

# Light Orthogonal Networks with Constant Geometric Dilation

Adrian Dumitrescu\*

Csaba D. Tóth†

December 30, 2006

## Abstract

An orthogonal network for a given set of  $n$  points in the plane is an axis-aligned planar straight line graph that connects all input points. We show that for any set of  $n$  points in the plane, there is an orthogonal network that (i) is *short* having a total edge length of  $O(|T|)$ , where  $|T|$  denotes the length of a minimum Euclidean spanning tree for the point set; (ii) is *small* having  $O(n)$  vertices and edges; and (iii) has *constant geometric dilation*, which means that for any two points  $u$  and  $v$  in the network, the shortest path in the network between  $u$  and  $v$  is at most constant times longer than the (Euclidean) distance between  $u$  and  $v$ .

## 1 Introduction

A typical problem in the theory of metric embeddings asks for a mapping from one metric space to another that distorts the distances between point pairs as little as possible. In this paper, we address the following problem about geometric dilation: Given a finite set  $S$  of points in the plane, find a small plane graph  $G(S)$  containing  $S$  so that the distortion between the  $L_2$  distance and the Euclidean shortest path distance between any two points (on edges or at vertices) of  $G(S)$  is bounded by a constant.

A special case of this problem received frantic attention in the late 80s and early 90s in the context of geometric spanners [6, 8, 16, 18] (see [14] for a survey). One of the latest results, due to Bose *et al.* [5], goes as follows: For any set  $S$  of  $n$  points in the plane, there is a plane graph  $H$  with four properties: (i) the vertex set of  $H$  is  $S$ , (ii)  $H$  has  $O(1)$  maximum degree, (iii) the total length of the edges of  $H$  is  $O(|T_S|)$ , where  $|T_S|$  is the length of the minimum Euclidean spanning tree for  $S$ , and (iv) for any two vertices  $u, v \in S$  the (Euclidean) shortest path along  $H$  is at most  $O(1)$  times longer than the distance between  $u$  and  $v$ . The last property is also referred to as constant *vertex-dilation*. Note that the graph  $H$  is sparse and the bound  $O(|T_S|)$  is the best possible, since  $H$  has to be connected at least. Intuitively, this graph  $H$  corresponds to a road network that has constant *detour* (precise definition is below) between any two of  $n$  given cities. However, there may be pairs of points along the roads (halfway between cities) with arbitrarily large detour. In the current paper, we further extend the results in [5] and construct a graph  $G$  of constant geometric dilation, that is, constant detour between *any* two points of the graph (not just between vertices).

Let us first define the (geometric) dilation formally (see also [11, 12]). Let  $G$  be an embedded planar graph whose edges are curves. Let  $G \subseteq \mathbb{R}^2$  also denote the set of points covered by the edges and vertices of the embedded graph  $G$ . The *detour* between two points  $u, v \in G$  (on edges or vertices of  $G$ ) is the ratio between the length  $d_G(p, q)$  of a Euclidean shortest path connecting  $u$  and  $v$  in  $G$  and their Euclidean

---

\*Department of Computer Science, University of Wisconsin-Milwaukee, WI 53201-0784, USA. E-mail: ad@cs.uwm.edu  
Research supported by NSF CAREER grant CCF-0444188.

†Department of Mathematics, MIT, Cambridge, MA 02139, USA. E-mail: toth@math.mit.edu

distance  $|uv|$ . The supremum value of detours over all pairs of points, denoted  $\delta(G)$ , is called the *geometric dilation* of  $G$ :

$$\delta(G) := \sup_{u,v \in G} \frac{d_G(u,v)}{|uv|}.$$

In contrast, the *vertex-dilation* is  $\max_{u,v \in V(G)} d_G(p,q)/|pq|$ , where  $V(G)$  is the vertex set of  $G$ . For instance, the dilation of a rectangle of aspect ratio  $t \geq 1$  is  $t + 1$ , while its vertex-dilation is only  $\frac{t+1}{\sqrt{t^2+1}}$ .

For a set  $S$  of  $n$  points in the plane, we construct an *orthogonal network*  $G = G(S)$ , which is a planar straight line graph with  $S \subseteq V(G)$  and with axis-parallel edges.  $G(S)$  has constant geometric dilation and retains all the good properties listed above for  $H$ . We use only  $O(n)$  Steiner points, thus  $|V(G)| = O(n)$ . The length of our network, that is, the total length of the edges of  $G$ , is  $|G| = O(|T_S|)$ .

**Theorem 1** *For every set  $S$  of  $n$  points in the plane, there is an orthogonal network  $G$  such that (i) its geometric dilation is at most  $c_1$ ; (ii) it has at most  $c_2 n$  vertices; (iii) its length is at most  $c_3 |T_S|$ . Here  $c_1$ ,  $c_2$ , and  $c_3$  are absolute constants.*

These constants are probably too large for designing a real-life orthogonal road network with small dilation for a given set of sites. Our first priority was proving that such constants exist, rather than optimizing them.

## Related previous results

**Geometric spanners and vertex-dilation.** Planar straight line graphs with constant vertex-dilation were thoroughly studied in the context of geometric spanners, motivated by VLSI design problems [14, 19]. Chew [7] proved that the vertex-dilation of the rectilinear Delaunay triangulation of  $n$  points in the plane is at most  $\sqrt{10}$ ; Dobkin *et al.* [10] gave a constant bound on vertex-dilation of the Euclidean Delaunay triangulation. Das and Joseph [8] found a large class of geometric graphs with this property, characterized by a certain *diamond property* similar to our concept of *lofty PLSGs* (see Def. 1). A lot of work has been done on finding "good" spanners: sparse and light graphs with constant vertex-dilation. Quite a few papers [1, 3, 6, 18] present algorithms that compute, for a set  $S$  of  $n$  points in the plane, a graph  $G$  with vertex set  $S$  that has constant vertex-dilation,  $O(n)$  edges, and  $O(|T_S|)$  length. Das *et al.* [9] generalized the result to  $d$ -space. Some of these algorithms run in  $O(n \log n)$  time, some compute graphs that are planar or have bounded maximal degree. Recently, Bose *et al.* [5] were able to combine all these properties. However, none of these papers provide any upper bound on the resulting *geometric dilation*. Aronov *et al.* [2] gave a tight worst-case bound on the vertex-dilation in terms of the number of edges of the graph used to connect  $n$  points.

**Geometric dilation of planar point sets.** The problem of embedding a given planar point set in a network of small geometric dilation, as well as the problem of computing or estimating the dilation of planar networks has only recently received attention. First attempts were made in designing efficient algorithms to compute the dilation of a polygonal curve [13, 17]. Ebberts-Baumann *et al.* [12] proved that every finite point set can be embedded in a plane graph (with curved edges) of geometric dilation at most 1.678, and Dumitrescu *et al.* [11] showed that some point sets require geometric dilation strictly more than  $\pi/2 \approx 1.5707$ : at least  $(1 + 10^{-11})\pi/2$ , to be precise.

