An Upper Bound on the Number of Extreme Shortest Paths in Arbitrary Dimensions

Florian Barth
Universität Stuttgart, Germany

Stefan Funke
Universität Stuttgart, Germany

Claudius Proissl
Universität Stuttgart, Germany

Abstract

Graphs with multiple edge costs arise naturally in the route planning domain when apart from travel time other criteria like fuel consumption or positive height difference are also objectives to be minimized. In such a scenario, this paper investigates the number of *extreme shortest paths* between a given source-target pair s, t. We show that for a fixed but arbitrary number of cost types $d \geq 1$ the number of extreme shortest paths is in $n^{O\left(\log^{d-1} n\right)}$ in graphs G with n nodes. This is a generalization of known upper bounds for $d = 2$ and $d = 3$.

2012 ACM Subject Classification Theory of computation → Shortest paths

Keywords and phrases Parametric Shortest Paths, Extreme Shortest Paths

Digital Object Identifier 10.4230/LIPIcs.ESA.2022.14

Funding This work was in part supported by the Deutsche Forschungsgemeinschaft (DFG) within the priority program 1894: Volunteered Geographic Information: Interpretation, Visualization and Social Computing.

Acknowledgements We thank Kshitij Gajjar and Jaikumar Radhakrishnan for their very helpful input and for identifying an error in a previous version of this paper.

1 Introduction

Given a finite set of points $P \subset \mathbb{R}^d$, what is the number of vertices (or extreme points) of the convex hull of P? This is a frequently asked question, for instance, in the area of Multi-objective Linear Programming [3] or in probability theory [4]. In this work, we examine this question for the case when P is the set of cost vectors of paths in a graph $G(V, E)$ with multiple edge costs.

Multi-objective path computation has obvious applications in the transportation domain, where the cost values (also called metrics) might correspond to quantities like travel time, fuel consumption, or positive height difference.

Figure 1 shows example cost vectors with two metrics. The red points are non-dominated (or Pareto-optimal) and, thus, may be the solution to constrained minimization problems. However, optimizing over all non-dominated cost vectors often turns out to be too expensive as, for instance, in the constrained shortest path problem [12, 14]. A typical strategy in such cases is to restrict the set of possible solutions to the non-dominated extreme points of the convex hull (circled in blue). The extreme points have the property that for any convex combination of the metrics there is at least one extreme point optimal for it. Therefore, extreme points are interesting on their own and one can hope that restricting the search to extreme points will lead to a good approximation of the optimal solution.

Let P_{st} be the set of cost vectors of all simple paths from a node s to a node t in a graph G. While it is easy to see that the number of non-dominated points in P_{st} can be exponential in the size of G, there is little known about the complexity of the extreme points in P_{st}. In this work we tackle the problem of counting the extreme points in P_{st}, which we call *extreme shortest paths*. We show that the number of extreme shortest paths in P_{st} is in $n^{O\left(\log^{d-1} n\right)}$, where n is the number of nodes in G and d is the fixed number of metrics. Thus, complexity-wise there is indeed a considerable gap between extreme points and non-dominated points in P_{st}.

There are multiple ways to model multi-objective shortest path problems. A well established one is the parametric shortest path problem, which is typically formulated for the two-metric case. We have decided to use the variant developed in the context of personalized route planning [10, 8] as it fits very well to practical application scenarios. Complexity-wise there is no difference between these two models.

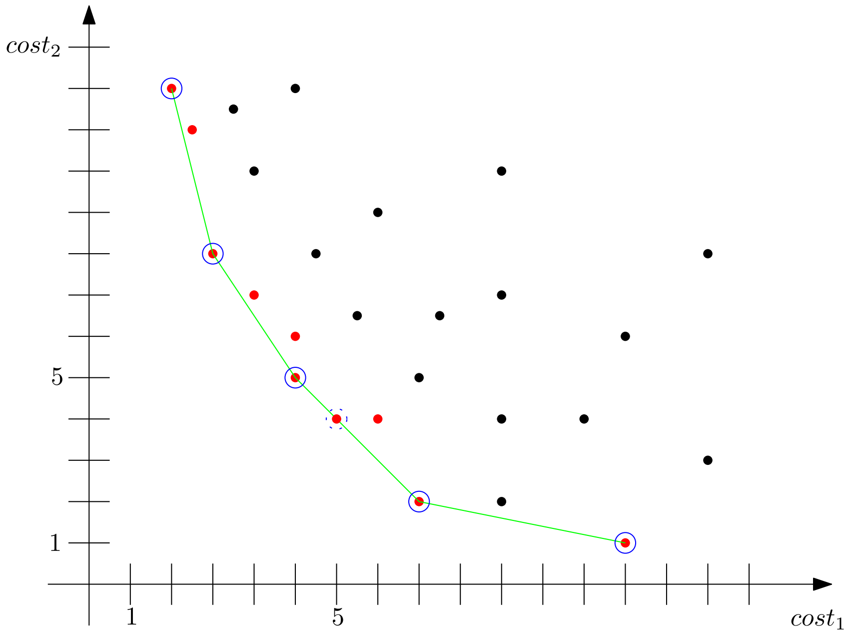

Figure 1 Paths in cost space (black and red dots); Pareto-optimal paths (red); (lower left part of) the boundary of the convex hull of all Pareto-optimal paths in green; extreme shortest paths/extreme points of the CH circled in blue; shortest (but not extreme) path dot-circled in blue.

