

Pareto envelopes in simple polygons ¹

VICTOR CHEPOI, KARIM NOUIOUA, EDOUARD THIEL, AND YANN VAXÈS

Laboratoire d'Informatique Fondamentale de Marseille,
Université de la Méditerranée, Faculté des Sciences de
Luminy, F-13288 Marseille Cedex 9, France,
`{chepoi,nouioua,thiel,vaxes}@lif.univ-mrs.fr`

Abstract. For a set T of n points in a metric space (X, d) , a point $y \in X$ is *dominated* by a point $x \in X$ if $d(x, t) \leq d(y, t)$ for all $t \in T$ and there exists $t' \in T$ such that $d(x, t') < d(y, t')$. The set of non-dominated points of X is called the *Pareto envelope* of T . H. Kuhn (1973) established that in Euclidean spaces, the Pareto envelopes and the convex hulls coincide. Chalmet et al. (1981) characterized the Pareto envelopes in the rectilinear plane (\mathbb{R}^2, d_1) and constructed them in $O(n \log n)$ time. In this note, we investigate the Pareto envelopes of point-sets in simple polygons P endowed with geodesic d_2 - or d_1 -metrics (i.e., Euclidean and Manhattan metrics). We show that Kuhn's characterization extends to Pareto envelopes in simple polygons with d_2 -metric, while that of Chalmet et al. extends to simple rectilinear polygons with d_1 -metric. These characterizations provide efficient algorithms for construction of these Pareto envelopes.

Keywords. Pareto envelope, dominance, geodesic l_1 - and l_2 -distance, simple polygon, algorithm.

Corresponding author: Victor Chepoi

¹An abstract of this paper appeared in: Proc. 24th European Workshop on Computational Geometry (EWCG'08), Nancy. This research was partly supported by the ANR grant BLAN06-1-138894 (projet OPTI-COMB)

1 Introduction

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. They host such remarkable points as center, barycenter, and median as well as the optimal solutions of some *NP*-hard problems like the Steiner tree, the p -median, and the p -center problems. H. Kuhn [15] noticed that $\text{conv}(T)$ can be described in truly distance terms: a point $p \in \mathbb{R}^m$ belongs to $\text{conv}(T)$ if and only if the vector of Euclidean distances of p to the points of T is not dominated by the distance vector of any other point of \mathbb{R}^m . Inspired by this characterization of $\text{conv}(T)$, one can define analogous geometric objects by replacing the Euclidean distance d_2 by any other distance d on \mathbb{R}^m , or by replacing \mathbb{R}^m by a polygonal or a polyhedral domain endowed with an intrinsic distance. This leads to the following general concept of Pareto envelope. Given a set T of n points in a metric space (X, d) , a point $y \in X$ is *dominated* by a point $x \in X$ if $d(x, t) \leq d(y, t)$ for all $t \in T$ and there exists $t' \in T$ such that $d(x, t') < d(y, t')$ [2, 19, 24]. The set of non-dominated points of X is called the *Pareto envelope* of T and is denoted by $\mathcal{P}_d(T)$.

Pareto envelopes have been investigated in several papers under the name of “sets of efficient points”. Thisse, Ward, and Wendell [19] proved that $\mathcal{P}_{d_2}(T) = \text{conv}(T)$ holds for all distances induced by round norms [21] (i.e., with strictly convex unit balls). The investigation of Pareto envelopes for particular polyhedral norms has been initiated by Wendell, Hurter, Lowe [24] and continued by Chalmet, Francis, Kolen [2] and Durier, Michelot [7, 8]. The main result of [2] is the following nice characterization of Pareto envelopes in the Manhattan plane:

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

where $I_{d_1}(t_i, t_j)$ is the smallest axis-parallel rectangle with diagonal $[t_i, t_j]$. This result was used in [2] to establish the correctness of an optimal $O(n \log n)$ sweeping-line algorithm for constructing $\mathcal{P}_{d_1}(T)$ in \mathbb{R}^2 . Subsequently, Pelegrin and Fernandez [16] described an algorithm for constructing Pareto envelopes in the plane endowed with a polygonal norm. Recently, Chepoi and Nouioua [5] characterized $\mathcal{P}_{d_1}(T)$ in (\mathbb{R}^3, d_1) and showed that the characterization of Chalmet et al. [2] holds for $\mathcal{P}_{d_\infty}(T)$ in (\mathbb{R}^m, d_∞) . They also presented efficient algorithms for constructing $\mathcal{P}_{d_1}(T)$ and $\mathcal{P}_{d_\infty}(T)$ in \mathbb{R}^3 . In several papers [2, 7, 8, 16, 19, 23, 24], Pareto envelopes are used to reduce the search space of some optimization problems by showing that they host all or at least one optimal solution(s) of the respective problems. For example, Wendell and Hurter [23] establish this type of results for the *Weber problem* (the weighted version of the median problem) while Hansen, Perreux, and Thisse [12] proved a similar result for the *NP*-hard *multifacility location problem*. In [6], we showed that Pareto envelopes in the l_1 -plane contain at least one minimum Manhattan network and we used this to design a factor 2 approximation algorithm for *minimum Manhattan network problem*. For other results in this vein, see [19].