**Related problems.** A somewhat related problem is the *Manhattan network* problem. A plane graph whose vertex-dilation is 1 under the  $L_1$  metric is called a Manhattan network [4, 15]. For our purpose such networks might be too expensive: Take, for instance,  $n$  equidistant points on the boundary of an axis-aligned unit square; the minimum Manhattan network is  $\Omega(n)$  times longer than the MST, and it contains  $\Omega(n^2)$  Steiner vertices; while the unit square itself has  $O(|T_S|)$  length,  $n$  vertices, and geometric dilation 2.

## 2 Reduction to axis-aligned polygons

**Notation on planar straight line graphs and polygons.** A *planar straight line graph* (PSLG) is a finite graph together with a planar embedding, where the vertices are distinct points and the edges are straight line segments, any pair of which being either disjoint or having a common endpoint. The complement  $\mathbb{R}^2 \setminus G$  of a PSLG  $G$  may have several components, which are the *faces* of  $G$ . Since  $G$  is finite, exactly one face extends to infinity, while all other faces are *bounded*. The portion of  $G$  that lies on the boundary of a face  $f$  is the PSLG  $\partial f$ . If  $f$  is a simply connected region, then the PSLG  $\partial f$  is a *weakly simple polygon*, for convenience called *polygon* in this paper. A polygon  $P$  and its interior jointly form the polygonal domain  $\text{dom}(P) \subset \mathbb{R}^2$ . A subdivision of a polygon  $P$  is a PSLG  $G$  with  $P \subset G \subset \text{dom}(P)$ .

The *length* of a PSLG  $G$ , denoted  $|G|$ , is the total length of the edges of  $G$ . The *perimeter* of a (weakly simple) polygon  $P$  is the length of a shortest closed path that visits all vertices of  $P$  along the boundary. Since this closed path can traverse some edges twice, the perimeter of  $P$  is less than  $2|P|$ .

### 2.1 Our algorithm in a nutshell

We construct an orthogonal network for a given set  $S$  of  $n$  points in the plane (Fig. 1(a)). First, we reduce the problem to a polygon subdivision problem. We construct a constant factor approximation  $T_n$  of a minimum axis-aligned Steiner tree (MAST) of  $S$ .  $T_n$  retains a key property of a MAST, which we call *loftiness*. Intuitively, a PSLG  $G$  is lofty if nearby parallel edges do not form "narrow channels." Such narrow channels are undesirable because the detour between closest points on opposite sides of a channel is too large. We enclose  $T_n$  in an appropriate axis-aligned bounding square  $B$ , add a segment connecting  $T_n$  and  $B$  and thus obtain a *lofty* weakly simple polygon  $P$  (Fig. 1(b)). It suffices to subdivide  $P$  into polygonal faces of constant geometric dilation such that the total length and the number of vertices increase by at most constant factors.




Figure 1: The three main steps of our algorithm. (a) A point set; (b) a rectilinear Steiner tree  $T_n$  and a bounding square  $B$ ; (c) a mountain subdivision; and (d) a refined subdivision into polygons of constant geometric dilation.

We augment  $P$  with new edges and vertices in two phases. The first phase decomposes a lofty axis-aligned polygon into lofty *pocketed mountain* polygons; see Def. 3 and Fig. 1(c). The advantage of mountains is that it is easy to approximate their dilation in terms of the detours of horizontal and vertical point pairs (see Lemma 2). In the second phase, we greedily decompose each pocketed mountain polygon into pocketed mountains of constant dilation in a top-down swepline algorithm: Whenever the portion of a mountain above the swepline has "critical" vertical or horizontal dilation, we insert new edges that separate this area and an adjacent buffer zone from the rest of the mountain (Fig. 1(d)). The buffer zones serve

to make sure that the detour is bounded by a constant for points lying on the newly inserted edges.

## 2.2 Reduction to axis-aligned subdivisions

Let  $P$  be a polygon. The *internal dilation* of  $P$  is  $\delta_{\text{int}}(P) = \sup d_G(u, v) / |uv|$  over all point pairs  $u, v \in P$  such that the line segment  $uv$  lies in  $\text{dom}(P)$ . To prove a constant bound on the geometric dilation in Theorem 1 part (i), it will suffice to bound the internal dilation of all polygonal faces of the final network. For this, recall a result of Ebberts-Baumann *et al.* [12] which says that the dilation of a plane graph  $G$  is attained for a pair  $u, v$  of *visible points* (where  $u, v \in G$  but the relative interior of the segment  $uv$  is disjoint from  $G$ ). In our final graph  $G$ , any pair of visible points lie on the boundary of a bounded (polygonal) face of  $G$ . Thus if the internal dilation of every bounded face is at most  $c_1$ , then the dilation of  $G$  is also at most  $c_1$ .

**Theorem 2** *For every set  $S$  of  $n$  points in the plane, there is an axis-aligned subdivision  $G$  of a bounding square of  $S$  such that (i) the internal dilation of every bounded face of  $G$  is at most  $c_1$ ; (ii)  $G$  has at most  $c_2 n$  vertices; and (iii)  $|G| \leq c_3 |T_S|$ .*

It is now easy to see that Theorem 1 follows from Theorem 2, since conditions (ii) and (iii) are the same.

## 2.3 Reduction to lofty axis-aligned polygons

Given a set  $S$  of  $n$  points in the plane, we first construct a Steiner spanning tree  $T_n$  with  $S \subseteq V(T_n)$ . Ideally,  $T_n$  should be the *minimum axis-aligned Steiner tree (MAST)* of  $S$ , which has at most  $2n - 1$  vertices and whose length is at most  $\sqrt{2}|T_S|$ . Since the MAST problem is NP-complete [8], we construct  $T_n$  as an approximation of a MAST that retains three important properties: it has at most  $2n - 1$  vertices,  $\sqrt{2}|T_S|$  length, and is 2-lofty as defined below.

**Definition 1** *Given an axis-aligned PSLG  $G$  and a parameter  $\kappa \geq 1$ , a  $\kappa$ -narrow channel is an axis-aligned rectangle  $r$  of aspect ratio at least  $\kappa$  such that (a) the two longer sides of  $r$  are contained in two parallel edges of  $G$  (but neither of these sides contains any vertex of  $G$ ); (b) the interior of  $r$  is disjoint from  $G$ . (See Fig. 2(a).)*

*An axis-aligned PSLG  $G$  is  $\kappa$ -lofty, for  $\kappa \geq 1$ , if it does not admit any  $\kappa$ -narrow channel.*

By definition, if  $\kappa_1 < \kappa_2$ , and  $G$  is  $\kappa_1$ -lofty, then it is also  $\kappa_2$ -lofty. Note that a MAST  $T$  is  $\kappa$ -lofty for any  $\kappa > 2$ : If there were a  $\kappa$ -narrow channel  $r$  with  $\kappa > 2$  for an MAST  $T$ , then one can construct a shorter axis-aligned Steiner tree by replacing a portion  $T$  along a longer side of  $r$  with the two shorter sides of  $r$  (see Fig. 2(a-b)).

