length ε is at most $\phi\varepsilon$. This is true for every random variable that is described by a probability density function that is bounded from above by ϕ . Hence instead of allowing the adversary to choose an interval of length $1/\phi$ for each weight w_i , we could also allow him to choose a density function $f_i : [0, 1] \rightarrow [0, \phi]$ according to which w_i is chosen independently of the other weights. This includes as a special case the uniform distribution in an interval of length $1/\phi$ but it also allows different types of random noise. Observe that we have restricted the density functions to $[0, 1]$ to normalize the weights.

In the following we will use the term *ϕ -perturbed random variable* to refer to a random variable described by a density $f : \mathbb{R} \rightarrow [0, \phi]$. If we replace all occurrences of Lemma 8 in the proof of Theorem 7 by the following lemma then Theorem 7 follows also for general ϕ -perturbed weights from $[0, 1]$.

Lemma 15. *Let X be a ϕ -perturbed random variable that is described by a density function $f : [0, 1] \rightarrow [0, \phi]$. For any interval I of length ε , $\Pr[X \in I] \leq \phi\varepsilon$.*

Proof. The lemma follows by the following simple calculation:

$$\Pr[X \in I] = \int_I f(x) dx \leq \int_I \phi dx = \phi \varepsilon. \quad \square$$

Next we state an even more general version of Theorem 7. The first generalization compared to Theorem 7 is that an arbitrary set $\mathcal{S} \subseteq \{0, 1\}^n$ of solutions is given. In the case of the knapsack problem, every vector from $\{0, 1\}^n$ is a solution, i.e., $\mathcal{S} = \{0, 1\}^n$. The second generalization is that the adversarial objective function p does not have to be linear. In fact, it can be an arbitrary function that maps every solution to some real value. The third generalization is that we extend the range of the ϕ -perturbed weights from $[0, 1]$ to $[-1, 1]$.

Theorem 16. *Let $\mathcal{S} \subseteq \{0, 1\}^n$ and $p : \mathcal{S} \rightarrow \mathbb{R}$ be arbitrary. Let w_1, \dots, w_n be arbitrary ϕ -perturbed numbers from the interval $[-1, 1]$. Then the expected number of solutions $x \in \mathcal{S}$ that are Pareto-optimal with respect to the objective functions $p(x)$ and $w^T x$ is $O(n^2 \phi)$. This upper bound holds regardless of whether the objective functions are to be maximized or minimized.*

We will not prove Theorem 16, but let us remark that its proof is very similar to the proof of Theorem 7. In fact we never used in the proof that $\mathcal{S} = \{0, 1\}^n$ and that p is linear. The fact that all weights w_i are positive was only used to argue that there must be an index i with $x_i^* = 0$ and $\hat{x}_i = 1$. For general w_i , it could also be the other way round. Handling this issue is the only modification of the proof that is not completely straightforward.

To illustrate the power of Theorem 16, let us discuss its implications on graph problems. For a given graph with m edges e_1, \dots, e_m , one can identify every vector $x \in \{0, 1\}^m$ with a subset of edges $E' = \{e_i \mid x_i = 1\}$. Then x is the so-called incidence vector of the edge set E' . If, for example, there is a source vertex s and a target vertex v given, one could choose the set \mathcal{S} of feasible solutions as the set of all incidence vectors of paths from s to v in the given graph. This way, Theorem 16 implies that the expected number of Pareto-optimal s - v -paths in the bicriteria shortest-path problem is $O(m^2 \phi)$. Similarly, one could choose \mathcal{S} as the set of incidence vectors of all spanning trees of a given graph. Then the result implies that in expectation there are only $O(m^2 \phi)$ Pareto-optimal spanning trees in the bicriteria spanning tree problem. In the traveling salesman problem (TSP) we are given an undirected graph with edge weights and the goal is to find a shortest tour (i.e., Hamiltonian cycle) that visits every vertex exactly once. As for the bicriteria shortest path problem, Theorem 16 implies that in expectation there are only $O(m^2 \phi)$ Pareto-optimal tours in the bicriteria version of the TSP.

For the Bicriteria Bellman-Ford algorithm we obtain the following corollary.

