Minimum weight convex Steiner partitions*

Adrian Dumitrescu† Csaba D. Tóth‡

May 24, 2009

Abstract

New tight bounds are presented on the minimum length of planar straight line graphs connecting n given points in the plane and having convex faces. Specifically, we show that the minimum length of a convex Steiner partition for n points in the plane is at most $O(\log n/\log\log n)$ times longer than a Euclidean minimum spanning tree (EMST), and this bound is the best possible. Without Steiner points, the corresponding bound is known to be $\Theta(\log n)$, attained for n vertices of a pseudo-triangle. We also show that the minimum length convex Steiner partition of n points along a pseudo-triangle is at most $O(\log\log n)$ times longer than an EMST, and this bound is also the best possible. Our methods are constructive and lead to $O(n\log n)$ time algorithms for computing convex Steiner partitions having $O(n)$ Steiner points and weight within the above worst-case bounds in both cases.

1 Introduction

Geometric spanner networks for n given points in Euclidean plane have been studied extensively, particularly in the last 20 years [15, 35]. Some desirable properties of such spanners are constant stretch factor, constant degree, weight proportional to the Euclidean minimum spanning tree, and others. A *convex Steiner partition* for a set S of points in the plane is a planar straight line graph G where the vertex set of G contains S and the boundary of every face is a convex polygon; specifically, every bounded face is convex, and the unbounded face is the complement of a convex polygon. A *convex (non-Steiner) partition* has the additional property that every vertex of G is a point in S (i.e., there are no Steiner vertices). Similarly, one can consider *Steiner triangulations* and *triangulations* for a point set S. Clearly, every triangulation for S is a convex partition for S, and every Steiner triangulation for S is a Steiner convex partition for S. The *weight* of a planar straight line graph or network is the total Euclidean length of its edges. We denote by $W = W(S)$ the weight of a Euclidean minimum spanning tree (EMST) for a finite point set S. Every spanning network for S is at least as heavy as the Euclidean minimum Steiner tree of S, whose weight is known to be at least $W/2$ [19].[1]

A convex (non-Steiner) partition or triangulation for n points is a planar graph on n vertices, hence it has $O(n)$ edges. Since the weight of each edge is at most W, the weight of a convex partition or triangulation is trivially bounded by $O(Wn)$. This naïve bound is tight apart from a constant factor: Kirkpatrick [24]

*A preliminary version of this paper appeared in the *Proceedings of the 19th ACM-SIAM Symposium on Discrete Algorithms*, ACM Press, 2008, pp. 581–590.

†Department of Computer Science, University of Wisconsin-Milwaukee, WI 53201-0784, USA. E-mail: `ad@cs.uwm.edu`.
Supported in part by NSF CAREER grant CCF-0444188.

‡Department of Mathematics, University of Calgary, AB, T2N 1N4, Canada. E-mail: `cdtoth@ucalgary.ca`

[1]Gilbert and Pollak [19] conjectured that the Steiner ratio, i.e., the ratio of the weights of the EMST and the minimum Steiner tree, is at most $2/\sqrt{3}$. The validity of the proof of this conjecture given by Du and Hwang [11] has recently been questioned by O. de Wet [41].

exhibited sets of n points, shown in Figure 1(a), where the minimum weight of any convex partition is $\Omega(Wn)$.

Clarkson [10] proved that any set of n points in the plane admits a Steiner triangulation of weight $O(W \log n)$, and Eppstein [14] showed that this bound is the best possible. His construction consists of 4 vertices of a square and $n-4$ points evenly distributed along a circle placed in the interior of the square, see Figure 1(b). Both lower bound constructions consist of points along a reflex chain and a few other points off this chain. These constructions can be realized with a vertex set of $O(1)$ *pseudo-triangle*s, that is, polygons with exactly three convex vertices called *corners*, see Figure 1(c).

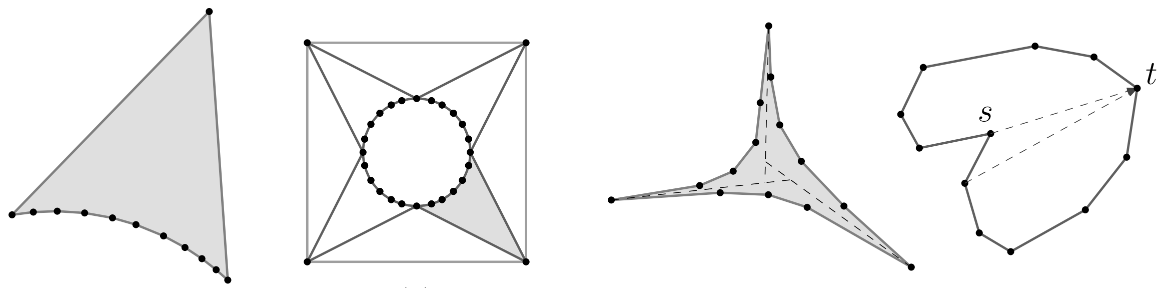

Figure 1: (a) n points for which any triangulation or convex partition has weight $\Omega(Wn)$. (b) $n+4$ points for which any Steiner triangulation has weight $\Omega(W \log n)$, and any convex Steiner partition has length $\Omega(W \log\log n)$. (c) A pseudo-triangle with the bisectors of its three corner-angles. (d) A nonconvex face where compass routing fails to route from s to t.

Our contribution. (i) Given n points along a pseudo-triangle, we prove that they admit a convex *Steiner* partition of weight $O(W \log\log n)$, and this bound is the best possible (Theorem 1). Recall that the weight of a minimum Steiner triangulation may be as large as $\Omega(W \log n)$ for n vertices of a pseudo-triangle.

(ii) Given n points in the plane, we prove that they admit a convex *Steiner* partition of weight $O(W \log n / \log\log n)$, and this bound is the best possible (Theorem 2). This is a $(\log\log n)$-factor improvement over the corresponding bound for minimum weight Steiner triangulations.

We prove our upper bounds constructively, using $O(n)$ Steiner points in both cases. Our methods lead to $O(n \log n)$ time algorithms in the real RAM model for computing convex Steiner partitions within these bounds.

Networks with constant stretch factor. A *spanner network* for set S of points in the plane is a connected planar straight line graph G whose vertex set contains S. The *vertex dilation* (also known as *stretch factor*) of G is the maximum ratio between the length of the shortest path in G and the Euclidean distance for any pair of points in S. In contrast, the *geometric dilation* of G is the maximum ratio between the length of the shortest path in G and the Euclidean distance for any two points, at vertices *or on edges* of G.

For the $\Omega(W \log n / \log\log n)$ lower bound, listed above in (ii), on the length of a minimum Steiner partition of a general point set, it is crucial that the unbounded face in a convex partition is required to be the complement of a convex set. If we drop this condition (but still require that all *bounded* faces be convex), then the minimum weight network would be the minimum Steiner tree, which has only a single, unbounded face. We show that the $\Omega(W \log n / \log\log n)$ lower bound holds even if we replace the convexity condition on the unbounded face by another condition on the stretch factor. We present (Theorem 3) n-element point sets for which any spanner network whose stretch factor is $o(n)$ and whose *bounded* faces are convex, has weight $\Omega(W \log n / \log \log n)$. The condition on the stretch factor rules out the EMST in some cases, for instance, for n equally spaced points along a circle.

It is known that given n points in the plane, there is a plane spanner network that (i) is *short*, namely, it has weight $O(W)$; (ii) is *small*, namely, it has $O(n)$ vertices and edges; and (iii) has constant vertex dilation [4, 31]. Recently, we have shown [12] that property (iii) can be strengthened by requiring constant geometric dilation rather than constant vertex dilation: There is a network plane spanner network with properties (i), (ii), and with constant geometric dilation. We observed that some of the faces of such networks are often non-convex. Theorem 3 shows that non-convex faces are in general unavoidable in a spanner network with properties (i)–(iii).

Motivation. *Location-based routing,* also known as *geographic routing*, has been studied extensively and is considered one of the most promising routing protocols in interconnection networks, sensor networks, and mobile networks [16, 23]. Even abstract networks are often embedded in Euclidean space (with "virtual coordinates") to be able to use location-based routing [28, 37, 39]. In geographic routing, each node stores its geographic coordinates, and makes routing decisions based on the coordinates of the packet destination, eliminating expensive routing tables. Variants differ in the number of bits carried along with a packet and in the amount of information stored at each node. A minimalist model is *compass routing* [27], where (i) the packets carry no routing information other than the target's coordinates, (ii) each node knows only its own coordinates and the directions of the adjacent edges, and (iii) they route each packet on a link whose direction is closest to the current direction of the target. Compass routing works on the Delaunay triangulation of a point set, as it was shown in [6]. As mentioned earlier, the Delaunay triangulation of n points in the plane may be $\Omega(n)$ times heavier than an EMST [24]. However, Bern *et al.* [2] showed that for every set S of n points in the plane, there is a super set S', $S \subseteq S'$, with $|S'| = O(n)$, such that the weight of the Delaunay triangulation of S' is $O(W \log n)$. This $O(W \log n)$ bound is the best possible due to Eppstein's lower bound on minimum weight Steiner triangulations [14].

A randomized version of compass routing guarantees delivery on any convex network [3]. The deterministic compass routing, however, can easily run into a loop: it fails for some source-destinations pairs in nonplanar graphs, and in planar straight line graphs with some bounded faces non-convex (Figure 1(d)). Moreover, any deterministic memoryless routing protocol fails on some convex network [5]. It is worth noting that many competing protocols, such as face routing [7], allow $O(1)$ memory carried along with each packet and they guarantee delivery on any plane graph; with $O(1)$ memory one can also guarantee that each packet is routed along a path at most constant-times longer than the Euclidean distance between source and destination [6, 21].

Related results. In this paper, we prove tight worst-case upper bounds on the minimum weight convex Steiner partition for n points in the plane in terms of n and W, and present polynomial-time algorithms for computing convex Steiner partitions within these bounds. No polynomial exact or approximation algorithm is known for computing the *minimum* convex Steiner partition for n given points. The problem is not known to be NP-hard, either, although we suspect that this is the case.

The related minimum weight triangulation problem has seen many developments in the last few years. Mulzer and Rote [34] proved that computing the minimum weight triangulation is NP-hard. Remy and Steger [40] gave a quasi-polynomial time approximation scheme for minimum weight triangulation. Levcopoulos and Krznaric [29] proposed $O(n \log n)$ time algorithms for constant-factor approximation of the minimum weight triangulation and the minimum weight convex partition, based on earlier work by Plaisted and Hong [38]. Eppstein [14] gave an $O(n \log n)$ time algorithm for computing a 316-factor approximation of the minimum weight Steiner triangulation of n points. Gudmundsson and Levcopoulos recently gave a tight $O(W \log n)$ upper bound on the minimum weight *pseudo-triangulation* of n points in the plane [20]. In a minimal pseudo-triangulation, every vertex is incident to a reflex angle, and so this is, in some sense, the opposite of convex partition.

The minimum weight convex partition problem restricted to the interior of a simple polygon can be solved exactly. Gilbert [18] and Klincsek [25] independently gave $O(n^3)$ time algorithms for computing a minimum weight convex partition of a simple polygon with n vertices using dynamic programming. Levcopoulos and Lingas [30] showed that the minimum weight of a convex Steiner partition of the interior of a simple polygon P with n vertices, m of which are reflex, is always $O(|P| \log n)$ and sometimes $\Omega(|P| \log m / \log \log m)$. Here $|P|$ denotes the perimeter of P. Our Theorem 4 in Section 3 shows that this lower bound is tight: Every simple polygon P with m reflex vertices admits a convex Steiner partition of weight $O(|P| \log m / \log \log m)$.