It is not difficult to devise a constant-factor approximation to the MAST that is also 2-lofty. Start with an arbitrary input point  $p_1 \in S$  and let  $T_1 = \{p_1\}$  be a singleton graph. For every  $i = 2, 3, \dots, n$ , construct an axis-aligned Steiner tree  $T_i$  on  $i$  points of  $S$  by extending  $T_{i-1}$ . If  $T_{i-1}$  is available, compute the minimum  $L_1$  distance from  $T_{i-1}$  to remaining points and connect  $T_{i-1}$  to a closest point using at most two axis-parallel edges (forming an L-shape) and at most one Steiner point (the closest point in  $T_{i-1}$  or the joint of the L-shape). By Prim's result [20], the axis-parallel Steiner tree  $T_n$  is not longer than the minimum rectilinear spanning tree (which has no Steiner points but the edge length is measured in  $L_1$  norm); which in turn is at most  $\sqrt{2}$  times longer than the minimum spanning tree  $T_S$ .

The above approximation  $T_n$  is also 2-lofty: Assume that the two longer sides of a 2-narrow channel  $r$  lie along two parallel edges  $e_1$  and  $e_2$  of  $T_n$ . Refer to Fig. 2. We may assume that  $e_1$  was created prior to  $e_2$ , and  $e_2$  connects  $q \in S$  to  $T_i$ . Since the aspect ratio of  $r$  is 2, the  $L_1$  distance between  $q$  and  $e_1$  is less than  $|e_2|$ . So  $e_2$  is not part of a shortest axis-parallel path from  $q$  to  $T_i$ : a contradiction.




Figure 2: (a-b) If an axis-aligned Steiner tree  $T$  is not 2-lofty, then it is not minimal. (c) This argument does not work if the two longer sides of  $r$  contains some vertices of  $T$ .

Let  $B'$  be the minimum axis-aligned bounding box of  $S$ , and let  $B$  be a bounding square of  $B'$  of side length  $2|T_S|$  which extends  $B'$  by at least  $|T_S|/2$  in each direction. Let now  $P = P(B)$  be the PSLG formed by the union of  $B$ ,  $T_n$ , and an axis-parallel segment connecting a vertex of  $T_n$  on  $B'$  to the closest point in  $B$ . Note that  $P$  is also 2-lofty, and we have  $|P| \leq (4 \cdot 2 + 1 + \sqrt{2})|T_S| = (9 + \sqrt{2})|T_S|$ . (The perimeter of  $P$ , however, is at most  $(9 + \sqrt{2})|T_S| + (1 + \sqrt{2})|T_S| = (10 + 2\sqrt{2})|T_S|$ .)  $P$  has at most  $2n - 1 + 4 + 1 = 2n + 4$  vertices:  $T_n$  has at most  $2n - 1$  vertices and there are 5 more vertices on the bounding box  $B$ . Note that  $P$  has exactly one bounded face which is simply connected and which lies in the bounding square  $B$ . The following theorem immediately implies Theorem 2.

**Theorem 3** *Every 2-lofty axis-aligned polygon  $P$  with  $n$  vertices has an axis-aligned subdivision  $G$  such that (i) the internal dilation of every face of  $G$  is at most  $c_1$ ; (ii)  $G$  has at most  $a_2n$  vertices; and (iii)  $|G| \leq a_3|P|$ . Here  $c_1$ ,  $a_2$ , and  $a_3$  are absolute constants.*

### 3 Subdividing axis-aligned lofty polygons

In this section, we prove Theorem 3 and present an algorithm that constructs an axis-aligned subdivision  $G$  for an input 3-lofty axis-aligned polygon  $P$  with  $n$  vertices. This algorithm has two phases: First we decompose  $P$  into 3-lofty pocketed mountains in Subsection 3.1. In the second phase, we decompose 3-lofty pocketed mountains into axis-aligned polygons of bounded internal dilation (in Subsections 3.2 and 3.3).

In both phases, we add new edges and vertices to  $P$ . We charge every new vertex to old vertices (that is, vertices of  $P$ ) such that each vertex of  $P$  is charged at most  $a_2$  times. Similarly, we charge the length of every new edge to portions of edges of  $P$  of the same length such that each point of  $G$  is charged at most  $a_3$  times.

#### 3.1 Subdividing lofty polygons into lofty pocketed mountains

We partition a 3-lofty axis-aligned polygon into 3-lofty *pocketed mountain* polygons defined below. We start with the definition of *mountain* polygons and attach *pockets* to them later.

**Definition 2** (see Fig. 4) *A vertical mountain (alternatively, histogram) is an axis-aligned polygon  $P$  that has a special horizontal base side  $b$ , such that for every point  $u \in P$  there is a vertical segment  $uv \subset \text{dom}(P)$  that connects  $u$  to a point  $v \in b$ . Horizontal mountains (with a vertical base) are defined analogously.*

Our algorithm is a modified version of a standard algorithm that subdivides an axis-aligned polygon  $P$  into mountains. For completeness, we first present this standard algorithm. Its input is  $P$  and a *base edge*  $b$ .

Rotate  $P$  to make  $b$  horizontal. Let  $M(b)$  be the boundary polygon of the set of all points  $x \in \mathbb{R}^2$  for which  $\exists y \in b$  such that  $xy$  is vertical and  $xy \subset \text{int}(P)$ . Clearly,  $M(b)$  is a mountain. If  $P \neq M(b)$ , then  $P$  decomposes into  $M(b)$  and other faces, each of which has a unique edge that is adjacent to  $M(b)$  but is not contained in edges of  $P$ . Recurse on each face, except for  $M(b)$ , independently, setting the base to be the edge adjacent to  $M(b)$ . (See Fig. 3.)




Figure 3: The progress of the standard subdivision algorithm into mountains.

Unfortunately this standard subdivision scheme may produce mountains  $M(b)$  that are not 3-lofty. A narrow channel may be located either outside of  $\text{dom}(M(b))$  along a vertical edge of  $M(b)$ , or in  $\text{dom}(M(b))$  between two vertical edges of  $M(b)$ . To eliminate all narrow channels, we extend the faces of the graph  $G$  to fill adjacent narrow channels. Intuitively, we attach "pockets" to the mountains.

