Pareto envelopes in $\mathbb{R}^3$ under l_1 and l_∞ distance functions*

[Extended Abstract]

Victor Chepoi and Karim Nouioua

ABSTRACT

Given a vector objective function $\mathbf{f} = (f_1, \ldots, f_n)$ defined on a set X, a point $y \in X$ is *dominated* by a point $x \in X$ if $f_i(x) \le f_i(y)$ for all $i \in \{1, \ldots, n\}$ and there exists an index $j \in \{1, \ldots, n\}$ such that $f_j(x) < f_j(y)$. The non-dominated points of X are called the *Pareto optima* of $\mathbf{f}$. H. Kuhn (1967,1973) applied the concept of Pareto optimality to distance functions and characterized the convex hull $\mathrm{conv}(T)$ of any set $T = \{t_1, \ldots, t_n\}$ of $\mathbb{R}^m$ as the set of all Pareto optima of the vector function $\mathbf{d}_2(x) = (d_2(x, t_1), \ldots, d_2(x, t_n))$, where $d_2(x, y)$ is the Euclidean distance between $x, y \in \mathbb{R}^m$. Motivated by this result, given a metric space (X, d) and a set $T = \{t_1, \ldots, t_n\}$ of X, we call the set $\mathcal{P}_d(T)$ of all Pareto optima of the function $\mathbf{d}(x) = (d(x, t_1), \ldots, d(x, t_n))$ the *Pareto envelope* of T.

In this paper, we investigate the Pareto envelopes in $\mathbb{R}^m$ endowed with l_1- or l_∞-distances. We characterize $\mathcal{P}_{d_\infty}(T)$ in all dimensions and $\mathcal{P}_{d_1}(T)$ in $\mathbb{R}^3$. Using these characterizations, we design efficient algorithms for constructing these envelopes in $\mathbb{R}^3$, in particular, an optimal $O(n \log n)$-time algorithm for $\mathcal{P}_{d_1}(T)$ and an $O(n \log^2 n)$-time algorithm for $\mathcal{P}_{d_\infty}(T)$.

Keywords

Pareto envelope, dominance, l_1- and l_∞-distance, algorithm

1. INTRODUCTION

1.1 Pareto envelopes

Convex hulls, in particular convex hulls in 2- and 3-dimensional spaces, are used in various applications and represent a basic object of investigations in computational geometry. Efficient algorithms for computing convex hulls are described in the textbooks [4, 11, 26]. The convex hull of an n point set $T = \{t_1, \ldots, t_n\}$ of $\mathbb{R}^m$ also hosts such remarkable points like the center, the barycenter, the Fermat-Torricelli point (the median) of T as well as the optimal solutions of several NP-hard problems like the Steiner tree problem, the p-median, or the p-center problem for T. All such problems consist in minimizing a certain function depending of the Euclidean distances to (or between) the terminals of T. This leads H. Kuhn [22, 23] to characterize $\mathrm{conv}(T)$ in truly distance terms: a point $x \in \mathbb{R}^m$ belongs to $\mathrm{conv}(T)$ if and only if the Euclidean distance vector $\mathbf{d}_2(x) = (d_2(x, t_1), d_2(x, t_2), \ldots, d_2(x, t_n))$ of p is not dominated by the distance vector $\mathbf{d}_2(y)$ of any other point $y \in \mathbb{R}^m$.

*LIF, Faculté des Sciences de Luminy, Université de la Méditerranée, F-13288 Marseille Cedex 9, France, {chepoi,nouioua}@lif.univ-mrs.fr.
This research was partly supported by the ANR grant BLAN06-1-138894 (projet OPTICOMB)

Kuhn's characterization of $\mathrm{conv}(T)$ invites to investigate analogous geometric objects obtained by replacing the Euclidean distance d_2 by any other distance d on $\mathbb{R}^m$, in particular by a distance induced by a norm. For this, we recall the fundamental notion from multi-objective optimization, that of *Pareto optimality* [21, 28]. Given a vector objective function $\mathbf{f} = (f_1, \ldots, f_n)$ defined on a set X, a point $y \in X$ is *dominated* by a point $x \in X$ if $f_i(x) \le f_i(y)$ for all $i \in \{1, \ldots, n\}$ and there exists an index $j \in \{1, \ldots, n\}$ such that $f_j(x) < f_j(y)$. The non-dominated (also named *efficient*) points of X are called the *Pareto optima* with respect to the vector function $\mathbf{f}$. Now, given a metric d on $\mathbb{R}^m (=: X)$ and a set of n terminals $T = \{t_1, \ldots, t_n\}$, we consider the vector objective function $\mathbf{d}(x) = (d(t_1, x), \ldots, d(t_n, x))$ and Pareto optimality with respect to this particular function. We write $p \succ_T q$ (or simply $p \succ q$) if p dominates q, and $p \sim_T q$ if p and q have the same distance vector. In view of Kuhn's result, we call the set of Pareto optima of $\mathbf{d}$ the *Pareto envelope* of T and denote it by $\mathcal{P}_d(T)$. The set of all points $p \in \mathcal{P}_d(T)$, such that $\mathbf{d}(p) \ne \mathbf{d}(q)$ for any other point q, constitute the *strict part* of the Pareto envelope and is denoted by $\mathcal{P}^0_d(T)$. Notice that strict Pareto envelopes share the monotonicity property of usual convex hulls (and of all closure operators): $T' \subseteq T$ implies $\mathcal{P}^0_d(T') \subseteq \mathcal{P}^0_d(T)$. In this paper, we investigate Pareto and strict Pareto envelopes for l_1 and l_∞ distances; for illustration of these envelopes see Figures 1 and 2.

1.2 l_1 and l_∞ distances

Recall that, given two points $p = (p^1, \ldots, p^m)$ and $q = (q^1, \ldots, q^m)$ of $\mathbb{R}^m$, the l_1-*distance* between p and q is $d_1(p, q) = \sum_{i=1}^m |p^i - q^i|$ and the l_∞-*distance* between p and q is $d_\infty(p, q) = \max\{|p^i - q^i| : i = 1, \ldots, m\}$. For a distance d on $\mathbb{R}^m$ denote by $B_d(p, r)$ the closed *ball* of radius r centered at p, i.e., $B_d(p, r) = \{q \in \mathbb{R}^m : d(p, q) \le r\}$. Denote also by $I_d(p, q) = \{x \in \mathbb{R}^m : d(p, x) + d(x, q) = d(p, q)\}$ the *interval* between the points p and q. The l_1- and l_∞-distances are particular cases of distances on $\mathbb{R}^m$ induced by polyhedral norms, because any ball $B_{d_\infty}(p, r)$ is a cube and any ball $B_{d_1}(p, r)$ is a cross-polytope (regular octahedron). The interval $I_{d_1}(p, q)$ is the axis-parallel box having the segment $[p, q]$ as a diagonal. On the other hand, $I_{d_\infty}(p, q)$ is an octahedron obtained as an intersection of two particular cones, one with vertex at p and another at q; see [6] for details. Finally notice that there is an isometry from the l_1-plane to the l_∞-plane: it suffices to rotate the plane by 45° and then shrink it by a factor $\frac{1}{\sqrt{2}}$. Therefore all results about Pareto envelopes in the l_1-plane immediately apply to the l_∞-plane and vice versa. This is not longer true for dimensions larger than 2.

1.3 Related work

As we mentioned already, Pareto optimality was first applied to distance vectors by H. Kuhn [22, 23], who established the equality $\mathcal{P}_{d_2}(T) = \text{conv}(T)$ (in fact, any point q outside $\text{conv}(T)$ is dominated by its unique closest point q' in $\text{conv}(T)$). In [30], Thisse, Ward, and Wendell proved that this equality is true for all distances induced by round norms. In general, neither $\mathcal{P}_d(T) \subseteq \text{conv}(T)$ nor $\text{conv}(T) \subseteq \mathcal{P}_d(T)$ hold for all normed spaces, in particular these inclusions are not true for l_1- and l_∞-norms. Nevertheless, Wendell and Hurter [34] established that the strict Pareto envelope $\mathcal{P}^0_d(T)$ is always included in $\text{conv}(T)$.

The investigation of Pareto envelopes for particular norms has been initiated by Wendell, Hurter, Lowe [35] and continued by Chalmet, Francis, Kolen [7] and Durier, Michelot [14, 15]. The main result of [7] is the following nice characterization of Pareto envelopes in the l_1-plane:

$$\mathcal{P}_{d_1}(T) = \cap_{i=1}^n (\cup_{j=1}^n I_{d_1}(t_i, t_j)) =: \Upsilon_{d_1}(T). \tag{1}$$

This result is used in [7] to establish the correctness of an optimal $O(n \log n)$ sweeping-line algorithm for constructing $\mathcal{P}_{d_1}(T)$ in the l_1-plane (notice that the first algorithm for this problem was proposed in [35] and has complexity $O(n^2)$). For polyhedral norms, Durier and Michelot [14, 15, 16] introduce the notion of *elementary convex sets* as closed convex sets C of $\mathbb{R}^m$ which can be written as the intersection of cones generated by facets of the unit ball and centered at terminals. They proved that for any polyhedral norm, $\mathbb{R}^m$ is partitioned into elementary convex sets and that the Pareto envelope is a connected union of a finite number of such cells. Durier [14] presents two rules to test if a point p belongs to the Pareto envelope, nevertheless no characterizations or algorithmic procedures for constructing Pareto envelopes in polyhedral norms have not been proposed in [14, 15]. Consequently, Pelegrin and Fernandez [25] describe an algorithm for constructing Pareto envelopes in the plane endowed with a polygonal norm (its complexity depends on the number of extremal vertices of the unit ball and the number of terminals).

In several papers, Pareto envelopes are used to reduce the search space in some optimization problems (notice that in this literature, Pareto envelopes are named *efficient sets* [7, 14, 15, 16, 25, 30, 34, 35]) by showing that Pareto envelopes host all or at least one optimal solution(s) of respective problems. For example, Wendell and Hurter [34] establish this type of results for *Weber problem* (the weighted version of the median problem), while Hansen, Perreur, and Thisse [18] prove a similar result for the NP-hard *multifacility location problem*. In [9], we show that Pareto envelopes in the l_1-plane contains at least one minimum Manhattan network and we use this to design a factor 2 approximation algorithm for *minimum Manhattan network problem*. For other results in this vein, see [15, 30]. Notice also that distance problems induced by polyhedral norms have been investigated from algorithmic point of view by Widmayer, Wu, and Wong in [36]; for results and additional references on fixed-orientation computational geometry see the book by Fink and Wood [17]. Voronoi diagrams with polyhedral metrics in high dimensions have been investigated by Boissonnat, Sharir, Tagansky, and Yvinec [5].

1.4 Our results

In this paper, we investigate Pareto and strict Pareto envelopes in l_1- and l_∞-norms. The departing point of our research was the paper by Chalmet, Francis, and Kolen [7], in particular, the equality (1) in the l_1-plane. Due to the isometry between the l_1- and l_∞-distances in the plane, a similar result holds in the l_∞-plane. We establish below that, in fact, the equality $\mathcal{P}_{d_\infty}(T) = \Upsilon_{d_\infty}(T)$ holds in l_∞-spaces of arbitrary dimension. Surprisingly, our proof is much simpler than the proof given in [7] for the plane. We also characterize the strict Pareto envelope $\mathcal{P}^0_{d_\infty}(T)$ as the intersection of $2m$ unions of cones, each union consisting of cones centered at the terminals and oriented in the same coordinate direction. We also establish a relationship between Pareto envelopes in l_∞-spaces and the injective hulls of finite metric spaces, introduced and investigated in [13, 19].

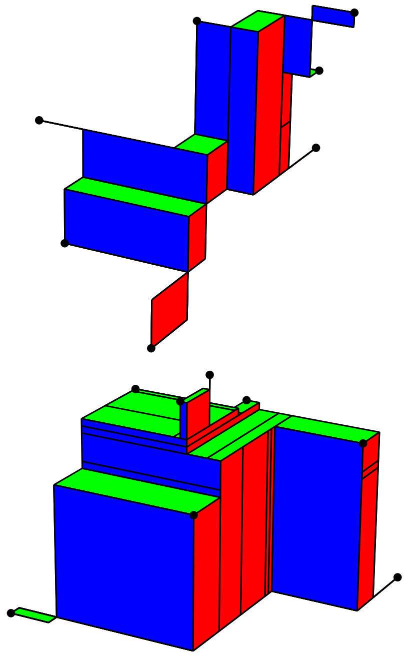

Figure 1: Examples of $\mathcal{P}_{d_1}(T)$.