The minimum *number* convex partition problem asks for the minimum number of faces in a convex partition of n points. Knauer and Spillner [26] recently showed that any planar n-element point set admits a convex partition with at most $\frac{15n-24}{11}$ faces (improving an earlier bound of $\frac{10n-18}{7}$ by Neumann-Lara *et al.* [36]); García-López and Nicolás [17] gave a lower bound construction of $\frac{12n}{11} - 2$ for $n \geq 4$. Knauer and Spillner [26] also gave a polynomial time $\frac{30}{11}$-approximation for the minimum number convex partition problem. No corresponding results are known for the minimum number convex *Steiner* partition problem. Restricted to simple polygons, these problems have efficient solutions. Keil and Snoeyink [22] gave an $O(n^3)$ time algorithm for computing the minimum number convex partition of a simple polygon with n vertices; this problem is NP-hard for polygons with holes [32]. Chazelle and Dobkin [9] gave an $O(n^3)$ time algorithm for the minimum number convex Steiner partition of a simple polygon with n vertices.

Definitions and notations. If A is a finite set, let $\#A$ denote the cardinality of A. For a polygonal curve γ, let $|\gamma|$ denote the length (or weight) of γ. For a polygon P, let $|P|$ denote the perimeter of P. A *convex* chain is a polygonal chain whose vertices are consecutive vertices of some convex polygon. A *reflex* chain is a convex chain appearing on the boundary of a nonconvex polygon with all internal angles of the chain larger than π. Let γ be a convex (or reflex) chain with endpoints a and b; see Figure 2(a). The *turning angle* of γ is the angle in $[0, 2\pi]$ of the rays along the first and last segments of γ. The *width* of γ is the width of the smallest parallel strip that contains γ and is parallel to the segment ab connecting the endpoints of γ. Denote by $h(\gamma)$ the half-plane containing γ determined by the line other than ab bounding this strip. A *pseudo-triangle* is a simple polygon with exactly three vertices with interior angles less than π, called *corners*; see Figure 1(c).

2 Vertex sets of pseudo-triangles and reflex chains

In this section, we prove a tight bound on the minimum weight convex Steiner partition for a vertex set of a pseudo-triangle. The corners partition the pseudo-triangle into three *reflex chains* (some of which may consist of a single segment). Each reflex chain is a polygonal path of reflex internal angles and total turning angle of at most π. If all three reflex chains of a pseudo-triangle P are straight line segments then P is a triangle; if two reflex chains are straight line segments then P is a *one-chain pseudo-triangle*; see Figure 1(a).

Theorem 1

(i) *For every set S of n points lying along a pseudo-triangle, a convex Steiner partition of weight $O(W \log \log n)$ with $n + O(\sqrt{n})$ Steiner points can be computed in $O(n)$ time.*

(ii) *For every n, there is a set S of n points along a pseudo-triangle such that any convex Steiner partition for S has weight $\Omega(W \log \log n)$.*

Upper bound. We construct a convex Steiner partition for S by augmenting the vertex set S with new (Steiner) vertices and new edges as follows. Refer to Figure 1(c). Add the 3 edges of the convex hull $\text{conv}(S)$ of S to the network. Each edge has weight at most W, hence their total weight is $O(W)$. Since the three polygons between the convex hull of the pseudo-triangle and the three reflex chains are convex, it suffices to partition the interior of the pseudo-triangle. First we reduce the problem to one-chain pseudo-triangles. Note that the weight of a reflex chain of P is less than twice the diameter of P, hence $|P| = O(W)$.

Lemma 1 *The interior of every pseudo-triangle P with n vertices can be partitioned into three one-chain pseudo-triangles and a (possibly degenerate) triangle along three line segments of total weight $O(|P|)$.*

Proof. The bisector of an interior angle of P at a corner separates the two adjacent reflex chains of P. The bisectors of the three corners bound a (possibly degenerate) triangle Δ lying in the interior of P; see Figure 1(c). Consider the line segment along each bisector between a corner and its intersections with the other two bisectors. These segments jointly partition the interior of P into three one-chain pseudo-triangles (each adjacent to a reflex chain of P) and Δ. Each segment is shorter than a diagonal of P, hence its weight is at most $|P|/2$. The total weight of the three segments is $O(|P|)$. □

Lemma 2 *The width of a convex chain γ of turning angle α, $0 < \alpha \leq \frac{\pi}{2}$, is at most $\frac{\alpha}{4}|\gamma|$.*

Proof. Draw a circle in which ab is a chord of inscribed angle $\pi - \alpha$. Note that γ must lie in the circular disk cap bounded by the chord ab, otherwise its turning angle is more than α. The width of this cap is $(|ab|/2)\sin(\alpha/2) < \alpha|\gamma|/4$, hence this is also an upper bound on the width of γ. □

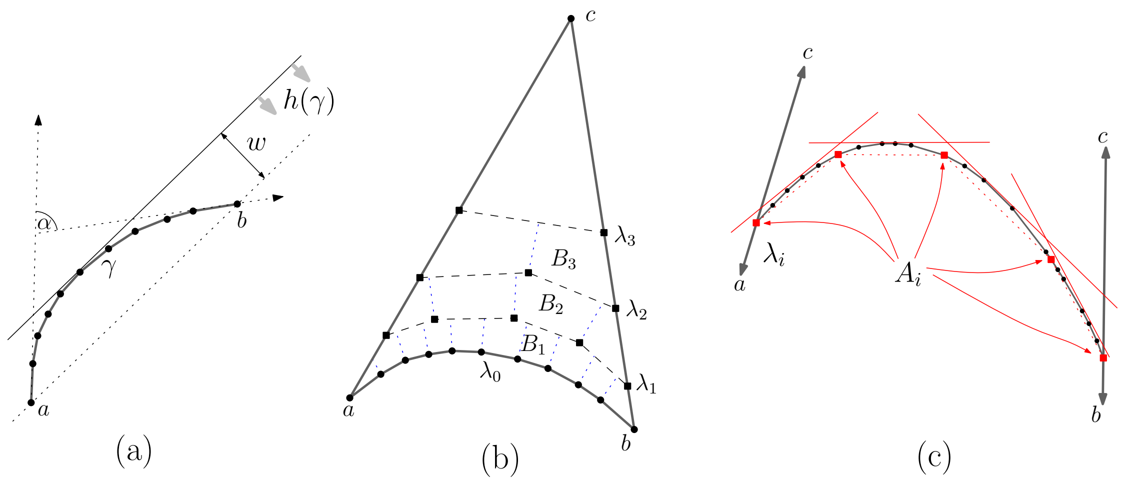

Figure 2: (a) The a reflex chain of width w and turning angle α. (b) A one-chain pseudo-triangle with corners a, b, and c, and with reflex chains λ_0, λ_1, λ_2, and λ_3. (c) In step i, we choose a vertex set A_i, halfplanes $h(\gamma)$ for $\gamma \in \Gamma_i$.

Lemma 3 *The interior of every one-chain pseudo-triangle P with n vertices has a convex Steiner partition of weight at most $O(|P| \log \log n)$ with at most $n+O(\sqrt{n})$ Steiner points. Such a partition can be computed in $O(n)$ time.*

Proof. Denote the corners of a one-chain pseudo-triangle P by a, b, and c such that the reflex chain λ_0 between a and b has $n-3$ internal vertices, and the other two reflex chains are the line segments ac and bc; see Figure 2(b). We partition the interior $\mathrm{int}(\mathrm{P})$ of P by pairwise non-crossing reflex chains λ_i, $i = 1, 2, \ldots, t = O(\log\log n)$. Each reflex chain connects segments ac and bc, and λ_i lies in the interior of the one-chain pseudo-triangle bounded by λ_{i-1}, ac, and bc. It is clear that the weight of each reflex chain is at most $|ac| + |bc|$. The last reflex chain, λ_t, will be a single segment, and so the portion of $\mathrm{int}(\mathrm{P})$ bounded by λ_t, ac, and bc is a triangle. Once these polygonal chains are constructed, we still need to subdivide the faces $B_i \subset \mathrm{int}(\mathrm{P})$ between consecutive reflex chains, λ_{i-1} and λ_i, into convex faces. Note that the vertices of λ_i are *convex* vertices of B_i, and the internal vertices of λ_{i-1} are *reflex* vertices of B_i. Partition B_i by rays emanating from each reflex vertex of B_i and subdividing the reflex angles into convex angles. We will choose the chains λ_i and the rays that partition B_i such that the portion of the rays lying in B_i have total weight proportional to the weight of λ_i for every $i = 1, 2, \ldots t$.

We define the reflex chains λ_i, $i = 0, 1, \ldots, t$, inductively. λ_0 is the given reflex chain of the one-chain pseudo-triangle P. Assume that λ_i is given, and we need to construct λ_{i+1}. We will partition λ_i into a set Γ_i of subchains ordered along λ_i such that the width of each subchain $\gamma \subset \lambda_i$, where $\gamma \in \Gamma_i$, is at most the average weight of the segments in γ. For every i, let m_i denote the number of internal vertices of λ_i, and let $\alpha_i < \pi$ denote its turning angle.

Initially, we have $i = 0$; λ_0 is the reflex chain of P, and $m_0 = n - 3$. As long as $m_i > 0$, construct the chain λ_{i+1} as follows. We choose a subset A_i of vertices of λ_i, and partition λ_i into a set Γ_i of subchains, each lying between consecutive vertices in A_i. Choose the splitting points A_i by the following simple algorithm: Put the endpoint $\lambda_i \cap bc$ of λ_i into A_i. Traverse λ_i from $\lambda_i \cap bc$ to $\lambda_i \cap ac$. Denoting by $\gamma_v \subset \lambda_i$ the subchain between the previous vertex of A_i and a vertex $v \in \lambda_i$, put v into A_i if γ_v is the maximal subchain with at most $\lfloor \sqrt{m_i} \rfloor$ segments and turning angle of at most $\alpha_i/\sqrt{m_i}$. Finally, put the endpoint $\lambda_i \cap ac$ into A_i (this process is schematically shown in Figure 2(c)). The number of subchains of λ_i created in this way is $|\Gamma_i| \le 2\sqrt{m_i}$.

Let C_i be the intersection of the halfplanes $h(\gamma)$ for $\gamma \in \Gamma_i$. Let λ_{i+1} be the portion of the boundary of C_i lying in P. Finally, for each vertex v of each subchain $\gamma \in \Gamma_i$, partition B_i by the ray emanating from v and perpendicular to $h(\gamma)$. In particular, we draw two rays at the common endpoint of any two consecutive subchains in Γ_i. This completes the construction of the chain λ_{i+1} and the partition of B_i into convex faces.

The weight of each ray drawn at a vertex of a subchain $\gamma \in \Gamma_i$ is at most the width of γ, which is at most $\alpha_i|\gamma|/(4\sqrt{m_i}) \le \frac{\pi}{4}|\gamma|/\sqrt{m_i}$ by Lemma 2. By construction, each $\gamma \in \Gamma_i$ has at most $\sqrt{m_i}$ vertices. Hence the total weight of the rays that partition B_{i+1} into convex faces is

$$\sum_{\gamma\in\Gamma_i} (\lfloor\sqrt{m_i}\rfloor + 1) \cdot \frac{\pi}{4} \cdot \frac{|\gamma|}{\sqrt{m_i}} \le \frac{\pi}{2} \cdot \sum_{\gamma\in\Gamma_i} |\gamma| = \frac{\pi}{2}|\lambda_i| \le \frac{\pi}{2}|P|.$$

Since $m_{i+1} \le |\Gamma_i| \le 2\sqrt{m_i}$, we also have $t = O(\log\log n)$. Summing over all $i = 0, 1, \ldots, t-1$, the total weight of the resulting convex Steiner partition of P is $O(|P| \cdot t) = O(|P| \log\log n)$, as required. The number of Steiner points is upper bounded by

$$\sum_{i=0}^{t}(2 + m_i) = 2(t+1) + \sum_{i=0}^{t} m_i \le O(\log\log n) + \sum_{i=0}^{t} 2^{1-1/2^i} \cdot (n-3)^{1/2^i} = n + O(\sqrt{n}).$$

Each polygonal chain λ_{i+1} is constructed in a single traversal of λ_i, in $O(m_i)$ time. Between the chains λ_i and λ_{i+1}, every segment splitting a reflex angle at an internal vertex of λ_i hits ac, bc, or one of two possible edges of λ_{i+1}, so each ray can be computed in $O(1)$ time. Hence, the total runtime of the partitioning algorithm is $O(\sum_{i=0}^{t} m_i) = O(n)$. □

Lemma 4 *The interior of every pseudo-triangle P with n vertices has a convex Steiner partition of weight $O(|P|\log\log n)$ with $n+O(\sqrt{n})$ Steiner points. Such a partition can be computed in $O(n)$ time.*

Proof. Partition the pseudo-triangle P by Lemma 1 into three one-chain pseudo-triangles and a (possibly degenerate) small interior triangle, as in Figure 1(c). This partition has weight $O(|P|)$ and at most 3 Steiner points. By Lemma 4, the interior of each pseudo-triangle can be partitioned into convex faces using $n+O(\sqrt{n})$ Steiner points and edges of total weight $O(|P|\log\log n)$ in $O(n)$ time. □

This completes the proof of the upper bound in part (i) of Theorem 1.