**Definition 3** (see Fig. 4) A vertical (horizontal) pocketed mountain is a polygon obtained from a vertical (horizontal) mountain  $M$  by replacing some segments  $s$  along  $M$  by a 3-path  $p_s$  such that  $s \cup p_s$  forms a rectangle  $r_s$  (a pocket) lying outside  $\text{dom}(M)$ , where the side of  $r_s$  orthogonal to  $s$  has length at most  $|s|/2$ .




Figure 4: A mountain polygon (left) and a pocketed mountain polygon (right).

**Lemma 1** Every axis-aligned 3-lofty polygon  $P$  with  $n$  vertices admits an orthogonal subdivision  $G$ , where: every face of  $G$  is a 3-lofty pocketed mountain;  $|G| \leq 3|P|$ ; and  $G$  has at most  $66n$  vertices.

**Proof.** We describe a recursive algorithm, whose input is a polygon  $P$  and a base segment  $b$  contained in  $P$ , which computes a subdivision of  $P$  into 3-lofty pocketed mountains. Initially  $b$  is an arbitrary horizontal edge of  $P$ . Let  $M(b)$  be the boundary polygon of the set of all points  $x \in \mathbb{R}^2$  for which  $\exists y \in b$  such that  $xy$  is vertical and  $xy \subset \text{int}(P)$ . See Fig. 5(a-b)). The graph  $G = P \cup M(b)$  is a subdivision of  $P$ , in which  $M(b)$  is a face. If  $P = M(b)$ , then the algorithm terminates, otherwise it modifies  $G$  in several steps to eliminate all  $\kappa$ -narrow channels,  $\kappa \geq 3$ . The main tool is the following *pocketing* subroutine, which extends a face  $f$  of  $G$  by attaching to it adjacent narrow channels.

*pocketing*( $G, f$ ) (see Fig. 5). Input: A PSLG  $G$  and a face  $f$ .

As long as  $G$  has a  $\kappa$ -narrow channel not contained in  $\text{dom}(f)$  with  $\kappa \geq 3$  and one of its long sides lying along  $\partial f$ , do:

Let  $r$  be such a narrow channel with maximal  $\kappa$ . Let  $r'$  be the rectangle obtained from  $r$  by removing two rectangles of aspect ratio 2 along its top and bottom sides. Delete the long side of  $r'$  that lies along  $\partial f$  and insert the two short sides of  $r'$  into  $G$ .

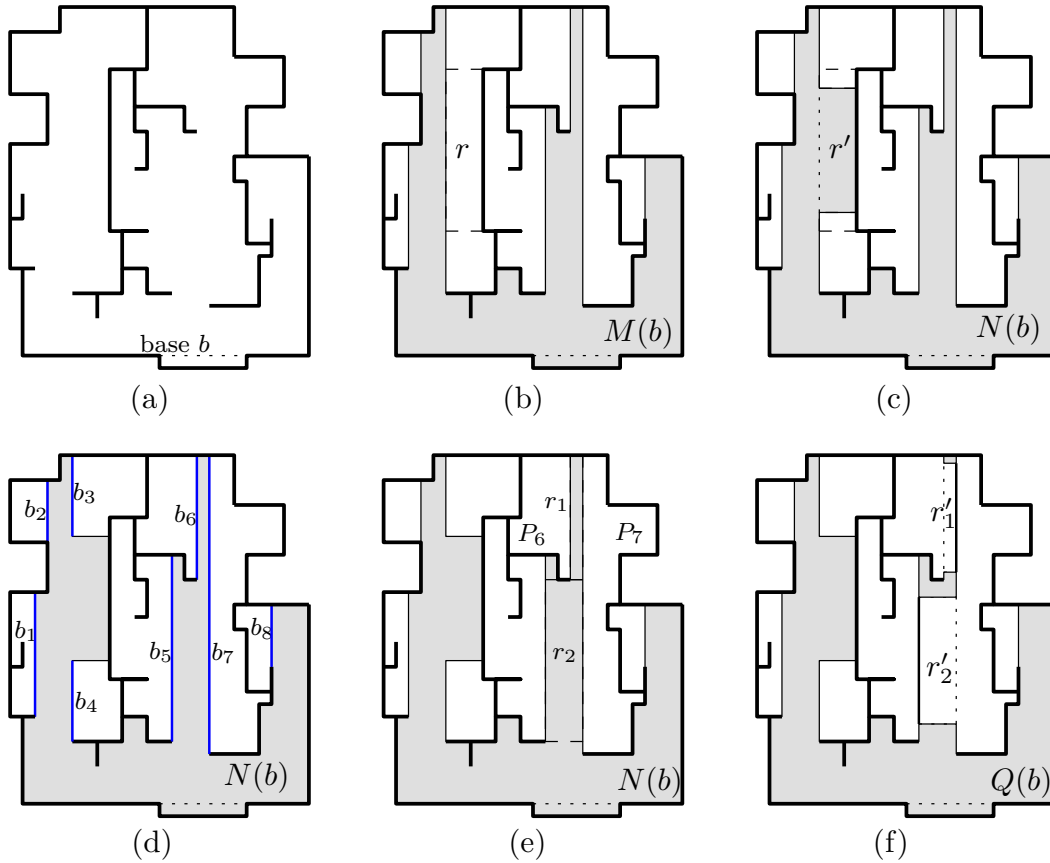


Figure 5: (a) A polygon  $P$  with a (pocketed) base  $b$ . (b) The polygon  $M(b)$  is adjacent to a narrow channel  $r$ . (c) Subroutine *pocketing*( $G, M(b)$ ) extends  $M(b)$  to a polygon  $N(b)$ . (d) The base of every face in  $\mathcal{P}(b)$  is the vertical edge adjacent to  $N(b)$ . (e) Narrow channels  $r_1$  and  $r_2$  in  $\text{dom}(N(b))$ . (f) Subroutines *pocketing*( $G, P_6$ ) and *pocketing*( $G, P_7$ ) splits  $N(b)$  into several polygons  $Q(b)$ .

Apply *pocketing*( $G, M(b)$ ) (see Fig. 5(b-c)). Since  $P$  is 3-lofty, any narrow channel of  $G$  in the exterior of  $\text{dom}(M(b))$  must lie between a vertical edge of  $M(b)$  and a vertical edge of  $P$ . Note that the

pockets are disjoint. Each step of the pocketing subroutine adds at most four new vertices (the four corners of  $r'$ ) and two new edges (the two horizontal sides of  $r'$ ). The removal of two small rectangles from  $r$  along its horizontal sides guarantees that the new horizontal edges do not form 3-narrow channels with any other edge.