Corollary 17. *Consider an instance of the bicriteria shortest-path problem with arbitrary costs and non-negative ϕ -perturbed weights from the interval $[0, 1]$. Let n and m denote the number of vertices and edges, respectively. Then the expected running time of the Bicriteria Bellman-Ford algorithm is $O(nm^3 \phi)$.*

Proof. We can use Theorem 16 to bound the expected size of each list L_v^i that occurs throughout the algorithm by $O(m^2 \phi)$, where m denotes the number of edges in the graph. Using linearity of expectation and Theorem 6 yields that the expected running time is

$$\Theta \left(\sum_{i=1}^M \left(\mathbf{E} \left[|L_{u(R_i)}^{i-1}| \right] + \mathbf{E} \left[|L_{v(R_i)}^{i-1}| \right] \right) \right).$$

Using that the expected length of each list is $O(m^2 \phi)$ and $M = (n - 1) \cdot m$ implies the claimed bound. \square

Let us finally remark that Theorem 16 can also be adapted to the setting where the set \mathcal{S} of feasible solutions is an arbitrary subset of $\{0, \dots, k\}^n$ for some $k \in \mathbb{N}$. Then the expected number of Pareto-optimal solutions is $O(n^2 k^2 \phi)$. This is useful to model, for example, the bounded knapsack problem, in which a number of identical copies of every item is given.

2.3 Multiobjective Optimization Problems

Even though Theorem 16 is quite general, it still has one severe restriction: it only applies to optimization problems with two objective functions. The extension to optimization problems with more than two objectives is rather challenging and requires different methods. In this section, we summarize the main results.

In Theorem 16 one of the objective functions is assumed to be arbitrary while the other is linear with ϕ -perturbed coefficients. We consider now optimization problems with one arbitrary objective function and d linear objective functions with ϕ -perturbed coefficients. Röglin and Teng (2009) were the first to study this model. They proved an upper bound of $O((n^2 \phi)^{f(d)})$ for the expected number of Pareto-optimal solutions where f is a rapidly growing function (roughly $2^d d!$). This has been improved by Moitra and O'Donnell (2012) to $O(n^{2d} \phi^{d(d+1)/2})$. Brunsch and Röglin (2015) improved the upper bound further to $O(n^{2d} \phi^d)$ under the assumption that all density functions are unimodal, where a function $f : \mathbb{R} \rightarrow \mathbb{R}$ is called unimodal if there exists some $x \in \mathbb{R}$ such that f is monotonically increasing on $(-\infty, x]$ and monotonically decreasing on $[x, \infty)$.

The c th moment of a random variable X is the expected value $\mathbf{E}[X^c]$ if it exists. Brunsch and Röglin (2015) also prove upper bounds on the moments of the number of Pareto-optimal solutions. In particular they show that for any constant c the c th moment is bounded from above by $O((n^{2d} \phi^{d(d+1)/2})^c)$ and $O((n^{2d} \phi^d)^c)$ for general and unimodal densities, respectively. Upper bounds for the moments give rise to non-trivial tail bounds. Consider the case $d = 1$. Then the c th moment is bounded from above by $b_c (n^2 \phi)^c$ for some constant b_c depending on c . Applying Markov's inequality to the c th moment yields for every $\alpha \geq 1$

$$\Pr[|\mathcal{P}| \geq \alpha \cdot (n^2 \phi)] = \Pr[|\mathcal{P}|^c \geq \alpha^c (n^2 \phi)^c] = \Pr\left[|\mathcal{P}|^c \geq \frac{\alpha^c}{b_c} \cdot b_c (n^2 \phi)^c\right] \leq \frac{b_c}{\alpha^c},$$

while applying Markov's inequality directly to $|\mathcal{P}|$ yields only a bound of (roughly) $1/\alpha$. Upper bounds for the moments are also important for another reason: If the running time of an algorithm depends polynomially but not linearly on the number of Pareto-optimal solutions (like the running time of the Nemhauser-Ullmann algorithm for more than two objective functions), then Theorem 16 cannot be used to derive any bound on its expected running time. This is because a bound on $\mathbf{E}[|\mathcal{P}|]$ does not imply any bound on, for example, $\mathbf{E}[|\mathcal{P}|^2]$. Only with the result of Brunsch and Röglin about the moments of $|\mathcal{P}|$ a polynomial bound follows for the expected running time of these algorithms.