Lower bound construction with points along a reflex chain. We prove that the weight of every convex Steiner partition for a set S of $n+4$ points arranged as indicated in Figure 1(b) is at least $\Omega(W\log\log n)$. Since this construction can be tiled with 5 convex faces and 4 congruent pseudo-triangles as shown in the figure, we also obtain an $\Omega(W\log\log n)$ lower bound for the minimum Steiner partition of n vertices of a pseudo-triangle.

Consider an integer n such that $\log\log n$ is even. Let S_0 be a set of n evenly spaced points on a circle of unit radius centered at the origin o; let S be the union of S_0 and the 4 vertices of a square Q at points $(\pm 2, \pm 2)$; see Figure 1(b). Clearly $Q = \mathrm{conv}(S)$. Denoting by W_n the length of an EMST for S, observe that $\lim_{n\to\infty} W_n = 2\pi + 4(2\sqrt{2}-1) \approx 13.60$. For sufficiently large n, an EMST consists of $n-1$ edges of $\mathrm{conv}(S_0)$ and 4 edges, each of length $2\sqrt{2}-1$, connecting the vertices of Q to $\mathrm{conv}(S_0)$.

Let G be a convex Steiner partition for S. For $r>0$, denote by $C(r)$ (resp. $D(r)$) the circle (resp. disk) of radius r centered at o. For $0<r<R$, let $K(r,R) = D(R)\setminus D(r)$ be the annulus between the circles $C(R)$ and $C(r)$. We construct inductively a sequence of $k = (\log\log n)/2$ concentric circles of radii $1 = r_0 < r_1 < \ldots < r_k \le 3/2$ and show (in Lemma 7) that the length of the portion of G lying in each annulus $K(r_i, r_{i+1})$ is $\Omega(1)$. This immediately implies $|G| = \Omega(\log\log n)$.

We say that a set of m points on a circle C is *dense* if every arc of C of size (measured by the angle of apex at the center of C) at least $4\pi/m$ contains at least one of the points. Along each circle $C(r_i)$, we choose a dense set $A_i \subset G\cap C(r_i)$, which consists of some intersection points of the circle $C(r_i)$ with vertices or edges of G. Initially, let $A_0 = S_0$ be the a dense set of $m_0 = n$ points of S along $C(r_0)$. We next describe how we choose the radius r_i and the point set A_i for $i = 1,2,\ldots,k$.

We will choose the widths $\varepsilon_i = r_{i+1} - r_i$ of the annuli to satisfy the following two conditions:

1. ε_i should be small enough so that the circle $C(r_{i+1})$ intersects G in a large dense set of points.
2. ε_i should be large enough so that the length of the portion of G in $K(r_i, r_{i+1})$ is at least $\Omega(1)$.

We choose the radii r_i so that they satisfy the recurrence $r_{i+1} = r_i + 1/(9m_i)$ for $i = 1,2,\ldots,k$. The following lemma ensures that we find a dense set of at least $\sqrt{m_i}$ points in $C(r_{i+1})\cap G$ for every $i = 0,1,\ldots,k-1$, if $r_{i+1} < 3/2$. This shows that

$$m_i \ge n^{2^{-i}}, \tag{1}$$

and we repeat the argument $k = (\log\log n)/2$ times, where $m_i \ge 10$ holds for every $0 \le i \le k$.

Lemma 5 *For any two points $p,q \in C(r_i)$ with $\angle poq \ge 1/\sqrt{m_i}$, the chords of the circle $C(r_{i+1})$ whose midpoints are p and q, respectively, are disjoint.*

Proof. Consider two points $p,q\in C(r_i)$ and let $\alpha = \angle poq$. Recall that $r_{i+1} = r_i + \frac{1}{9m_i}$. If the chords of the circle $C(r_{i+1})$ with midpoints p and q, respectively, share a common endpoint, then

$$\cos\alpha = \frac{r_i}{r_{i+1}} = \frac{r_i}{r_i + \frac{1}{9m_i}}.$$

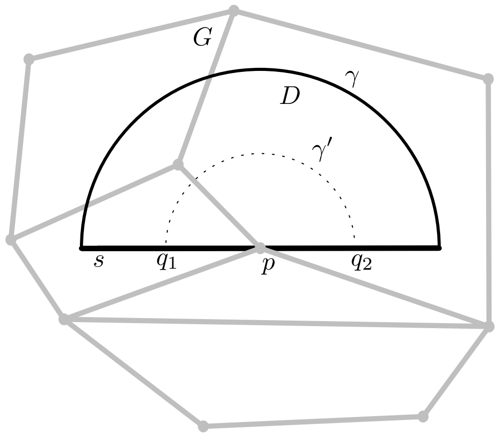

Figure 3: A half-disk D whose boundary (black) consists of a segment s and a half-circle γ, and a convex partition G (grey).

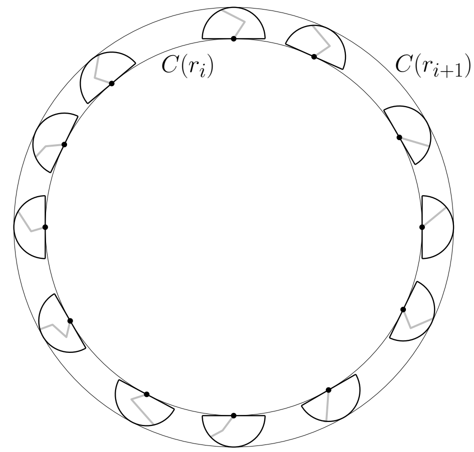

Figure 4: The graph G contains a portion of length at least ε_i in each half-disk lying in the annulus $K(r_i, r_{i+1})$.

From the Taylor expansion of cosine, for $\alpha \in \left(0, \frac{\pi}{2}\right]$ we have

$$1 - \frac{\alpha^2}{2} < \cos\alpha < 1 - \frac{\alpha^2}{2} + \frac{\alpha^4}{4!} \leq 1 - \frac{\alpha^2}{3}. \tag{2}$$

Hence, the chords with midpoints at p and q are disjoint if

$$1 - \frac{\alpha^2}{3} \leq \frac{r_i}{r_i + \frac{1}{9m_i}}. \tag{3}$$

The inequality holds for $\alpha = 1/\sqrt{m_i}$ and $r_i \geq 1$. $\square$

The following simple lemma is crucial for enforcing the convexity of faces in the partition.

Lemma 6 *Let D be a closed half-disk of radius r centered at p and bounded by a diameter segment s and a half-circle γ; see Figure 3. Let G be network such that p lies at a vertex or on an edge of G, and the half-disk D is covered by convex faces of G. Then the weight of the portion of G lying in D is at least r.*

Proof. It is enough to show that every halfcircle centered at p and lying in D intersects the network G, since then the lower bound on the weight follows by integration. Assume to the contrary, that there is a halfcircle γ' centered at p lying in D but disjoint from G. Denote the endpoints of γ' by $q_1, q_2 \in s$. The points q_1 and q_2 are on opposite sides of p and lie in the interior of some faces of G. They must lie in distinct faces of G, otherwise any point in the segment q_1q_2, including p, would be in the interior of the same face by convexity. However, they must lie in the same face of G, since they can be connected by curve γ' disjoint from G. This is a contradiction, hence every halfcircle centered at p and lying in D intersects G. $\square$

We next choose a set of points $A_i \subseteq C(r_i) \cap G$ inductively for every $i = 1, 2, \ldots k$. Partition $C(r_i)$ into $2\lfloor\sqrt{m_i/2}\rfloor^2 \leq m_i/2$ congruent arcs. By the induction hypothesis, A_i is a dense set along $C(r_i)$, and so there is a point of A_i in each arc. Let $A'_i \subset A_i$ consist of exactly one point of A_i from every $\lfloor\sqrt{m_i/2}\rfloor$-th arc. The cardinality of A'_i is $\#A'_i = 2\lfloor\sqrt{m_i/2}\rfloor \geq \sqrt{m_i}$ if $m_i \geq 2$. Draw the chords of circle $C(r_{i+1})$ with midpoints at points in A'_i. Note that each chord is tangent to $C(r_i)$. The size of an arc between consecutive points in A_i' is at least

$$2\pi\frac{\lfloor\sqrt{m_i/2}\rfloor-1}{2\lfloor\sqrt{m_i/2}\rfloor^2}\geq\frac{1}{\sqrt{m_i}},$$

if $m_i > 2$. By Lemma 5, these chords of $C(r_{i+1})$ are pairwise disjoint, and determine disjoint caps of $C(r_{i+1})$. Recall that we assumed that disk $D(r_{i+1})$ is covered by bounded faces whenever $r_{i+1} < 3/2$. By Lemma 6, every cap contains a point in $G \cap C(r_{i+1})$. Let A_{i+1} consists of one point of $G \cap C(r_{i+1})$ from each cap. The cardinality of A_{i+1} is $m_{i+1} = \#A_i' \geq \sqrt{m_i}$. Moreover, the point set A_{i+1} is dense, since every arc of size least $4\pi/m_{i+1}$ contains at least one entire cap, hence at least one point of A_{i+1}.

Lemma 7 *The total weight of the portion of G that lies in $K(r_i, r_{i+1})$ is at least $1/50$ for every $i = 0, 1, \dots, k-1$.*

Proof. Recall that the width of $K(r_i, r_{i+1})$ is $r_{i+1} - r_i = \frac{1}{9m_i}$. Partition $C(r_i)$ into $\lfloor m_i/2\rfloor$ congruent arcs. Since A_i is dense, each arc of size $4\pi/m_i$ contains a point of A_i. Pick one point of A_i from every other arc, and let A_i'' denote the resulting set of at least $\lfloor m_i/4\rfloor$ points. Any two consecutive points in A_i'' are separated by an arc of size at least $8\pi/m_i$. For every point $p \in A_i''$, construct a half-disk centered at p with radius ε_i and bounded by the tangent to $C(r_i)$ at p; see Figure 4. These half-disks are pairwise disjoint and lie in the annulus $K(r_i, r_{i+1})$. By Lemma 6, the weight of the portion of G lying in each half-disk of radius ε_i is at least ε_i. Hence, the length of $G \cap K(r_i, r_{i+1})$ is at least

$$\#A_i''\cdot\varepsilon_i \geq \left\lfloor\frac{m_i}{4}\right\rfloor\cdot\frac{1}{9m_i}\geq\frac{1}{50}.$$

□

A standard calculation using (1) gives that $m_k \geq \log n$ holds for sufficiently large n. Note also that we assumed in Lemma 7 that each annulus $K(r_i, r_{i+1})$ is covered by bounded (convex) faces of G. Since the sequence of m_i's is non-decreasing, we have

$$\sum_{i=0}^{k-1}\varepsilon_i \leq \frac{k}{9\log n}\leq\frac{1}{2},\ \text{ therefore }\ r_k = 1+\sum_{i=0}^{k-1}\varepsilon_i\leq\frac{3}{2}.$$

The disk of radius $3/2$ centered at the origin lies in Q. Since the annuli $K(r_i, r_{i+1})$, for $0 \leq i \leq k-1$, have pairwise disjoint interiors, and each annulus is covered by bounded faces of G, we can apply Lemma 7, and the total weight of G is at least $k/50 = \Omega(W \log\log n)$. This completes the proof of part (ii) of Theorem 1.