Related Work

For the comparison with related work let us first define the well-studied parametric shortest path problem as in [9]. Given an acyclic, directed graph $G\,(V,\, E)$ the weights w_e of the edges $e \in E$ are linear functions of the form

$$w_e\,(\lambda) := a_e\lambda + b_e.$$

The cost of a path p is then given by $C\,(p,\, \lambda) := \sum_{e\in p} w_e\,(\lambda)$. Clearly, if we compute the shortest path from a node $s \in V$ to a node $t \in V$ for different values of λ, we may get different paths. Let Π_{st} bet the set of all paths from s to t. Then

$$C\,(\lambda) := \min_{p\in\Pi_{st}} C\,(p,\, \lambda)$$

is the shortest path cost function, which is a concave, piece-wise linear function. The maximum possible number of pieces of $C(\lambda)$ with respect to the size of G is called the *parametric shortest path complexity*. Complexity-wise, counting the pieces of $C(\lambda)$ is equivalent to counting extreme shortest paths from s to t. Therefore, all results regarding the parametric shortest path complexity are closely related to our work.

For the two-metric case, it is well known that the number of pieces in $C(\lambda)$ is upper bounded by $n^{O(\log n)}$, where n is the number of nodes in G [11]. This upper bound is tight [5, 6, 15], even for planar graphs [9].

Gajjar and Radhakrishnan [9] extend the parametric shortest path problem to three dimensions by setting

$$w_e(\lambda := (\lambda_1, \lambda_2, \lambda_3)) := a_e\lambda_1 + b_e\lambda_2 + c_e\lambda_3.$$

They show that in this case the number of extreme shortest paths is in $n^{(\log n)^2+O(\log n)}$.

We are not aware of any results regarding the parametric shortest path complexity beyond three dimensions. Parts of our way (especially Section 3.2) to prove the general upper bound are inspired by the proof for three dimensions in [9].

Both, Pareto-optimal paths as well as extreme shortest paths have been instrumented to create alternative route recommendations. The former approach, pursued e.g. in [7, 13], unfortunately only seems to be viable on rather small graphs due to the too rapidly growing number of Pareto-optimal paths. Restricting to extreme shortest paths, though, as in [8], has been shown to be feasible in different practical application scenarios [2, 1].

2 Preliminaries

In this section, we introduce the notions used in Section 3 and show some basic, well known properties.

For a set $f \subseteq \mathbb{R}^d$, we define its dimension $\dim(f)$ to be the maximum number of affinely independent points in f minus one. For a finite set $P \subset \mathbb{R}^d$ we denote the number of elements in P with $|P|$.

▶ **Definition 1.** *The d-metric preference space $\mathcal{P}_d$ is defined as follows.*

$$\mathcal{P}_d := \{(\alpha_1, \alpha_2, \ldots, \alpha_d) \in \mathbb{R}^d_{\geq 0} \mid \sum_{i=1}^{d} \alpha_i = 1\}$$

Note that the d-metric preference space is a $(d-1)$-dimensional simplex in d dimensions. Given a finite set of points $P \subset \mathbb{R}^d$ with $v \in P$, let $f_P(v) \subseteq \mathcal{P}_d$ be the set of preferences for which αv^T is minimal, where αv^T is the dot product of the (row) vectors α, v. Or more formally,

$$f_P(v) := \{\alpha \in \mathcal{P}_d \mid \alpha(v - v')^T \leq 0 \ \forall v' \in P\}. \tag{1}$$

▶ **Definition 2.** *Given a finite set of points $P \subset \mathbb{R}^d$, a subset $P' \subseteq P$ is a* preference cover *(PC) of P if and only if*

$$\bigcup_{v \in P'} f_P(v) = \mathcal{P}_d. \tag{2}$$

The following lemma states that there is a special, minimal PC for each finite point set P.

▶ **Lemma 3.** *Given a finite set of points $P \subset \mathbb{R}^d$ and let $X := \{v \in P \mid \dim(f_P(v)) = d-1\}$, then X is a PC of P and it holds*

$$|X| = \min_{P' \text{ is PC of } P} |P'|.$$

Proof. Let $\mathcal{M} := \min_{P' \text{ is PC of } P} |P'|$. As the number of points in P is finite and because the inequality in (1) is not strict, it holds

$$\bigcup_{v \in X} f_P(v) = \mathcal{P}_d.$$

Thus, X is a PC of P and $|X| \geq \mathcal{M}$. Moreover, for each $v \in X$ one can find an $\alpha \in \mathcal{P}_d$ with

$$\alpha v^T < \alpha v'^T \; \forall v' \in P \setminus \{v\}.$$

It follows that $|X| \leq \mathcal{M}$ and, thus, $|X| = \mathcal{M}$. ◀

The following definition is motivated by Lemma 3.

▶ **Definition 4.** *Given a finite set of points $P \subset \mathbb{R}^d$, we call the subset*

$$\mathcal{M}(P) := \{v \in P \mid \dim(f_P(v)) = d-1\}$$

minimum preference cover (MPC) of P.

▶ **Definition 5.** *Given a finite set of points $P \subset \mathbb{R}^d$, the d-metric* preference space subdivision *(PSS) $\mathcal{S}_d(P)$ (or simply $\mathcal{S}(P)$) is the arrangement induced by the set $\{f_P(v) \mid v \in P\}$. We write $f_P^r \in \mathcal{S}_d(P)$ for an r-dimensional facet of $\mathcal{S}_d(P)$.*

Figure 2 shows an example 3-metric PSS. The following lemma motivates the term *subdivision*.

▶ **Lemma 6.** *Given a finite set of points $P \subset \mathbb{R}^d$ with $d > 1$, for any two $v, v' \in \mathcal{M}(P)$ we have* $\dim(f_P(v) \cap f_P(v')) < d-1$. *Furthermore, for any $v \in P$ there is a $v' \in \mathcal{M}(P)$ with $f_P(v) \subseteq f_P(v')$.*

Proof. Given two points $v, v' \in \mathcal{M}(P)$, let H be the hyperplane described by $\alpha(v-v')^T = 0$, $\alpha \in \mathbb{R}^d$. For the first part of the lemma we have to show that $\dim(H \cap \mathcal{P}_d) \leq d-2$. This is the case if $\mathcal{P}_d \nsubseteq H$. The vector $v - v'$ must have at least one positive and one negative entry (otherwise, one vector would dominate the other). Thus, H cannot be parallel to $\sum_{i \leq d} \alpha_i = 1$ and $\mathcal{P}_d \nsubseteq H$ follows.