Improving earlier work of Brunsch et al. (2014), Brunsch (2014) shows lower bounds for the expected number of Pareto-optimal solutions of $\Omega(n^2 \phi)$ for $d = 1$ and $\Omega(n^{d-1.5} \phi^d)$ for $d \geq 2$. Hence the upper bound in Theorem 16 for the bicriteria case is asymptotically tight.

A ϕ -perturbed number is non-zero with probability 1. This implies that each of the d objective functions depends on all the variables. This limits the expressibility of the model because there are many examples of problems in which some objective function depends only on a certain subset of the variables. Brunsch and Röglin (2015) discuss this subtle issue in more detail and they also give concrete examples. To circumvent this problem, they introduce *zero-preserving perturbations*. In their model, the adversary can decide for each coefficient whether it should be a ϕ -perturbed number or is deterministically set to zero. For this model they prove upper bounds of $O(n^{O(d^3)} \phi^d)$

and $O((n\phi)^{O(d^3)})$ for unimodal and general ϕ -perturbed coefficients, respectively, for the expected number of Pareto-optimal solutions.

3 Smoothed Complexity of Binary Optimization Problems

The results on the expected number of Pareto-optimal solutions imply that the knapsack problem can be solved in expected polynomial time on instances with ϕ -perturbed weights or ϕ -perturbed profits (Corollary 14). A natural question is whether or not similar results also hold for other NP-hard optimization problems. Does, for example, the TSP admit an algorithm with expected polynomial running time if all distances are ϕ -perturbed? Instead of studying each problem separately, we will now present a general characterization due to Beier and Vöcking (2006) which combinatorial optimization problems can be solved efficiently on instances with ϕ -perturbed numbers.

While most smoothed analyses in the literature focus on the analysis of specific algorithms, this section instead considers *problems* in the sense of complexity theory. We will study *linear binary optimization problems*. In an instance of such a problem Π , a linear objective function $c^T x = c_1 x_1 + \dots + c_n x_n$ is to be minimized or maximized over an arbitrary set $\mathcal{S} \subseteq \{0, 1\}^n$ of feasible solutions. The problem Π could, for example, be the TSP and the coefficients c_i could be the edge lengths. (See also the discussion in Section 2.2 on how graph problems can be encoded as binary optimization problems.) One could also encode the knapsack problem as a linear binary optimization problem. Then \mathcal{S} contains all subsets of items whose total weight does not exceed the capacity.

We will study the *smoothed complexity* of linear binary optimization problems, by which we mean the complexity of instances in which the coefficients c_1, \dots, c_n are ϕ -perturbed numbers from the interval $[-1, 1]$. We will assume without loss of generality that the objective function $c^T x$ is to be minimized. Since ϕ -perturbed numbers have infinite encoding length with probability 1, we have to discuss the machine model that we will use in the following. One could change the input model and assume that the ϕ -perturbed coefficients are discretized by rounding them after a polynomial number, say n^2 , of bits. The effect of this rounding is so small that it does not influence our results. We will, however, not make this assumption explicit and use, for the sake of simplicity, the continuous random variables in our probabilistic analysis. When defining the input size we will not take the encoding length of the coefficients c_i into account. Instead we assume that the coefficients c_1, \dots, c_n contribute in total only n to the input length.

To state the main result, let us recall two definitions from computational complexity. We call a linear binary optimization problem *strongly NP-hard* if it is already NP-hard when restricted to instances with integer coefficients c_i in which the largest absolute value $C := \max_i |c_i|$ of any of the coefficients is bounded by a polynomial in the input length. The TSP is, for example, strongly NP-hard because it is already NP-hard when all edges have length either 1 or 2. The knapsack problem, on the other hand, is not strongly NP-hard because instances in which all profits are integers and polynomially bounded in the input size can be solved by dynamic programming in polynomial time.

A language L belongs to the complexity class ZPP (zero-error probabilistic polynomial time) if there exists a randomized algorithm A that decides for each input x in expected polynomial time whether or not x belongs to L . That is, A always produces the correct answer but its running time is a random variable whose expected value is bounded polynomially for every input x . Let us point out that the expectation is only with respect to the random decisions of the algorithm and not with respect to a randomly chosen input. It is yet unclear whether or not $P=ZPP$. In any case,

languages that belong to ZPP are generally considered to be easy to decide and NP-hard problems are believed to not lie in ZPP.