3 General point sets in the plane

In this section we derive asymptotically tight bounds on the minimum length of a convex Steiner partition for n points in the plane.

Theorem 2

(i) *For n points in the plane, there is a convex Steiner partition with $O(W \log n/\log\log n)$ weight and $O(n)$ Steiner points. Such a partition can be computed in $O(n \log n)$ time.*

(ii) *For every n, there is a set S of n points in the plane such that any convex Steiner partition for S has weight $\Omega(W \log n/\log\log n)$.*

We also extend the lower bound construction and show that there are n-element point sets for which any spanner network with all bounded faces convex and stretch factor $o(n)$ has weight $\Omega(W \log n / \log\log n)$.

Theorem 3 *For every n, there is an n-element point set S in the plane such that any spanner network for S, whose bounded faces are convex and whose stretch factor is $o(n)$, must have weight $\Omega(W \log n / \log\log n)$.*

Our upper bound in Theorem 2 is based on first partitioning the convex hull of an n-element point set into polygons, and then partitioning the polygons into convex faces. A convex Steiner partition of a polygon P is a planar straight line graph G, where the boundary of every bounded face of G is convex and all edges of P are covered by edges of G (in particular, there is no constraint on the unbounded face of G). For convex Steiner partitions of polygons, we prove the following.

Theorem 4 *Every polygon P with n vertices, m of which are reflex, admits a convex Steiner partition with $O(|P| \log m / \log\log m)$ weight and $O(n)$ Steiner points. Such a partition can be computed in $O(n \log n)$ time.*

The upper bound for the weight is the best possible, as it was shown by Levcopoulos and Lingas [30]. Interestingly, in our lower bound construction in Theorem 2, the EMST is a path which can be completed to a simple polygon by adding one edge. The resulting simple polygon P is, in fact, somewhat reminiscent of the lower bound construction presented in [30].

3.1 Lower bound construction

Proof of Theorem 2(ii). For every $n > 4$, we describe a set of at most n points in the plane for which any convex Steiner partition has weight $\Omega(W \log n / \log\log n)$. Refer to Figure 5. Let $k \in \mathbb{N}$ be the maximal integer such that $k^k < n$. Since $n > 4$, we have $k \geq 2$ and $k = \Theta(\log n / \log\log n)$. Consider a circular arc subtended by an inscribed angle $\alpha = \pi/(4k)$; that is, the central angle is $2\alpha = \pi/(2k)$. Let β be a polygonal path connecting $k+1$ points evenly distributed along this arc. We construct inductively, in k steps, a polygonal path γ_k, whose vertex set will be the point set S. The initial polygonal path is a straight line segment $\gamma_0 = ab$. In step $i = 1, 2, \ldots, k$, construct a polygonal path γ_i by replacing each segment s of γ_{i-1} by a scaled copy of the polygonal path β above segment s, as illustrated in Figure 5. The polygonal path γ_i consists of k^i segments of equal length. Let S_i denote the set of vertices of γ_i, for $i = 0, 1, \ldots, k$. Since γ_i is a refinement of γ_{i-1}, we have $S_{i-1} \subset S_i$. By induction, we have $\#S_i = k^i + 1$. Our point set for the lower bound construction is $S = S_k$. Observe that by the definition of k, $\#S = \#S_k \leq n$.

Every edge of γ_i makes an angle of at most $i\alpha$ with the x-axis. So every edge of the final curve γ_k makes an angle at most $k\alpha \leq \pi/4$ with the x-axis. Hence, an EMST of S_i is the path γ_i for each i (e.g., by Prim's algorithm). We give an upper bound on the weight of γ_k. In each step, the increase in length of the path is bounded by a constant factor:

$$\frac{|\gamma_i|}{|\gamma_{i-1}|} < \frac{2\alpha}{2\sin\alpha} \leq \frac{\alpha}{\alpha - \alpha^3/6} < \frac{1}{1-\alpha/6} = \frac{24k}{24k-\pi} = 1 + \frac{\pi}{24k-\pi} < 1 + \frac{1}{k},$$

for $k \geq 2$. The weight of γ_k is $|\gamma_k| \leq |\gamma_0| \cdot (1+1/k)^k < |ab| \cdot e = O(|ab|)$, where e stands for the base of the natural logarithm.

Next, we show that the weight of any convex Steiner partition of S is $\Omega(|\gamma_k| \cdot k) = \Omega(W \log n / \log\log n)$. By Lemma 6, it is enough to construct a set of pairwise disjoint half-disks of total radii $\Omega(|ab| \cdot k)$ such that each half-disk is centered at a point of S and is contained in the convex hull of S. Specifically, we construct disjoint half-disks of radius $|ab|/(32k^i)$ centered at at least half of the points in $S_i \setminus S_{i-1}$, for $i = 1, 2, \ldots, k-1$ (but not for the last level, $i = k$). So the total radii of these disks will be

$$\sum_{i=1}^{k-1} \frac{\#(S_i \setminus S_{i-1})}{2} \cdot \frac{|ab|}{32k^i} = \frac{|ab|}{64} \sum_{i=1}^{k-1} \frac{k^i - k^{i-1}}{k^i} = \Omega(|ab| \cdot k) = \Omega\left(\frac{W \log n}{\log\log n}\right).$$

The length of each segment of the path γ_i is at least $|ab|/k^i$ because in each step, we replace a segment by a sequence of k segments of larger total length. If the diameter of the polygonal path β is ℓ, then its width is

$$\frac{1-\cos\alpha}{2\sin\alpha} \cdot \ell = \frac{2\sin^2\frac{\alpha}{2}}{4\sin\frac{\alpha}{2}\cos\frac{\alpha}{2}} \cdot \ell = \tan\frac{\alpha}{2} \cdot \frac{\ell}{2} \geq \frac{\alpha}{4} \cdot \ell = \frac{\pi}{16k} \cdot \ell \geq \frac{\ell}{8k}.$$

Since the vertices of a scaled copy of the path $\beta \subset \gamma_i$ lie along a circular arc, at least $k/2$ of its $k-1$ internal vertices lie at a distance at least $\ell/(16k)$ from the diameter of β (that is, at least half-width distance from the chord of the circular arc). For each copy of the path β along the paths γ_i for $i = 1, 2, \ldots, k-1$ (except for the last level, $i = k$), at each such vertex $v \in \beta$, $v \in S_i \setminus S_{i-1}$, place a half-disk of radius $|ab|/(32k^i) \leq \ell/(32k)$ centered at v and bounded by the tangent line of the circular arc at v. For a fixed i, let D_i denote this set of congruent half-disks.

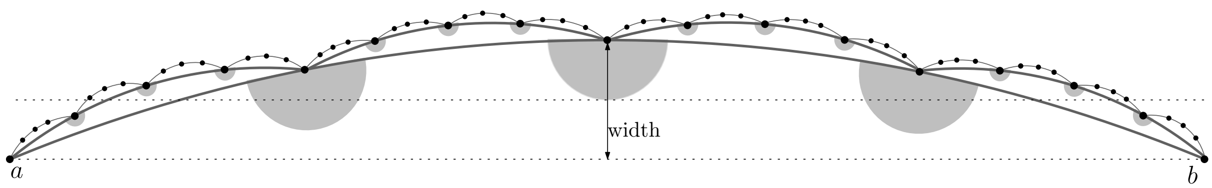

Figure 5: Lower bound construction for $k = 2 : S = S_3$. Each segment is replaced with a sequence of 4 segments. We place half-disks of radius $|ab|/(32k^i)$ centered at at least half of the points in $S_i \setminus S_{i-1}$, for $i = 1, 2, \ldots, k-1$.

It remains to show that any half-disk in $\cup_{i=1}^{k-1} D_i$ lies in the convex hull of S, and that the half-disks in $\cup_{i=1}^{k-1} D_i$ are pairwise disjoint. Note that the half-disks in D_i do not lie entirely below γ_i, and so γ_i does not separate half-disks in D_i and D_{i+1}. We show that the half-disks in D_i (i) are pairwise disjoint; (ii) lie above curve γ_{i-1}, at distance at least $|ab|/(32k^i)$ from γ_{i-1}; (iii) lie below γ_{i+1} and their portions above γ_i remains within distance at most $|ab|/(64k^{i+2})$ from γ_i. (i) It is easy to see that the half-disks in D_i, centered at vertices of γ_i, are pairwise disjoint: The half-disks in D_i have radius $|ab|/(32k^i)$. Consecutive vertices of γ_i in the same copy of β are at least $|ab|/k^i$ distance apart, and so the x-coordinates of any two vertices of γ_i (even in different copies of β) are at least $\cos(\pi/4)|ab|/k^i$ distance apart. (ii) By construction, the centers of the half-disks in D_i are above γ_{i-1}, at distance at least $|ab|/(16k^i)$ from γ_{i-1}. Hence every half disk in D_i is above γ_{i-1}, at distance at least $|ab|/(32k^i)$ from γ_{i-1}. (iii) Consider a circle passing through the vertices of a copy of β along γ_i. Each edge of β subtends an inscribed angle of $\alpha/k = \pi/(4k^2)$; and the diameter of a half-disk centered at a vertex of β is tangent to this circle. Hence, at the center of each half-disk in D_i, the incident edges of γ_i make an angle of $\pi/(4k^2)$ with the diameter of the half-disk. The portions of the half-disk above γ_i remain within distance at most $\sin(\pi/(4k^2)) \cdot |ab|/(32k^i) \leq |ab|/(64k^{i+2})$ from γ_i. Therefore the half-disks in D_i lie below γ_{i+1}, yet they are disjoint from any half-disk in D_{i+1}. This completes the proof of Theorem 2(ii). □

Proof of Theorem 3. We modify the previous construction as follows. Let S be a set of at most $2n$ points: a set S_1 of n points distributed evenly on a circle of radius 2 centered at the origin o, and a set S_2 of at most n points making our construction in Figure 5, starting from a horizontal segment $\gamma_0 = ab$ centered at o, where $|ab| = 2$. In particular, S_2 is contained in a circle of unit radius centered at o.

An EMST of S consists of a polygonal path of $n-1$ edges connecting consecutive points along the circle of radius 2, a polygonal path of $n-1$ edges connecting consecutive points along γ_k, and a unit segment connecting the two. Denoting by W_n the length of an EMST for S, we have $\lim_{n\to\infty} W_n \leq 2e+4\pi+1 = O(1)$. Consider a spanner network G for S with stretch factor $o(n)$ and all *bounded* faces convex. By the condition on the stretch factor, the shortest path between any two consecutive vertices of $\mathrm{conv}(S)$ has length $o(n)\cdot 4\pi/n = o(1)$. Hence the unbounded face is disjoint from the unit disk if n is sufficiently large. That is, all faces of the network that cover the unit disk are convex, and our previous argument shows that the network has weight $\Omega(\log n/\log\log n) = \Omega(W\log n/\log\log n)$, as claimed. $\square$

Remark. Consider again the point set S for the lower bound construction in Figure 5. Let P be the simple polygon obtained by connecting the two endpoints a and b of γ_k. Note that P is x-monotone, and $|P| \leq 2W$. We have given a lower bound on the total weight of the portions of a convex Steiner partition of S that lie in some pairwise disjoint half-disks. Observe that these half-disks are disjoint from the edges of P and lie in the interior of P. It follows that any Steiner partition of P has weight $\Omega(|P|\log n/\log\log n)$. In our construction, the convex vertices of P are a, b, and the vertices $S_k \setminus S_{k-1}$ introduced in the last level. The number of reflex vertices is $m = \Theta(n/k) = \Theta(n\log\log n/\log n)$, and so $\log m/\log\log m = \Theta(\log n/\log\log n)$. This provides an alternative construction for the lower bound of $\Omega(|P|\log m/\log\log m)$ on the minimum weight Steiner partition of a simple polygon P with m reflex vertices, first established by Levcopoulos and Lingas [30].