Now we come to the second part of the lemma. If $\dim(f_P(v)) = d-1$, then $v \in \mathcal{M}(P)$ and we are finished. Thus, we may assume that $0 < \dim(f_P(v)) < d-1$. From Lemma 3 we know that

$$f_P(v) = \{\alpha \in \mathbb{R}^d \mid \alpha(v-v')^T \leq 0 \; \forall \; v' \in \mathcal{M}(P)\}.$$

Therefore, if $\dim(f_P(v)) < d-1$, there must be a point $v' \in \mathcal{M}(P)$ and a hyperplane H' described by $\alpha(v-v')^T = 0$ with $f_P(v) \subseteq H'$. It follows that $f_P(v) \subseteq f_P(v')$. ◀

In fact, a PSS looks similar to a Voronoi diagram.

▶ **Definition 7.** *Given a finite set of points $P \subset \mathbb{R}^d$, then $\varphi_d^r(P)$ is the number of r-dimensional facets in the PSS $\mathcal{S}(P)$.*

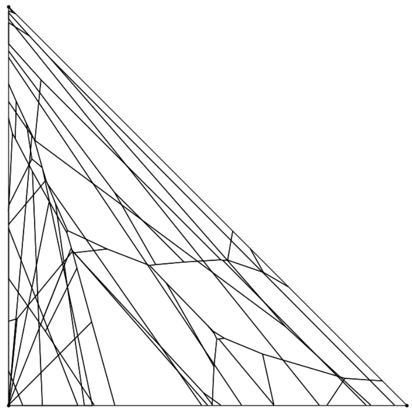

Figure 2 Example 3-metric preference space subdivision $\mathcal{S}_3$.

Throughout this work, we assume that a graph $G(V, E)$ with directed edges, without multi-edges and with n nodes is given. A path $\pi := u_1 e_1 u_2 e_2 \dots e_l u_{l+1}$ in G is an alternating sequence of nodes and edges that starts and ends with a node. Furthermore, for each edge $e_i \in \pi$ it holds $e_i = (u_i, u_{i+1})$. We say that a path π is simple if it contains each node $u \in V$ at most once.

A d-metric (or d-dimensional) cost function $c : E \to \mathbb{R}^d_{\geq 0}$ maps edges to (non-negative) cost vectors. We extend c to paths π as follows.

$$c(\pi) := \sum_{e \in \pi} c(e)$$

We define $\mathcal{C}_d$ to be the set of all d-metric cost functions of E.

Let Π be a set of paths (or: path set) in G and c a cost function for E. Then we define the point set $P(\Pi, c)$ as follows.

$$P(\Pi, c) := \{c(\pi) \mid \pi \in \Pi\}$$

With $\Pi_{st}(l)$ we denote all simple paths from $s \in V$ to $t \in V$ (with s and t being arbitrary but fixed nodes) with at most $\lceil l \rceil$ edges (l is a real number as we need to divide it by two in a recursion later on). We call $\Pi_{st}(l)$ a *complete* path set.

▶ **Definition 8.** *Given a path set Π, we define $\varphi^r_d(\Pi)$ with $1 \leq r \leq d-1$ as*

$$\varphi^r_d(\Pi) := \max_{c \in \mathcal{C}_d} \varphi^r_d(P(\Pi, c)).$$

Note that, by Lemma 6, it holds $\varphi^{d-1}_d(\Pi) = \max_{c \in \mathcal{C}_d} |\mathcal{M}(P(\Pi, c))|$.

▶ **Definition 9.** *Given a path set Π and a cost function $c \in \mathcal{C}_d$, an element $v \in P(\Pi, c) =: P$ is an* extreme shortest path *with respect to Π and c if and only if $v \in \mathcal{M}(P)$. Furthermore, for a preference $\alpha \in \mathcal{P}_d$ we call a cost vector $v \in \mathcal{M}(P)$ α-shortest path with respect to P if and only if $\alpha \in f_P(v)$.*

There are a few possibly confusing things to clarify here. First, we call the elements in $\mathcal{M}(P)$ extreme shortest paths even though they are mere cost vectors. The reason is that there is a bijective relationship between elements in P and Π as long as no two paths in Π have the same cost vector. However, if there are multiple paths in Π with the same cost vector we prefer to count them only once. This is the reason why we focus on the set P instead of the set Π.

Second, if a cost vector v is an extreme shortest path depends on the path set Π. For instance, if $\Pi := \{\pi\}$, then any cost vector $v = c(\pi)$ is an extreme shortest path. Thus, if we say that v is an extreme shortest path, this is always with respect to some path set Π and cost function c that we either mention explicitly or that should be clear from the context.

Third, a cell of a d-metric PSS $\mathcal{S}_d(P)$ is a $(d-1)$-dimensional and not a d-dimensional facet. The reason is that the preference space $\mathcal{P}_d$ lives in d dimensions but is itself a $(d-1)$-dimensional object (because the preferences sum up to one).

This work is about finding an upper bound for the number of extreme shortest paths in $P(\Pi_{st}(n), c)$ in arbitrary dimensions d and for arbitrary cost functions $c \in \mathcal{C}_d$.

Comparison of Personalized Route Planning and Parametric Shortest Path Problem

In the parametric shortest path problem, as defined for three dimensions in [9], the edge costs have the form

$$c\left(e \in E,\, \lambda \in \mathbb{R}^d\right) := \lambda c(e)^T \tag{3}$$

with $c:\ E \to \mathbb{R}^d$. In contrast to the definition in [9], λ_1 is typically set to one. However, note that the shortest path problem is homogeneous in the sense that if we scale all edge costs by a factor $\delta > 0$ we also scale the extreme shortest paths by δ. Thus, as argued in [9], (3) can always be scaled to either $\lambda_1 = 1$ or $\lambda_1 = -1$. The complexity of these two versions therefore can only differ by a factor of two.