We characterize the Pareto envelopes in $\mathbb{R}^3$ with l_1-norm using the equality $\mathcal{P}_{d_1}(T) = \mathcal{I}(T) = \mathcal{M}(T)$, where the sets $\mathcal{I}(T)$ and $\mathcal{M}(T)$ are defined in the following way. $\mathcal{I}(T)$ is the intersection of three polyhedra $\mathcal{I}^1(T)$, $\mathcal{I}^2(T)$, and $\mathcal{I}^3(T)$, where $\mathcal{I}^i(T)$ $(i \in \{1,2,3\})$ is the Cartesian product of the Pareto envelope of the orthogonal projection of T onto the coordinate plane $H^i := \{x \in \mathbb{R}^3 : x^i = 0\}$ with the line orthogonal to this plane. On the other hand, $\mathcal{M}(T)$ is a cubical complex obtained from the median closure of the set T in $\mathbb{R}^3$. Recall that the *median* $m(x, y, z)$ of three points x, y, and z is a point having as its ith coordinate the median of ith coordinates of x, y, and z. Then the *median closure* $\text{med}(T)$ of T

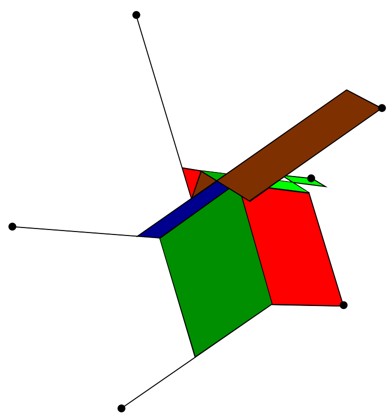

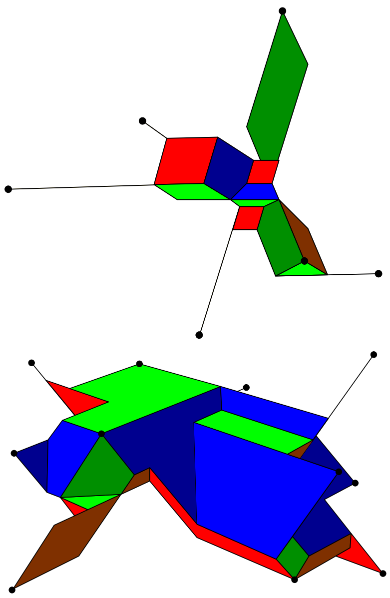

Figure 2: Examples of $\mathcal{P}_{d_\infty}(T)$.

is the smallest set containing T which together with any three points also contains their median. Median structures arise in discrete mathematics, theoretical computer science, geometry of metric spaces of non-positive curvature, and theory of abstract convexity; for an overview, see the book [31] and the papers [2, 3, 8, 20]. $\mathcal{M}(T)$ is the cubical complex obtained by replacing the graphical cubes of med(T) by solid boxes. The equality $\mathcal{P}_{d_1}(T) = \mathcal{M}(T)$ also holds in the l_1-plane, however neither of the equalities $\mathcal{P}_{d_1}(T) = \mathcal{I}(T)$ or $\mathcal{P}_{d_1}(T) = \mathcal{M}(T)$ is no longer true in dimensions larger than 3.

The characterizations of $\mathcal{P}_{d_\infty}(T), \mathcal{P}^0_{d_\infty}(T)$, and $\mathcal{P}_{d_1}(T)$ lead to efficient algorithms of their construction. For example, $\mathcal{P}_{d_1}(T) = \mathcal{I}(T)$ allows to construct $\mathcal{P}_{d_1}(T)$ in $\mathbb{R}^3$ in optimal $O(n \log n)$ time. For $\mathcal{P}^0_{d_\infty}(T)$ in $\mathbb{R}^3$, we present an $O(n \log^2 n)$-time algorithm. More related results and a detailed treatment of Pareto envelopes are given in the doctoral thesis of the second author [24]. The thesis also contains a description of the implementation of algorithms for constructing Pareto envelopes in $\mathbb{R}^2$ and $\mathbb{R}^3$.

The rest of this paper is organized in the following way. In Section 2, we present some preliminary results. Section 3 presents the characterization of $\mathcal{P}_{d_\infty}(T)$ and $\mathcal{P}^0_{d_\infty}(T)$, while Section 4 that of $\mathcal{P}_{d_1}(T)$. Section 5 describes the algorithms for constructing $\mathcal{P}_{d_\infty}(T)$ and $\mathcal{P}_{d_1}(T)$ in $\mathbb{R}^3$.

2. PRELIMINARY RESULTS

In this section, we establish some results which are true for Pareto envelopes in all metric spaces or all normed spaces. For a metric space (X, d) and a finite subset $T = \{t_1, \ldots, t_n\}$ of X, set $\Upsilon_d(T) = \cap_{i=1}^n(\cup_{j=1}^n I_d(t_i, t_j))$.

LEMMA 1. $\Upsilon_d(T) \subseteq \mathcal{P}_d(T)$.

PROOF. Suppose by way of contradiction that there exist two points $p, q \in \mathbb{R}^m$ such that $p \succ q$ and $q \in \Upsilon_d(T)$. Then there exists a terminal t_i so that $d(p, t_i) < d(q, t_i)$. Since q belongs to $\Upsilon_d(T)$, it also belongs to $\cup_{j=1}^n I_d(t_i, t_j)$, thus there exists a terminal t_k such that $q \in I_d(t_i, t_k)$. This yields $d(t_i, t_k) = d(t_i, q) + d(q, t_k)$. Notice also that $d(p, t_k) \le d(q, t_k)$ because $p \succ q$. All this implies that $d(p, t_i) + d(p, t_k) < d(q, t_i) + d(q, t_k) = d(t_i, t_k)$, contrary to the triangle inequality. □

The following result shows that the operator $\mathcal{P}^0_d$ is monotone, thus justifying the name "Pareto envelope".

LEMMA 2. *If $T' \subset T$, then $\mathcal{P}^0_d(T') \subseteq \mathcal{P}^0_d(T)$.*

PROOF. Suppose by way of contradiction that $q \in \mathcal{P}^0_d(T') \setminus \mathcal{P}^0_d(T)$. Then there exists a point p such that either $p \succ_T q$ or $p \sim_T q$. Since $T' \subset T$, in the second case we obtain $p \sim_{T'} q$, while in the first case we have $p \succ q$ or $p \sim_{T'} q$, a contradiction. □

For $p \in X$ and $t_i \in T$, set $r_i^p := d(p, t_i)$. The points of $\mathcal{P}^0_d(T)$ can be easily characterized in the following way:

LEMMA 3. *A point $p \in X$ belongs to $\mathcal{P}^0_d(T)$ if and only if $\cap_{i=1}^n B(t_i, r_i^p) = \{p\}$.*

Now, consider Pareto optimality for distances induced by norms. In this case, $d(p, x)$ is a convex function of variable x, therefore the following useful fact holds:

LEMMA 4. *If $p, q \in \mathbb{R}^m$ and $p \succ_T q$, then $p' \succ_T q$ for any point $p' \in [p, q]$.*

3. ENVELOPES $\mathcal{P}_{d_\infty}(T)$ AND $\mathcal{P}^0_{d_\infty}(T)$

3.1 $\mathcal{P}_{d_\infty}(T)$

Let $B_{d_\infty}(p, 1)$ be the unit ball centered at a point p (as we notices already, $B_{d_\infty}(p, 1)$ is an axis-parallel cube). For each coordinate $i \in \{1, \ldots, m\}$, let $F^i_+(p) = \{x \in B_{d_\infty}(p, 1) : x^i = p^i + 1\}$ and $F^i_-(p) = \{x \in B_{d_\infty}(p, 1) : x^i = p^i - 1\}$ be the two opposite facets of the cube $B_{d_\infty}(p, 1)$ which are orthogonal to the ith coordinate line. Denote by $C^i_+(p)$ and $C^i_-(p)$ the cones having their apex at the point p and induced by $F^i_+(p)$ and $F^i_-(p)$. They can be equally defined as

$$C^i_+(p) = \{x \in \mathbb{R}^m : x^i \ge p^i \text{ and } d_\infty(p, x) = |p^i - x^i|\},$$
$$C^i_-(p) = \{x \in \mathbb{R}^m : x^i \le p^i \text{ and } d_\infty(p, x) = |p^i - x^i|\}.$$

Using these cones, we recall the construction of intervals in $(\mathbb{R}^m, d_\infty)$ resulting from a general description of intervals in all normed spaces given by Boltyanski and Soltan:

LEMMA 5. *[6] For two points $p, q \in \mathbb{R}^m$, if $d_\infty(p, q) = p^i - q^i$ $(p^i \ge q^i)$, then $I_{d_\infty}(p, q) = C^i_-(p) \cap C^i_+(q)$. In particular, $x \in I_{d_\infty}(p, q)$ if and only if $p \in C^i_+(x)$ and $q \in C^i_-(x)$.*

A metric space (X, d) is called *hyperconvex* (or *injective*) [1, 19] provided that any family of closed balls $B(x_i, r_i)$ with centers x_i and radii r_i, $i \in I$, satisfying $d(x_i, x_j) \le r_i + r_j$ for all $i, j \in I$ has a nonempty intersection, that is, (X, d) is a *Menger-convex space* (i.e., for any two distinct points x and y there exists another point z between x and y) such that the closed balls have the (infinite) *Helly property* (which means that every family of closed balls that pairwise intersect has a nonempty intersection). The metric space $(\mathbb{R}^m, d_\infty)$ is hyperconvex because the closed balls are axis-parallel cubes, and they obviously satisfy the Helly property.

We continue with the characterization of $\mathcal{P}_{d_\infty}(T)$ in $\mathbb{R}^m$ which generalizes the result of [7]. Since the proof uses only the Helly property for balls, a similar characterization holds in all hyperconvex metric spaces (X, d).

THEOREM 1. *$\mathcal{P}_{d_\infty}(T) = \Upsilon_{d_\infty}(T)$ for any $T \subset \mathbb{R}^m$.*

PROOF. The inclusion $\Upsilon_{d_\infty}(T) \subseteq \mathcal{P}_{d_\infty}(T)$ follows from Lemma 1. It remains to show that $q \notin \Upsilon_{d_\infty}(T)$ implies $q \notin \mathcal{P}_{d_\infty}(T)$. Let $T = \{t_1, \ldots, t_n\}$. Since the point q does not belong to $\Upsilon_{d_\infty}(T)$, there exists a terminal t_i such that $q \notin \cup_{j=1}^n I_{d_\infty}(t_i,t_j)$, i.e. $d_\infty(t_i,t_j) < d_\infty(t_i,q)+d_\infty(q,t_j)$ for each terminal t_j. This means that $\varepsilon := \min\{d_\infty(t_i,q)+d_\infty(q,t_j) - d_\infty(t_i,t_j)\} > 0$. For each $t_k \in T$, set $r_k := d_\infty(t_k,q)$ if $t_k \neq t_i$ and $r_k := d(t_k,q) - \varepsilon$ if $t_k = t_i$. Let $B = \cap_{k=1}^n B_{d_\infty}(t_k,r_k)$. We assert that B is non-empty and that the point q is dominated by any point from this intersection. Since q belongs to the ball $B_{d_\infty}(t_k,r_k)$ for each $k \neq i$, we conclude that the balls $B_{d_\infty}(t_j,r_j)$ and $B(t'_j,r_{j'})$ intersect for all indices $j,j' \neq i$. For each t_j, the balls $B_{d_\infty}(t_i,r_i)$ and $B_{d_\infty}(t_j,r_j)$ equally intersect because $r_i+r_j = d_\infty(t_i,q)-\varepsilon+d_\infty(t_j,q) = d_\infty(t_i,q)+d_\infty(q,t_j)-\varepsilon$ and r_i+r_j is larger or equal than $d_\infty(t_i,t_j)$ because $d_\infty(t_i,q)+d_\infty(q,t_j)-\varepsilon \geq d_\infty(t_i,q)+d_\infty(q,t_j)-(d_\infty(t_i,q)+d_\infty(q,t_j)-d_\infty(t_i,t_j)) \geq d_\infty(t_i,t_j)$. Since $(\mathbb{R}^m,d_\infty)$ is Menger-convex and its balls satisfy the Helly property and since the balls defining the intersection B pairwise intersect, we deduce that $B \neq \varnothing$. But q does not belong to the ball $B_{d_\infty}(t_i,r_i)$, yielding $q \notin B$, whence q is dominated by any point p of B. □

A set S is called *empty* if $S \cap T = \varnothing$. Using Lemma 5, we can re-phrase Theorem 1 in the following way.