Theorem 18. *Let Π be a linear binary optimization problem that is strongly NP-hard. Then there does not exist an algorithm for Π whose expected running time is polynomially bounded in N and ϕ for instances with ϕ -perturbed coefficients from $[-1, 1]$, where N denotes the input length, unless $NP \subseteq ZPP$.*

The main idea of the proof of this theorem can be summarized as follows: An algorithm A for Π with expected running time polynomial in N and ϕ can be used to solve worst-case instances of Π with polynomially bounded numbers optimally in expected polynomial time. Given such a worst-case instance, one could add a small amount of random noise to all the numbers and then solve the resulting instance with A in expected time polynomial in N and ϕ . If this random noise is small enough ($\phi = \Theta(C)$) then it does not change the optimal solution. This way, we obtain an algorithm that solves worst-case instances with polynomially bounded numbers in expected polynomial time, implying that $NP \subseteq ZPP$.

The previous theorem shows that ϕ -perturbed instances of strongly NP-hard optimization problems are not easier to solve than worst-case instances. Hence, these problems stay hard also in the model of smoothed analysis. One consequence of this result is that there is no hope that the TSP can be solved efficiently when the edge lengths are randomly perturbed. This is in clear contrast to the knapsack problem, which is easy to solve on randomly perturbed inputs. We will now state a more general positive result. We say that a linear binary optimization problem Π can be solved in *pseudo-linear time* if there exists an algorithm whose running time on instances with integer coefficients is bounded from above by $p(N) \cdot C$, where p denotes a polynomial, N denotes the input length, and C denotes the largest absolute value of any of the coefficients.

Theorem 19. *A linear binary optimization problem Π that can be solved in pseudo-linear time in the worst case can be solved in expected polynomial time (with respect to the input length and ϕ) on instances with ϕ -perturbed numbers from $[-1, 1]$.*

Let A_p be an algorithm that solves integral instances of Π in pseudo-linear time. In the proof of Theorem 19, the algorithm A_p is used to construct an algorithm A that solves instances with ϕ -perturbed numbers in expected polynomial time. Algorithm A first rounds all ϕ -perturbed coefficients after some number b of bits after the binary point. Then it uses the algorithm A_p to solve the rounded instance. One can prove that for $b = \Theta(\log n)$ rounding all coefficients does not change the optimal solution with high probability. This is based on the observation that in instances with ϕ -perturbed numbers usually the best solution is significantly better than the second best solution and hence it stays optimal even after rounding all coefficients (see Exercise 8). For $b = \Theta(\log n)$ the running time of A_p to solve the rounded instance optimally is polynomial. This yields an algorithm that always runs in polynomial time and solves ϕ -perturbed instances of Π with high probability correctly. It is possible to adapt this approach to obtain an algorithm that always computes the optimal solution and whose expected running time is polynomial.

4 Conclusions

We have proven bounds on the expected number of Pareto-optimal solutions and we have studied the complexity of linear binary optimization problems in the framework of smoothed analysis. Our results are in many cases consistent with empirical observations. The knapsack problem is, for

example, easy to solve in applications and has few Pareto-optimal solutions while solving large-scale TSP instances optimally is computationally still expensive despite a lot of progress that has been made in the last decades and great speedups in the common solvers.

The models that we considered in this chapter are very general, in particular because the set \mathcal{S} of feasible solutions can be arbitrarily chosen, both in Section 2 and in Section 3. However, this generality is also a drawback of our results because the adversary is still rather powerful and can exactly determine the combinatorial structure of the problem. Often problems are easier in applications than in the worst case because the instances obey certain structural properties. Depending on the problem and application, input graphs might be planar or have small degree, distances might satisfy the triangle inequality etc. Such structural properties are not considered in our general model. Hence, often it is advisable to look in more detail into the instances that are really relevant in applications instead of only assuming that some coefficients are random.