Let  $N(b)$  denote the face of the resulting subdivision  $G$  that contains  $M(b)$  (that is,  $N(b)$  contains  $M(b)$  and some adjacent pockets). Denote by  $\mathcal{P}(b)$  the set of all faces of  $G$  except for  $N(b)$ . In each face  $P' \in \mathcal{P}(b)$ , choose a base  $b'$ , which is the unique edge adjacent to  $N(b)$  but not contained in  $P'$  (Fig. 5(d)). Apply subroutine `pocketing`( $G, P'$ ) for every  $P' \in \mathcal{P}(b)$  successively (Fig. 5(e-f)). This destroys narrow channels lying in  $\text{dom}(N(b))$  by attaching pockets to the base sides of the faces in  $\mathcal{P}(b)$ . It may also split the face  $N(b)$  into several faces: Let  $Q(b)$  denote the polygon(s) obtained from  $N(b)$  after this procedure. The graph  $G$  is a subdivision of polygon  $P$  into faces, where the faces of  $Q(b)$  are pocketed mountains, and every base  $b'$  of other faces  $P' \in \mathcal{P}(b)$  may have pockets attached. Apply this subdivision algorithm recursively with input  $(P', b')$  for every face  $P' \in \mathcal{P}(b)$ , independently. This completes the description of the algorithm.

**Charging scheme for length:** We charge every edge created by our algorithm to a portion of the perimeter of  $P$ . Recall that each step of the pocketing subroutines removes a long side of a rectangle  $r'$  of aspect ratio at least 2 and inserts its two short sides. Clearly, this operation does not increase the length of the graph. Assume that the edge set of  $P$  is  $E_0 = E(P)$ . Let  $E_1$  be the set of new edges constructed when building polygons  $M(b)$  for all bases  $b$  in our algorithm. It is enough to charge the total length of edges in  $E_1$  to the perimeter of  $P$ . Consider a step where the base  $b$  is horizontal, and the mountain  $M(b)$  extends vertically above  $b$ . Charge the length of each edge  $e \in E_1 \cap M(b)$  to the portion of the perimeter of  $P$  that is horizontally visible from  $e$ , and has the same length as  $e$ . (Note that the shorter edges of rectangles  $r'$  arising in pocketing subroutines are never charged.) Every point along the perimeter is charged at most once. Hence,  $|G| \leq |P| + 2|P| = 3|P|$ .

**Charging scheme for vertices:** We count the number vertices created during the decomposition of polygon  $P$  with  $n$  vertices. Every edge in  $E_1$  is incident to a reflex vertex of  $P$ ; and every reflex vertex  $v \in V(P)$  is incident to at most one edge of  $E_1$  because if  $v$  is incident to a new edge of some mountain  $M(b)$ , then  $v$  becomes a convex vertex in the recursive steps. Hence, we have  $|E_1| \leq n$ . Each edge of  $E_1$  is incident to a vertex of  $P$  and a potentially new vertex, so the construction of polygons  $M(b)$  increases the number of vertices by  $n$ . Each step of the pocketing subroutines increases the number of vertices by 4 (the corners of a rectangle  $r'$ ). Next, we deduce an upper bound on the number of these steps. First consider the pockets created in subroutines `pocketing`( $G, M(b)$ ) for all bases  $b$ . Every such pocket lies between an edge of  $E_0$  and a parallel edge of  $E_1$ , and every pair of parallel edges in  $E_0 \times E_1$  corresponds to at most one pocket. If we draw a curve in each pocket that connects the two corresponding edges of  $E_0$  and  $E_1$ , we obtain a planar bipartite graph on vertex set  $E_0 \cup E_1$ , which has less than  $2|E_0 \cup E_1|$  edges by Euler's polyhedron theorem. Since  $|E_0| + |E_1| \leq n + n = 2n$ , the number of pockets is less than  $4n$ . These pockets also split some edges of  $E_1$  into several pieces; denote the set of these pieces by  $E_2$ . Each pocket partitions an edge of  $E_1$  into two pieces, so we have  $|E_2| \leq |E_1| + 4n \leq 5n$ . Now consider the pockets created in subroutines `pocketing`( $G, P'$ ) for all  $P' \in \mathcal{P}(b)$ . Each such pocket lies between an edge of  $E_2$  and a parallel edge of  $E_0 \cup E_2$ . By a similar argument, the number of these pockets is at most  $2(|E_0| + |E_2|) \leq 2(n + 5n) = 12n$ . The subdivision  $G$  of the input polygon  $P$  has at most  $n + n + 4(4n + 12n) = 66n$  vertices.  $\square$

### 3.2 Subdividing lofty mountains into polygons of constant geometric dilation

In this subsection, we present an algorithm to subdivide a 3-lofty mountain polygon into polygons of constant internal dilation. We extend this algorithm in the next subsection to 3-lofty *pocketed* mountains. The

advantage of using mountains is that one can approximate their internal dilation in terms of special detours of axis-parallel segments. Consider a mountain  $M$  with a horizontal base side  $b$ . For every horizontal segment  $uv$  that lies in the polygonal domain  $\text{dom}(M)$  and  $u, v \in M$ , we denote by  $d_M^*(u, v)$  the length of the (upper) arc between  $u$  and  $v$  along the perimeter of  $M$  that does not contain the base side.

**Lemma 2** *The internal dilation of every vertical mountain  $M$  is upper bounded by  $\max(\delta_H(M) + 1, \delta_V(M))$ , where*

- $\delta_H(M) = \max_{uv} d_M^*(u, v)/|uv|$  over all horizontal  $uv$ , with  $u, v \in M$ ;
- $\delta_V(M) = |M|/(2|\lambda(M)|)$ , where  $\lambda(M)$  is the shortest vertical segment with endpoints on  $\partial M$  and whose interior lies in  $\text{int}(M)$ .

**Proof.** Consider two points  $p, q \in M$  for which the internal dilation of  $M$  is attained. That is,  $d_M(p, q)/|pq|$  is maximal over all segments  $pq$  that lie in the polygonal domain  $\text{dom}(M)$  and  $p, q \in M$ . We distinguish two cases: (1) either  $p$  or  $q$  lies in the base side, (2) neither  $p$  nor  $q$  lies in the base side.