COROLLARY 1. *A point $p \in \mathbb{R}^m$ belongs to $\mathcal{P}_{d_\infty}(T)$ iff whenever a terminal belongs to the interior of a cone $C_+^i(p)$ or $C_-^i(p)$ then the opposite cone is non-empty and whenever a terminal belongs to the boundary of several such cones, then at least one of respective opposite cones is non-empty.*

3.2 $\mathcal{P}^0_{d_\infty}(T)$

To characterize the strict Pareto envelope $\mathcal{P}^0_{d_\infty}(T)$ we introduce two sets $\mathcal{C}(T)$ and $\mathcal{U}(T)$. For a point $p \in \mathbb{R}^m$, let $\mathcal{C}(p)$ denote the union of all non-empty cones of the form $C_+^i(p)$ or $C_-^i(p)$. Then $\mathcal{C}(T) = \cap\{\mathcal{C}(p) : p \in \mathbb{R}^m\}$. For each index $i \in \{1,\ldots,m\}$, set $\mathcal{U}_+^i(T) = \cup_{j=1}^n C_+^i(t_j)$ and $\mathcal{U}_-^i(T) = \cup_{j=1}^n C_-^i(t_j)$. Then the set $\mathcal{U}(T)$ is defined as the intersection of these $2m$ sets: $\mathcal{U}(T) = \cap_{i=1}^m(\mathcal{U}_+^i(T) \cap \mathcal{U}_-^i(T))$.

THEOREM 2. *$\mathcal{P}^0_{d_\infty}(T) = \mathcal{U}(T) = \mathcal{C}(T)$ for any $T \subset \mathbb{R}^m$.*

PROOF. First we show that $\mathcal{U}(T) \subseteq \mathcal{C}(T)$. Suppose by way of contradiction that q belongs to $\mathcal{U}(T) \setminus \mathcal{C}(T)$. Then there exists a point p such that $q \notin \mathcal{C}(p)$. We can suppose without loss of generality that there exists an index $i \in \{1,\ldots,m\}$ so that $q \in C_+^i(p)$ and $C_+^i(p) \cap T = \varnothing$. In this case, the cone $C_+^i(q)$ is equally empty. Then $q \notin \mathcal{U}_-^i(T)$, because all cones $C_-^i(t_j)$ containing the point q have their apex in the cone $C_+^i(q)$ and the latter cone is empty. This contradicts the assumption that q belongs to $\mathcal{U}(T)$. To establish that $\mathcal{C}(T) \subseteq \mathcal{U}(T)$, we show if $p \notin \mathcal{U}(T)$, then $p \notin \mathcal{C}(T)$. Indeed, if $p \notin \mathcal{U}(T)$, then necessarily at least one of the $2m$ cones centered at p, say the cone $C_+^i(p)$, is empty. Then, decreasing appropriately the ith coordinate of the point p, we can derive another point p' so that $C_+^i(p') \cap T = \varnothing$. Since p belongs to the interior of the cone $C_+^i(p')$, we deduce that $p \notin \mathcal{C}(T)$.

To establish that $\mathcal{U}(T) \subseteq \mathcal{P}^0_{d_\infty}(T)$, by Lemma 3 it suffices to show that if $p \in \mathcal{U}(T)$, then the intersection $B := \cap_{i=1}^n B_{d_\infty}(t_i,r_i^p)$ consists of the point p only. Since $p \in \mathcal{U}(T)$, for each $i \in \{1,\ldots,m\}$ there exist two terminals t_j,t_k so that $t_j \in C_+^i(p)$ and $t_k \in C_-^i(p)$. Then $p \in I_{d_\infty}(t_j,t_k)$, thus $p \in F_-^i(t_j) \cap F_+^i(t_k) \subset \{x \in \mathbb{R}^m : x^i = p^i\}$. Since the intersection of the m hyperplanes $\{x \in \mathbb{R}^m : x^i = p^i\}$ coincides with the point p, we conclude that the required intersection B coincides with p as well, establishing that $\mathcal{U}(T) \subseteq \mathcal{P}^0_{d_\infty}(T)$.

Finally, we show that $\mathcal{P}^0_{d_\infty}(T) \subseteq \mathcal{U}(T)$. For this, we prove that if $p \notin \mathcal{U}(T)$, then $p \notin \mathcal{P}^0_{d_\infty}(T)$. Indeed, since $p \notin \mathcal{U}(T)$, we can suppose that there exists $i \in \{1,\ldots,m\}$ so that $p \notin \mathcal{U}_i^-(T)$. This means that the cone $C_+^i(p)$ is empty. This implies that the intersection of the closed ball $B_{d_\infty}(t_j,r_j^p)$ with the line L parallel to the ith axis and passing via p is a segment containing p and a point $q \in C_-^i(p) \setminus \{p\}$. Hence the intersection of the balls $B_{d_\infty}(t_1,r_1^p),\ldots,B_{d_\infty}(t_n,r_n^p)$ with L is also a segment containing p and yet another point $q' \in C_-^i(p) \setminus \{p\}$. Lemma 3 shows that $p \notin \mathcal{P}^0_{d_\infty}(T)$, thus the inclusion $\mathcal{P}^0_{d_\infty}(T) \subseteq \mathcal{U}(T)$ is verified. □

For each $i = 1,\ldots,m$, denote by $\mathcal{L}^i$ the set of all lines of $\mathbb{R}^m$ passing via the origin o and contained in the cones $C_+^i(o)$ and $C_-^i(o)$. Then the sets $\mathcal{U}_+^i(T)$ and $\mathcal{U}_-^i(T)$ and their intersection $\mathcal{U}_+^i(T) \cap \mathcal{U}_-^i(T)$ are $\mathcal{L}^i$-convex (recall, that for a set of lines $\mathcal{L}$, a set S is *$\mathcal{L}$-convex* [17] if the intersection of S with any line parallel to a line of $\mathcal{L}$ is convex). Since the intersection $S' \cap S''$ of an $\mathcal{L}'$-convex set S' with an $\mathcal{L}''$-convex set S'' is $\mathcal{L}' \cap \mathcal{L}''$-convex, from Theorem 2 we conclude that $\mathcal{P}^0_{d_\infty}(T)$ is $\mathcal{L}^*$-convex, where $\mathcal{L}^*$ consists of the 2^{m-1} lines $l^1, l^2, \ldots$ defined by the opposite corners of the unit l_∞-ball centered at the origin (they are also the lines defining the extremal rays of the cones from each of $2m$ families).

3.3 Pareto envelopes and injective hulls

In this subsection, we establish a close relationship between $\mathcal{P}_{d_\infty}(T), \mathcal{P}^0_{d_\infty}(T)$ and the injective hulls of metric spaces defined by Isbell [19], Dress [13], and Chrobak and Larmore [10]. Given a finite metric space (X,d) with $X = \{x_1,\ldots,x_n\}$, denote by $\mathcal{K}_d(X)$ the set of all points u of $\mathbb{R}^n$ so that $|u^i - u^j| \leq d(p_i,p_j) \leq u^i + u^j$. $\mathcal{K}_d(X)$ is an unbounded convex polyhedron of $\mathbb{R}^n_+$. This polyhedron endowed with the distance d_∞ is called the *extension* of the metric space (X,d) because the map $\phi : X \mapsto \mathbb{R}^n$ defined by $\phi(x) = (d(x,x_1),\ldots,d(x,x_n))$ is an isometric embedding (i.e., $d(x_i,x_j) = d_\infty(\phi(x_i),\phi(x_j))$ for all $i,j \in \{1,\ldots,n\}$). For $u,v \in \mathcal{K}_d(X)$, let $u \preccurlyeq v$ if $u^i \leq v^i$ for all $i \in \{1,\ldots,n\}$. The set of all minimal points of $\mathcal{K}_d(X)$ with respect to the partial order $\preccurlyeq$ is denoted by $\mathcal{IH}(X,d)$. It was shown in [13] that $\mathcal{IH}(X,d)$ consists exactly of all compact faces of the polyhedron $\mathcal{K}_d(X)$. The set $\mathcal{IH}(X,d)$ endowed with the metric d_∞ is called the *injective hull,* the *tight span,* or the *hyperconvex hull* of the metric space (X,d) [10, 13, 19]. This metric space is hyperconvex [13, 19], in fact, it is the smallest hyperconvex space into which (X,d) embeds isometrically. By a classical result of Aronszajn and Panitchpakdi [1], hyperconvex spaces are exactly the injective metric spaces. A map f from a metric space (Y,d') to a metric space (Z,d'') is *non-expansive* if $d''(f(x),f(y)) \leq d'(x,y)$ for all $x,y \in Y$. A metric space (Z,d'') is *injective* if, for each metric space (Y,d') and each subspace $Y' \subset Y$, if there exists a non-expansive map f' from (Y',d') to (Z,d''), then f' extends to a non-expansive map f from (Y,d') to (Z,d'').

Now, we relate the Pareto envelopes $\mathcal{P}_{d_\infty}(T)$ and $\mathcal{P}^0_{d_\infty}(T)$ with the injective hull $\mathcal{T}(T,d_\infty)$ of the metric space (T,d_∞). Since the metric space $(\mathbb{R}^m,d_\infty)$ is injective, it contains the injective hull $\mathcal{IH}(T,d_\infty)$ of (T,d_∞) as a subspace. The identity map $id : \mathcal{IH}(T,d_\infty) \mapsto \mathcal{IH}(T,d_\infty)$ is non-expansive, therefore id can be extended to a non-expansive map α from $(\mathbb{R}^m,d_\infty)$ to $(\mathcal{IH}(T,d_\infty),d_\infty)$.

PROPOSITION 1. *For any finite subset T of $\mathbb{R}^m$, we have $\mathcal{P}^0_{d_\infty}(T) \subseteq \mathcal{IH}(T,d_\infty) \subseteq \mathcal{P}_{d_\infty}(T)$. Moreover, $\mathcal{P}^0_{d_\infty}(T) = \{\alpha^{-1}(u) : u \in \mathcal{IH}(T,d_\infty)$ and $|\alpha^{-1}(u)| = 1\}$.*

PROOF. First notice that $\mathcal{IH}(T, d_\infty) \subseteq \mathcal{P}_{d_\infty}(T)$: indeed according to a result of [13, 19], a point $p \in \mathcal{K}_d(X)$ belongs to $\mathcal{IH}(T, d_\infty)$ if and only if for any terminal $t_i \in T$ there exists a terminal $t_j \in T$ such that p lies between (in the d_∞-metric) $\phi(t_i)$ and $\phi(t_j)$. Thus $\mathcal{IH}(T, d_\infty) \subseteq \Upsilon_{d_\infty}(T)$. Since $\Upsilon_{d_\infty}(T) = \mathcal{P}_{d_\infty}(T)$ by Theorem 1, we obtain the required inclusion. For each point $p \in \mathbb{R}^m$, we have $d_\infty(\alpha(p), \alpha(t_i)) \leq d_\infty(p, t_i)$ for any terminal $t_i \in T$. If at least one of these inequalities is strict, then $\alpha(p) \succ p$ and therefore $p \notin \mathcal{P}_{d_\infty}(T)$. Otherwise, we obtain $d_\infty(\alpha(p), t_i) = d_\infty(p, t_i)$ for all terminals $t_i \in T$, thus p and $\alpha(p)$ have the same distance vector. Since $\mathcal{I}_{d_\infty}(T) \subseteq \mathcal{P}_{d_\infty}(T)$, we deduce that the point p also belongs to $\mathcal{P}_{d_\infty}(T)$. Hence $p \in \mathcal{P}^0_{d_\infty}(T)$ if and only if $\alpha(p) = p$, i.e., if $p \in \mathcal{I}_{d_\infty}(T)$ and $\alpha^{-1}(p) = \{p\}$. $\square$

Proposition 3.3 can be used to visualize the injective hulls, in particular, the injective hulls of metric subspaces of $(\mathbb{R}^3, d_\infty)$; see Fig. 2 for an illustration. In general, the injective hull of a finite metric space (X, d) may have quite a sophisticated shape. It was shown in [13] that if $|X| = 3$, then the injective hull consists of three segments glued along a common end-point, if $|X| = 4$, then it consists of a rectangle and four segments glued to the corners of the rectangle and that if $|X| = 5$, then $\mathcal{IH}(X, d)$ has one of three canonical forms (each of them is two-dimensional). More recently, Sturmfels and Yu [29] established that there exist 339 canonical forms of the injective hulls of metric spaces on 6 points.