3.2 Upper bound—reduction to flat poygons

Let S be a set of n points in the plane. We show that S admits a convex Steiner partition of weight $O(W\log n/\log\log n)$ by reducing the partition problem for points to the corresponding partition problem for polygons. We proceed as follows. Compute the convex hull and an EMST of S. The EMST partitions the interior of the convex hull into weakly simple polygons [33, Section 10.2]. The total perimeter of these polygons is at most $4W$. Since the maximum degree of an EMST is at most 6, the total number of vertices of these polygons is at most $6n$. We construct a convex Steiner partition for each of these polygons, and then Theorem 4 completes the proof of Theorem 2(i).

We now present the proof of Theorem 4. Given a simple polygon P with n vertices, m of which are reflex, we construct a convex Steiner partition of P in three stages. The first stage (Lemma 8) partitions the interior of P into convex polygons and *monotone polygons* (defined below) by introducing new edges of total weight $O(|P|)$. The second stage (Lemma 10) partitions the interior of each x-monotone polygon M into convex polygons and monotone *1-flat polygons* (defined below) introducing new edges of total weight $O(|M|)$. The third stage (Lemma 11) partitions every monotone 1-flat polygon F into convex faces by adding new edges of total weight $O(|F|\log m/\log\log m)$. We proceed with the details.

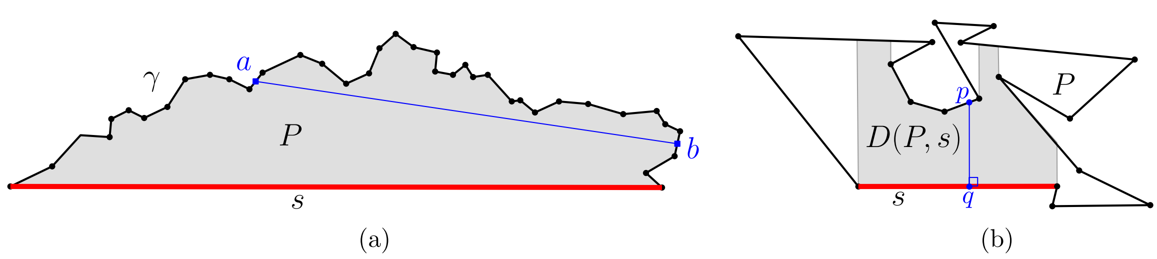

Figure 6: (a) A 1-flat polygon P. (b) A polygon P and a domain $D(P,s)$ spanned by a side s of P.

Definitions. A *diagonal* of a polygon P is a line segment connecting two vertices of P through the interior of P. A *chord* of a polygon P is a line segment that connects two points on the boundary of P (at vertices or on edges of P) and whose relative interior is disjoint from the exterior of P. For a parameter $\varepsilon > 0$, a polygon P is *ε-flat* if it is bounded by a *base* side s and a polygonal curve γ connecting the endpoints of s, and lying in one of the closed halfplanes determined by s, such that for any chord ab of P with $a, b \in \gamma$, the portion of γ between a and b has weight at most $(1+\varepsilon)|ab|$; see Figure 6 (i). In particular, we have $|\gamma| \leq (1+\varepsilon)|s|$.

Let d be a *direction*, represented by a directed or undirected line or line segment (e.g., a coordinate axis). A polygonal path γ is *d-monotone* if the intersection of γ with every line orthogonal to d is connected (that is, the intersection is a point, a line segment, or the empty set). A polygon P is *d-monotone* if it is bounded by a base side s parallel to d and a d-monotone path γ connecting the endpoints of s. We sometimes omit the direction d, and call a polygon *monotone* if it is d-monotone in some direction d. (Note that our definition of monotone polygon is slightly different from the standard one [1]).

Denote by $\overline{P}$ the closed polygonal domain bounded by P; that is, $\overline{P}$ is the closure of $\mathrm{int(P)}$. For a polygon P and a side s, let $D(P,s)$ be the set of all points $p \in \overline{P}$ such that the line segment pq lies in $\overline{P}$, $q \in s$ and pq is orthogonal to s; refer to Fig. 6(ii). Let the polygon $M(P,s)$ be the boundary of the domain $D(P,s)$. Note that $M(P,s)$ is s-monotone.

Lemma 8 *Every polygon P with n vertices admits a partition into convex faces and monotone polygons, such that the partition has $O(|P|)$ weight, $O(n)$ Steiner points, and can be computed in $O(n \log n)$ time. Furthermore, every reflex vertex of a monotone face is a reflex vertex of P.*

Proof. We partition P recursively. The input of each recursive step is a pair (Q,s), where Q is a simple polygon and s is a side of Q. For a polygon Q, let $v(Q)$ denote the number of vertices of Q. Initially, we start with the pair (P, s_0), where s_0 is an arbitrary side of P, and $v(P) = n$.

One recursion step works as follows (refer to Figure 7): We are given a pair (Q,s). Stop if Q is s-monotone or convex. Otherwise compute the s-monotone polygon $M(Q,s)$ spanned by s in Q. The edges of $M(Q,s)$ partition Q into polygonal faces. Denote the resulting polygons outside of $M(Q,s)$ by $Q_1, Q_2, \ldots, Q_t$, for some $t \geq 1$. The polygon $M(Q,s)$ has exactly one side adjacent to each Q_i, which we denote by s_i. We have $v(Q_i) \leq v(Q)$ for $i = 1, 2, \ldots, t$.

Case 1 If [$t = 1$ and $v(Q_1) < v(Q)$] or [$t \geq 2$], then partition Q along all segments s_i and recurse on (Q_i, s_i) for $i = 1, \ldots, t$.

Case 2 If $t = 1$ and $v(Q_1) = v(Q)$, then find a chord f of Q_1 parallel to s_1 such that $f \not\subseteq s_1$, the endpoints of f lie on the edges of Q_1 adjacent to s_1, and f contains a vertex of Q_1; see Figure 7(c). Partition Q along f into a convex quadrilateral Q_0 adjacent to s; and polygons $Q'_1, Q'_2, \ldots, Q'_{t'}$, for some $t' \geq 1$. Each Q'_i has exactly one side s'_i along f. Recurse on (Q'_i, s'_i) for $i = 1, \ldots, t'$.

First, we estimate the weight of the partition. Consider one step of the recursion. In Case 1, Q is partitioned along the sides of $M(Q,s)$ perpendicular to s, and the monotone polygon $M(Q,s)$ is discarded from further consideration. The partitioning edges s_i, $i = 1, 2, \ldots, t$, become the base sides in the subproblems (Q_i, s_i). Charge the weight of each s_i to the common boundary of $M(Q_i, s_i)$ and the input polygon P. The weight of this portion of the boundary of $M(Q_i, s_i)$ is at least s_i, and will not be charged again—since $M(Q_i, s_i)$ is discarded in the next step of the recursion. In Case 2, Q is partitioned along a chord f perpendicular to s, and a convex quadrilateral Q_0, which is strictly larger than $M(Q,s)$, is discarded from further consideration. The edges s'_i, $i = 1, 2, \ldots, t'$, along f become the base sides of the subproblems. Charge the weight of each s'_i to the common boundary of $M(Q'_i, s'_i)$ and the input polygon P as above. Over all, each portion of the boundary of P is charged at most once. Hence the total weight of the new edges is at most $|P|$, and the weight of the entire network (including P) is at most $2|P|$.

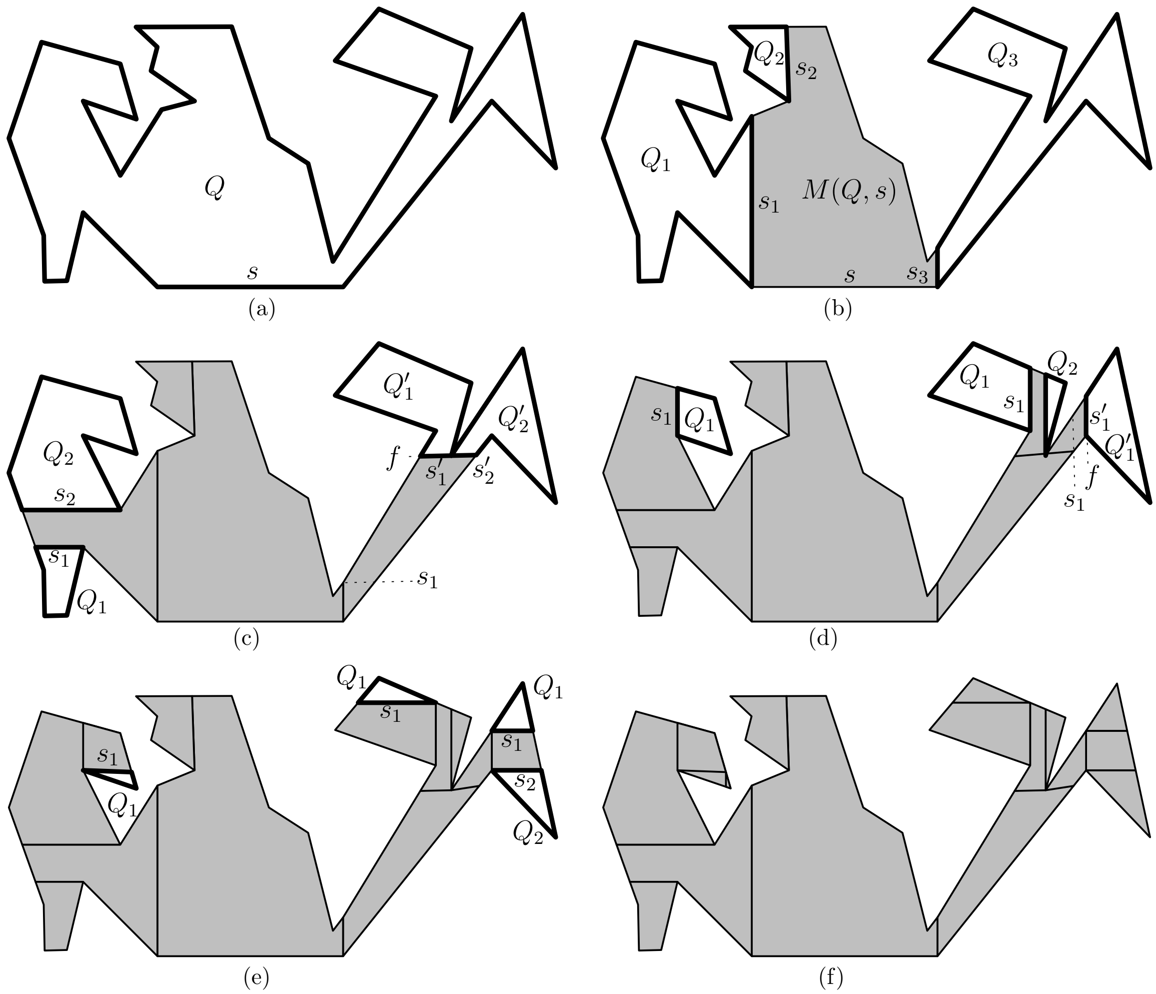

Figure 7: Partitioning a polygon P into convex and monotone polygons.

Next, we estimate the number of Steiner points. For a polygon Q, let $\mathrm{count}(Q) = 2v(Q) - 6$. Consider a step of the recursion that produces $t \geq 1$ subproblems (Q_i, s_i), $i = 1, 2, \ldots, t$. In Case 1, at most t new Steiner points are created (at most one endpoint of each s_i). In Case 2, at most 2 Steiner points are created (namely, the endpoints of the chord f). So in both cases, at most $2t$ Steiner points are created. We claim that

$$\sum_{i=1}^{t} \mathrm{count}(Q_i) \leq \mathrm{count}(Q) - t.$$

That is, the total count decreases by at least t. Initially, $\mathrm{count}(P) = 2n-6$. We have $\mathrm{count}(Q) \geq 2 \cdot 3 - 6 = 0$ for any subproblem (Q, s) throughout the recursion. This gives an upper bound of $2(2n-6) = 4n-12$ on the number of new Steiner points created altogether.