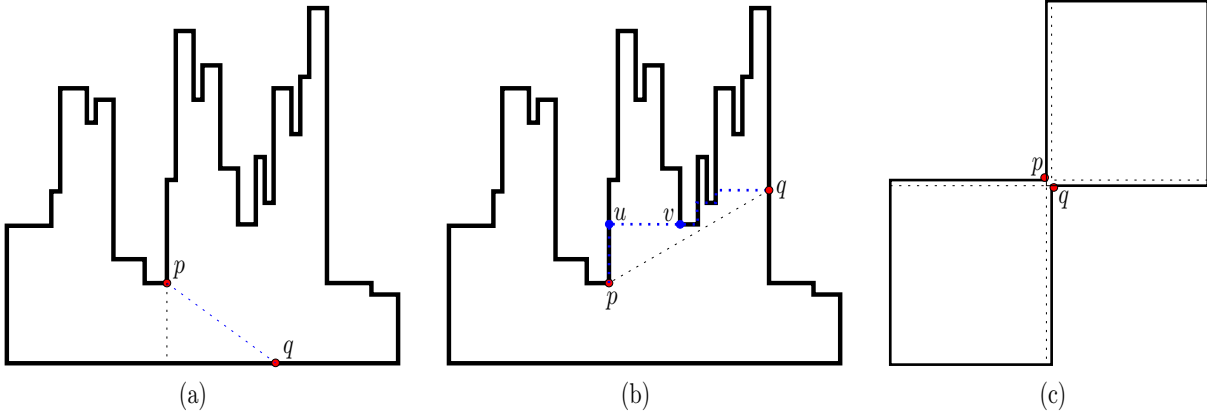


Figure 6: (a-b) Approximating the internal dilation of a mountain in two cases. (c) For an  $x$ - and  $y$ -monotone axis-aligned polygon  $P$ , the internal dilation can be arbitrarily large even though  $\delta_H(P)$  and  $\delta_V(P)$  are bounded.

If  $p \in b$ , then  $q \notin b$  and so  $|pq|$  is at least as long as  $\lambda(M)$ . Since  $d_M(p, q)$  is less than  $|M|/2$ , we have  $d_M(p, q)/|pq| \leq \delta_V(M)$ . Assume that  $p, q \notin b$ . Denote by  $|pq|_H$  (resp.,  $|pq|_V$ ) the length of the horizontal (resp., vertical) component of segment  $pq$ . Let  $\pi(p, q)$  be the staircase path between  $p$  and  $q$  whose vertical segments all lie in  $M$  (see Fig. 6). Clearly, we have  $|\pi(p, q)| = |pq|_H + |pq|_V$ . The graph distance  $d_M(p, q)$  is at most the sum of the graph distances of the portions of  $\pi(p, q)$ , and for every portion above a horizontal segment  $uv$ , it is  $d_M(u, v) \leq \delta_H(M)|uv|$ . Hence,  $d_M(p, q) \leq \delta_H(M)|pq|_H + |pq|_V < (\delta_H(M) + 1)|pq|$ .  $\square$

Note that the internal dilation of arbitrary axis-aligned polygons cannot be bounded in terms of detours of only horizontal and vertical point pairs. Fig. 6(c) shows that an  $x$ - and  $y$ -monotone polygon  $P$  can have arbitrarily large dilation even though the ratio  $d_P(u, v)/|uv|$  is at most 3 for any horizontal or vertical segment  $uv$ .

In the remainder of this subsection, we present and analyze an algorithm for subdividing 3-lofty mountains into axis-aligned polygons with constant internal dilation. Our algorithm greedily chooses polygons for which the dilation bound of Lemma 2 is above a constant threshold. We prove the following.

**Lemma 3** *Every 3-lofty mountain  $M$  with  $n$  vertices admits an orthogonal subdivision  $G$ , where: the internal dilation of every face of  $G$  is at most 45;  $|G| \leq 2|P|$ ; and  $G$  has at most  $52n$  vertices.*

**Proof.** We are given a vertical mountain  $M$  that lies above the  $x$ -axis, with the base side  $b$  on the  $x$ -axis. For every horizontal segment  $s$ , we define a *padding*, which is a rectangle of aspect ratio 3 whose top longer side is  $s$ . Perturb the  $y$ -coordinates of horizontal edge of  $M$  by a tiny  $\varepsilon > 0$ , if necessary, so that no two horizontal edges of  $M$  are collinear. This simplifies our notation when we scan  $M$  in a sweep-line algorithm. Let  $H$  denote the set of maximal horizontal segments  $uv$  where  $u, v \in M$  and  $uv \subset \text{dom}(M)$ . We subdivide  $M$  recursively into 3-lofty mountains as follows.

Move a horizontal sweep line  $\ell$  from the top of  $M$  down, and scan the segments of  $H$  lying along  $\ell$ . We subdivide  $M$  if either of the following two events occurs.

1. If the padding of the segment  $uv$  intersects the base side  $b$ , then insert two vertical edges connecting  $u$  and  $v$  to the base side, and apply the `pocketing` subroutine to the face containing  $uv$ .
2. If  $d_M^*(u, v) = 7|uv|$ , then insert the lower, left, and right edges of the padding of  $uv$ , and apply the `pocketing` subroutine to the face containing  $uv$ .

Recurse on each face of the resulting subdivision of  $M$  that lies in the closed halfplane below  $\ell$ .

An illustration of the subdivision algorithm is shown in Fig. 7. □

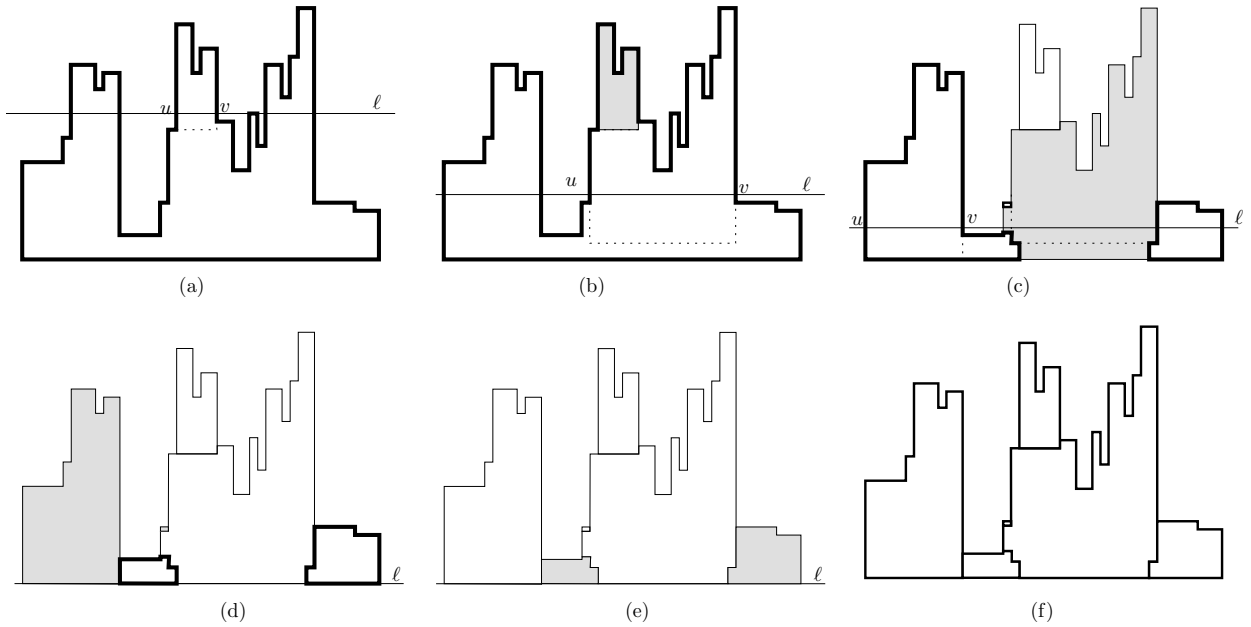


Figure 7: The progress of our subdivision algorithm for mountain polygons.