An illustrative experimental study of the multiobjective shortest path problem is due to Müller-Hannemann and (2006). They consider a graph that is obtained from the daily train schedule of the German railway network and observe that the number of Pareto-optimal train connections in view of travel time, fare, and number of train changes is very small (for no pair of nodes there were more than 8 Pareto-optimal connections in the experiments). This is much smaller than suggested by Theorem 16. One possible explanation is that in this and many other applications, the objective functions are not independent but to some degree correlated, which might reduce the number of Pareto-optimal solutions. It would be interesting to find a formal model for correlated objective functions that explains the extremely small number of Pareto-optimal solutions observed in this setting.

Notes

The Bicriteria Bellman-Ford algorithm was described by Corley and Moon (1985). The analysis of its running time presented in this chapter can also be found in Beier (2004). Beier and Vöcking (2004) initiated the study of the number of Pareto-optimal solutions in the framework of smoothed analysis. The proof of Theorem 7 in this chapter follows an improved and simplified analysis due to Beier et al. (2007). This analysis also generalizes the original work of Beier and Vöcking to integer optimization problems. The bound stated in Beier et al. (2007) is $O(n^2 k^2 \log(k) \phi)$ if $\mathcal{S} \subseteq \{0, \dots, k\}^n$. It has been improved to $O(n^2 k^2 \phi)$ by Röglin and Rösner (2017).

The results in Section 3 can be found in Beier and Vöcking (2006). Theorems 18 and 19 do not give a complete characterization of the smoothed complexity of linear binary optimization problems because Theorem 19 does only apply to pseudo-linear and not to general pseudo-polynomial algorithms. Beier and Vöcking circumvent this problem by introducing a notion of polynomial smoothed complexity that is not based on expected running times (similar to polynomial average-case complexity). Later Röglin and Teng (2009) showed that all problems that can be solved in pseudo-polynomial time in the worst case can be solved in expected polynomial time on ϕ -perturbed instances, which completes the characterization.

Exercises

1. Implement the Nemhauser-Ullmann algorithm so that your implementation achieves a running time of $\Theta(\sum_{i=0}^{n-1} |\mathcal{P}_i|)$.
2. Find an instance of the knapsack problem with $|\mathcal{P}_{i+1}| < |\mathcal{P}_i|$ for some i .

3. Construct instances for the bicriteria shortest path problem with an exponential number of Pareto-optimal s - v -paths for some vertices s and v .
4. Prove that the Bicriteria Bellman-Ford algorithm is correct, i.e., that after termination the list L_v equals for every vertex $v \in V$ the set of Pareto-optimal s - v -paths
5. A famous algorithm for the single criterion all-pairs shortest path problem is the Floyd-Warshall algorithm. Adapt this algorithm to the bicriteria all-pairs shortest path problem (given a graph G with costs and weights, compute for each pair (u, v) of vertices the set of Pareto-optimal u - v -paths in G). State a bound on its running time in the same fashion as Theorem 6. What is the expected running time if the weights are ϕ -perturbed?
6. The concept of zero-preserving perturbations could also be applied to the bicriteria case with one adversarial objective function and one linear objective function with ϕ -perturbed coefficients. Show that, in contrast to the multiobjective case, for bicriteria optimization problems it does not increase the expressibility. For this, show that zero-preserving perturbations for the bicriteria case can be simulated by ϕ -perturbed coefficients if the set \mathcal{S} of feasible solutions is adapted appropriately. Why does this simulation not work for problems with three or more objectives?
7. Prove that the expected number of Pareto-optimal points among n points drawn independently and uniformly at random from the unit square is $O(\log n)$.
8. Given an instance \mathcal{I} of some linear binary optimization problem Π with a set $\mathcal{S} \subseteq \{0, 1\}^n$ of feasible solutions, the winner gap is defined as

$$\Delta = c^T x^{**} - c^T x^*,$$

where

$$x^* = \arg \min\{c^T x \mid x \in \mathcal{S}\} \quad \text{and} \quad x^{**} = \arg \min\{c^T x \mid x \in \mathcal{S} \setminus \{x^*\}\}$$

denote the best and second best solution of \mathcal{I} , respectively. Let \mathcal{I} be an instance of Π with ϕ -perturbed coefficients c_1, \dots, c_n . Prove that, for every $\varepsilon > 0$,

$$\Pr[\Delta \leq \varepsilon] \leq 2n\phi\varepsilon.$$

Hint: This statement follows by similar arguments as Lemma 10.