In this note, we characterize and efficiently construct the Pareto envelopes of sets in simple polygons endowed with the geodesic d_2 and d_1 -distances. Distance problems for simple polygons constitute a classical subject in computational geometry; [10, 13, 17, 18, 20] is a




Figure 1: Examples of $\mathcal{P}_{d_2}(T)$ and $\mathcal{P}_{d_1}(T)$

small sample of papers devoted to this subject. We show that, like in Euclidean spaces, Pareto envelopes of finite sets in simple polygons with d_2 -distance coincide with their geodesic convex hulls and therefore can be constructed using an algorithm of Toussaint [20]. On the other hand, we show that Pareto envelopes in simple rectilinear polygons can be characterized using equality (1). This characterization is used to design an efficient algorithm for constructing these envelopes.

We conclude this section with some definitions. Let (X, d) be a metric space. The *interval* $I(x, y)$ between two points $x, y \in X$ consists of all points *between* x and y : $I(x, y) := \{u \in X : d(x, u) + d(u, y) = d(x, y)\}$. A set M of X is *convex* if $I(x, y) \subseteq M$ for all $x, y \in M$. The *convex hull* $\text{conv}(S)$ of a set $S \subset X$ is the smallest convex set containing S .

2 Simple polygons

In this section, P is a simple polygon with m sides endowed with the geodesic d_2 -metric. For two points $x, y \in P$, $\gamma(x, y)$ is the unique geodesic path inside P between x and y , and $d_2(x, y)$ is the length of this path. For a set of n points $T \subset P$, we denote by $\text{conv}(T)$ and $\mathcal{P}_{d_2}(T)$ the *geodesic convex hull* and the Pareto envelope of T . Since two points of a simple polygon P are connected by a unique geodesic, (P, d_2) is a metric space of global non-positive curvature, i.e., a *CAT(0)-space* [1]. CAT(0) spaces are characterized in several ways (in particular, by uniqueness of geodesic paths, convexity of the distance function, etc.) and have many important properties, placing them in the center of modern geometry; for results and definitions the reader can consult the book [1]. Below we will show that $\mathcal{P}_d(T) \subseteq \text{conv}(T)$ holds for any finite subset of a CAT(0)-space (X, d) and we conjecture that in fact $\mathcal{P}_d(T) = \text{conv}(T)$ holds.

2.1 $\mathcal{P}_{d_2}(T) = \text{conv}(T)$

We aim to establish the following result:

Proposition 1 $\mathcal{P}_{d_2}(T) = \text{conv}(T)$. Consequently, $\mathcal{P}_{d_2}(T)$ can be constructed in $O(m + n \log m)$ -time.

The inclusion $\mathcal{P}_{d_2}(T) \subseteq \text{conv}(T)$ follows from the following more general result which holds for all CAT(0) metric spaces:

Lemma 2.1 $\mathcal{P}_d(T) \subseteq \text{conv}(T)$ for any finite set of a CAT(0) metric space (X, d) .

Proof. Let $x \notin \text{conv}(T)$. By Proposition 2.4(1) of [1] there exists a unique point $\pi(x)$ (the metric projection of x) such that $d(x, \pi(x)) = \inf_{y \in \text{conv}(T)} d(x, y)$. As in the case of Euclidean spaces, $\pi(x)$ can be viewed as the orthogonal projection of x on $\text{conv}(T)$, because by Proposition 2.4(3) the Alexandrov angle α at $\pi(x)$ between the geodesics $\gamma(x, \pi(x))$ and $\gamma(y, \pi(x))$ is at least $\pi/2$ for any point $y \in \text{conv}(T), y \neq \pi(x)$. By law of cosines which holds in CAT(0) spaces (page 163 of [1]), if $a = d(x, \pi(x)), b = d(y, \pi(x))$, and $c = d(x, y)$, then $c^2 \geq a^2 + b^2 - 2ab \cos \alpha \geq a^2 + b^2 > b^2$ for any $y \in \text{conv}(T), y \neq \pi(x)$. Hence $d(x, y) > d(\pi(x), y)$, i.e., x is dominated by $\pi(x)$ yielding $x \notin \mathcal{P}_d(T)$. Since x is an arbitrary point outside $\text{conv}(T)$, we conclude that $\mathcal{P}_d(T) \subseteq \text{conv}(T)$. \square