4. ENVELOPE $\mathcal{P}_{d_1}(T)$ IN $\mathbb{R}^3$

In this section, we establish that for any set of terminals T of $\mathbb{R}^3$, the Pareto envelope $\mathcal{P}_{d_1}(T)$ coincides with the cubical polyhedron $\mathcal{M}(T)$ induced by the median closure $\text{med}(T)$ of T and with the cubical polyhedron $\mathcal{I}(T)$ defined below.

For a set T of $\mathbb{R}^3$ and $i \in \{1, 2, 3\}$, denote by T^i the orthogonal projection of T on the hyperplane $H^i = \{x \in \mathbb{R}^3 : x^i = 0\}$. Let $\mathcal{P}_{d_1}(T^i)$ be the Pareto envelope of the set T^i in the plane H^i. Finally, let $\mathcal{I}^i(T)$ denote the Cartesian product of $\mathcal{P}_{d_1}(T^i)$ with the line l^i orthogonal to the plane H^i. Then $\mathcal{I}(T)$ is defined as the intersection of the three "cylindrical" sets $\mathcal{I}^1(T), \mathcal{I}^2(T)$, and $\mathcal{I}^3(T)$; for an illustration, see Fig. 3. In the same way, we can define $\mathcal{I}(T)$ for any set T of $\mathbb{R}^m$.

The *median* of three points $u, v, w \in \mathbb{R}^m$ is a point $m(u, v, w) = m \in \mathbb{R}^m$ such that m^i is the median value of the triplet u^i, v^i, and w^i for all $i \in \{1, \ldots, m\}$. The median $m(u, v, w)$ is the unique point of $\mathbb{R}^m$ which belongs to the intervals $I_{d_1}(u, v)$, $I_{d_1}(u, w)$, and $I_{d_1}(v, w)$. A subset $S \subset \mathbb{R}^m$ is called *median* (or median stable) iff $m(u, v, w) \in S$ for each triplet $u, v, w \in S$. Since the intersection of median stable sets is median stable, for a set T of $\mathbb{R}^m$ one can always define the smallest median stable set $\text{med}(T)$ containing T (the *median closure* of T). If T is finite, then $\text{med}(T)$ is finite as well [31]. Moreover, if $T \subset \mathbb{R}^2$, then $\text{med}(T)$ consists of T and the medians of all triplets of T. If $T \subset \mathbb{R}^3$, then to obtain $\text{med}(T)$, one extra-iteration suffices [33]. With $\text{med}(T)$, one can associate a median graph $G(T)$ [2, 31] having $\text{med}(T)$ as a vertex set and $u, v \in \text{med}(T)$ define an edge of G iff $I_{d_1}(u, v) \cap \text{med}(T) = \{u, v\}$. There is a standard way to derive a cubical complex from the median graph $G(T)$: replace every graphic cube by a solid box of the same dimension. The resulting cubical polyhedron (which is a median space itself) is denoted by $||\text{med}(T)||$. Endowed with intrinsic l_1, l_2, or l_∞ metric, $||\text{med}(T)||$ has many nice characteristic properties; for more details on this construction and its properties, see [2, 8, 31]. In particular, $||\text{med}(T)||$ is a median isometric subspace

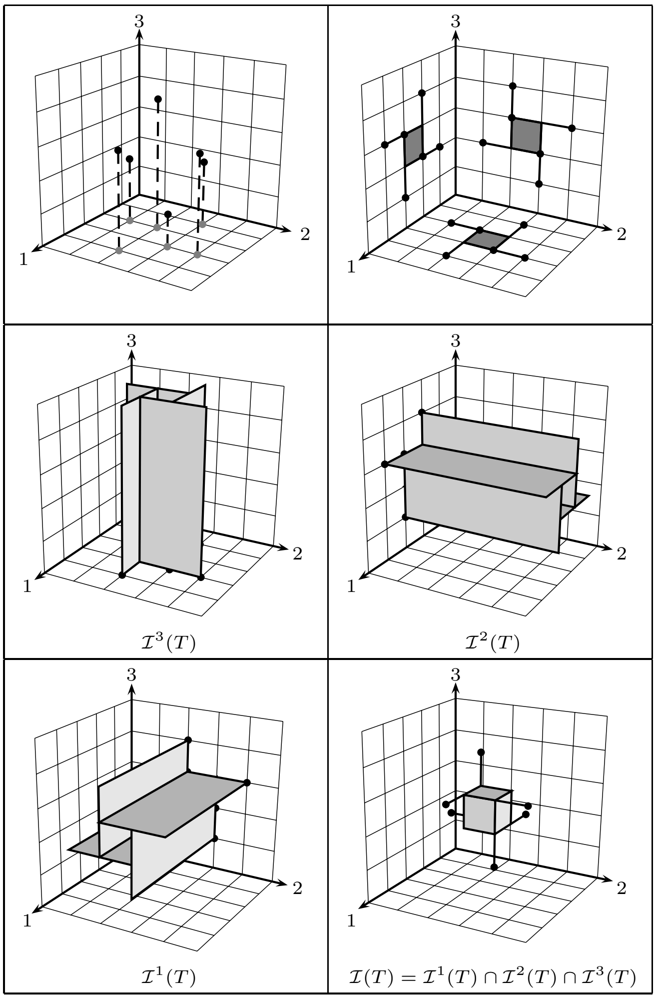

Figure 3: The set $\mathcal{I}(T)$ in $\mathbb{R}^3$

of $(\mathbb{R}^m, d_1)$. We will slightly modify this construction to obtain $\mathcal{M}(T)$.

Given $T \subset \mathbb{R}^m$, denote by $\Gamma(T)$ the grid obtained by drawing for each terminal t_i the m axis-parallel planes $H^1(t_i), \ldots, H^m(t_i)$ passing via t_i. It can be easily seen that the vertex-set of $\Gamma(T)$ contains $\text{med}(T)$. For $T \subset \mathbb{R}^3$, the cell complex $\mathcal{M}(T)$ consists of certain boxes, rectangles, and edges of the grid $\Gamma(T)$. If we consider a 3-cube F of the underlying median graph $G(T)$, then all its vertices are corners of a solid box $Q(F)$ of the grid $\Gamma(T)$. We replace the graphic cube F by $Q(F)$ and add it to $\mathcal{M}(T)$ (notice that $Q(F)$ also belongs to the polyhedron $||\text{med}(T)||$). If F is a square of $G(T)$, then the four vertices of this square are either corners of a rectangular face $R(F)$ of Γ or they define two opposite sides of a box $Q(F)$ of $\Gamma(T)$. In the first case, we replace F by the rectangle $R(F)$ and in the second case we replace F by the box $Q(F)$; in both cases, we insert them in $\mathcal{M}(T)$. Finally, if an edge $e = xy$ of $G(T)$ is an edge of $\Gamma(T)$, then we add the segment $[x, y]$ to $\mathcal{M}(T)$. Otherwise, if e is a diagonal of a 2-dimensional or 3-dimensional box $Q(e)$ of Γ, then we add $Q(e)$ to $\mathcal{M}(T)$. In analogy to the cubical complex $||\text{med}(T)||$, it can be shown that the cell complex $\mathcal{M}(T)$ is a median subspace of $\mathbb{R}^3$; moreover, $\mathcal{M}(T)$ is l_1-isometric, in the sense that any two points $p, q \in \mathcal{M}(T)$ can be connected inside $\mathcal{M}(T)$ by a path of length $d_1(p, q)$; for more details see the van de Vel's book [31]. The rest of this section is devoted to the proof of the following result (due to space limitations, some details are deferred to the full version; see also [24]).

THEOREM 3. *$\mathcal{P}_{d_1}(T) = \mathcal{M}(T) = \mathcal{I}(T)$ for any $T \subset \mathbb{R}^3$.*

A similar characterization is not longer true in higher dimensions, in particular in $\mathbb{R}^4$. Let $p = (1,1,1,1)$ and $q = (0,0,0,0)$ be two opposite vertices of the unit 4-cube $Q_4 = [0,1]^4$. Let T consist of all vertices t of Q_4 such that $d_1(p,t) \le d_1(q,t)$. For this set of terminals, med(T) consists of all corners of the cube Q_4, therefore $\mathcal{M}(T) = [0,1]^4$. The set $\mathcal{I}(T)$ is also the whole cube because the Pareto envelopes of the projections of T on four coordinate hyperplanes are all 3-dimensional cubes. However, the point q does not belongs to $\mathcal{P}_{d_1}(T)$, because q is dominated by p. Thus $\mathcal{P}_{d_1}(T) \ne \mathcal{M}(T)$ and $\mathcal{P}_{d_1}(T) \ne \mathcal{I}(T)$.

Notice also that, analogously to Theorem 3, one can characterize the strict envelope $\mathcal{P}^0_{d_1}(T)$. Then $\mathcal{I}^0(T)$ is defined as $\mathcal{I}(T)$, only instead of Pareto envelopes $\mathcal{P}_{d_1}(T^i)$ we should consider the strict envelopes $\mathcal{P}^0_{d_1}(T^i)$, $i = 1,2,3$. On the other hand, $\mathcal{M}^0(T)$ is the union of the boxes of $\mathcal{M}(T)$ for which all corners belong to med(T). Then one can show that $\mathcal{P}^0_{d_1}(T) = \mathcal{M}^0(T) = \mathcal{I}^0(T)$ holds.

4.1 Auxiliary results

In this subsection we establish the notation and some auxiliary results used in the proof of Theorem 3. Throughout this section, by a hyperplane of $\mathbb{R}^m$ we will mean an axis parallel hyperplane. First we recall the notion of "gate" which for l_1-spaces plays the same role as the metric projection for Euclidean spaces. A subset S of $(\mathbb{R}^m, d_1)$ is called *gated* [31] if any point $p \notin S$ contains a (necessarily unique) *gate* in S, i.e., a point $p' \in S$ such that $p' \in I_{d_1}(p,q)$ for any $q \in S$. Closed half-spaces of $\mathbb{R}^m$ defined by axis-parallel hyperplanes and their intersections are basic examples of gated sets of $(\mathbb{R}^m, d_{l_1})$; cf. [31]. Axis-parallel boxes and orthants are particular instances of such sets. It can be easily shown that if a gated set Π contains the set T, then any point p outside Π is dominated by its gate in Π. This gives the first "approximation" for $\mathcal{P}_{d_1}(T)$.

LEMMA 6. *$\mathcal{P}_{d_1}(T)$ is contained in the smallest axis-parallel box $\Pi(T)$ spanned by T.*

For an axis-parallel box Π of $\mathbb{R}^m$ (not necessarily full-dimensional) and a corner p of Π, the set of all points q of $\mathbb{R}^m$ such that the gate of q in Π is the point p, is called the *orthant* of p with respect to Π and is denoted by Ort(Π, p). Then $p \in I_{d_1}(q,x)$ for any $q \in$ Ort(Π,p) and $x \in \Pi$. Notice also that the orthant Ort(Π,p) is gated and p is the gate in Ort(Π,p) of all points $x \in \Pi$. Geometrically, in $\mathbb{R}^3$, Ort(Π,p) is either a closed halfspace if Π is 1-dimensional, or a closed wedge (the intersection of two closed halfspaces) if Π is 2-dimensional, or the intersection of three closed pairwise crossing halfspaces if Π is 3-dimensional.

LEMMA 7. *Let $p,q \in \mathbb{R}^m$ and set $\Pi := I_{d_1}(p,q)$. If $p \succ q$, then* Ort$(\Pi,q) \cap \mathcal{P}_{d_1}(T) = \varnothing$.

PROOF. The proof uses the following easily verified assertion: *if a,b,δ are real numbers such that $a \ge b$ and $\delta \le 0$, then* $|a| - |b| \ge |a+\delta| - |b+\delta|$.

Pick $q' \in$ Ort(Π,q) and suppose without loss of generality that $q'^i \le q^i \le p^i$ for each index i. We assert that q' is dominated by the point p' obtained from p by translation $\overrightarrow{qq'}$ (i.e., $p'^i = p^i + (q'^i - q^i)$ for each i). The auxiliary assertion yields $|p^i - t^i| - |q^i - t^i| \ge |p^i - t^i + (q'^i - q^i)| - |q^i - t^i + (q'^i - q^i)| \ge |p'^i - t^i| - |q'^i - t^i|$. Hence, for each terminal $t \in T$ we have $d_1(p,t) - d_1(q,t) = \sum_{i=1}^m |p^i - t^i| - \sum_{i=1}^m |q^i - t^i| = \sum_{i=1}^m (|p^i - t^i| - |q^i - t^i|) \ge \sum_{i=1}^m (|p'^i - t^i| - |q'^i - t^i|) \ge \sum_{i=1}^m |p'^i - t^i| - \sum_{i=1}^m |q'^i - t^i| \ge d_1(p',t) - d_1(q',t)$. Since $p \succ q$, we infer that $p' \succ q'$. □

From this result one can easily conclude that $\mathcal{P}_{d_1}(T)$ is ortho-convex in the sense that the intersection of $\mathcal{P}_{d_1}(T)$ with any line parallel to a coordinate axis is convex. This generalizes an analogous property of planar l_1-envelopes established by [7, 35].