In the personalized route planning model the edge costs are defined as

$$c(e \in E,\, \alpha \in \mathcal{P}_d) := \alpha c(e)^T \tag{4}$$

with $c:\ E \to \mathbb{R}^d_{\geq 0}$. The differences between (4) and (3) are that the preferences $\alpha \in \mathcal{P}_d$ and cost vectors are non-negative and the preferences sum up to one. We can handle the normalization with the same scaling argument as above. The non-negativity makes sure that there are no negative cost cycles in G regardless of α. This issue is treated differently in the parametric shortest path problem by requiring G to by acyclic. In the end, it is a matter of taste which requirement to choose. With minor adjustments our proofs also work for the case when G is acyclic and the costs and preferences are allowed to be negative.

3 A General Upper Bound on the Number of Extreme Shortest Paths

In this section we prove the following theorem.

▶ **Theorem 10.** *For any fixed but arbitrary $d \geq 1$ and any cost function $c \in \mathcal{C}_d$ the number of extreme shortest paths in $P(\Pi_{st}(n), c)$ is upper bounded by $n^{O\left(\log^{d-1} n\right)}$.*

▶ **Definition 11.** *We define $\varphi_d^r(n, l)$ to be the maximum of $\varphi_d^r(\Pi_{st}(l))$ over all possible graphs $G(V, E)$ with n nodes and all possible node pairs $s, t \in V$.*

Thus, Theorem 10 is equivalent to the statement $\varphi_d^{d-1}(n, n) \in n^{O\left(\log^{d-1} n\right)}$. We prove Theorem 10 with a recursion of the form $\varphi_d^{d-1}(n, n) \leq f\left(\varphi_{d-1}^{d-2}(n, n)\right)$. Thus, we upper bound $\varphi_d^{d-1}(n, n)$ with a recursion in the dimension d, which is an idea also used in the proof of the upper bound for $d = 3$ in [9].

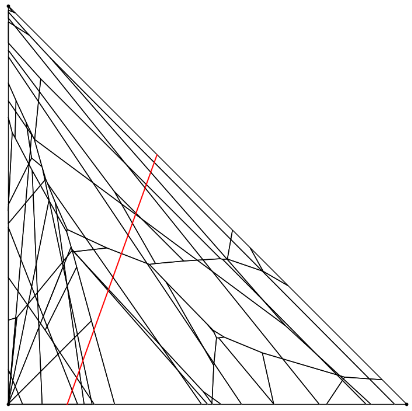

Figure 3 Example preference space subdivision with red hyperplane intersection.

We obtain the function f as follows. First, we show that the intersection of a d-dimensional PSS and a hyperplane can have at most the complexity of a $(d-1)$-dimensional PSS (Lemma 12). In fact, this observation is based on a similar result shown in [9] in the context of parametric shortest paths. Second, we use this insight to upper bound $\varphi_d^{d-2}(n, l)$ for special path sets (Lemma 13 and 14), which then allows us to construct a second recursion of the form $\varphi_d^{d-2}(n, l) \leq g\left(\varphi_d^{d-2}\left(n, \frac{l}{2}\right)\right)$. This leads to an upper bound for $\varphi_d^{d-2}(n, l)$ based on $\varphi_{d-1}^{d-2}(n, l)$ (Lemma 15). The proof of Theorem 10 then uses Lemma 15 together with an observation we discuss in Section 3.1 to construct the function f. We consider Lemma 13 and 14 in Section 3.3 as our main contributions as these ingredients allow us to generalize the ideas shown in [9] to arbitrary values d.

3.1 A First Upper Bound on the Number of Cells

The number of $(d-2)$-dimensional facets in a PSS is an upper bound of the number of $(d-1)$-dimensional facets in the same PSS for the following reasons. Given any finite point set $P \subset \mathbb{R}^d$ with $d > 1$. Each facet f_P^{d-2} of $\mathcal{S}_d(P)$ supports at most two $(d-1)$-dimensional facets. Moreover, every facet f_P^{d-1} is supported by at least d $(d-2)$-dimensional facets. This is true because the preference space $\mathcal{P}_d$ is bounded itself and, thus, there is no unbounded cell in $\mathcal{S}_d(P)$. Therefore, we have

$$\varphi_d^{d-1}(P) \leq \frac{2}{d}\varphi_d^{d-2}(P). \tag{5}$$

For the case $d = 1$, our task of counting extreme shortest paths is simple. It holds

$$\varphi_1^0(P) \leq 1 \tag{6}$$

for any finite point set P because $\mathcal{M}(P) = \{\min_{v \in P} v\}$.

3.2 Bounding the Complexity of PSS Intersections

In this section we show that the complexity of the intersection of a hyperplane with a d-metric PSS is upper bounded by the complexity of a $(d-1)$-metric PSS. This observation is a crucial ingredient of our proof of Theorem 10 as it allows us to construct recursive upper bounds in the dimension d. The authors of [9] prove a similar statement in the context of parametric shortest paths. As their setting and notation slightly differ from ours, we decided to give a proof of the following lemma.

▶ **Lemma 12.** *Let $d > 1$ and let H_d be a hyperplane in d dimensions with $\mathcal{P}_d \nsubseteq H_d$ that intersects the preference space $\mathcal{P}_d$. Then for any path set Π and cost function $c \in \mathcal{C}_d$ the set $Y := \{f^{d-1}_{P(\Pi,\, c)} \cap H_d \mid \dim\left(f^{d-1}_{P(\Pi,\, c)} \cap H_d\right) = d-2\}$ has at most $\varphi^{d-2}_{d-1}(\Pi)$ elements.*

Proof. We fix an arbitrary cost function $c \in \mathcal{C}_d$ and define $P := P(\Pi,\, c)$. See Figure 3 for an example of the intersection $H_d \cap \mathcal{P}_d$. The set $Z := \{f \cap H_d \mid f \in \mathcal{S}_d(P)\}$ looks like a $(d-1)$-metric PSS. However, it is unclear whether the complexity bounds for $(d-1)$-metric PSS also apply for sets like Z, which live in d dimensions. This lemma says that the answer is yes in the case of extreme shortest paths.