Now we show the converse inclusion $\text{conv}(T) \subseteq \mathcal{P}_{d_2}(T)$. Pick $q \in \text{conv}(T)$. If q belongs to the boundary of $\text{conv}(T)$, then q belongs to the geodesic path $\gamma(t, t')$ between two vertices t, t' of $\text{conv}(T)$. Since $t, t' \in T$, if q is dominated by some point p , then $d_2(p, t) \leq d_2(q, t)$ and $d_2(p, t') \leq d_2(q, t')$. Since $q \in \gamma(t, t')$, this is possible only if these inequalities hold as equalities, thus $p \in \gamma(t, t')$, yielding $p = q$. Thus $q \in \mathcal{P}_{d_2}(T)$ in this case. Now, suppose that q belongs to the interior of the simple polygon $\text{conv}(T)$. Suppose by way of contradiction that q is dominated by some point $p' \in P$; see Fig. 2 for an illustration. By Lemma 1 of [17] the distance function d_2 on P is convex. This means that for any point $t \in T$, as p varies along the geodesic $\gamma(p', q)$, $d_2(t, p)$ is a convex function of p . Since q belongs to the interior of $\text{conv}(T)$, one can select a point $p \in \gamma(p', q) \cap \text{conv}(T)$ which still dominates q and is visible from q (i.e., $[p, q] \subseteq P$). Denote by q' the intersection of the boundary of $\text{conv}(T)$ with the ray with origin p which passes via the point q . By definition of q' , we infer that $q \in [p, q'] = \gamma(p, q')$. Pick any point $t \in T$. By second part of Lemma 1 of [17], $d_2(t, q) < \max\{d_2(t, q'), d_2(t, p)\}$. Since $d_2(t, p) \leq d_2(t, q)$ by the choice of p , we obtain that $d_2(t, q) < d_2(t, q')$. Since this inequality holds for all points of T , we conclude that q and p both dominate the boundary point q' , a contradiction with the fact that $q' \in \mathcal{P}_{d_2}(T)$. This establishes the inclusion $\text{conv}(T) \subseteq \mathcal{P}_{d_2}(T)$ and concludes the proof of Proposition 1.




Figure 2: To the proof of the inclusion $\text{conv}(T) \subseteq \mathcal{P}_{d_2}(T)$

G. Toussaint [20] presented an $O(m+n \log m)$ -time algorithm for constructing the geodesic convex hull of an n -point set T of a simple polygon P with m sides. Together with Proposition 1 this shows that $\mathcal{P}_{d_2}(T)$ can be constructed within the same time bounds.

3 Simple rectilinear polygons

In this section, P is a simple rectilinear polygon (i.e., a simple polygon having all edges axis-parallel) with m edges endowed with the geodesic d_1 -metric. A *rectilinear path* is a polygonal chain consisting of axis-parallel segments lying inside P . The length of a rectilinear path in the d_1 -metric equals the sum of the length of its constituent segments. For two points $x, y \in P$, the *geodesic d_1 -distance* $d_1(x, y)$ is the length of a minimum length rectilinear path (i.e., rectilinear geodesic) connecting x and y . An axis-parallel segment c is a *cut segment* of P if it connects two edges of P and lies entirely in P . One basic property of resulting metric space (P, d_1) is that its axis-parallel cuts and the two subpolygons defined by such cuts are convex and gated [4]. A subset M of a metric space (X, d) is called *gated* [22] provided every point $v \in X$ admits a *gate* in M , i.e., a point $g(v, M) \in M$ such that $g(v, M) \in I(v, u)$ for all $u \in M$.

3.1 Characterization

We extend the characterization of [2] to Pareto envelopes $\mathcal{P}_{d_1}(T)$ in rectilinear polygons:

Proposition 2 $\mathcal{P}_{d_1}(T) = \bigcap_{i=1}^n (\bigcup_{j=1}^n I(t_i, t_j))$.

Proof. One direction of the proof is obvious: if p is not Pareto, then there exists $p' \in P$ such that $d_1(p', t_j) \leq d_1(p, t_j)$ for all j and $d_1(p', t_i) < d_1(p, t_i)$ for some i . Hence, $d_1(t_i, p') + d_1(p', t_j) < d_1(t_i, p) + d_1(p, t_j)$ for all j and thus $p \notin \bigcup_{j=1}^n I(t_i, t_j)$.

To prove the converse, let $p \in \mathcal{P}_{d_1}(T)$ but $p \notin \bigcup_{j=1}^n I(t_i, t_j)$ for some t_i . Let $c_v = [q', q'']$ and $c_h = [p', p'']$ be the maximal vertical and horizontal cuts which pass through the point p . Denote by P_1, P_2, P_3 and P_4 the subpolygons of P defined by these cuts. Let $P_1 \cap P_3 = P_2 \cap P_4 = \{p\}$ and $t_i \in P_1$. Obviously, P_1, \dots, P_4 are gated. Note that p is the gate in P_1 of any point of P_3 . As $p \notin \bigcup_{j=1}^n I(t_i, t_j)$, we conclude that $P_3 \cap T = \emptyset$. Set $P_{j,k} := P_j \cup P_k$, where $j, k \in \{1, 2, 3, 4\}$ and $j \neq k$. Note that the four subpolygons $P_{j, j+1(\text{mod } 4)}$ are gated sets of P .