Given two sets $T, S \subset \mathbb{R}^m$, $|T| = |S|$, we say that S is obtained from T via a coordinate-wise isotonic map if $t^k_i < t^k_j$ iff $s^k_i < s^k_j$ and $t^k_i = t^k_j$ iff $s^k_i = s^k_j$. In particular, consider the map φ which associate to each terminal $t_i \in T$ the vertex $\varphi(t_i)$ of the uniform grid $\mathbb{Z}^m$ so that the kth coordinate of $\varphi(t_i)$ is the rank of t^k_i in the sorted list of the pairwise distinct values of kth coordinates of the terminals of T. Notice that $\varphi(T)$ can be constructed in $O(mn \log n)$ time. Then φ can be extended accordingly to a map from the grid $\Gamma(T)$ to the uniform grid defined by $\mathbb{Z}^m$.

LEMMA 8. $\mathcal{P}_{d_1}(\varphi(T)) = \varphi(\mathcal{P}_{d_1}(T))$.

PROOF. It suffices to show that if $q \notin \mathcal{P}_{d_1}(T)$, then $\varphi(q) \notin \varphi(\mathcal{P}_{d_1}(T))$ (the converse inclusion can be proven in the same way). Consider a translation of $\mathbb{R}^m$ which identifies the points q and $\varphi(q)$. Denote by S the image of $\varphi(T)$ under this translation. Let $p \succ_T q$. By Lemma 4 we can select the point p enough close to q so that the box $\Pi := I_{d_1}(p,q)$ satisfies the condition: for each $t \in (T \cup S)$, there exists a corner c of Π such that $t \in$ Ort(Π,c). We assert that $p \succ_S q$. Pick $t_i \in T$ and let s_i be its image in S. By definition of φ and Π, we conclude that there exists a corner c of Π such that $t_i, s_i \in$ Ort(Π,c). Since $d_1(p,x) - d_1(q,x) = d_1(p,c) - d_1(q,c)$ for any point $x \in$ Ort(Π,s), we deduce that $d_1(p,t_i) - d_1(q,t_i) = d_1(p,s_i) - d_1(q,s_i)$. Thus $p \succ_S q$, and therefore the pre-image of p under the translation will dominate the point $\varphi(q)$. □

We call a set of terminals T *uniform* if $\varphi(T) = T$. In view of Lemma 8, further we can assume without loss of generality that T is uniform. Then med$(T) \subset \mathbb{Z}^m$ and $||$med$(T)||$, $\mathcal{M}(T)$, and $\mathcal{I}(T)$ are all three cubical complexes. We call a finite cubical complex $\mathcal{C}$ with vertices in $\Gamma(T)$ *conformal* if a face Π of $\Gamma(T)$ belongs to $\mathcal{C}$ whenever all corners of Π belongs to $\mathcal{C}$. The cubical complexes $||$med$(T)||$ and $\mathcal{M}(T)$ are conformal by definition. Now, since $\mathcal{P}_{d_1}(T) = \mathcal{M}(T)$ holds in $\mathbb{R}^2$ by Proposition 2, we can easily deduce that $\mathcal{I}(T)$ in $\mathbb{R}^3$ is conformal as well. In fact $\mathcal{I}(T)$, $\mathcal{M}(T)$, and $\mathcal{P}_{d_1}(T)$ (see the next result) are conformal in all dimensions.

LEMMA 9. *For $T \subset \mathbb{R}^m$, the Pareto envelope $\mathcal{P}_{d_1}(T)$ is a conformal cubical complex of $\Gamma(T)$.*

PROOF. Let Π be a cell of $\Gamma(T)$ whose all corners belong to $\mathcal{P}_{d_1}(T)$. Suppose by way of contradiction that a point $q \in \Pi$ is dominated by a point p. Let $\Pi' := I_{d_1}(p,q)$. At least one corner c of Π belongs to Ort(Π', q). By Lemma 7, $c \notin \mathcal{P}_{d_1}(T)$, contrary to our choice of Π. □

The following result (given here without proof) presents a discrete analogous of local domination provided by Lemma 4.

LEMMA 10. *If T is an uniform set of $\mathbb{R}^3$, then a vertex $q \in \Gamma(T)$ belongs to $\mathcal{P}_{d_1}(T)$ if and only if q is not dominated by opposite to q corners of the 2-dimensional faces of $\Gamma(T)$ incident to q.*

We continue with the following characterization of med(T).

Lemma 11. *[31] For $T \subset \mathbb{R}^m$, a point p belongs to med(T) if and only if for each pair of halfspaces H_1, H_2 of $\mathbb{R}^m$ containing p, we have $T \cap (H_1 \cap H_2) \neq \varnothing$.*

Proposition 2. *$\mathcal{P}_{d_1}(T) = \mathcal{M}(T) = \Upsilon_{d_1}(T)$ for $T \subset \mathbb{R}^2$.*

Proof. The equality $\mathcal{P}_{d_1}(T) = \Upsilon_{d_1}(T)$ is a result of [7]; see also Theorem 1. The inclusion $\mathcal{M}(T) \subseteq \mathcal{P}_{d_1}(T)$ will be established below for $\mathbb{R}^3$. It remains to show that $\mathcal{P}_{d_1}(T) \subseteq \mathcal{M}(T)$. Since the complexes $\Upsilon_{d_1}(T)$ and $\mathcal{P}_{d_1}(T)$ are conformal, to show that $\Upsilon_{d_1}(T) = \mathcal{P}_{d_1}(T) \subseteq \mathcal{M}(T)$ it suffices to show that every $p \in \Upsilon_{d_1}(T) \setminus \text{med}(T)$ which is a vertex of $\Gamma(T)$ belongs to $\mathcal{M}(T)$. Then the vertical line l^v and the horizontal line l^h passing via p each contain a terminal t_i and t_j, respectively. Suppose that p is below t_i and to the left of t_j. The quadrant $\mathcal{Q}_1 = \{q \in \mathbb{R}^2 : q^1 \geq p^1 \text{ and } q^2 \geq p^2\}$ contains the terminals t_i and t_j. If the quadrant $\mathcal{Q}_3$ opposite to $\mathcal{Q}_1$ is nonempty as well, then for any terminal t_k located in $\mathcal{Q}_3$ we have $m(t_i, t_j, t_k) = p$, thus $p \in \text{med}(T)$. Otherwise, if $\mathcal{Q}_3$ is empty, then Corollary 1 implies that the terminals located in $\mathcal{Q}_1$ all belong to its boundary. Thus, all terminals of T are located in the second and fourth quadrants. Let Π denote the rectangle of $\Gamma(T)$ located in the first quadrant and having the corners p, a, b, c. The definition of $\Gamma(T)$ and the emptiness of interiors of $\mathcal{Q}_1$ and $\mathcal{Q}_3$ implies that the second quadrant contains a terminal t_l and the fourth quadrant contains a terminal t_k on the same horizontal and vertical lines as b. Then $a = m(t_l, t_i, t_j)$ and $c = m(t_j, t_k, t_i)$. Since all terminals are located in $\mathcal{Q}_2 \cup \mathcal{Q}_4$, the corners p and b are not medians of triplets of terminals. Thus Π and its corners belong to $\mathcal{M}(T)$. □

4.2 $\mathcal{P}_{d_1}(T) = \mathcal{I}(T)$

First we prove that $\mathcal{P}_{d_1}(T) \subseteq \mathcal{I}(T)$ (this inclusion holds in all dimensions). For this, it suffices to show that $q \notin \mathcal{P}_{d_1}(T)$ whenever there exists an index $i \in \{1, 2, 3\}$ such that the orthogonal projection of q on H^i does not belong to $\mathcal{P}_{d_1}(T^i)$. Suppose without loss of generality that $i = 3$, i.e., $q' = (q^1, q^2) \notin \mathcal{P}_{d_1}(T^3)$. Suppose that q' is dominated by a point $p' = (p^1, p^2)$ of H^i. We assert that q is dominated by the point $p = (p^1, p^2, q^3)$. Since $p^3 = q^3$, for any terminal $t_i \in T$ we have $d_1(p, t_i) - d_1(q, t_i) = |t_i^1 - p^1| + |t_i^2 - p^2| - |t_i^1 - q^1| - |t_i^2 - q^2|$. Therefore $p \succ_T q$ if and only if $p' \succ_{T^3} q'$, yielding $q \notin \mathcal{P}_{d_1}(T)$.

Conversely, let $q \notin \mathcal{P}_{d_1}(T)$. We assert that $q \notin \mathcal{I}(T)$. Since $\mathcal{P}_{d_1}(T)$ and $\mathcal{I}(T)$ are conformal cubical complexes, we can suppose that q is a vertex of the (uniform) grid $\Gamma(T)$. By Lemma 10, q is dominated by a point p such that p and q are opposite corners of a 2-dimensional face of $\Gamma(T)$. Thus p and q differ in exactly two coordinates, say $p^3 = q^3$. Let $q' = (q^1, q^2)$ and $p' = (p^1, p^2)$. For any terminal t_i we have $|t_i^3 - p^3| = |t_i^3 - q^3|$, therefore, if we denote $t_i' = (t_i^1, t_i^2)$, we obtain $d_1(p, t_i) - d_1(q, t_i) = |t_i^1 - p^1| + |t_i^2 - p^2| - |t_i^1 - q^1| - |t_i^2 - q^2| = d_1(p', t_i') - d_1(q', t_i')$. Since $p \succ_T q$, we conclude that $p' \succ_{T^3} q'$, thus $q' \notin \mathcal{P}_{d_1}(T^3)$. This shows that $q \notin \mathcal{I}(T)$.

4.3 $\mathcal{M}(T) \subseteq \mathcal{P}_{d_1}(T)$

First notice that $\mathcal{I}(T) = \mathcal{P}_{d_1}(T)$ is a median stable set. Indeed, the three sets $\mathcal{P}_{d_1}(T^i)$ as well as their Cartesian products with lines are median stable. Since the intersection of median stable sets is median stable, we conclude that $\mathcal{I}(T)$ is median. Thus $\text{med}(T) \subset \mathcal{I}(T) = \mathcal{P}_{d_1}(T)$. Now, we will show that any cell Π of $\mathcal{M}(T)$ also belongs to $\mathcal{P}_{d_1}(T)$. If all corners of Π belong to med(T), then since $\text{med}(T) \subset \mathcal{P}_{d_1}(T)$ and the cubical complex $\mathcal{P}_{d_1}(T)$ is conformal, we deduce that Π belongs to $\mathcal{P}_{d_1}(T)$. On the other hand, if $\Pi \cap \text{med}(T) = \{a, b, a', b'\}$ and $ab, a'b'$ are opposite edges of Π, then from Lemma 11 we deduce that $T \subset \text{Ort}(\Pi, a) \cup \text{Ort}(\Pi, b) \cup \text{Ort}(\Pi, a') \cup \text{Ort}(\Pi, b')$ and each of these orthants contains at least one terminal. Pick any two terminals t_i and t_j belonging to opposite orthants. Then $\Pi \subseteq I_{d_1}(t_i, t_j) \subseteq \cup_{k=1}^n I_{d_1}(t_i, t_k)$, thus $\Pi \subseteq \mathcal{P}_{d_1}(T)$, because $\Upsilon_{d_1}(T) \subseteq \mathcal{P}_{d_1}(T)$ by Lemma 1. Finally suppose that only the opposite corners a, b of Π belong to med(T). Again, Lemma 11 implies that $T \subset \text{Ort}(\Pi, a) \cup \text{Ort}(\Pi, b)$ and each of these orthants contains terminals. As in previous case, we conclude that $\Pi \subseteq \Upsilon_{d_1}(T) \subseteq \mathcal{P}_{d_1}(T)$.