Our strategy to prove Lemma 12 looks as follows.

1. We show that for each $v \in P$ there is at most one $f \in Y$ with $\dim(f_P(v) \cap f) = d-2$.
2. We define mappings $\eta : \mathbb{R}^d \to \mathbb{R}^{d-1}$ and $\beta : H_d \cap \mathcal{P}_d \to \mathcal{P}_{d-1}$ such that for each $\alpha \in H_d \cap \mathcal{P}_d$ and each $v \in P$ it holds $\alpha v^T = \beta(\alpha)\, \eta(v)^T$.

We prove these points later on and first show that they are sufficient to prove the lemma. Let $\eta(P)$ and $\beta(f_P)$ be the mapped sets P and f_P. It follows from the second point that for any $v \in P$ with $\dim(f_P(v) \cap H_d) = d-2$ it holds

$$\beta(f_P(v) \cap H_d) \subseteq f_{\eta(P)}(\eta(v)).$$

Therefore, for each $f \in Y$ there is also an extreme shortest path $v' \in \mathcal{M}(\eta(P))$ with $\beta(f_P(v)) \subseteq f_{\eta(P)}(\eta(v'))$ (see Lemma 6).

Let $v' \in \mathcal{M}(\eta(P))$ be any extreme shortest path with respect to $\eta(P)$. Then the set $W := \{f \in Y \mid \beta(f) \subseteq f_{\eta(P)}(v')\}$ contains at most one element. Otherwise, we could take any cost vector $v \in P$ with $\eta(v) = v'$ and show, using point two, that all elements of W are subsets of $f_P(v)$. This contradicts the first point.

In summary, for each element $f \in Y$ we find an extreme shortest path $v' \in \mathcal{M}(\eta(P))$ with $\beta(f) \subseteq f_{\eta(P)}(v')$. Furthermore, for each extreme shortest path $v' \in \mathcal{M}(\eta(P))$ we find at most one $f \in Y$ with $\beta(f) \subseteq f_{\eta(P)}(v')$. Thus, it is possible to injectively map Y to $\mathcal{M}(\eta(P))$. As $|\mathcal{M}(\eta(P))| \le \varphi^{d-2}_{d-1}(\Pi)$ by definition and because we chose the cost function c arbitrarily, this finishes the proof.

It remains to show that the two points are correct. Assume that we find a cost vector $v \in P$ and two elements $f_1,\, f_2 \in Y$ with $f_1 \subseteq f_P(v)$ and $f_2 \subseteq f_P(v)$. We know from Lemma 6 that there is a $v' \in \mathcal{M}(P)$ with $f_P(v) \subseteq f_P(v')$. Let $f_3 := f_P(v') \cap H_d$. Clearly, $f_1 \subseteq f_3$ and $f_2 \subseteq f_3$. As $f_1 \neq f_2$ at least one of them cannot be equal to f_3. W.l.o.g. let $f_2 \subset f_3$.

By definition of Y, we find a shortest path $v_2 \in \mathcal{M}(P)$ with $f_2 = f_P(v_2) \cap H_d$. As $f_2 \subset H_d$ and $f_3 \subset H_d$ it follows that H_d is described by $\alpha(v' - v_2)^T = 0$. But then $f_P(v_2) \cap H_d = f_P(v') \cap H_d = f_3$, which is a contradiction to $f_2 \subset f_3$.

We come to the second point. W.l.o.g. let the intersection $H_d \cap \mathcal{P}_d$ be describable by an equation of the form

$$\alpha_1 = x_2 \cdot \alpha_2 + x_3 \cdot \alpha_3 + \cdots + x_{d-1} \cdot \alpha_{d-1} + b =: f(\alpha)$$

with $x_i, b \in \mathbb{R}$. We map each point $v := (v_1,\, v_2,\, \ldots,\, v_d) \in P$ to

$$\eta(v) := (v_2 + v_1 \cdot \tilde{v}_{2,1} + v_d \cdot \tilde{v}_{2,d},\, v_3 + v_1 \cdot \tilde{v}_{3,1} + v_d \cdot \tilde{v}_{3,d}, \ldots,\, v_d + v_1 \cdot \tilde{v}_{d,1} + v_d \cdot \tilde{v}_{d,d})$$

with $\tilde{v}_{i,1} = x_i + b$, $\tilde{v}_{i,d} = -(x_i + b)$ and $x_d = 0$. Note that η reduces the number of dimensions by one. Let $\eta(P)$ be the set of mapped points of P. By definition, $\mathcal{S}_{d-1}(\eta(P))$ has at most $\varphi^{d-2}_{d-1}(n,\, l)$ $(d-2)$-dimensional facets. Thus, there are at most $\varphi^{d-2}_{d-1}(n,\, l)$ extreme shortest paths in $\eta(P)$.

Let the map $\beta : H_d \cap \mathcal{P}_d \to \mathcal{P}_{d-1}$ be defined as follows.

$$\beta(\alpha) = (\alpha_2, \alpha_3, \ldots, \alpha_d + \alpha_1)$$

For any $v \in P$ and $\alpha \in H_d \cap \mathcal{P}_d$ we have

$$\begin{aligned}\alpha \cdot v^T = &\left(f(\alpha), \alpha_2, \ldots, 1 - f(\alpha) - \sum_{1<i<d} \alpha_i \right) \cdot (v_1, \ldots, v_d)^T \\ =&f(\alpha)(v_1 - v_d) + v_d + \sum_{1<i<d} \alpha_i (v_i - v_d) \\ =&b(v_1 - v_d) + (\alpha_1 + \alpha_d) v_d + \sum_{1<i<d} \alpha_i x_i (v_1 - v_d) + \sum_{1<i<d} \alpha_i v_i \\ =&(\alpha_1 + \alpha_d)(v_d + v_1 \tilde{v}_{d,1} + v_d \tilde{v}_{d,d}) + \sum_{1<i<d} \alpha_i (v_i + v_1 \tilde{v}_{i,1} + v_d \tilde{v}_{i,d}) \\ =&\beta(\alpha) \eta(v)^T.\end{aligned}$$