First, suppose that p is the gate of t_i in one of the cuts c_h or c_v , say in c_h ; see Fig. 3. Then p is also the gate of t_i in $P_{3,4}$. Indeed, the gate q of t_i in $P_{3,4}$ belongs to c_h . Since p is the gate of t_i in c_h , we have $p \in I(t_i, q)$, hence $p = q$. From the choice of p and t_i we conclude that $P_{3,4} \cap T = \emptyset$. Let g_1, \dots, g_n be the gates of the points t_1, \dots, t_n of T in c_v . First, assume that these gates are all different from p . Then all g_1, \dots, g_n belong to the segment $[q', p] \subset c_v$ which separates P_1 and P_2 . Let g_k be a closest to p point of the set g_1, \dots, g_n . Then $g_k \in I(p, g_j)$ and, since $g_j \in I(p, t_j)$, we infer that $g_k \in \bigcap_{j=1}^n I(p, t_j)$, thus g_k dominates p , contradiction that p is Pareto.

Now assume that p is the gate of some point $t_j \neq t_i$ in c_v . If $t_j \in P_2$, then p is the gate of t_j in $P_{1,4}$, contrary to $p \notin I(t_i, t_j)$. Thus $t_j \in P_1$. Let u and w be the gates of t_i and t_j in c_v and




Figure 3: First case in the proof of Proposition 2

c_h . Pick some rectilinear geodesics $\gamma(t_i, u)$, $\gamma(t_j, w)$, and $\gamma(t_i, t_j)$ between the pairs t_i, u ; t_j, w , and t_i, t_j , respectively; for an illustration, see Fig. 4. Since $p \notin I(t_i, t_j)$, the geodesic $\gamma(t_i, t_j)$ cannot share common points with both segments $[u, p]$ and $[p, w]$. Let $\gamma(t_i, t_j) \cap [u, p] = \emptyset$. Let u' be a closest to u point of $\gamma(t_i, u) \cap \gamma(t_i, t_j)$. Necessarily $u' \neq u$. Let w' be a closest to w point of $\gamma(t_j, w) \cap \gamma(t_i, t_j)$. Since P is a simple polygon, the region of the plane bounded by $[u, p]$, $[p, w]$, the part of $\gamma(t_i, u)$ between u, u' , the part of $\gamma(t_i, t_j)$ between u', w' , and the part of $\gamma(t_j, w)$ between w', w , is contained in P . Let $[u'', u]$ be the last link in the subpath of $\gamma(t_i, u)$ between u' and u . Then for some $\delta > 0$, the segment $[v', v'']$ belongs to P , where $v' \in [u'', u]$, $v'' \in [p, w]$ and $d(u, v') = d(p, v'') = \delta$. This contradicts that p is the gate of t_i in c_h .

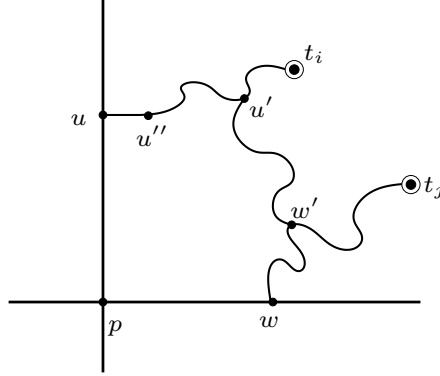


Figure 4: Second case in the proof of Proposition 2

Now, suppose that the gates q and z of t_i in c_v and c_h are different from p . Let u be the furthest from t_i point of $I(t_i, q) \cap I(t_i, z)$. Pick the rectilinear geodesics $\gamma(u, q)$ and $\gamma(u, z)$ between u, q and u, z . Let $[q', q]$ and $[z', z]$ be the last links of these paths. Let q'' be the point of c_h with the same x -coordinate as q' . Let z'' be the point of c_v with the same y -coordinate

as z' . Since P is a simple polygon, the region between $[q, p], [z, p]$ and $\gamma(u, q), \gamma(u, z)$ belongs to P . Moreover, since $q, z \in I(p, t_i)$ and $I(p, t_i)$ is convex, this region necessarily belongs to $I(p, t_i)$. In particular, both rectangles $R' = [q', q, p, q'']$ and $R'' = [z', z, p, z'']$ belong to $I(p, t_i)$. As we already stated, all points of T are outside P_3 . Let g_v be the closest to p gate in c_v of a point of $T \cap P_2$. Let also g_h be the closest to p gate in c_h of a point of $T \cap P_4$. Since $p \notin \cup_{j=1}^n I(t_i, t_j)$, we conclude that g_v and g_h are different from p . Let $0 < \delta < \min\{d(p, g_v), d(p, g_h), d(z', z), d(q', q)\}$. Consider a point $p' \in R' \cap R''$ whose coordinates differ by δ from those of p . Since $p' \in I(p, t_i)$, we obtain that $d(p', t_i) = d(p, t_i) - 2\delta$. For any other t_j we have $d(p', t_j) \leq d(p, t_j)$. This contradicts that p is Pareto. \square

A subset S of P is *ortho-convex* [11] if the intersection of S with any axis-parallel cut of P is connected.