4.4 $\mathcal{I}(T) \subseteq \mathcal{M}(T)$

Since $\mathcal{I}(T)$ and $\mathcal{M}(T)$ are conformal, it suffices to show that if a vertex q of the uniform grid $\Gamma(T)$ belongs to $\mathcal{I}(T) = \mathcal{P}_{d_1}(T)$, then q also belongs to $\mathcal{M}(T)$. Since $\text{med}(T) \subset \mathcal{M}(T)$, we can suppose that $q \notin \text{med}(T)$. Then the projection of q on at least one of the coordinate planes, say on H^3, does not belong to med(T^3), otherwise it can be easily shown that $q \in \text{med}(T)$. Let $q' = (q^1, q^2)$. Since $q \in \mathcal{I}(T)$, we must have $q' \in \mathcal{P}_{d_1}(T^3)$. Hence $q' \in \mathcal{M}(T^3)$ by Proposition 2, i.e., q' belongs to a square Π_0 of $\Gamma(T^3)$ having either all corners in med(T^3) or only two opposite corners in med(T^3). The first case is impossible because $q' \notin \text{med}(T^3)$. Therefore, only two opposite corners s_1 and s_2 (both different from q') of Π_0 belong to med(T^3). Let s_3 and s_4 be the two other corners of Π_0 (see Fig. 4(a)). Let Π be a 3-dimensional cell of $\Gamma(T)$ having q as a corner and whose orthogonal projection on the plane H^3 is the cell Π_0. Denote by a, b, c, d, e, f, and g the seven remaining corners of Π, as indicated in Fig. 4(a). Lemma 11 implies that $T^3 \subset \text{Ort}(\Pi', s_1) \cup \text{Ort}(\Pi', s_2)$, because s_3 and s_4 are not median points of T^3. Since T^3 is the projection of T on the plane H^3, we deduce that $T \subset \text{Ort}(\Pi, a) \cup \text{Ort}(\Pi, b) \cup \text{Ort}(\Pi, c) \cup \text{Ort}(\Pi, d)$ and that four other orthants are empty. Lemma 11 yields that the vertices e, f, g, and q do not belong to med(T), because each of them belongs to an empty wedge. The definition of $\Gamma(T)$ implies that the hyperplane passing via a and c contains a terminal, which we denote by t_a. We will assume without loss of generality that t_a belongs to $\text{Ort}(\Pi, a)$. Using this information about $\text{Ort}(\Pi, a)$, we will proceed with a case analysis depending of the presence or absence of terminals in the orthants defined by the corners b, c, and d of Π.

If each of the four orthants defined by a, b, c, and d is nonempty, then a, b, c, and d belong to med(T), moreover they are the only corners of Π from med(T). Thus $q \in \mathcal{M}(T)$, because $\Pi \subseteq \mathcal{M}(T)$ by the definition of $\mathcal{M}(T)$. If only two orthants are non-empty, since $t_a \in \text{Ort}(\Pi, a)$, from the definition of $\Gamma(T)$ we infer that $\text{Ort}(\Pi, d)$ is the second nonempty orthant. Since these two orthants are opposite, we conclude that a and d are the only corners of Π belonging to med(T), whence $q \in \Pi \subseteq \mathcal{M}(T)$. Now suppose that there are exactly three non-empty orthants. If these are the orthants of a, b, and c, then a dominates q because the terminals from $\text{Ort}(\Pi, b) \cup \text{Ort}(\Pi, c)$ are equidistant from a and q, while the terminals from $\text{Ort}(\Pi, a)$ are closer to a than to q. This contradicts our choice of q. Analogously, if the non-empty orthants are those defined by a, c, and d, then we show that q is dominated by the point c. Finally, suppose that the non-empty orthants are the orthants defined by a, b, and d. Since Π is a cell of the grid $\Gamma(T)$, there exists a terminal t_1 in the hyperplane passing via e and d, a terminal t_2 in the hyperplane passing via d and g, a terminal t_3 in the hyperplane H passing via d and b, and, finally, a terminal t_b located in the orthant $\text{Ort}(\Pi, b)$. Since all terminals are located in three orthants defined by a, b, and d, we conclude that $t_1, t_2 \in \mathrm{Ort}(\Pi, d)$ and $t_3 \in \mathrm{Ort}(\Pi, b) \cup \mathrm{Ort}(\Pi, d)$. If the intersection $\mathrm{Ort}(\Pi, d) \cap H$ contains some terminal t', then from Lemma 11 and the location of the terminals t', t_a, and t_b we infer that the points b and d must belong to $\mathrm{med}(T)$. In this case, q belongs to $\mathcal{M}(T)$ because the 2-dimensional cell defined by f, b, q and d has exactly two corners in $\mathrm{med}(T)$. If $\mathrm{Ort}(\Pi, d) \cap H$ does not contain terminals, then t_3 is located in $\mathrm{Ort}(\Pi, b) \cap H$ (see Fig. 4(b)). In this case, let Π' denote the 3-cell located immediately below the cell Π. This cell necessarily exists because the terminals t_1 and t_2 are below the plane H. Denote by b', d', f', and q' the corners of Π' which are incident in $\Gamma(T)$ to b, d, f, and q, respectively. Since $\mathrm{Ort}(\Pi, d) \cap H \cap T = \varnothing$, we infer that $\mathrm{Ort}(\Pi, d) \cap T = \mathrm{Ort}(\Pi', d') \cap T \neq \varnothing$. If the orthant $\mathrm{Ort}(\Pi', b')$ is non-empty, then the location of terminals yields that q is dominated by b'. Otherwise, we conclude that b and d' are the only corners of $\Pi' \cap \mathrm{med}(T)$, because $T \subset \mathrm{Ort}(\Pi', d') \cup \mathrm{Ort}(\Pi', b)$. In this case, the definition of $\mathcal{M}(T)$ yields $q \in \Pi' \subseteq \mathcal{M}(T)$. This establishes the inclusion $\mathcal{I}(T) \subseteq \mathcal{M}(T)$ and concludes the proof of Theorem 3.

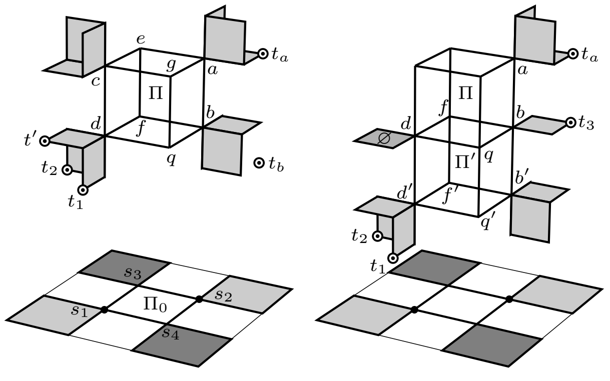

(a) (b)

Figure 4: To the proof of $\mathcal{I}(T) \subseteq \mathcal{M}(T)$.

5. ALGORITHMS

5.1 $O(n \log n)$ algorithm for $\mathcal{P}_{d_1}(T)$ in $\mathbb{R}^3$

The equality $\mathcal{P}_{d_1}(T) = \mathcal{I}(T)$ from Theorem 3 paves the way to an optimal $O(n \log n)$-time algorithm for constructing Pareto envelopes $\mathcal{P}_{d_1}(T)$ in $\mathbb{R}^3$ (its optimality follows from [7], which shows that the sorting problem reduces in linear time to the construction of $\mathcal{P}_{d_1}(T)$ in the plane). As a by-product of our algorithm, we obtain an $O(n \log n)$ algorithm for computing the median closure $\mathrm{med}(T)$ and its geometric realization $||\mathrm{med}(T)||$. Let $T \subset \mathbb{R}^3$. After sorting the terminals of T by each coordinate, we can perform in $O(n \log n)$ time the coordinatwise-isotonic transform φ, which maps T to an uniform set $\varphi(T)$ of terminals. Since from cubical complex $\mathcal{P}_{d_1}(\varphi(T))$ we can easily reconstruct the cell complex $\mathcal{P}_{d_1}(T)$, further, for simplicity, we will assume without loss of generality that T is uniform, i.e., that $T = \varphi(T)$. The algorithm, called `ParetoR3L1`, consists of the following three phases:

1. Compute $\mathcal{P}_{d_1}(T^1), \mathcal{P}_{d_1}(T^2)$, and $\mathcal{P}_{d_1}(T^3)$.
2. Construct the faces of the sets $\mathcal{I}^1(T)$, $\mathcal{I}^2(T)$, $\mathcal{I}^3(T)$.
3. Restrict each face of $\mathcal{I}^1(T)$, $\mathcal{I}^2(T)$, $\mathcal{I}^3(T)$ to $\mathcal{P}_{d_1}(T)$.

`Algorithm ParetoR3L1`(T)

Phase 1. In this phase, using the algorithm of [7] we construct the 2-dimensional Pareto envelopes $\mathcal{P}_{d_1}(T^1), \mathcal{P}_{d_1}(T^2)$, and $\mathcal{P}_{d_1}(T^3)$. Since $\mathcal{P}_{d_1}(T^i)$ are ortho-convex, the intersection of any $\mathcal{P}_{d_1}(T^i)$ with a horizontal or vertical line is empty or a segment. We call the intersection of $\mathcal{P}_{d_1}(T^i)$ with a horizontal line passing via a terminal a *horizontal cut* (*vertical cuts* are defined in a similar way). The horizontal lines passing via the terminals of T^i partition the plane into horizontal slabs. The intersection of each slab with $\mathcal{P}_{d_1}(T^i)$ is a rectangle, which we call a *horizontal strip.* The vertical slabs and strips are defined in similar way. We modify the algorithm of [7] so that, together with the boundary of $\mathcal{P}_{d_1}(T^i)$, for each horizontal cut $c = [a, b]$ it returns the ranks of the abscises of a and b in the sorted list of abscises of the terminals. Analogously, for each horizontal strip, the algorithm returns the ranks of the abscises of its vertical sides. The same computations are done for vertical cuts and strips. The boundary of $\mathcal{P}_{d_1}(T^i)$ is subdivided into a linear number of segments which are either vertical sides of horizontal strips or horizontal sides of vertical strips of $\mathcal{P}$. We call such segments *boundary segments* of $\mathcal{P}$. For each vertical boundary segment s, we keep the rank of the horizontal strip defining it as well at the rank of the abscise of its end-vertices. We call the resulting arrays, the *cut-strip representation* of $\mathcal{P}$. Obviously, its size is $O(n)$.

Phase 2. The sets $\mathcal{I}^1(T), \mathcal{I}^2(T), \mathcal{I}^3(T)$ are constructed using the cut-strip representation of $\mathcal{P}_{d_1}(T^1), \mathcal{P}_{d_1}(T^2), \mathcal{P}_{d_1}(T^3)$ by taking the Cartesian product of each boundary segment s of the respective envelope with the line perpendicular to the plane containing it, thus getting unbounded strips.

Phase 3. For each unbounded strip f constructed in previous phase, we compute its intersection with $\mathcal{I}(T)$. Suppose without loss of generality that f is a face of $\mathcal{I}^1(T)$, i.e., f is the Cartesian product of the line l^1 with the vertical side s of a horizontal strip of $\mathcal{P}_{d_1}(T^1)$. By definition of $\mathcal{I}(T)$, to compute the restriction of f to $\mathcal{I}(T)$, it suffices to compute its intersection with $\mathcal{I}^2(T)$ and $\mathcal{I}^3(T)$. The algorithm first computes the intersection $f' = f \cap \mathcal{I}^3(T)$ and then the intersection $f'' = f' \cap \mathcal{I}^2(T) = (f \cap \mathcal{I}^3(T)) \cap \mathcal{I}^2(T)$, which is the final "contribution" of f to $\mathcal{I}(T)$. To compute the intersection of f with $\mathcal{I}^3(T)$, it suffices to find the cut of the envelope $\mathcal{P}_{d_1}(T^3)$ which is contained in the projection of the unbounded strip f to the plane H^3. The minimum and the maximum ranks of this cut defines the rectangle $f' = f \cap \mathcal{I}^3(T)$; see Fig. 5(a). To compute the intersection of f' with $\mathcal{I}^2(T)$, first we find the horizontal slab of H^2 containing the projection of f' on this plane. This is easy because the rank of the strip defined by this slab coincides with the rank of the horizontal strip of $\mathcal{P}_{d_1}(T^1)$ defining the boundary segment s. Using the minimum and the maximum ranks of this strip of $\mathcal{P}_{d_1}(T^2)$ and the coordinates of f', we compute the rectangle f''; see Fig. 5(b). Notice that using the cut-strip representation of the three 2-dimensional Pareto envelopes, f' and f'' can be computed in constant time.