◀

3.3 Decomposing Path Sets

In this section we first look at path sets Π_{sut} that can be written as the pairwise concatenation of two path sets Π_{su} and Π_{ut} that end/start at a common node u. We will see that in such a scenario the PSS $\mathcal{S}(P(\Pi_{sut}, c))$ is the overlay of $\mathcal{S}(P(\Pi_{su}, c))$ and $\mathcal{S}(P(\Pi_{ut}, c))$ for any cost function c (see Figure 4 for an example). Thus, we can upper bound $\varphi_d^{d-2}(\Pi_{sut})$ with the help of Lemma 12.

▶ **Lemma 13.** *Let Π_{su} and Π_{ut} be two path sets such that each path in Π_{su} ends at node $u \in V$ and each path in Π_{ut} starts at node u. Furthermore, let $\Pi_{sut} := \{\pi_1\pi_2 \mid \pi_1 \in \Pi_{su}, \pi_2 \in \Pi_{ut}\}$ be the pairwise concatenation of Π_{su} and Π_{ut}. Then it holds*

$$\varphi_d^{d-2}(\Pi_{sut}) \leq \varphi_{d-1}^{d-2}(\Pi_{ut}) \cdot \varphi_d^{d-2}(\Pi_{su}) + \varphi_{d-1}^{d-2}(\Pi_{su}) \cdot \varphi_d^{d-2}(\Pi_{ut}).$$

Proof. We fix an arbitrary cost function $c \in \mathcal{C}_d$ and prove the inequality for c. Let $P_{sut} := P(\Pi_{sut}, c)$, $P_{su} := P(\Pi_{su}, c)$ and $P_{ut} := P(\Pi_{ut}, c)$. Given any $\alpha \in \mathcal{P}_d$, let v_{su} and v_{ut} be the α-shortest paths in P_{su} and P_{ut}. Then with $v_{sut} := v_{su} + v_{ut}$ it holds $v_{sut} \in P_{sut}$ and v_{sut} is the α-shortest path in P_{sut}. This is true because all paths in Π_{sut} go via node u. Thus, $\mathcal{S}(P_{sut})$ is the overlay of $\mathcal{S}(P_{su})$ and $\mathcal{S}(P_{ut})$.

Every facet $f_{P_{su}}^{d-2}$ is part of a hyperplane H_d. If we let H_d intersect the preference subdivision $\mathcal{S}_d(P_{ut})$ we know from Lemma 12 that this intersection contains no more than $\varphi_{d-1}^{d-2}(\Pi_{ut})$ $(d-2)$-dimensional facets. Thus, in the overlay of $\mathcal{S}(P_{su})$ and $\mathcal{S}(P_{ut})$ the facet $f_{P_{su}}^{d-2}$ can be split into at most $\varphi_{d-1}^{d-2}(\Pi_{ut})$ $(d-2)$-dimensional facets as well. An analogous statement holds for the facets $f_{P_{ut}}^{d-2}$, which finishes the proof. ◀

The following lemma addresses the problem of upper bounding $\varphi_d^{d-2}(\Pi')$ if the path set Π' is the union of multiple path sets.

▶ **Lemma 14.** *Given k path sets $\Pi_1, \Pi_2, \ldots, \Pi_k$ and let $\Pi' := \bigcup_{1\leq i\leq k} \Pi_i$, then it holds for every $d > 1$*

$$\varphi_d^{d-2}(\Pi') \leq 2 \cdot k \cdot \varphi' \cdot \sum_{1\leq i\leq k} \varphi_d^{d-2}(\Pi_i),$$

with $\varphi' := \max_{1\leq i\leq k} \varphi_{d-1}^{d-2}(\Pi_i)$.

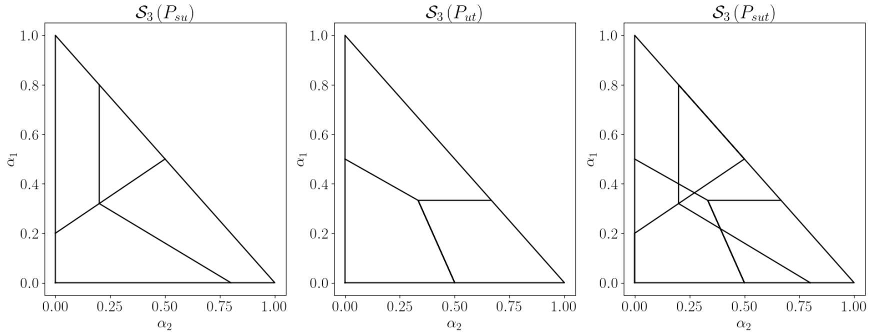

Figure 4 This figure illustrates the meaning of Lemma 13. With the two path sets Π_{su} and Π_{ut} the PSS $\mathcal{S}(P_{sut})$, as defined in Lemma 13, is the overlay of $\mathcal{S}(P_{su})$ and $\mathcal{S}(P_{ut})$. An edge $f^1_{P_{su}}$ can cause multiple edges in the PSS $\mathcal{S}_3(P_{sut})$ by intersecting the PSS $\mathcal{S}_3(P_{ut})$. Lemma 13 says that $f^1_{P_{su}}$ can be split into at most $\varphi^1_2(\Pi_{ut})$ many edges in $\mathcal{S}_3(P_{sut})$.

Proof. We fix an arbitrary cost function $c \in \mathcal{C}_d$ and define $P' := P(\Pi', c)$ and $P_i := P(\Pi_i, c)$ for all $1 \le i \le k$.