We now justify the above claim. For a subproblem (Q, s), let $V(Q, s)$ denote the set of vertices of Q with the exception of the endpoints of s. Clearly, $\#V(Q, s) = v(Q) - 2$. As above, consider a step of the recursion that produces subproblems (Q_i, s_i), $i = 1, 2, \ldots, t$. The sets $V(Q_i, s_i)$ are pairwise disjoint and $\bigcup_{i=1}^{t} V(Q_i, s_i) \subseteq V(Q, s)$. Therefore, $\sum_{i=1}^{t}(v(Q_i) - 2) \le v(Q) - 2$, and $\sum_{i=1}^{t} \mathrm{count}(Q_i) \le \mathrm{count}(Q) - 2(t-1)$. For $t \ge 2$, we have $t \le 2(t-1)$, so our claim is established in this case. Now assume that $t = 1$ (i.e., a single subproblem (Q_1, s_1) is produced). In Case 1, we assumed that $v(Q_1) < v(Q)$. In Case 2, a vertex of Q incident to the chord f belongs to $V(Q, s)$ but lies at the base side of the subproblem, hence $v(Q_1) < v(Q)$. This implies that for $t = 1$, we also have $\mathrm{count}(Q_1) \le \mathrm{count}(Q) - 2$ in both cases.

It is easy to implement the algorithm in $O(n \log n)$ time. Assume w.l.o.g. that s_0 is horizontal. Then every partitioning segment is axis-parallel, and is incident to a vertex of P or a Steiner point. Using a ray-shooting data structure for the polygon P [8], compute all axis-aligned rays from every vertex of P in advance; and sort the heads of rays along each edge of P. During the partition algorithm, we insert new rays from every new Steiner point. Shoot a ray from the endpoints of s_i orthogonally to s_i into the interior of Q_i (in Case 1) and from the endpoints of f orthogonally to f in the interior of Q'_i (in Case 2). The rays allow computing each polygon $M(Q, s)$ in time proportional to its number of vertices. The sorted list of ray heads along the edges allows finding the segment f incident to a vertex of P. $\Box$

Consider now a x-monotone polygon P with a horizontal base s. Let $\mathcal{C}$ denote the set of all chords of P whose endpoints are not in the interior of s; and let $\mathcal{H} \subset \mathcal{C}$ denote the set of horizontal chords. For every $ab \in \mathcal{C}$, let $\mu(ab)$ denote the portion of the boundary of P between a and b that does not contain s.

Lemma 9 *Let P be a x-monotone polygon with a horizontal base s. If there is a constant $\kappa > 0$ such that $|\mu(ab)| \le (1+\kappa)|ab|$ for all horizontal chords $ab \in \mathcal{H}$, then $|\mu(ab)| \le (1+\kappa)\sqrt{2}|ab|$ for all chords $ab \in \mathcal{C}$.*

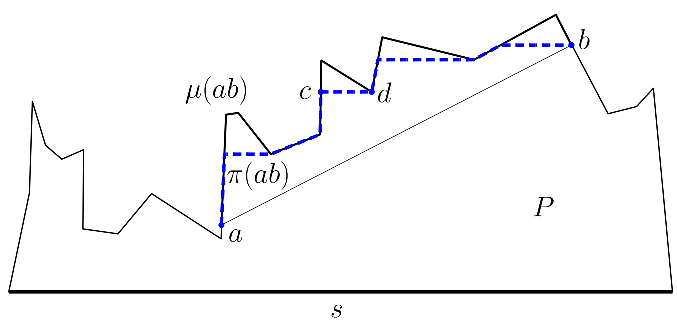

Figure 8: A chord ab in a monotone polygon P with a horizontal base s, the boundary path $\mu(ab)$ (bold), and the x-monotone path $\pi(ab)$ (dashed).

Proof. Let $ab \in \mathcal{C}$ be a chord of P; see Figure 8. Assume that ab is not horizontal, otherwise $\mu(ab) \le (1+\kappa)|ab|$. Assume without loss of generality that a has smaller y-coordinate than b. Let $\pi(ab)$ be an x-monotone path between a and b such that every segment along $\pi(ab)$ is either horizontal or lies along an edge of P. Since the weight of $\pi(ab)$ is at most that of an axis-aligned staircase path between a and b, we have $|\pi(ab)| \le \sqrt{2}|ab|$. For each horizontal portion of $\pi(ab)$, say, $cd \subset \pi(ab)$, we have $\mu(cd) \le (1+\kappa)|cd|$. Hence $|\mu(ab)| \le (1+\kappa)|\pi(ab)| \le (1+\kappa)\sqrt{2}|ab|$. $\Box$

Lemma 10 *Every x-monotone polygon P with a horizontal base s and n vertices admits a Steiner partition into convex polygons and 1-flat x-monotone faces such that the partition has $O(|P|)$ weight, $O(n)$ Steiner points, and it can be computed in $O(n \log n)$ time. Furthermore, every reflex vertex of a face is a reflex vertex of P.*

Proof. We sweep a horizontal line ℓ top-down over P, and insert horizontal chords along ℓ when certain events occur. Let Q denote the polygonal face in the current partition adjacent to the base s. Initially, let

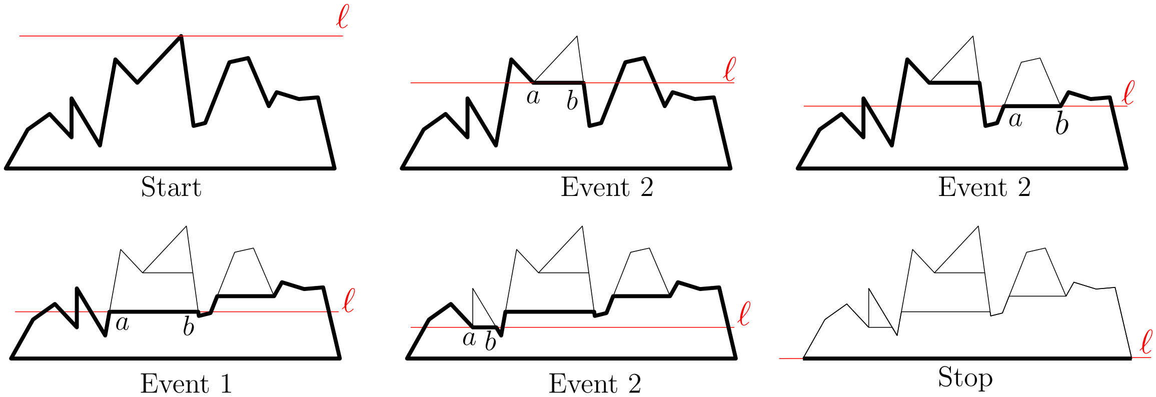

Figure 9: Partitioning an x-monotone polygon with horizontal base s into convex faces and 1-flat x-monotone faces.

$Q = P$. The face Q is always an x-monotone polygon with base s, and each insertion of a horizontal chord along ℓ cuts off a polygon from Q. The algorithm is designed so that each polygon cut off from Q is either convex or 1-flat with a horizontal base along the current position of ℓ. For every chord $ab \in \mathcal{H}$ lying along the sweep-line ℓ, we define $\mu_Q(ab)$ to be the portion of the boundary of Q between a and b that does not contain s.

Sweep a horizontal line ℓ from the top vertex down until ℓ reaches the base s. If any of the following two events occurs for a chord $ab \in \mathcal{H}$ lying along ℓ, partition Q along ab into two faces, and let Q be the face below ab.

Event 1. *$\mu_Q(uv)$ contains at least three vertices of Q in its interior and $|\mu_Q(ab)| = \sqrt{2}|ab|$.*

Event 2. *a or b is a reflex vertex of Q and $|\mu_Q(ab)| \geq \sqrt{2}|ab|$.*

In each step, we inserted a chord ab and partitioned Q along chords ab into two faces. The face above ab is either convex and its common boundary with the new face Q is at least $\sqrt{2}$-times heavier than $|ab|$; or it is 1-flat and x-monotone with base ab by Lemma 9, with $\kappa = \sqrt{2} - 1$. This face is discarded from further consideration, and the remaining face (adjacent to s) is partitioned recursively. In each step, we discard a portion of weight at least $\sqrt{2}|ab|$ from the boundary of Q, and introduce a new boundary segment of weight $|ab|$. If we charge the weight of each chord ab uniformly to the portion of the polygon P that it replaces, then each point along the boundary of P is charged at most $\sum_{i=1}^{\infty}(\sqrt{2})^{-i} = \frac{1}{\sqrt{2}-1}$ times. So the total length of the partition (including the weight of P) is at most $(1 + \frac{1}{\sqrt{2}-1})|P| = O(|P|)$.

Next, we estimate the number of Steiner points. Let $\text{count}(Q)$ denote the number of vertices of Q plus the number of vertices of Q with an acute interior angle. Initially, when $Q = P$, we have $\text{count}(Q) \leq 2n$, since there are no more than n acute angles. We claim that every partition step decreases $\text{count}(Q)$ by at least one, and introduces at most two Steiner points. This implies that the number of new (Steiner) vertices cannot exceed $4n$. In Event 2, at least three vertices are removed from Q, and at most two new (Steiner) vertices are created, neither of which may have an acute interior angle. In Event 1, one or two convex vertices are removed and at most one new (Steiner) vertex is created, which does not have acute interior angle. If only one vertex is removed from Q, then the part of Q above ab is a triangle, and $\mu_Q(ab) \geq |ab|$ is possible only if the triangle has an acute angle at the vertex opposite to ab. This confirms our claim that $\text{count}(Q)$ strictly decreases in each step, and concludes the proof of the lemma. $\square$

3.3 Upper bound—partitioning a flat polygon

In this section, we partition a monotone 1-flat polygon recursively. The intermediate polygons in the recursion steps are not necessarily monotone or flat, however, they have a very special structure: A polygon P is a *clamp polygon* with *spine* ab if it is bounded by a convex chain (path) σ and a d-monotone polygonal path γ (for some direction d), with both paths connecting the same two points a and b; see Figure 10. The *width* of the clamp polygon is the width of the minimum strip that contains the polygon and is parallel to ab. Observe that every s-monotone polygon is a clamp polygon, where the convex chain is the side s.

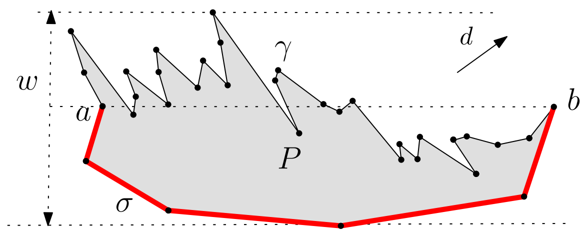

Figure 10: A clamp polygon P, bounded by a convex chain σ (in bold) and a d-monotone polygonal path γ.

Lemma 11 *Every 1-flat x-monotone polygon F with a horizontal base and n vertices, m of which are reflex, admits a convex Steiner partition with $O(|F| \log m / \log\log m)$ weight and $O(m)$ Steiner points. Such a partition can be computed in $O(n + m \log n)$ time.*

Proof. Consider a 1-flat x-monotone polygon F bounded by a horizontal side s_0 and an x-monotone polygonal path γ_0. Assume that F has n vertices, m of which are reflex. We partition F recursively. We describe a generic step of the recursion, where we are given a clamp polygon Q with r reflex vertices, bounded by a convex chain and an x-monotone polygonal path $\gamma \subseteq \gamma_0$, and we want to partition Q into convex faces and some clamp polygons Q_i, each having at most $r/2$ reflex vertices. We say that the *weight of the problem* associated with Q is $|\gamma|$. The polygonal path γ of Q will be partitioned among the subproblems Q_i; and so the total weight of the subproblems in each level of the recursion is at most $|\gamma_0|$. In a recursion step, we will introduce new edges of total weight $O(|\gamma_0|)$. Intuitively, it is enough to show that the "average depth" of the recursion is $O(\log m/ \log\log m)$, in order to establish a $O(W \log m/ \log\log m)$ bound on the weight of the partition.

