Available online at www.sciencedirect.com

ScienceDirect

Computational Geometry 40 (2008) 207–219

Computational Geometry
Theory and Applications

www.elsevier.com/locate/comgeo

Sparse geometric graphs with small dilation ☆

Boris Aronov [a], Mark de Berg [b], Otfried Cheong [c,*], Joachim Gudmundsson [d,1], Herman Haverkort [b], Michiel Smid [e], Antoine Vigneron [f]

[a] *Department of Computer and Information Science, Polytechnic University, Brooklyn, NY, USA*
[b] *Department of Mathematics and Computing Science, TU Eindhoven, Eindhoven, the Netherlands*
[c] *Division of Computer Science, KAIST, Daejeon, South Korea*
[d] *National ICT Australia Ltd, Australia*
[e] *School of Computer Science, Carleton University, Ottawa, Canada*
[f] *Unité Mathématiques et Informatique Appliquées, INRA, Jouy-en-Josas, France*

Received 13 February 2007; received in revised form 28 June 2007; accepted 16 July 2007

Available online 7 August 2007

Communicated by R. Fleischer

Abstract

Given a set S of n points in $\mathbb{R}^D$, and an integer k such that $0 \leqslant k < n$, we show that a geometric graph with vertex set S, at most $n-1+k$ edges, maximum degree five, and dilation $\mathrm{O}(n/(k+1))$ can be computed in time $\mathrm{O}(n \log n)$. For any k, we also construct planar n-point sets for which any geometric graph with $n-1+k$ edges has dilation $\Omega(n/(k+1))$; a slightly weaker statement holds if the points of S are required to be in convex position.

Keywords: Dilation; Spanner; Geometric network; Small-dilation spanning tree

1. Preliminaries and introduction

A *geometric network* is an undirected graph whose vertices are points in $\mathbb{R}^D$. Geometric networks, especially geometric networks of points in the plane, arise in many applications. Road networks, railway networks, computer networks—any collection of objects that have some connections between them can be modeled as a geometric network. A natural and widely studied type of geometric network is the *Euclidean network*, where the weight of an edge

☆ B.A. was supported in part by NSF ITR Grant CCR-00-81964 and by a grant from the US–Israel Binational Science Foundation. Part of the work was carried out while B.A. was visiting TU/e in February 2004 and in the summer of 2005. O.C. was supported by LG Electronics. M.d.B. was supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 639.023.301. M.S. was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). A.V. was supported by NUS research grant R-252-000-166-112.

* Corresponding author.
E-mail addresses: mdberg@win.tue.nl (M. de Berg), otfried@tclab.kaist.ac.kr (O. Cheong), joachim.gudmundsson@nicta.com.au (J. Gudmundsson), cs.herman@haverkort.net (H. Haverkort), michiel@scs.carleton.ca (M. Smid), antoine.vigneron@jouy.inra.fr (A. Vigneron).
URL: http://cis.poly.edu/~aronov (B. Aronov).

1 NICTA is funded through the Australian Government's Backing Australia's Ability initiative, in part through the Australian Research Council.

0925-7721/$ – see front matter
doi:10.1016/j.comgeo.2007.07.004

is simply the Euclidean distance between its two endpoints. Such networks for points in $\mathbb{R}^D$ form the topic of study of our paper.

When designing a network for a given set S of points, several criteria have to be taken into account. In particular, in many applications it is important to ensure a short connection between every two points in S. For this it would be ideal to have a direct connection between every two points; the network would then be a complete graph. In most applications, however, this is unacceptable due to the high costs. Thus the question arises: is it possible to construct a network that guarantees a reasonably short connection between every two points while not using too many edges? This leads to the well-studied concept of *spanners*, which we define next.

The weight of an edge $e = (u, v)$ in a Euclidean network $G = (S, E)$ on a set S of n points is the Euclidean distance between u and v, which we denote by $d(u, v)$. The *graph distance* $d_G(u, v)$ between two vertices $u, v \in S$ is the length of a shortest path in G connecting u to v. The *dilation* (or *stretch factor*) of G, denoted $\Delta(G)$, is the maximum factor by which the graph distance d_G differs from the Euclidean distance d, namely

$$\Delta(G) := \max_{\substack{u,v \in S \\ u \neq v}} \frac{d_G(u, v)}{d(u, v)}.$$

The network G is a *t-spanner* for S if $\Delta(G) \leqslant t$.

Spanners find applications in robotics, network topology design, distributed systems, design of parallel machines, and many other areas and have been a subject of considerable research [3]. Recently spanners found interesting practical applications in metric space searching [21,22] and broadcasting in communication networks [1,12,18]. The problem of constructing spanners has received considerable attention from both an experimental perspective [11,24] and theoretical perspective—see the surveys [9,13,25] and the book by Narasimhan and Smid [20].

The complete graph has dilation 1, which is optimal, but we already noted that the complete graph is generally too costly. The main challenge is therefore to design *sparse* networks that have small dilation. There are several possible measures of sparseness, for example, the total weight of the edges or the maximum degree of a vertex. The measure that we will focus on is the number of edges. Thus the main question we study is this: Given a set S of n points in $\mathbb{R}^D$, what is the best dilation one can achieve with a network on S that has few edges? Notice that the edges of the network are allowed to cross or overlap.