References

- Beier, René. 2004. *Probabilistic Analysis of Discrete Optimization Problems*. Ph.D. thesis, Universität des Saarlandes.
- Beier, René, and Vöcking, Berthold. 2004. Random Knapsack in Expected Polynomial Time. *Journal of Computer and System Sciences*, **69**(3), 306–329.
- Beier, René, and Vöcking, Berthold. 2006. Typical Properties of Winners and Losers in Discrete Optimization. *SIAM Journal on Computing*, **35**(4), 855–881.
- Beier, René, Röglin, Heiko, and Vöcking, Berthold. 2007. The Smoothed Number of Pareto Optimal Solutions in Bicriteria Integer Optimization. Pages 53–67 of: *Proceedings of the 12th International Conference on Integer Programming and Combinatorial Optimization (IPCO)*.

- Böckler, Fritz, Ehr Gott, Matthias, Morris, Christopher, and Mutzel, Petra. 2017. Output-sensitive complexity of multiobjective combinatorial optimization. *Journal of Multi-Criteria Decision Analysis*, **24**(1-2), 25–36.
- Brunsch, Tobias. 2014. *Smoothed Analysis of Selected Optimization Problems and Algorithms*. Ph.D. thesis, Universität Bonn.
- Brunsch, Tobias, and Röglin, Heiko. 2015. Improved Smoothed Analysis of Multiobjective Optimization. *Journal of the ACM*, **62**(1), 4:1–4:58.
- Brunsch, Tobias, Goyal, Navin, Rademacher, Luis, and Röglin, Heiko. 2014. Lower Bounds for the Average and Smoothed Number of Pareto-Optima. *Theory of Computing*, **10**, 237–256.
- Corley, H. William, and Moon, I. Douglas. 1985. Shortest paths in networks with vector weights. *Journal of Optimization Theory and Application*, **46**(1), 79–86.
- Diakonikolas, Ilias, and Yannakakis, Mihalis. 2007. Small Approximate Pareto Sets for Bi-objective Shortest Paths and Other Problems. Pages 74–88 of: *Proceedings of the 10th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*.
- Hansen, Pierre. 1979. Bicriterion Path Problems. Pages 109–127 of: *Multiple Criteria Decision Making: Theory and Applications*. Lecture Notes in Economics and Mathematical Systems, vol. 177.
- Hansen, Pierre. 1980. Bicriterion Path Problems. Pages 109–127 of: *Multiple Criteria Decision Making: Theory and Applications*. Lecture Notes in Economics and Mathematical Systems, vol. 177.
- Kung, H. T., Luccio, Fabrizio, and Preparata, Franco P. 1975. On Finding the Maxima of a Set of Vectors. *Journal of the ACM*, **22**(4), 469–476.
- Moitra, Ankur, and O’Donnell, Ryan. 2012. Pareto Optimal Solutions for Smoothed Analysts. *SIAM Journal on Computing*, **41**(5), 1266–1284.
- Müller-Hannemann, Matthias, and Weihe, Karsten. 2006. On the cardinality of the Pareto set in bicriteria shortest path problems. *Annals of Operations Research*, **147**(1), 269–286.
- Nemhauser, George L., and Ullmann, Zev. 1969. Discrete dynamic programming and capital allocation. *Management Science*, **15**(9), 494–505.
- Papadimitriou, Christos H., and Yannakakis, Mihalis. 2000. On the approximability of trade-offs and optimal access of Web sources. Pages 86–92 of: *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- Röglin, Heiko, and Rösner, Clemens. 2017. The Smoothed Number of Pareto-Optimal Solutions in Non-integer Bicriteria Optimization. Pages 543–555 of: *Proceedings of the 14th Annual Conference on Theory and Applications of Models of Computation (TAMC)*.
- Röglin, Heiko, and Teng, Shang-Hua. 2009. Smoothed Analysis of Multiobjective Optimization. Pages 681–690 of: *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- Vassilvitskii, Sergei, and Yannakakis, Mihalis. 2005. Efficiently computing succinct trade-off curves. *Theoretical Computer Science*, **348**(2–3), 334–356.