One step of the recursion. We are given a clamp polygon Q with $r \geq 1$ reflex vertices, a spine s, where Q is bounded by a convex chain σ and a polygonal path $\gamma \subseteq \gamma_0$. By rotating Q, if necessary, we may assume that the spine s is horizontal, and γ is d-monotone for some direction d (which is not necessarily horizontal). It is clear that any vertex of Q with maximal y-coordinate is a vertex of γ. We also assume that one of the vertices of Q with minimal y-coordinate is a vertex of γ. This property holds for the initial clamp polygon F, and we maintain the property for every clamp polygon Q during the recursion. It follows that $|\gamma| \geq 2w$, where w is the width of Q. Let $k = \lceil |\gamma|/w \rceil \geq 2$. We partition the polygonal path γ into at least k subpaths such that

(i) each subpath passes through at most r/k reflex vertices of Q;

(ii) each subpath passing through a reflex vertex of Q has weight at most $|\gamma|/k$; and

(iii) a furthest point of γ from the supporting line of the spine of Q is the endpoint of a subpath.

The partition is done by successively selecting the elements of a set $A \subset \gamma$ of *spliting points*. Refer to Figure 11. The splitting points are selected as follows: Place one endpoint of γ in A. Move a point p along γ continuously from this endpoint to the other. If the subpath of γ between the previous splitting point of A and $p \in \gamma$ satisfies any of the following conditions, then insert p into A: (1) p is the $(\lfloor r/k \rfloor + 1)^{\text{th}}$ reflex vertex of γ along the subpath; (2) the weight of the subpath is at least $|\gamma|/k$ and it passes through at least one reflex vertex; (3) the weight of the subpath exceeds $|\gamma|/k$ and p is its first reflex vertex; (4) p is a the furthest point of γ from the supporting line of the spine of Q, which may be either above or below the spine; or (5) p is the right endpoint of γ. The cardinality of A is at most $2k + 1$, with the two endpoints of γ, at most $k - 1$ additional splitting points of type (1), at most $k - 1$ additional splitting points of types (2) and (3), and the furthest point from the spine (4).

Next we partition the interior of Q into a set $\mathcal{F}$ of at least k polygons by using the points $A \subset \gamma$ and by drawing some new segments. If γ is x-monotone then drop a vertical segment from each point $a \in A$ to the convex chain σ. If γ is not x-monotone, however, additional splitting points may be necessary. Let γ_* be the *lower envelope* of γ (that is, the set of points $p \in \gamma$ such that there is no other point of γ with the same x-coordinate and a smaller y-coordinate). If γ is x-monotone, then $\gamma = \gamma_*$, otherwise $\gamma_* \subset \gamma$ consists of several components. If pq is a vertical segment that connects distinct components of γ_*, then the portion of γ between p and q is called a *pocket of* γ *bounded by* pq. For every point $a \in \gamma$ in a pocket bounded by pq, there is a line segment orthogonal to d that connects a to pq, since γ is d-monotone. Now, connect every splitting point $a \in A$ to the convex chain σ as follows. If $a \in \gamma_*$, then drop a vertical segment from a to the convex chain σ. If a is part of a pocket bounded by some vertical segment pq, then insert p and q as new splitting points, extend the vertical segment pq to σ, and connect a to pq by a segment orthogonal to d. Let A' denote the union of A and the sets $\{p, q\}$ for each $a \in A$ lying in a pocket bounded by pq.

There are at most $2k+1$ splitting points in A. Each vertical edge between γ and the convex chain σ has weight at most w. The segments orthogonal to d have weight at most $2w$ since F is 1-flat and the weight of a pocket bounded by pq is at most $2|pq|$, where $|pq| \le w$. So the total weight of the partition edges is at most

$$(2k + 1)(1 + 2)w = (6k + 3)w \le 8kw = 8\left\lceil\frac{|\gamma|}{w}\right\rceil w < 8\frac{|\gamma| + w}{w}w \le 12|\gamma|.$$

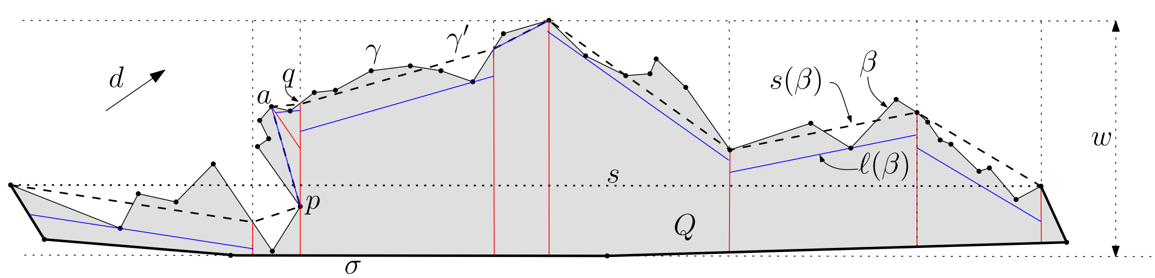

Figure 11: One level of partitions for a clamp polygon Q.

The splitting points in A' partition the path γ into a set Γ of subpaths, satisfying conditions (i)–(iii). By connecting consecutive points of A' along γ, we obtain a polygonal path γ' (dashed path in Figure 11).

So far, we have partitioned Q into a set $\mathcal{F}$ of polygons, each of which is bounded by a convex chain and a path in Γ. For each path $\beta \in \Gamma$, let $s(\beta) \subset \gamma'$ denote the segment connecting the endpoints of β. Shift segment $s(\beta)$ continuously to a position $\ell(\beta)$ such that it partitions the corresponding face in $\mathcal{F}$ into a convex face and a clamp polygon with spine $s(\beta)$; see Figure 11. Since $|\ell(\beta)| \le |s(\beta)|$, the total weight of the segments $\ell(\beta)$, for all $\beta \in \Gamma$, is $|\gamma'| \le |\gamma|$. The total weight of all new edges introduced in one recursion step is at most $12|\gamma| + |\gamma| = 13|\gamma|$. Partition all non-convex clamp polygons recursively until all faces are convex. This completes the description of our partition algorithm.

Analysis. Assume F is the initial 1-flat clamp polygon with n vertices, m of which are reflex, a base s_0 and a polygonal path γ_0. (We can assume m is large enough, when needed.) Obviously, we have $|\gamma_0| \le |F|$. At level j of the recursion, $j \ge 0$, we construct a polygonal path λ_j connecting the two endpoints of the spine s_0 as follows. At the root level, λ_0 is a straight line segment $\lambda_0 = s_0$. At level j, we construct λ_j by replacing some segments of λ_{j-1} with polygonal paths: If a segment $s \subset \lambda_{j-1}$ is the spine of a subproblem with at least one reflex vertex, then s is replaced by the polygonal path γ' (defined previously); if a segment $s \subset \lambda_{j-1}$ is the spine of a convex clamp polygon in a subproblem, then s is replaced by the (convex) chain γ of that clamp polygon. We can establish a piecewise linear homeomorphism $H_j : \lambda_{j-1} \to \lambda_j$. If a segment $s \subseteq \lambda_{j-1}$ is replaced by a polygonal path $\gamma' \subseteq \lambda_j$ of weight $|\gamma'|$, then H_j maps s to γ' while increasing each portion of s with the same factor. By transitivity, the composition $H_j \circ \mathcal{H}_{j-1} \circ \cdots \circ \mathcal{H}_1$ is a piecewise linear homeomorphism between s_0 and λ_j. We next show that the weights of the approximations λ_j monotonically increase ($|\lambda_{j-1}| \le |\lambda_j|$, for $j \ge 1$) but remain in the range $|s_0| \le |\lambda_j| \le |\gamma_0| \le 2|s_0|$.

Consider the recursion tree T. In T, the subproblem at a terminal node at level j is a *convex* clamp polygon, and so the corresponding portion of λ_j is a portion of the initial path γ_0. At the next level, $j+1$, only the clamp polygons adjacent to λ_j are further partitioned. Each non-terminal node v in T corresponds to a subproblem, where a spine s_v is replaced by a path γ'_v and v has at least $k_v \ge 2$ children (recall that one step of the recursion partitions $\gamma_v \subseteq \gamma_0$ into at least k_v subpaths). By the definition of k_v, we have $k_v \ge |\gamma_v|/w_v \ge |s_v|/w_v$. Observe that the path γ'_v reaches a vertex of γ_v at distance at least $w_v/2$ from the spine s_v, so its weight is at least

$$|\gamma'_v| \ge 2\sqrt{(|s_v|/2)^2 + (w_v/2)^2} = \sqrt{|s_v|^2 + w_v^2} \ge |s_v|\sqrt{1 + 1/k_v^2}.$$

Let nodes $v(i)$, $i = 0, 1, \ldots, t$, in T form a path from the root to a leaf, where $v(i)$ lies at level i of the recursion. If we follow a path $v(0)$, $v(1)$, …, $v(t)$ in the tree of the recursion, then the number of reflex vertices in the subproblems decreases by a factor of at least $k_{v(i)}$ for $i = 0, 1, 2 \ldots, t-2$. We have zero reflex vertices at a terminal node, and we may have fewer than $k_{v(t-1)}$ reflex vertices at the parent of a terminal node. At any other node $v(i) \in V(T)$, the number of reflex vertices decreases by a factor of at least $k_{v(i)}$. Since at the root of T corresponds to the initial problem with m reflex vertices, the product of the $k_{v(i)}$ values along the path is bounded by m:

$$\prod_{i=0}^{t-2} k_{v(i)} \le m.$$

In any chain $v(0), v(1), \ldots, v(t)$ from the root to a leaf in T, there are at most $4\log m/\log\log m$ nodes with $k_{v(i)} \ge (\log m)^{1/4}$, since $((\log m)^{1/4})^{4\log m/\log\log m} = 2^{(\log\log m)(\log m/\log\log m)} = 2^{\log m} = m$.

We next establish lower bounds on the weight of the approximation paths λ_j. In each step with $k_v \le (\log m)^{1/4}$, a segment s_v is replaced by a path γ'_v of weight at least $|s_v| \cdot \sqrt{1 + 1/(\log m)^{1/2}}$. In each step with $k > (\log m)^{1/4}$, we use the trivial lower bound that a segment s_v is replaced by a path γ'_v of weight at least $|s_v|$. Let $h \ge 0$ be an integer. We show that the total weight of those portions of the initial segment s_0 that undergo at least $h\lceil \log m/\log\log m\rceil$ steps that each expand the weight by a factor of at least $\sqrt{1 + 1/(\log m)^{1/2}}$ is at most $|s_0|/2^h$. This is immediate for $h = 0$; and if it did not hold for $h \ge 1$, then the weight of the final approximation path γ_0 would be

$$\begin{aligned} |\gamma_0| &\geq \frac{|s_0|}{2^h}\left(1+\frac{1}{\sqrt{\log m}}\right)^{\frac{h\log m}{2\log\log m}} \\ &= \frac{|s_0|}{2^h}\left(\left(1+\frac{1}{\sqrt{\log m}}\right)^{\sqrt{\log m}}\right)^{\frac{h\sqrt{\log m}}{2\log\log m}} \\ &\geq |s_0|\left(2^{\frac{\sqrt{\log m}}{2\log\log m}}\right)^h \geq 2^h|s_0| > 2|s_0|, \end{aligned}$$

for $m \geq 2^{256}$, using the fact that $(1+1/x)^x \geq 2$ for $x \geq 1$. This contradicts our assumption that F is 1-flat and so $|\gamma_0| \leq 2|s_0|$.