Lemma 3.1 $\mathcal{P}_{d_1}(T)$ is a closed ortho-convex set of P .

Proof. Proposition 2 shows that $\mathcal{P}_{d_1}(T)$ is closed. Since the intersection preserves ortho-convexity, by Proposition 2 it suffices to show that for any point $t_i \in T$ the union $\cup_{j=1}^n I(t_i, t_j)$ is orthogonally convex. Pick any horizontal or vertical cut c of P . If for some point $t_j \in T$ the intersection $c' := c \cap I(t_i, t_j)$ is non-empty, then one can easily show that $c' = [g(t_i, c), g(t_j, c)]$, where $g(t_i, c)$ and $g(t_j, c)$ are the gates of the points t_i, t_j in c . Since the intersection of c with $\cup_{j=1}^n I(t_i, t_j)$ is the union of such segments $[g(t_i, c), g(t_j, c)]$ and all these segments share the point $g(t_i, c)$, this union is necessarily a segment as well, thus we are done. \square

3.2 The algorithm

Now, we describe the algorithm for constructing the Pareto envelope $\mathcal{P}_{d_1}(T)$ for a set T of n terminals in a simple rectilinear polygon P with m vertices. The algorithm uses Chazelle's algorithm for computing all vertex-edge visible pairs of a simple polygon [3] and the optimal point location methods [9, 14]. Using Chazelle's algorithm, we derive a decomposition of the polygon P into rectangles, employing only horizontal cuts which pass through the vertices of P . Using the optimal point location methods [9, 14] we compute in total $O(n \log m)$ time which rectangles of the decomposition contain the terminals (notice that the induced subdivision is monotone, whence the point location structure can be built in linear time). At the next step, we sort by y all terminals from each rectangle. With these sorted lists, we refine the initial subdivision by dividing each rectangle containing terminals with the horizontal cuts passing via terminals. The dual graph of this decomposition \mathcal{D} is a tree \mathcal{T} : the nodes of the tree are the rectangles of \mathcal{D} , and two nodes in \mathcal{T} are adjacent iff the corresponding rectangles are bounded by the common cut. We suppose that \mathcal{T} is rooted at some rectangle. Any cut c of our subdivision divides the polygon P into two subpolygons P'_c and P''_c which correspond to two subtrees \mathcal{T}'_c and \mathcal{T}''_c of \mathcal{T} . It can be easily shown that if $P''_c \cap T = \emptyset$ (in this case we say that P''_c is T -empty), then $\mathcal{P}_{d_1}(T)$ is contained in $P'_c \cup c$ (any point of P''_c is dominated by its gate in c). By processing the tree \mathcal{T} , in linear time we can remove all T -empty subpolygons and their corresponding subtrees. We will denote the resulting polygon, subdivision, and tree

also by P, \mathcal{D} , and \mathcal{T} . The resulting decomposition \mathcal{D} and its tree \mathcal{T} can be constructed in time $O(m + n(\log n + \log m))$. If all terminals are vertices of P , then we avoid the application of point location methods and ranking of terminals, requiring only $O(n + m)$ time.

Given a non-root rectangle R , we denote by e'_R and e''_R the horizontal sides of R , so that e'_R separates R from the root of \mathcal{T} . The set of gates of all terminals in R can be partitioned into the subset G'_R of gates located on e'_R and the subset G''_R of gates located on e''_R . Let $g'_l(R), g'_r(R)$ be the leftmost and the rightmost points from G'_R and let $g''_l(R), g''_r(R)$ be the leftmost and the rightmost points from G''_R . We will show below how to compute the four *extremal gates* $g'_l(R), g'_r(R), g''_l(R)$ and $g''_r(R)$ for all rectangles $R \in \mathcal{D}$ in total linear time.

In the first stage, we consider only the gates of terminals that belong to R or to rectangles that are descendants of R in \mathcal{T} (called *lower terminals*) and, in the second stage, the gates of the remaining terminals (called *upper terminals*). At the beginning of the procedure, for each $R \in \mathcal{D}$, we assign to $g'_l(R), g'_r(R), g''_l(R), g''_r(R)$ the leftmost or the rightmost terminal of e'_R and e''_R (or a conventional null value if such terminals do not exist). This step requires $O(n + m)$ total time. The first stage is implemented by applying the following simple upward traversal of \mathcal{T} . If R is a leaf of \mathcal{T} , then $g'_l(R), g'_r(R), g''_l(R), g''_r(R)$ have been defined at the initialisation step. Now suppose that R is an inner node of \mathcal{T} . For each child rectangle R' of R we already computed the leftmost and the rightmost gates of all its lower terminals. Then we compute the gates in R of these four points of R' (since R and R' are adjacent rectangles this requires constant time). Each time we compute a gate g in R (suppose it lies on e'_R , the case of e''_R is identical) such that either (i) $g'_l(R)$ and $g'_r(R)$ are null, (ii) g is to the left of current $g'_l(R)$ or (iii) g is to right of current $g'_r(R)$, we update the value of $g'_l(R)$ and $g'_r(R)$ accordingly.