For example, in Fig. 5, f is defined by the vertical boundary segment s of the 3rd horizontal strip of $\mathcal{P}_{d_1}(T^1)$ and the end-vertices of s have rank 2. The projection of f on the plane H^3 is a ray defining the 2nd vertical cut of $\mathcal{P}_{d_1}(T^3)$. The end-vertices of this cut have ranks 1 and 4. Using these values, we compute the rectangle f', namely, f' is defined by the 3rd horizontal strip of $\mathcal{P}_{d_1}(T^1)$ and the ranks 1 and 4 of the 2nd vertical cut of $\mathcal{P}_{d_1}(T^3)$. The rectangle f' projects precisely on the 3rd horizontal strip of $\mathcal{P}_{d_1}(T^2)$. The vertical sides of this strip have ranks 2 and 4. Intersecting [1..4] and [2..4], we deduce that the vertical sides of f'' have ranks 2 and 4. Since the cut-strip representation is implemented using an array, all such computations are done in constant time per boundary segment s.

Notice that all rectangles f'' computed in **Phase 3** belong to the boundary of $\mathcal{I}(T)$. We assert that `ParetoR3L1` completely describe the boundary of $\mathcal{I}(T)$, i.e., that each boundary point p of $\mathcal{I}(T)$ belongs to a face returned by this algorithm. Since $\mathcal{I}(T)$ is the intersection of the polyhedra $\mathcal{I}^1(T)$, $\mathcal{I}^2(T)$, and $\mathcal{I}^3(T)$, the point p necessarily belongs to some unbounded strip f occurring at **Phase 2**, say to f from $\mathcal{I}^1(T)$. Since p belongs to the restriction of f to $\mathcal{I}(T)$, we see that p will belong to f', and then it will belong also to f''.

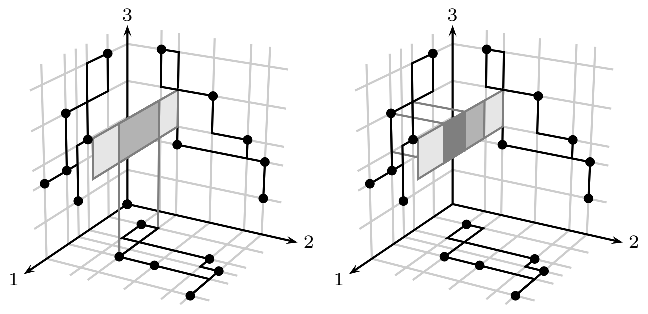

(a) (b)

Figure 5: Generating a face of $\mathcal{I}(T)$

The algorithm returns a subdivision of the boundary of $\mathcal{I}(T)$ into a linear number of rectangular faces, thus the complexity of $\mathcal{P}_{d_1}(T) = \mathcal{I}(T)$ is linear. Using a standard sweeping procedure, one can concatenate the sets of coplanar incident faces of $\mathcal{I}(T)$ into single faces. Notice also that, as described above, the **Phase 3** of the algorithm returns the faces of the Pareto envelope in an arbitrary order. To retrive the incidence between these rectangular faces, each time when a new face f'' is computed, we call **Phase 3** to the boundary segments of the strip used to derive f'' from f'.

The algorithm `ParetoR3L1` runs in total $O(n \log n)$ time because the envelopes $\mathcal{P}_{d_1}(T^1), \mathcal{P}_{d_1}(T^2)$, and $\mathcal{P}_{d_1}(T^3)$ can be constructed in $O(n \log n)$ [7] as well as their cut-strip representations. The (unbounded) faces of $\mathcal{I}^1(T)$, $\mathcal{I}^2(T)$, and $\mathcal{I}^3(T)$ are defined in $O(n)$ time. Finally, each of these unbounded faces induces a rectangular face of $\mathcal{I}(T)$, which is computed in constant time. Summarizing, we obtain the following result:

THEOREM 4. *The algorithm* `ParetoR3L1` *computes the faces of the boundary of $\mathcal{P}_{d_1}(T)$ in optimal $O(n \log n)$ time.*

5.2 $O(n \log^2 n)$ algorithm for $\mathcal{P}^0_{d_\infty}(T)$ in $\mathbb{R}^3$

5.2.1 Complexity of $\mathcal{P}^0_{d_\infty}(T)$

We establish that the boundary of the strict Pareto envelope $\mathcal{P}^0_{d_\infty}(T) = \cap_{i=1}^3 (\mathcal{U}^i_+(T) \cap \mathcal{U}^i_-(T))$ in $\mathbb{R}^3$ has linear complexity $\kappa(\mathcal{P}^0_{d_\infty}(T))$. Denote by $\partial\mathcal{U}^i_+(T)$ and $\partial\mathcal{U}^i_-(T)$ the boundaries of the unions of cones $\mathcal{U}^i_+(T)$ and $\mathcal{U}^i_-(T)$. The 6 surfaces $\partial\mathcal{U}^i_-(T)$, $\partial\mathcal{U}^i_+(T), i = 1,2,3$, subdivide $\mathbb{R}^3$ into cells which define a cellular complex $\Gamma^0_\infty(T)$. The equality established in Theorem 2 implies that $\mathcal{P}^0_{d_\infty}(T)$ is a subcomplex of $\Gamma^0_\infty(T)$. Moreover, any cell C of $\Gamma^0_\infty(T)$ belongs to $\mathcal{P}^0_{d_\infty}(T)$ if and only if any interior point of C does. Every edge e of $\Gamma^0_\infty(T)$ belongs to a line $l(e)$ which is either parallel to one of the lines l^1, l^2, l^3, l^4 or is parallel to one of the coordinate lines m^1, m^2, m^3 (m^i is perpendicular to the plane H^i). We will show below that in fact the complexity $\kappa(\Gamma^0_\infty(T))$ of the cell complex $\Gamma^0_\infty(T)$ is linear.

Denote by $\mathcal{F}$ the collection of all faces of the surfaces $\partial\mathcal{U}^i_+(T)$ and $\partial\mathcal{U}^i_-(T), i = 1,2,3$. Any face $F \in \mathcal{F}$ is a part of an unbounded face of some cone of one of the six families. To fix the notation, we will assume without loss of generality that F belongs to the face $F(t_j)$ of the downward cone $C^3_-(t_j)$ visible from the direction m^1. Then the boundary of F contains two incident at t_j edges which are parts of the extremal rays of $F(t_j)$ plus a polygonal chain $Q^3_-(F)$, each edge of which is either horizontal or parallel to an extremal ray of $F(t_j)$ (see Fig. 6(a)). Denote by $Q^3_+(F(t_j))$ the polygonal line obtained as the intersection of $F(t_j)$ with the surface $\partial\mathcal{U}^3_+(T)$ (see Fig. 6(b)). Analogously, for each $i = 1,2$, denote by $Q^i_-(F(t_j))$ and $Q^i_+(F(t_j))$ the intersections of $F(t_j)$ with the surfaces $\partial\mathcal{U}^i_-(T)$ and $\partial\mathcal{U}^i_+(T)$. Denote by $Q^i_-(F)$ and $Q^i_+(F)$ ($i = 1,2,3$) the upper envelopes on $F(t_j)$ of the chain $Q^3_-(F)$ and each of the chains $Q^i_-(F(t_j))$ and $Q^i_+(F(t_j))$, respectively (see Fig. 6(c)). Each of 6 resulting polygonal lines $Q^i_-(F), Q^i_+(F), i = 1,2,3$, is monotone with respect to the extremal rays of $F(t_j)$. They induce a subdivision of F into 2-dimensional cells, which are nothing else than the 2-dimensional faces of $\Gamma^0_\infty(T)$ belonging to F. The vertices and the pairwise intersections of the 6 polygonal chains $Q^i_-(F)$ and $Q^i_+(F)$, $i = 1,2,3$, taken over all $F \in \mathcal{F}$ define the vertices of $\Gamma^0_\infty(T)$. Finally, the resulting subdivision of the polygonal chains $Q^i_-(F)$ and $Q^i_+(F)$ induced by the intersection points define the edges of $\Gamma^0_\infty(T)$. For each of the polygonal lines $Q^i_-(F)$ and $Q^i_+(F)$, we call its vertex *convex* if this is a locally convex vertex in the intersection of F with the respective union of cones and *concave* otherwise.

THEOREM 5. $\kappa(\Gamma^0_\infty(T)) = \kappa(\mathcal{P}^0_{d_\infty}(T)) = O(n)$.

PROOF. Since every vertex of the cell complex $\Gamma^0_\infty(T)$ is incident to a constant number of edges and cells, it suffices to show that the cell complex $\Gamma^0_\infty(T)$ contains a linear number of vertices. Pick any face $F \in \mathcal{F}$. To fix the notation, we will assume as before that F belongs to the face $F(t_j)$ of the downward cone $C^3_-(t_j)$ visible from the direction m^1.

Claim 1: *$\partial\mathcal{U}^i_-(T)$ and $\partial\mathcal{U}^i_+(T)$ contain $O(n)$ vertices.*

Proof. The result follows from Theorem 2.1 of [27] about linearity of lower envelopes of particular bivariate functions or, as noticed by R. Klein, because the projection of $\partial\mathcal{U}^i_+(T)$ on H^i can be viewed as a planar Voronoi l_1-diagram with additive weights. We provide an alternative proof to give a better understanding of more complicated cases. Suppose $i = 3$, i.e., we consider the union of downward cones $C^3_-(t_j)$. As we noticed above, the edges of the polygonal line $Q^3_-(F)$ are either horizontal or parallel to the extremal rays of the face $F(t_j)$. Each vertex v of $Q^3_-(F)$ gives rise to a vertex of $\mathcal{U}^3_-(T)$ and vice versa. Each convex vertex v of $Q^3_-(F)$ is the intersection of F with an extremal ray R of some cone $C^3_-(t_k)$. We charge the terminal t_k for v. Notice that each t_k is in charge of at most 4 such vertices of $\mathcal{U}^3_-(T)$: indeed, since the intersection of $\mathcal{U}^3_-(T)$ with R is convex, R contains at most one vertex of $\mathcal{U}^3_-(T)$. On the other hand, each concave vertex of $Q^3_-(F)$ necessarily is incident to a convex vertex of this chain (see Fig. 6(a)). Summing up over all faces F of $\partial\mathcal{U}^3_-(T)$, our charging scheme implies that $\partial\mathcal{U}^3_-(T)$ contains $O(n)$ vertices.

Claim 2: *$\partial\mathcal{U}^i_-(T) \cap \partial\mathcal{U}^i_+(T)$ contains $O(n)$ vertices.*

Proof. Let $i = 3$. To count the number of vertices of $\partial\mathcal{U}^3_-(T) \cap \partial\mathcal{U}^3_+(T)$ located on the faces F of the surface $\partial\mathcal{U}^3_-(T)$

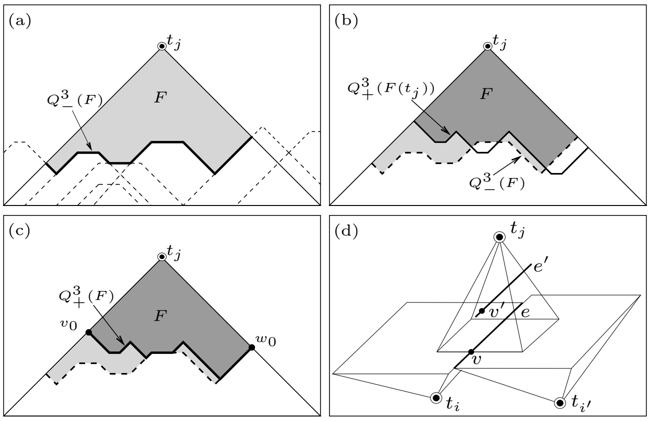

Figure 6: To the definition of $Q^i_-(F)$ and $Q^i_+(F)$

(see Fig. 6(c)), we will charge the terminals of T and the edges of $\partial\mathcal{U}^3_+(T)$ to pay for such vertices. Every vertex of $Q^3_-(F)$ is payed as in the proof of Claim 1. Analogously, every convex vertex of $Q^3_+(F) \setminus Q^3_-(F)$ is payed by a terminal. It remains to pay the concave vertices of $Q^3_+(F)$ and the points obtained as the intersection of two chains. The latter are incident to concave vertices, thus it suffices to count only the concave vertices. Every such vertex v belongs to an edge e of $\partial\mathcal{U}^3_+(T)$ parallel to the line m^1. Suppose that this edge e belongs to the intersection $C^3_+(t_i) \cap C^3_+(t_{i'})$. In general, e may contain convex vertices of other polygonal lines $Q^3_-(F')$. If this happens, and e contains several such vertices, then e traverses all respective downwards cones except the first and the last one. Then we assert that all cones traversed by e cannot be intersected by another m^1-edge e' of $\partial\mathcal{U}^3_+(T)$. Suppose by way of contradiction that the edge e traverses the cone $C^3_-(t_j)$ and the face F contains a vertex $v' := e' \cap F$ different from $v = e \cap F$. Suppose without loss of generality that v' is above v (otherwise, interchange v and v'). Consider the sections of the three cones in question with the horizontal plane $H^3(v)$ (see Fig. 6(d)). Since $H^3(v) \cap (C^3_+(t_i) \cup C^3_+(t_{i'}))$ contains the square $H_3(v) \cap C^3_-(t_j)$, necessarily each point of the part of $C^3_-(t_j)$ lying above $H^3(v)$ belongs either to the interior of $C^3_+(t_i)$ or to the interior of $C^3_+(t_{i'})$. This contradicts the fact that v' is a boundary point of $\mathcal{U}^i_-(T) \cap \mathcal{U}^i_+(T)$, establishing our assertion.