For any $h \geq 4$ and $j = h\lceil \log m/\log\log m\rceil$, we have $|\lambda_j \cap \gamma_0| \geq (1-2^{4-h})|\gamma_0|$, that is, the common portion of λ_j and γ_0 has weight at least $(1-2^{4-h})|\gamma_0|$. This implies that for any $h \geq 4$, at level $j = h\lceil \log m/\log\log m\rceil$, we are left with subproblems of total weight at most $|\gamma_0|/2^h$. At each level of the recursion, we introduce new edges whose total weight is proportional to the weight of the current subproblems in that level. Therefore, for $h \geq 4$, the total weight of the new edges introduced between levels $h\lceil \log m/\log\log m\rceil$ and $(h+1)\lceil \log m/\log\log m\rceil$ is at most $(1/2)^{(h-4)}O(|\gamma_0|)\cdot\lceil \log m/\log\log m\rceil$. This also holds for $h = 0,1,2,3$. The weight of the final convex Steiner partition of F is bounded by

$$\sum_{h=0}^{\infty}\left(\frac{1}{2}\right)^{h-4} O(|\gamma_0|)\cdot\left\lceil\frac{\log m}{\log\log m}\right\rceil = O\left(|\gamma_0|\cdot\frac{\log m}{\log\log m}\right) = O\left(|F|\cdot\frac{\log m}{\log\log m}\right).$$

The number of Steiner points in level j of the recursion is proportional to the total number splitting points created along γ_{j-1}. Each splitting point in a set A lies at a reflex vertex or between two reflex vertices of γ_0; and there are at most one splitting point of A between two consecutive reflex vertices of γ_0. Hence, there are at most $2m-1$ splitting points along paths λ_j for all $j \geq 0$, and the number of Steiner points is $O(m)$.

The runtime of constructing the partition is $O(n + m\log n)$. In an $O(n)$-time preprocessing step, traverse γ_0 and for each convex chain $\beta \subset \gamma_0$ between consecutive reflex vertices, compute the weight and build a binary search structure, which can report for a query value $q > 0$ a point $p \in \beta$ in $O(\log n)$ time such that the weight of the portion of β between its left endpoint and p equals q. In order to find the splitting points in a clamp polygon Q, we traverse the corresponding path $\gamma \subseteq \gamma_0$ once. We have seen that the number of reflex vertices traversed more than $h\lceil \log m/\log\log m\rceil$ times is $O(m/2^h)$, and so we traverse reflex vertices $O(m\log m/\log\log m)$ times in total. We can skip convex vertices, since we have computed the weight of γ_0 between consecutive reflex vertices. There is at most one splitting point in A between any two reflex vertices, each of which can be located in $O(\log n)$ time based on a binary search structure. We spend $O(m\log n)$ total time on finding points of A between reflex vertices. For connecting the splitting points to the convex chain σ of the corresponding clamp polygon Q, we use a ray shooting data structure [8] for F, which can be constructed in $O(n)$ time and permits $O(\log n)$ query time. Since each segment $s(\beta)$ either connects the endpoints of β or is tangent to a reflex vertex of β, we can compute $s(\beta)$ in a single traversal of the reflex vertices of β. The total time for partitioning F is $O(m\log n)$. This completes the proof of Lemma 11. $\square$

Lemma 11 was the last in the reduction chain, so the proof of Theorem 4 is now complete.

4 Conclusion

We deduced tight bounds on the weight of a minimum convex partition of a point set in terms of their Euclidean minimum spanning trees. The worst-case ratio of the minimum Steiner partition and an EMST is $\Theta(\log n/\log\log n)$ in general, and $\Theta(\log\log n)$ for the special case of pseudo-triangles. We conclude with some remaining unanswered questions.

1. Our partition for n points has $O(n)$ Steiner points. Does every set of n points admit a convex Steiner partition of weight $O(W\log n/\log\log n)$ and with only $o(n)$ Steiner points?
2. Does every set of n points admit a convex Steiner partition of weight $O(W\log n/\log\log n)$ and with $O(n)$ Steiner points, such that every face is *fat* (that is, the ratio of the radii of the minimum enclosing and maximum inscribed circles over all faces is bounded by a constant)? Networks with fat convex faces are of interest because they have constant geometric dilation; this follows from a result of [13].
3. What is the minimum size and weight, in terms of n and W, of a Steiner network that supports compass routing?

Acknowledgment. We are indebted to Nadia Benbernou, Erik Demaine, Martin Demaine, Mashhood Ishaque, and Diane Souvaine for valuable conversations on these matters.

References

[1] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf, *Computational Geometry: Algorithms and Applications* 3rd edition, Springer, 2008.

[2] M. W. Bern, D. Eppstein, and J. R. Gilbert, Provably good mesh generation, *J. Comput. Syst. Sci.* **48(3)** (1994), 384–409.

[3] P. Bose, A. Brodnik, S. Carlsson, E. D. Demaine, R. Fleischer, A. Lopez-Ortiz, P. Morin, and I. Munro, Online routing in convex subdivisions, *Intern. J. Comput. Geom. & Appl.* **12(4)** (2002), 283–295.

[4] P. Bose, J. Gudmundsson, and M. Smid, Constructing plane spanners of bounded degree and low weight, *Algorithmica* **42** (2005), 249–264.

[5] P. Bose and P. Morin, Competitive online routing in geometric graphs, *Theor. Comput. Sci.* **324(2-3)** (2004), 273–288.

[6] P. Bose and P. Morin, Online routing in triangulations, *SIAM J. Comput.* **33(4)** (2004), 937–951.

[7] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia, Routing with guaranteed delivery in ad hoc wireless networks, *Wireless Networks* **7** (2001), 609–616.

[8] B. Chazelle, H. Edelsbrunner, M. Grigni, L. J. Guibas, J. Hershberger, M. Sharir, and J. Snoeyink, Ray shooting in polygons using geodesic triangulations, *Algorithmica* **12** (1994), 54–68.

[9] B. Chazelle and D. P. Dobkin, Decomposing a polygon into its convex parts, in *Proc. 11th Sympos. Theory of Computing*, ACM Press, 1979, pp. 38–48.

[10] K. L. Clarkson, Approximation algorithms for planar traveling salesman tours and minimum-length triangulations, in *Proc. 2nd Sympos. on Discrete Algorithms*, ACM Press, 1991, pp. 17–23.

[11] D.-Z. Du and F. K. Hwang, A proof of Gilbert-Pollak's conjecture on the Steiner ratio, *Algorithmica*, **7** (1992), 121–135.

[12] A. Dumitrescu and C. D. Tóth, Light orthogonal networks with constant geometric dilation, *J. Discrete Algorithms* **7** (2009), 112–129.

[13] A. Ebbers-Baumann, A. Grüne, and R. Klein, On the geometric dilation of finite point sets, *Algorithmica* **44** (2006), 137–149.

[14] D. Eppstein, Approximating the minimum weight Steiner triangulation, *Discrete Comput. Geom.* **11** (1994), 163–191.

[15] D. Eppstein, Spanning trees and spanners, in *Handbook of Computational Geometry (J.-R. Sack and J. Urrutia, eds.)*, Elsevier, 2000, pp. 425–461.

[16] J. Gao, L. J. Guibas, J. Hershberger, L. Zhang, and A. Zhu, Geometric spanners for routing in mobile networks, *IEEE J. Selected Areas in Comm.* **23(1)** (2005), 174–185.

[17] J. García-López and C. M. Nicolás, Planar point sets with large minimum convex partitions, *Proc. European Workshop Comput. Geom.*, Delphi, 2006, pp. 51–54.

[18] P. D. Gilbert, New results in planar triangulations, Report R-850, Coordinated Science Laboratory, University of Illinois, Urbana, Illinois, 1979.

[19] E. N. Gilbert and H. O. Pollak, Steiner minimal trees, *SIAM J. Appl. Math.* **16(1)** (1968), 1–29.

[20] J. Gudmundsson and C. Levcopoulos, Minimum weight pseudo-triangulations, *Comput. Geom. Theory Appl.* **38(3)** (2007), 139-153.

[21] J. M Keil and C. A. Gutwin, Classes of graphs that approximatate the complete Euclidean graph, *Discrete Comput. Geom.* **7** (1992), 13–28.

[22] J. M. Keil and J. Snoeyink, On the time bound for convex decomposition of simple polygons, *Internat. J. Comput. Geom. Appl.*, **12** (2002), 181–192.

[23] Y-J. Kim, R. Govindan, B. Karp, and S. Shenker, Geographic routing made practical, in *Proc. 2nd Sympos. Networked Systems Design and Implementation* USENIX Assoc., 2005, pp. 217–230,

[24] D. G. Kirkpatrick, A note on Delaunay and optimal triangulations, *Inform. Proc. Lett.* **10(3)** (1980), 127–128.

[25] G. T. Klincsek, Minimal triangulations of polygonal domains, in *Submodular Functions and Optimization*, vol. 9 of Annals of Discrete Math., Elsevier, 1980, pp. 121–123.

[26] C. Knauer and A. Spillner, Approximation algorithms for the minimum convex partition problem, in *Proc. 10th Scandinavian Workshop in Algorithm Theory*, vol. 4059 of LNCS, Springer, 2006, pp. 232–241.

[27] E. Kranakis, H. Singh, and J. Urrutia. Compass routing on geometric networks, in *Proc. 11th Canadian Conf. Comput. Geom.*, Vancouver, BC, 1999, pp. 51–54.

[28] T. Leighton and A. Moitra, Some results on greedy embeddings in metric spaces, in *Proc. 49th Sympos. on Foundations of Comp. Sci.*, IEEE, 2008, pp. 337–346.

[29] C. Levcopoulos and D. Krznaric, Quasi-greedy triangulations approximating the minimum weight triangulation, *J. Algorithms* **27(2)** (1998), 303–338.

[30] C. Levcopoulos and A. Lingas, Bounds on the length of convex partitions of polygons, in *Foundations of Software Technology and Theoretical Computer Science*, vol. 181 of LNCS, 1984, Springer, Berlin, pp. 279–295.

[31] C. Levcopoulos and A. Lingas, There are planar graphs almost as good as the complete graphs and almost as cheap as minimum spanning trees, *Algorithmica* **8** (1992), 251-256.

[32] A. Lingas, The power of non-rectilinear holes, *Proc. 9th Internat. Colloq. Automata, Languages and Programming*, vol. 140 of LNCS, Springer, 1982, pp. 369–383.

[33] K. Mehlhorn and S. Näher, *LEDA: A Platform for Combinatorial and Geometric Computing*, Cambridge University Press, 1999.

[34] W. Mulzer and G. Rote, Minimum weight triangulation is NP-hard, *J. ACM* **55(2)** (2008), article 11.

[35] G. Narasimhan and M. Smid, *Geometric Spanner Networks*, Cambridge University Press, 2007.

[36] V. Neumann-Lara, E. Rivera-Campo, and J. Urrutia, A note on convex decompositions of a set of points in the plane, *Graphs and Combinatorics* **20(2)** (2004), 223–231.

[37] C. H. Papadimitriou and D. Ratajczak, On a conjecture related to geometric routing, *Theor. Comput. Sci.* **344** (2005), 3–14.

[38] D. A. Plaisted and J. Hong, A heuristic triangulation algorithm, *J. Algorithms* **8(3)** 1987, 405–437.

[39] A. Rao, C. H. Papadimitriou, S. Shenker, and I. Stoica, Geographic routing without location information, in *Proc. 9th Conf. on Mobile Computing and Networking*, 2003, ACM Press, pp. 96–108.

[40] J. Remy and A. Steger, A quasi-polynomial time approximation scheme for minimum weight triangulation, *Proc. 38th Sympos. Theory of Computing*, 2006, ACM Press, pp. 316–325. To appear in the *Journal of the ACM*.

[41] P. O. de Wet, Geometric Steiner minimal trees, PhD thesis, University of South Aftrica, 2008. `http://etd.unisa.ac.za/ETD-db/theses/available/etd-08052008-130058/unrestricted/thesis.pdf`.