Now, let us fix a facet $f^{d-2}_{P'} \in \mathcal{S}(P')$. The facet is part of a hyperplane H that is described by $\alpha(v_1 - v_2)^T = 0$ with $v_1, v_2 \in \mathcal{M}(P')$ and $v_1 \neq v_2$.

We first assume that v_1 and v_2 belong to the same set P_i. As $P_i \subseteq P'$, we have $v_1, v_2 \in \mathcal{M}(P_i)$ and $f_{P'}(v_1) \cap f_{P'}(v_2) \subseteq f_{P_i}(v_1) \cap f_{P_i}(v_2)$. Therefore, $\dim(f_{P_i}(v_1) \cap f_{P_i}(v_2)) = d-2$ and there is a facet $f^{d-2}_{P_i} = f_{P_i}(v_1) \cap f_{P_i}(v_2)$. Thus, we have at most $\varphi^{d-2}_d(\Pi_i)$ such pairs v_1, v_2 in $\mathcal{M}(P_i)$.

We now assume that v_1 and v_2 belong to different sets P_i and P_j. From $P_i \cup P_j \subseteq P'$ it follows that $f_{P'}(v_1) \cap f_{P'}(v_2) \subseteq f_{P_i}(v_1) \cap f_{P_j}(v_2)$. Thus, $\dim\left(f_{P_i}(v_1) \cap f_{P_j}(v_2)\right) \ge d-2$.

Therefore, there must be a facet $f^{d-2}_{P_i}$ or $f^{d-2}_{P_j}$ (or both) that intersects with $f_{P_i}(v_1) \cap f_{P_j}(v_2)$.

In summary, for each facet $f^{d-2}_{P'}$ we either find a corresponding facet $f^{d-2}_{P_i}$ with $f^{d-2}_{P'} \subseteq f^{d-2}_{P_i}$ (case one) or an intersection $f^{d-2}_{P_i} \cap f^{d-1}_{P_j}$ (or $f^{d-1}_{P_i} \cap f^{d-2}_{P_j}$) with $\dim\left(f^{d-2}_{P_i} \cap f^{d-1}_{P_j}\right) = d-2$ (case two). Clearly, we can upper bound the first case with $\sum_{1 \le i \le k} \varphi^{d-2}_d(\Pi_i)$.

The second case we handle with the help of Lemma 12. The intersection of the facet $f^{d-2}_{P_i}$ with any PSS $\mathcal{S}_d(P_j)$ can contain at most $\varphi^{d-2}_{d-1}(\Pi_j) \le \varphi'$ $(d-2)$-dimensional facets (Lemma 12). Moreover, there are at most two extreme shortest paths in P_i that are adjacent to $f^{d-2}_{P_i}$. Thus, in total the facet $f^{d-2}_{P_i}$ can contribute to at most $2 \cdot (k-1) \cdot \varphi'$ intersections of case two. Therefore, the number of facets $f^{d-2}_{P'}$ of case two can be upper bounded by $2 \cdot (k-1) \cdot \varphi' \cdot \sum_{1 \le i \le k} \varphi^{d-2}_d(\Pi_i)$.

With $\sum_{1 \le i \le k} \varphi^{d-2}_d(\Pi_i) + 2 \cdot (k-1) \cdot \varphi' \cdot \sum_{1 \le i \le k} \varphi^{d-2}_d(\Pi_i) \le 2 \cdot k \cdot \varphi' \cdot \sum_{1 \le i \le k} \varphi^{d-2}_d(\Pi_i)$ the statement follows. ◀

3.4 Proving the Upper Bound via Recursion

In the following lemma we describe a recursive upper bound for $\varphi^{d-2}_d(n, l)$ in the dimension d, which we then use to prove Theorem 10.

▶ **Lemma 15.** *$\varphi^{d-2}_d(n, l) \le d \cdot l^2 \cdot n^{2\log l} \cdot \varphi^{d-2}_{d-1}(n, l)^{2\log l}$ for $d > 1$.*

Proof. We prove Lemma 15 for an arbitrary complete path set $\Pi_{st}(l)$ and cost function $c \in \mathcal{C}_d$ and define $P_{st} := P(\Pi_{st}(l), c)$. We first introduce path sets $\Pi_{sut} := \{\pi \in \Pi_{st}(l) \mid u \in V$ divides π into subpaths of at most $\lceil \frac{l}{2} \rceil$ edges$\}$. Clearly, each path $\pi \in \Pi_{st}(l)$ is contained in at least one path set Π_{sut}. Thus, we can apply Lemma 14 and get

$$\varphi_d^{d-2}\left(\Pi_{st}\left(l\right)\right) \leq 2 \cdot n \cdot \varphi_{d-1}^{d-2}\left(n,\, l\right) \cdot \sum_{u \in V} \varphi_d^{d-2}\left(\Pi_{sut}\right), \tag{7}$$

as one can easily show that $\varphi_{d-1}^{d-2}\left(n,\, l\right) \geq \max_{u \in V} \varphi_{d-1}^{d-2}\left(\Pi_{sut}\right)$.

The path sets Π_{sut} can be written as the pairwise concatenation of the path sets $\Pi_{su}\left(\frac{l}{2}\right)$ and $\Pi_{ut}\left(\frac{l}{2}\right)$. Thus, they satisfy the requirements to apply Lemma 13 such that for each node u we get

$$\begin{aligned}\varphi_d^{d-2}\left(\Pi_{sut}\right) &\leq \varphi_{d-1}^{d-2}\left(\Pi_{ut}\left(\frac{l}{2}\right)\right) \cdot \varphi_d^{d-2}\left(\Pi_{su}\left(\frac{l}{2}\right)\right) + \varphi_{d-1}^{d-2}\left(\Pi_{su}\left(\frac{l}{2}\right)\right) \cdot \varphi_d^{d-2}\left(\Pi_{ut}\left(\frac{l}{2}\right)\right) \\ &\leq 2 \cdot \varphi_{d-1}^{d-2}\left(n,\, \frac{l}{2}\right) \cdot \varphi_d^{d-2}\left(n,\, \frac{l}{2}\right).\end{aligned} \tag{8}$$