The second stage for upper terminals is analogous, except that (i) the tree \mathcal{T} is traversed in a downward way and (ii) an extremal gate of the parent rectangle R' of R is taken into account for computing the extremal gates in R only if it is the gate of an upper terminal (with respect to R). Since, by construction of \mathcal{D} , any point of P belongs to at most two rectangles of \mathcal{D} , a gate in R' is defined by an upper terminal with respect to R if and only if it does not belong to R . If R is the root of \mathcal{T} , then there are no upper terminals and the gates $g'_l(R), g'_r(R), g''_l(R), g''_r(R)$ have been already correctly computed in Phase 1. Now suppose that R is an inner node of \mathcal{T} and that the extremal gates in the parent rectangle R' of R have been already computed. Then among such gates we select those which do not belong to R (as we noticed above, these gates are defined by upper terminals with respect to R) and compute their gates in R . Each time when we find a gate g in R (it lies on e'_R because R' is the parent of R) such that either (i) $g'_l(R)$ and $g'_r(R)$ are null, (ii) g is to the left of current $g'_l(R)$ or (iii) g is to the right of current $g'_r(R)$, then we update the values of $g'_l(R)$ and $g'_r(R)$ accordingly.

At the end of these two stages, we claim that the points $g'_l(R), g'_r(R), g''_l(R), g''_r(R)$ have been correctly computed for each rectangle R of \mathcal{D} . The total complexity of this procedure is linear in the number of terminals and rectangles, thus $O(n + m)$.

The correctness of the first stage is proved by induction. Pick $R \in \mathcal{D}$ and suppose that the procedure correctly computes the four extremal gates for its children. Suppose that there is a lower terminal t whose gate g_t in R lies on e'_R to the left of $g'_l(R)$. Suppose that t lies in

the subpolygon defined by the child rectangle R' . The choice of g_t implies that $t \notin R'$ and therefore the gate of t in R coincides with the gate in R of its gate in R' . From the algorithm and the induction hypothesis we conclude that the gate g'_t of t in R' cannot be the leftmost gate in R' , otherwise g'_t will be involved in the computation of $g'_l(R)$. Therefore g'_t lies to the right of $g'_l(R')$. Suppose that $g'_l(R')$ is the gate g'_v in R' of a terminal v . The procedure ensures that $g'_l(R)$ coincides or lies to the left of the gate g_v of g'_v in R . Since g_t is to the left of $g'_l(R)$, any two shortest paths between g'_t, g_t and g'_v, g_v have a common point z . By definition of g_t and g_v , the points z and g_t lie on a common shortest path between g'_t and g_v , while z and g_v lie on a common shortest path between g'_v and g_t , and we conclude that $g_t = g_v$, a contradiction.

The correctness proof of the second stage is analogous. We pick a rectangle $R \in \mathcal{D}$, its parent rectangle R' in \mathcal{T} and suppose that the procedure correctly computed the four extremal gates for R' . Then, since every shortest path between an upper terminal (with respect to R) and any point in R passes via R' , the gates in R of the upper terminals are exactly the gates in R of their gates in R' . Now, we proceed as in the proof of the first stage: we suppose that there is an upper terminal t whose gate g_t in R lies on e'_R to the left of $g'_l(R)$. The gate g'_t of t in R' cannot be the leftmost gate in R' , otherwise g'_t will be involved in the computation of $g'_l(R)$. Therefore g'_t lies to the right of $g'_l(R')$. Suppose that $g'_l(R')$ is the gate g'_v in R' of a terminal v . The procedure ensures that $g'_l(R)$ coincides or lies to the left of the gate g_v of g'_v in R . Since g_t is to the left of $g'_l(R)$, any two shortest paths between g'_t, g_t and g'_v, g_v have a common point z , leading to a similar contradiction as in the first case.

For each rectangle $R \in \mathcal{D}$, given the quadruple of gates $Q_R = \{g'_l(R), g'_r(R), g''_l(R), g''_r(R)\}$, at the next step we compute the Pareto envelope $\mathcal{P}_{d_1}(Q_R)$ of Q_R . It consists of a box B_R having its horizontal sides on the sides of R and two vertical segments which are incident either to two points of the quadruple lying on the same horizontal side of R or to two opposite points lying on different horizontal sides of R (one or both these segments can be degenerated); see Fig. 5 for the three occurring cases. In general, these segments do not necessarily belong to the final Pareto envelope $\mathcal{P}_{d_1}(T)$. On the other hand, as we will show below, B_R minus its horizontal sides is exactly the set $\mathcal{P}_{d_1}(T) \cap R^0$, where $R^0 := R \setminus (e'_R \cup e''_R)$ (clearly, the horizontal sides of B_R belong to $\mathcal{P}_{d_1}(T)$ as well because $\mathcal{P}_{d_1}(T)$ is closed). Now, if we consider any horizontal cut c , then we show that $\mathcal{P}_{d_1}(T) \cap c$ is the smallest segment $s_c \subseteq c$ spanned by the terminals and/or the horizontal sides of all boxes B_R located on c . Clearly, having at hand the four gates of each rectangle, the sets B_R and s_c can be determined in $O(n + m)$ time. To conclude, it remains to prove the correctness of the two last steps of the algorithm.