Thus, if the face F is intersected by several m^1-edges of $\mathcal{U}^3_+(T)$, then we can assign all vertices of $\partial\mathcal{U}^3_-(T) \cap \partial\mathcal{U}^3_+(T)$ obtained in such a way to respective incident m^1-edges. Otherwise, if F contains only one such convex vertex, then we charge this vertex to the terminal t_j defining F. Since by Claim 1 the number of m^1-edges of $\partial\mathcal{U}^3_+(T)$ is linear, and each m^1-edge and each terminal is in charge of at most one convex vertex of $Q^3_+(F)$, summing up over all faces of all cones, we deduce the linearity of $\partial\mathcal{U}^i_-(T) \cap \partial\mathcal{U}^i_+(T)$. This establishes Claim 2.

The proof of next claim uses the same arguments as the proof of Claim 2 and is deferred to the complete version:

Claim 3: $\partial\mathcal{U}^i_-(T) \cap \partial\mathcal{U}^j_+(T)$, $\partial\mathcal{U}^i_-(T) \cap \partial\mathcal{U}^j_-(T)$, and $\partial\mathcal{U}^i_+(T) \cap \partial\mathcal{U}^j_+(T)$ contain $O(n)$ vertices.

For each face F of $\mathcal{F}$ we have defined 6 polygonal chains $Q^i_-(F), Q^i_+(F), i = 1,2,3$. Each of these chains is convex with respect to the lines defining the extremal rays of F. Claims 1,2, and 3 show that the total number of vertices on all chains taken over all faces F of $\mathcal{F}$ is linear. Using the directional convexity of the chains of F, we conclude that the total number of the intersection points of any two from the 6 chains of F is also linear in the size of respective chains. Since every vertex of $\Gamma^0_\infty(T)$ is obtained in this way, the proof of the theorem is finished. □

5.2.2 The algorithm

Our algorithm first constructs the vertices and the edges of the cell complex $\Gamma^0_\infty(T)$ by using ray shooting queries. After deriving the cells of $\Gamma^0_\infty(T)$, for each cell, it tests if it belongs or not to $\mathcal{P}^0_{d_\infty}(T)$. The algorithm, called `ParetoR3LINF`, consists of the following five phases:

1. Compute the 6 boundaries $\partial\mathcal{U}^i_-(T)$ and $\partial\mathcal{U}^i_+(T)$.
2. Compute the vertices and the edges of the 6 polygonal chains $Q^i_-(F), Q^i_+(F)$ for each face $F \in \mathcal{F}$.
3. Compute the subdivision of any face $F \in \mathcal{F}$ induced by the 6 monotone chains $Q^i_-(F), Q^i_+(F)$.
4. Compute $\Gamma^0_\infty(T)$ by merging the incidence lists of the vertices of the subdivisions obtained in Phase 3.
5. Compute the cells of $\Gamma^0_\infty(T)$ which belong to $\mathcal{P}^0_{d_\infty}(T) = \cap^3_{i=1}(\mathcal{U}^i_+(T) \cap \mathcal{U}^i_-(T))$.

`Algorithm ParetoR3LINF(T)`

Phase 1. The surfaces $\partial\mathcal{U}^i_-(T)$ and $\partial\mathcal{U}^i_+(T), i = 1,2,3$, can be viewed as lower envelopes of bivariate functions satisfying the conditions of Theorem 2.1 of [27] and therefore can be computed in $O(n \log n)$-time as described in this paper. Each of the 6 resulting boundaries is preprocessed in order to support ray shootings in 14 directions defined by the lines l^1, l^2, l^3, l^4 and m^1, m^2, m^3 in fixed-oriented polyhedra as described in [12] (this can be done in total $O(n \log^2 n)$ randomized time).
Phase 2. For any face $F \in \mathcal{F}$, to compute the vertices and the edges of the 6 polygonal chains $Q^i_-(F), Q^i_+(F), i = 1,2,3$, it suffices to find the end-vertices of each chain, and then to construct the respective chain edge-by-edge by performing repetitive ray shootings in certain directions belonging to the face F. For example, to compute the chain $Q^3_+(F)$ as defined in Fig. 6(c), first we find its end-vertices v_0, w_0 by performing two ray shootings from the terminal t_j defining F in the directions of the extremal rays of F. Letting initially $v := v_0$, and, given the current vertex v of $Q^3_+(F)$, we perform ray shootings from v in a constant number of directions with respect to $\partial\mathcal{U}^3_-(T)$ and $\partial\mathcal{U}^3_+(T)$ to find the next neighbor of v in $Q^3_+(F)$, and so on until we reach w_0 (see Fig. 6(c)). Each ray shooting query can be answered in $O(\log^2 n)$ time using the algorithm described in Subsection 6.2 of [12]. Since in order to find a new vertex, we perform a constant number of ray shootings, the total number of such queries is linear.
Phase 3. By sweeping each pair of 6 monotone chains $Q^i_-(F), Q^i_+(F), i = 1,2,3$, we can find they pairwise intersections. These points together with the vertices of the chains define the vertex-set of the subdivision. To obtain the subdivision itself, each time when an intersection point is found, the respective crossing segments are subdivided accordingly.
Phase 4. To construct the cells of $\Gamma^0_\infty(T)$, it suffices to detect the coinciding vertices coming from different face-subdivisions and merge their incidence edge lists (for this, it suffices to sort the $O(n)$ vertices of all subdivisions obtained in **Phase 3** in lexicographic order).
Phase 5. To decide if a cell C of $\Gamma^0_\infty(T)$ belongs to $\mathcal{P}^0_{d_\infty}(T)$, we pick any interior point p of C and add C to $\mathcal{P}^0_{d_\infty}(T)$ if

and only if $p \in \cap_{i=1}^{3}(\mathcal{U}_{+}^{i}(T) \cap \mathcal{U}_{-}^{i}(T)) = \mathcal{P}_{d_\infty}^{0}(T)$. For this, we compute the intersection of each $\partial\mathcal{U}_{+}^{i}(T)$ and $\partial\mathcal{U}_{-}^{i}(T)$ with the line passing via p and parallel to l^1 (in fact, any other convexity direction l^j does the job). To find this intersection, we perform two ray shootings from p in the opposite directions defined by l^1. Since $\Gamma_{\infty}^{0}(T)$ contains $O(n)$ cells, this phase can be implemented in total $O(n \log^2 n)$ time. Concluding, we obtain the following result:

THEOREM 6. *The algorithm* `ParetoR3LINF` *computes the cells of* $\mathcal{P}_{d_\infty}^{0}(T)$ *in* $O(n \log^2 n)$ *randomized time.*

6. REFERENCES

[1] N. Aronszajn and P. Panitchpakdi, Extensions of uniformly continous transformations and hyperconvex metric spaces, *Pacif. J. Math.* **6** (1956), 405-439.

[2] H.-J. Bandelt and V. Chepoi, Metric graph theory and geometry: a survey. *Contemp. Math.* (to appear).

[3] H.-J. Bandelt and J. Hedlíková, Median algebras, *Discr. Math.* **45** (1983), 1-30.

[4] J.-D. Boissonnat and M. Yvinec, *Géométrie Algorithmique,* Ediscience International, Paris, 1995.

[5] J.-D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec, Voronoi diagrams in higher dimensions under certain polyhedral distance functions, *Discr. Comput. Geom.* **19**(1998), 485-519.

[6] V. Boltyanski, H. Martini, and P. S. Soltan, *Excursions into Combinatorial Geometry,* Springer, Berlin, 1997.

[7] G. Chalmet, L. Francis, and A. Kolen, Finding efficient solutions for rectilinear distance location problems efficiently, *Europ. J. Oper. Res.* **6** (1981), 117–124.

[8] V. Chepoi, Graphs of same CAT(0) complexes, *Adv. Appl. Math.* **24** (2000), 125–179.

[9] V. Chepoi, K. Nouioua, Y.Vaxès, A rounding algorithm for approximating minimum Manhattan networks, In *APPROX-RANDOM'2005,* pp. 40-51.

[10] M. Chrobak and L. L. Larmore, Generosity helps or an 11-competitive algorithm for three Servers. *J. Algorithms* **16** (1994), 234–263.

[11] M. de Berg, M. van Kreveld, M. Overmars, O. Schwarzkopf, and M. H. Overmars, *Computational Geometry : Algorithms and Applications*, 2000, Springer-Verlag, Berlin.

[12] M. de Berg, *Ray Shooting, Depth Orders and Hidden S urface Removal,* 1993, Springer-Verlag, Berlin.

[13] A. W. M. Dress, Trees, tight extensions of metric spaces, and the cohomological dimension of certain groups : a note on combinatorial properties of metric spaces. *Adv. Math.*, **53** (1984), 321–402.

[14] R. Durier, On Pareto optima, the Fermat-Werber problem, and polyhedral gauges. *Math. Progr.* **47** (1990), 65–79.

[15] R. Durier and C. Michelot, Geometrical properties of the Fermat-Weber problem. *Europ. J. Oper. Res.* **20** (1985), 332–343.

[16] R. Durier and C. Michelot, Sets of efficient points in normed space, *J. Math. Analys. Appl.* **117** (1986), 506–528.

[17] E. Fink and D. Wood, *Restricted-Orientation Convexity,* 2004, Springer-Verlag, Berlin.

[18] P. Hansen, T. Perreur, and J.F. Thisse, Location theory, dominance and convexity : some further results. *Oper. Res.* **28** (1980), 1241–1250.

[19] J.R Isbell, Six theorems about metric spaces. *Comment. Math. Helv.* **39** (1964), 65–74.

[20] J.R. Isbell, Median algebra. *Trans. Amer. Math. Soc.* **260** (1980), 319–362.

[21] R. L. Keeney and H. Raiffa, *Decisions with Multiple Objectives,* John Wiley, New York, 1976.

[22] H. W. Kuhn, A pair of dual nonlinear programs. In *Methods of Nonlinear Programming,* Ed. J. Abadie, pp. 37–64. 1967, North-Holland, Amsterdam.

[23] H. W. Kuhn, A note on Fermat's problem, *Math. Prog.* **4** (1973), pp. 98–107.

[24] K. Nouioua, *Enveloppes de Pareto et Réseaux de Manhattan: caractérisations et algorithmes,* Thèse de Doctorat en Informatique, Université de la Méditerranée, 2005.

[25] B. Pelegrin and F.R. Fernandez, Determination of efficient points in multiple-objective location problems, *Nav. Res. Logistics* **35** (1988), 697–705.

[26] F.P. Preparata and M.I. Shamos, *Computational Geometry: an Introduction,* 1983, Springer-Verlag, Berlin.

[27] J. T. Schwartz and M. Sharir, On the two-dimensional Davenport-Schinzel problem, *J. Symbolic Comput.* **10** (1990), 371-393.

[28] R. E. Steuer, *Multiple Criteria Optimization: Theory, Computation and Application,* Wiley, New York, 1986.

[29] B. Sturmfels and J. Yu, Classification of six-point metrics, *Electr. J. Combin.,* **11** (2004).

[30] J. F. Thisse, J. E. Ward, and R. E. Wendell, Some properties of location problems with block and round norms. *Oper. Res.* **32** (1984), 1309–1327.

[31] M. van de Vel, *Theory of Convex Structures*, Elsevier, Amsterdam, 1993.

[32] M. van de Vel, Collapsible polyhedra and median spaces, *Proc. Amer. Math. Soc.* **126** (1998), 2811–2818.

[33] M. van de Vel, Determination of $msd(L^n)$. *J. Algebr. Combin.* **9** (1999), 161–171.

[34] R. E. Wendell and A.P. Hurter, Location theory, dominance, and convexity. *Oper. Res.* **21** (1973), 314–320.

[35] R. E. Wendell, A.P. Hurter, and T.J. Lowe, Efficient points in location theory. *AIEE Trans.* **9** (1973), 238–246.

[36] P. Widmayer, Y. F. Wu, and C. K. Wong, On some distance problems in fixed orientations. *SIAM J. Comput.* **16** (1987), 728–746.