**Analysis.** We show first that every face in the resulting subdivision of  $M$  has bounded dilation. Initially, when the sweep line  $\ell$  passes through the top horizontal edge of  $M$ , we have  $d_M^*(u, v) = |uv|$ . The ratio  $d_M^*(u, v)/|uv|$  increases while  $u$  and  $v$  move along vertical edges of  $M$ , and it does not increase when  $\ell$  reaches a horizontal edge of  $M$ . The padding of  $uv$  moves continuously with the sweep line  $\ell$  while  $u$  and  $v$  move along vertical edges of  $M$ , but the padding may increase dramatically when  $\ell$  reaches a horizontal edge of  $M$ . Let us consider the case that  $\ell$  passes through segments two visibility  $uv, u'v'$ , and

a horizontal edge  $vu' \subset M$ . (Segment  $vu'$  is a single edge of  $M$ , since no two edges of  $M$  are collinear.) The paddings of  $uv$  and  $u'v'$  are disjoint from  $b$ , and the padding of the edge  $vu'$  is also disjoint from  $b$ , since  $M$  is 3-lofty. Therefore, if the padding of  $uv'$  intersects  $b$ , then the distance of  $\ell$  and  $b$  is at least  $\frac{1}{3} \max(|uv|, |vu'|, |u'v'|) \geq \frac{1}{9}|uv'|$ .

Assume that step 1 is applied. Denote by  $P$  the (mountain) polygon between the vertical segments incident to  $u$  and  $v$ ; and let  $P'$  be the polygon obtained from  $P$  by pocketing. The perimeter of  $P$  is at most  $(7 + 1 + \frac{2}{3})|uv| = \frac{26}{3}|uv|$ . The pocketing subroutine can increase the length of the new edges by a factor of at most  $\frac{5}{3}$ . So the perimeter of  $P'$  is still at most  $(7 + 1 + \frac{5}{3} \cdot \frac{2}{3})|uv| < 10|uv|$ . By construction, we have  $\delta_H(P) < 7 + \frac{2}{3}$  and with pockets  $\delta_H(P') < 7 + \frac{5}{3} \cdot \frac{2}{3} < 9$ . We next give an upper bound on  $\delta_V(P')$ . If  $uv$  does not contain any horizontal edge of  $M$ , then the length of  $\lambda(P)$  is at least  $|uv|/3$ , the height of the padding of  $uv$ . If  $uv$  contains a horizontal edge of  $M$ , then the length of  $\lambda(P)$  is at least  $|uv|/9$ . Hence,  $\delta_V(P') < 10|uv|/(2|uv|/9) < 45$ .

Next assume that step 2 is applied. Denote by  $P$ , the (mountain) polygon  $P$  formed by three sides of  $r(uv)$  and the portion of  $M$  above  $uv$ ; and let  $P'$  be the polygon obtained from  $P$  by pocketing. The perimeter of  $P$  is at most  $(7 + 1 + \frac{2}{3})|uv|$ , after pocketing it is still less than  $(7 + \frac{5}{3}(1 + \frac{2}{3}))|uv| < 10|uv|$ . By construction, we have  $\delta_H(P) = 7$  and pocketing can increase it to at most  $\delta_H(P') < 8$ . The length of every maximal vertical segment  $w_1w_2$ , with  $w_1, w_2 \in P$  is at least  $|uv|/3$ . Hence,  $\delta_V(P') < 15$ .

**Maintaining active polygons.** Our sweep-line algorithm for a 3-lofty mountain produces one face of dilation at most 45: This face is discarded from further consideration, and we run the algorithm on the remaining *active* faces, each of which is a 3-lofty mountain. Let  $C_i$  denote the set of (pairwise disjoint) active mountains after  $i$  steps, assuming that in each step  $i$  only one active mountain  $M \in C_i$  is processed.

**Charging scheme for length:** We account for the total perimeter of the active polygons during our algorithm. Consider the effect of a sweep-line algorithm on a mountain  $M \in C_i$ . If step 1 is applied, then at least  $2|uv|$  portion of the perimeter of  $M$  is discarded from the active polygons, and at most  $\frac{2}{3}|uv|$  long new perimeter is created along the faces of  $C_{i+1}$ . If step 2 is applied, then at least  $7|uv|$  portion of the perimeter of  $M$  is discarded, and at most  $\frac{5}{3}|uv|$  long perimeter is created. It follows that the total length of the new edges is at most half of the perimeter of  $M$ , which is  $2|M|$ . The length of the resulting subdivision is at most  $2|M|$ .

**Charging scheme for vertices:** We classify the set of vertices of the active polygons in  $C_i$ : Let  $W_i$  denote the set of vertices of the mountains in  $C_i$  that are created by pocketing; let  $V_i$  denote the vertices that are either vertices of the input mountain  $M$  or created in step 2 as the lower corners of the padding of  $uv$ ; finally let  $U_i$  denote the set of vertices along the base  $b$  created in step 1. Each subdivision step discards at least 3 vertices of  $V_i$  and creates at most two new vertices in  $V_{i+1} \cup U_{i+1}$ . Hence the total number of vertices in  $V_i$  strictly decreases in each step. There are at most  $n$  subdivision steps, and so at most  $2n$  vertices are created in  $V_i \cup U_i$ .

It remains to bound the number of vertices created by pocketing. Let  $E_0$  be the set of edges of the input mountain  $M$ , and let  $E_1$  be the edges inserted by step 1 (the vertical edges connecting  $u$  and  $v$  to the base side) or step 2 (the left, right and lower sides of the padding of  $uv$ ). Since there are at most  $n$  subdivision steps,  $|E_1| \leq 3n$ . Each pocket corresponds to two edges of  $E_0 \cup E_1$ , and each pair corresponds to at most one pocket. If we connect the two edges of  $E_0 \cup E_1$  in the pocket by a curve, we obtain a planar graph on  $|E_0| + |E_1| \leq 4n$  vertices. The number of pockets corresponds to the number of edges of this planar graph, which is at most  $3(|E_0| + |E_1|) - 6 < 12n$  by Euler's polyhedron theorem. Hence pocketing creates at most  $48n$  vertices. Altogether, at most  $51n$  vertices are created during the algorithm.

### 3.3 Subdividing lofty pocketed mountains into polygons of constant geometric dilation

In Subsection 3.1, we have subdivided a polygon into pocketed mountains, and in Subsection 3.2 we have subdivided mountains into polygons of constant dilation. It remains to show how to subdivide a *pocketed* mountain.