Lemma 3.2 $\mathcal{P}_{d_1}(T) \cap R^0 = B_R \cap R^0$.

Proof. Pick any point $p \in B_R \cap R^0$. To show that p is Pareto, by Proposition 2 it suffices to show that $p \in \cup_{j=1}^n I(t_i, t_j)$ for any $t_i \in T$. Suppose that the gate g of t_i in R belongs to e'_R . Consider the points $g''_l(R)$ and $g''_r(R)$ and suppose that they are the gates of the terminals t_j and t_k , respectively. First suppose that the vertical projection p' of p on e''_R belongs to the segment $[g''_l(R), g''_r(R)]$. Now if the projection of g on e''_R is to the left of that of p , then $p \in I(t_i, t_k)$, otherwise $p \in I(t_i, t_j)$. On the other hand, if $p' \notin [g''_l(R), g''_r(R)]$, then we can suppose without loss of generality that p' is to the right of $g''_r(R)$. Since $p \in B_R$, the vertical




Figure 5: The three shapes of $\mathcal{P}_{d_1}(Q_R)$

projection of p on e'_R is to the left of $g'_l(R)$. Then again $p \in I(t_i, t_j)$. This establishes that $B_R \cap R^0 \subseteq \mathcal{P}_{d_1}(T)$. Conversely, we assert that any point $q \in R^0 \setminus B_R$ is not Pareto. Suppose without loss of generality that q is located to the left of B_R and that $g''_l(R)$ is a corner of B_R . Let t_i be the terminal defining the gate $g'_r(R)$. Then it can be easily seen that $q \notin I(t_i, t_j)$ for any other terminal t_j . Indeed, this is evidently true if the gate of t_j in R belongs to e'_R . On the other hand, if this gate belongs to e''_R , then this gate is to the right of g''_l , thus to the right of q , showing that $q \notin I(t_i, t_j)$. This concludes the proof. \square

Lemma 3.3 $\mathcal{P}_{d_1}(T) \cap c = s_c$.

Proof. Since by Lemma 3.1 the set $\mathcal{P}_{d_1}(T)$ is closed and ortho-convex, the definition of s_c implies that $s_c \subset \mathcal{P}_{d_1}(T)$. Let $s_c = [u, v]$, where u is to the left of v . Pick any point $q \in c$ to the left of u . Suppose by way of contradiction that $q \in \mathcal{P}_{d_1}(T)$. By Proposition 2, for any $t_i \in T$ there exists $t_j \in T$ such that $q \in I(t_i, t_j)$. This implies that q is comprised between the gates of the points t_i and t_j on c , in particular, one of these gates coincides with q or is to the left of q . Let t be the terminal having the leftmost gate g_t in c . Suppose that $I(t, g_t)$ crosses the rectangle R having one horizontal side on c , say e'_R . Then necessarily the gate of t in R is the point $g''_l(R)$. The definition of s_c and the choice of q implies that $g''_l(R)$ is not a corner of the box B_R . Let u' be the corner of B_R incident to $g''_l(R)$ in the Pareto envelope $\mathcal{P}_{d_1}(Q_R)$. From the definition of u we infer that $u \in e'_R$ and that u is a gate of some terminal. Then one can easily see that $[u', u]$ is a vertical side of B_R . By the definition of the four gates of Q_R and the form of $\mathcal{P}_{d_1}(Q_R)$ we deduce that either $u' = g''_r(R)$ or $u' = g'_l(R)$. In the first case all gates of G'_R are located to the right of u . Thus, if t' is the terminal defining the gate $g''_r(R)$, then q does not belong to the interval between t' and any terminal t_j having a gate in G'_R . Since q does not belong to $I(t', t_j)$ for any terminal having a gate in G'_R , we infer that $q \notin \cup_{j=1}^n I(t', t_j)$. On the other hand, if $u' = g'_l(R)$, then necessarily $g''_r(R)$ exists and is located to the right of u' . By definition all gates from G'_R are located between u and $g'_r(R)$, thus to the right of q . This implies again that $q \notin \cup_{j=1}^n I(t', t_j)$, where t' is the terminal defining the gate $g''_r(R)$. Hence, in both cases Proposition 2 yields $q \notin \mathcal{P}_{d_1}(T)$, establishing our assertion. \square

Summarizing the results of this section, we obtain the following algorithm and theorem for computing $\mathcal{P}_{d_1}(T)$:

Algorithm Pareto-RP(P, T)**Input:** A simple rectilinear polygon P and a set $T \subset P$ of n terminals.**Output:** The Pareto envelope $\mathcal{P}_{d_1}(T)$.