If we combine (7) and (8), we get

$$\begin{aligned}\varphi_d^{d-2}\left(\Pi_{st}\left(l\right)\right) &\leq 4 \cdot n^2 \cdot \varphi_{d-1}^{d-2}\left(n,\, l\right) \cdot \varphi_{d-1}^{d-2}\left(n,\, \frac{l}{2}\right) \cdot \varphi_d^{d-2}\left(n,\, \frac{l}{2}\right) \\ &\leq 4 \cdot n^2 \cdot \varphi_{d-1}^{d-2}\left(n,\, l\right)^2 \cdot \varphi_d^{d-2}\left(n,\, \frac{l}{2}\right).\end{aligned} \tag{9}$$

Using this recursion in l we obtain with $\varphi_d^{d-2}\left(n,\, 1\right) = d$

$$\begin{aligned}\varphi_d^{d-2}\left(\Pi_{st}\left(l\right)\right) &\leq d \cdot 4^{\log l} \cdot n^{2\log l} \cdot \varphi_{d-1}^{d-2}\left(n,\, l\right)^{2\log l} \\ &= d \cdot l^2 \cdot n^{2\log l} \cdot \varphi_{d-1}^{d-2}\left(n,\, l\right)^{2\log l}.\end{aligned} \tag{10}$$

◀

Proof of Theorem 10. We need to show that $\varphi_d^{d-1}\left(n,\, n\right) \in n^{O\left(\log^{d-1} n\right)}$ for a fixed but arbitrary number of cost types $d \geq 1$. From Lemma 15 we know that $\varphi_d^{d-2}\left(n,\, n\right) \leq d \cdot n^{2+2\log n} \cdot \varphi_{d-1}^{d-2}\left(n,\, n\right)^{2\log n}$. With (5) we get $\varphi_d^{d-1}\left(n,\, n\right) \leq 2 \cdot n^{2+2\log n} \cdot \varphi_{d-1}^{d-2}\left(n,\, n\right)^{2\log n}$. With $\varphi_1^0\left(n,\, n\right) = 1$ this leads to $\varphi_d^{d-1}\left(n,\, n\right) \in n^{O\left(\left(2 \cdot \log n\right)^{d-1}\right)}$, which is, for a fixed d, equal to $\varphi_d^{d-1}\left(n,\, n\right) \in n^{O\left(\log^{d-1} n\right)}$. ◀

4 Conclusions

In this paper we showed that the number of extreme shortest paths in a graph G is upper bounded by $n^{O\left(\log^{d-1} n\right)}$, where n is the number of nodes and d is the fixed but arbitrary number of edge costs in G. This is a generalization of previous results in the context of parametric shortest paths for two and three dimensions.

An open question is whether one can also generalize the matching two-dimensional lower bounds shown in [5].

References

1. Florian Barth, Stefan Funke, and Claudius Proissl. Preference-based trajectory clustering: An application of geometric hitting sets. In *32nd International Symposium on Algorithms and Computation (ISAAC 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.
2. Florian Barth, Stefan Funke, and Sabine Storandt. Alternative multicriteria routes. In *2019 Proceedings of the Meeting on Algorithm Engineering and Experiments (ALENEX)*, pages 66–80, 2019. doi:10.1137/1.9781611975499.6.

3 HP Benson and E Sun. Outcome space partition of the weight set in multiobjective linear programming. *Journal of Optimization Theory and Applications*, 105(1):17–36, 2000.

4 Jon Louis Bentley, Hsiang-Tsung Kung, Mario Schkolnick, and Clark D. Thompson. On the average number of maxima in a set of vectors and applications. *J. ACM*, 25(4):536–543, October 1978. doi:10.1145/322092.322095.

5 Patricia J Carstensen. Complexity of some parametric integer and network programming problems. *Mathematical Programming*, 26(1):64–75, 1983.

6 Patricia June Carstensen. *The complexity of some problems in parametric linear and combinatorial programming.* PhD thesis, University of Michigan, 1983.

7 Daniel Delling and Dorothea Wagner. Pareto paths with SHARC. In *Proc. 8th International Symposium on Experimental Algorithms (SEA)*, volume 5526 of *Lecture Notes in Computer Science*, pages 125–136. Springer, 2009.

8 Stefan Funke and Sabine Storandt. Personalized route planning in road networks. In *Proc. 23rd SIGSPATIAL International Conference on Advances in Geographic Information Systems, Bellevue, WA, USA*, pages 45:1–45:10. ACM, 2015.

9 Kshitij Gajjar and Jaikumar Radhakrishnan. Parametric shortest paths in planar graphs. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 876–895. IEEE, 2019.

10 Robert Geisberger, Moritz Kobitzsch, and Peter Sanders. Route planning with flexible objective functions. In *Workshop on Algorithms and Experimentation*, ALENEX '10, pages 124–137, USA, 2010. Society for Industrial and Applied Mathematics.

11 Daniel Mier Gusfield. *Sensitivity analysis for combinatorial optimization.* PhD thesis, University of California, Berkeley, 1980.

12 Gabriel Y Handler and Israel Zang. A dual algorithm for the constrained shortest path problem. *Networks*, 10(4):293–309, 1980.

13 Hans-Peter Kriegel, Matthias Renz, and Matthias Schubert. Route skyline queries: A multi-preference path planning approach. In *26th International Conference on Data Engineering (ICDE)*, pages 261–272. IEEE Computer Society, 2010.

14 Kurt Mehlhorn and Mark Ziegelmann. Resource constrained shortest paths. In *European Symposium on Algorithms*, pages 326–337. Springer, 2000.

15 Ketan Mulmuley and Pradyut Shah. A lower bound for the shortest path problem. In *Proceedings 15th Annual IEEE Conference on Computational Complexity*, pages 14–21. IEEE, 2000.