**Lemma 4** *Every 3-lofty pocketed mountain  $P$  with  $n$  vertices admits an orthogonal subdivision  $G$ , where: the internal dilation of every face of  $G$  is at most 75;  $|G| \leq 5|P|$ ; and  $G$  has at most  $102n$  vertices.*

**Proof.** We are given a pocketed mountain  $P$  with  $n$  vertices corresponding to a mountain  $M$  and a set of disjoint pockets (see Def. 3). Recall that each pocket  $r$  has a common side  $s(r)$  with  $M$ , and its sides orthogonal to  $s(r)$  have length at most  $|s(r)|/2$ . Note also that  $M$  is shorter and has fewer vertices than  $P$ .

Subdivide  $M$  into polygons of bounded dilation as described in Subsection 3.2, and let  $G$  be the resulting network. Run the pocketing subroutine for the graph  $G \cup P$  and each face of  $G$  (Subsection 3.1). This may attach all or some portions of each pocket to faces of  $G$ . Consider a maximal portion  $r'$  of a pocket  $r$  that has not been attached to any face of  $G$ . If the aspect ratio of  $r'$  is at most 3, then it is a rectangular face of dilation at most 4. If the aspect ratio of  $r'$  is  $t \geq 3$ , then  $G$  must have at least  $\lfloor t/3 \rfloor$  vertices along  $s(r')$ . Subdivide  $r'$  by segments orthogonal to  $s(r')$  into rectangles of aspect ratio at most 3. In this step, the number of vertices is at most doubled, and the length of  $G \cup P$  increases by at most a factor of  $\frac{7}{6}$ . Pocketing can increase the perimeter of a face of  $G$  by a factor of at most  $\frac{5}{3}$ ; and it can also increase the dilation by at most the same factor.  $\square$

## 4 Conclusion

We have shown that any set of  $n$  points in the plane can be embedded in a planar network with  $O(1)$  geometric dilation, has  $O(n)$  vertices, and  $O(|T_S|)$  length. In addition, we also guaranteed that all edges are axis-parallel (orthogonal network), so the maximum degree is 4. Some questions remain open. The constants hidden in the asymptotic notation are at the moment prohibitively large for practical applications. An obvious question is to improve the constants to a level close to that of low vertex-dilation networks. We do not know if Steiner points are necessary (if the edges are not required to be axis-parallel). Another natural question is whether it is possible to construct a network with the additional property that all bounded faces are convex (e.g., all bounded faces are rectangles in case of an orthogonal network). We conjecture the answer to be negative. Specifically, we suspect that a set  $S$  of  $n$  points uniformly distributed on two concentric circles of radii 1 and 2 has the property that every network with  $O(n)$  vertices,  $O(|T_S|)$  length, and  $O(1)$  dilation must have a nonconvex bounded face.

**Acknowledgments.** We thank Ansgar Grüne and Minghui Jiang for interesting discussions on the topic. We are also grateful to an anonymous reviewer for several useful comments and observations.

## References

- [1] I. Althöfer, G. Das, D. P. Dobkin, D. Joseph, and J. Soares, On sparse spanners of weighted graphs, *Discrete Comput. Geom.* **9** (1993), 81-100.
- [2] B. Aronov, M de Berg, O. Cheong, J. Gudmundsson, H. J. Haverkort, and A. Vigneron, Sparse geometric graphs with small dilation, in *Proc. 16th ISAAC*, vol. 3827 of LNCS, Springer, 2005, pp. 50–59.
- [3] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid, Euclidean spanners: short, thin, and lanky, in *Proc. 27th STOC*, 1995, ACM Press, pp. 489-498.

- [4] M. Benkert, A. Wolff, F. Widmann, and T. Shirabe, The minimum Manhattan network problem: approximations and exact solutions, *Comput. Geom. Theory Appl.* **35** (2006), 188–208.
- [5] P. Bose, J. Gudmundsson, and M. Smid, Constructing plane spanners of bounded degree and low weight, *Algorithmica* **42** (2005), 249–264.
- [6] B. Chandra, G. Das, G. Narasimhan, J. Soares, New sparseness results on graph spanners, *Int. J. Comput. Geometry Appl.* **5** (1995), 125–144.
- [7] L. P. Chew, There are planar graphs almost as good as the complete graph, *J. Computer Sys. Sci.* **39** (1989), 205–219.
- [8] G. Das and D. Joseph, Which triangulations approximate the complete graph? in *Proc. Int. Symp. on Optimal Algorithms*, vol 401 of LNCS, Springer, 1989, pp. 168–192.
- [9] G. Das, G. Narasimhan, and J. S. Salowe, A new way to weigh malnourished Euclidean graphs, in *Proc. 6th SODA*, ACM Press, 1995, pp. 215–222.
- [10] D. P. Dobkin, S. J. Friedman, and K. J. Supowit, Delaunay graphs are almost as good as complete graphs, *Discrete Comput. Geom.* **5** (1990), 399–407.
- [11] A. Dumitrescu, A. Ebberts-Baumann, A. Grüne, R. Klein, and G. Rote, On the geometric dilation of closed curves, graphs, and point sets, *Comput. Geom. Theory Appl.* **36** 2006, 16–38.
- [12] A. Ebberts-Baumann, A. Grüne, and R. Klein, On the geometric dilation of finite point sets, *Algorithmica* **44** (2006), 137–149.
- [13] A. Ebberts-Baumann, R. Klein, E. Langetepe, A. Lingas, A fast algorithm for approximating the detour of a polygonal chain, *Comput. Geom. Theory Appl.* **27** (2004), 123–134.
- [14] D. Eppstein, Spanning trees and spanners, in *Handbook of Computational Geometry* (J. R. Sack and J. Urrutia, eds), North-Holland, Amsterdam, 2000, pp. 425–461.
- [15] J. Gudmundsson, C. Levcopoulos, and G. Narasimhan, Approximating a minimum Manhattan network, *Nordic J. of Computing* **8** (2001), 216–229.
- [16] M. Keil and C. A. Gutwin, Classes of graphs which approximate the complete Euclidean graph, *Discrete Comput. Geom.* **7** (1992), 13–28.
- [17] S. Langerman, P. Morin, and M. Soss, Computing the maximum detour and spanning ratio of planar chains, trees and cycles, in *Proc. 19th STACS*, vol. 2285 of LNCS, Springer, 2002, pp 250–261.
- [18] 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.
- [19] J. MacGregor Smith & P. Winter, Computing in Euclidean geometry, in *Computational Geometry and Topological Network Design*, World Scientific, 1992, pp. 287–385.
- [20] R. C. Prim, Shortest connection networks and some generalizations, *Bell System Technical Journal* **36** (1957), 1389–1401.