1. Compute the subdivision \mathcal{D} into rectangles of the polygon P .
2. Refine the subdivision \mathcal{D} by dividing each rectangle containing terminals with the horizontal cuts passing via these terminals.
3. For each rectangle R of \mathcal{D} , compute the quadruple of gates $Q_R = \{g'_l(R), g'_r(R), g''_l(R), g''_r(R)\}$.
4. For each rectangle R of \mathcal{D} , compute the Pareto envelope $\mathcal{P}_{d_1}(Q_R)$ and the box B_R spanned by $R_0 \cap \mathcal{P}_{d_1}(Q_R)$.
5. For each horizontal cut c of \mathcal{D} , compute the segment s_c spanned by the terminals and/or the horizontal sides of all boxes B_R located on c .
6. Output the union of all boxes B_R and segments s_c for all rectangles R and horizontal cuts c of the subdivision \mathcal{D} .

Theorem 3.4 *The Pareto envelope of n terminals located in a simple rectilinear polygon P with m edges can be constructed in time $O(n + m(\log n + \log m))$ ($O(n + m)$ if all terminals are vertices of P).*

Acknowledgement. We would like to acknowledge the referee for finding an error in the first version of the paper.

References

- [1] M. Bridson, A. Haefliger, Metric Spaces of Non-Positive Curvature (Springer, Berlin, 1999).
- [2] G. Chalmet, L. Francis, and A. Kolen, Finding efficient solutions for rectilinear distance location problems efficiently, *Europ. J. Oper. Res.* 6 (1981) 117–124.
- [3] B. Chazelle, Triangulating a simple polygon in linear time, *Discrete Comput. Geom.* 6 (1991) 485–524.
- [4] V. Chepoi and F.F. Dragan, Computing a median point of a simple rectilinear polygon, *Inform. Process. Lett.* 49 (1994) 281–285.
- [5] V. Chepoi and K. Nouioua, Pareto envelopes in \mathbb{R}^3 under l_1 and l_∞ -distance functions, in: *Symposium on Computational Geometry 2007*, pp. 284–293.
- [6] V. Chepoi, K. Nouioua, Y.Vaxès, A rounding algorithm for approximating minimum Manhattan networks, *Theoret. Comput. Sci.* 390 (2008) 56–69 (extended abstract appeared in: *APPROX-RANDOM'2005*, pp. 40–51).

- [7] R. Durier, On Pareto optima, the Fermat-Weber problem, and polyhedral gauges, *Math. Progr.* 47 (1990) 65–79.
- [8] R. Durier and C. Michelot, Sets of efficient points in normed space, *J. Math. Anal. Appl.* 117 (1986) 506–528.
- [9] H. Edelsbrunner, L.J. Guibas, and J. Stolfi, Optimal point location in a monotone subdivision, *SIAM J. Comput.* 15 (1985) 317–340.
- [10] L. J. Guibas, J. Hershberger, D. Leven, M. Sharir, and R. E. Tarjan, Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons, *Algorithmica* 2 (1987) 209–233.
- [11] E. Fink and D. Wood, *Restricted-Orientation Convexity* (Springer-Verlag, Berlin, 2004).
- [12] P. Hansen, T. Perreux, and J.F. Thisse, Location theory, dominance and convexity : some further results. *Oper. Res.* 28 (1980) 1241–1250.
- [13] J. Hershberger and S. Suri, Matrix searching with the shortest-path metric, *SIAM J. Comput.* 26 (1997) 1612–1634.
- [14] D.G. Kirkpatrick, Optimal search in planar subdivisions, *SIAM J. Comput.* 12 (1983) 28–35.
- [15] H. W. Kuhn, A note on Fermat’s problem, *Math. Prog.* 4 (1973) 98–107.
- [16] B. Pelegrin and F.R. Fernandez, Determination of efficient points in multiple-objective location problems, *Nav. Res. Logistics* 35 (1988) 697–705.
- [17] R. Pollack, M. Sharir, and G. Rote, Computing the geodesic center of a simple polygon, *Discr. Comput. Geom.* 4 (1989) 611–626.
- [18] S. Suri, Computing geodesic furthest neighbors in simple polygons, *J. Comput. Syst. Sci.* 39 (1989) 220–235.
- [19] 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.
- [20] G.T. Toussaint, Computing geodesic properties inside a simple polygon, *Revue D’Intelligence Artificielle* 3(1989) 9–42.
- [21] A.C. Thompson, *Minkowski Geometry* (Encyclopedia of Mathematics and Applications, Vol. 63, Cambridge University Press, 1996).
- [22] M. van de Vel, *Theory of Convex Structures* (Elsevier, Amsterdam, 1993).
- [23] R. E. Wendell and A.P. Hurter, Location theory, dominance, and convexity. *Oper. Res.* 21 (1973) 314–320.
- [24] R. E. Wendell, A.P. Hurter, and T.J. Lowe, Efficient points in location theory. *AIEE Trans.* 9 (1973) 238–246.