This question has already received ample attention. For example, there are several algorithms [2,17,23,26] that compute a $(1 + \varepsilon)$-spanner for S, for any given constant $\varepsilon > 0$. The number of edges in these spanners is $\mathrm{O}(n)$. Although the number of edges is linear in n, it can still be rather large due to the hidden constants in the O-notation that depend on ε and the dimension D. Das and Heffernan [5] showed how to compute in $\mathrm{O}(n \log n)$ time, for any constant $\varepsilon' > 0$, a t-spanner with $(1 + \varepsilon')n$ edges and degree at most three, where t only depends on ε' and D.

Any spanner must have at least $n - 1$ edges, for otherwise the graph would not be connected, and the dilation would be infinite. In this paper, we are interested in the case when the number of edges is close to $n - 1$, not just linear in n. This leads us to define the quantity $\Delta(S, k)$:

$$\Delta(S, k) := \min_{\substack{V(G)=S \\ |E(G)|=n-1+k}} \Delta(G).$$

Thus $\Delta(S, k)$ is the minimum dilation one can achieve with a network on S that has $n - 1 + k$ edges.

Klein and Kutz [15] showed recently that given a set of points S, a dilation t and a number $k > 0$, it is NP-hard to decide whether $\Delta(S, k) \leqslant t$. Cheong et al. [4] showed that even the special case $k = 0$ is still NP-hard: given a set of points S and a real value $t > 1$, it is NP-hard to decide whether a spanning tree of S with dilation at most t exists. But for the special case when the tree is required to be a star, Eppstein and Wortman [10] gave an algorithm to compute minimum dilation stars of a set of points in the plane in $\mathrm{O}(n \log n)$ expected time.

In this paper we study the *worst-case behavior* of the function $\Delta(S, k)$: what is the best dilation one can guarantee for *any* set S of n points if one is allowed to use $n - 1 + k$ edges? In other words, we study the quantity

$$\delta(n, k) := \sup_{\substack{S \subset \mathbb{R}^D \\ |S|=n}} \Delta(S, k).$$

For the special case when the set S is in $\mathbb{R}^2$ and is required to be in convex position[2] we define the quantity $\delta_C(n,k)$ analogously.

The result of Das and Heffernan [5] mentioned above implies that, for any constant $\varepsilon' > 0$, $\delta(n, \varepsilon' n)$ is bounded by a constant. We are interested in what can be achieved for smaller values of k, in particular, for $0 \leqslant k < n$.

In the above definitions we have placed the perhaps unnecessary restriction that the graph may use no other vertices besides the points of S. We will also consider networks whose vertex sets are supersets of S. In particular, we define a *Steiner tree* on S as a tree T with $S \subset V(T)$. The vertices in $V(T) \setminus S$ are called *Steiner points*. Note that T may have any number of vertices, the only restriction is on its topology.

Our results. We first show that any Steiner tree on a set S of n equally spaced points on a circle has dilation at least n/π. We prove in a similar way that $\delta(n,0) \geqslant \frac{2}{\pi}n - 1$. Eppstein [9] gave a simpler proof of a similar bound for $\delta(n,0)$; we improve this bound by a constant factor.

We then continue with the case $0 < k < n$. Here we give an example of a set S of n points in the plane for which any network with $n-1+k$ edges has dilation at least $\frac{2}{\pi}\lfloor n/(k+1)\rfloor - 1$, proving that $\delta(n,k) \geqslant \frac{2}{\pi}\lfloor n/(k+1)\rfloor - 1$. We also prove that for points in convex position the dilation can be almost as large as in the general case, namely $\delta_C(n,k) = \Omega(n/((k+1)\log n))$.

Next we study upper bounds. We describe an $\mathrm{O}(n \log n)$ time algorithm that computes for a given set S of n points in $\mathbb{R}^D$ and a parameter $0 \leqslant k < n$ a network of at most $n-1+k$ edges, maximum degree five, and dilation $\mathrm{O}(n/(k+1))$. Combined with our lower bounds, this implies that $\delta(n,k) = \Theta(n/(k+1))$. In particular, our bounds apply to the case $k = o(n)$, which was left open by Das and Heffernan [5]. Notice that if $k \geqslant n$, then we have $1 \leqslant \delta(n,k) \leqslant \delta(n,n-1)$, and thus $\delta(n,k) = \Theta(1)$. This means that $\delta(n,k) = \Theta(1 + n/(k+1))$ holds for any $k \geqslant 0$.

Our lower bounds use rather special point sets and it may be the case that more 'regular' point sets admit networks of smaller dilation. Therefore we also study the special case when S is a point set with so-called *bounded spread*. The *spread* of a set S of n points is the ratio between the longest and shortest pairwise distance in S. We show that any set S with spread s admits a network of at most $n-1+k$ edges with dilation $\mathrm{O}(s/(k+1)^{1/D})$. This bound is asymptotically tight for $s = \mathrm{O}(n/(k+1)^{1-1/D})$ (otherwise $\mathrm{O}(n/(k+1))$ is a better bound). This also leads to tight bounds for the case when S is a $n^{1/D} \times \cdots \times n^{1/D}$ grid.

Notation and Terminology. Hereafter S will always denote a set of points in $\mathbb{R}^D$. Whenever it causes no confusion we do not distinguish an edge $e = (u,v)$ in the network under consideration and the line segment uv.

2. Lower bounds

In this section we prove lower bounds on the dilation that can be achieved with $n-1+k$ edges, for $0 \leqslant k < n$, by constructing high-dilation point sets in $\mathbb{R}^2$. Of course, these lower bounds apply as well in higher dimensions.

2.1. Steiner trees

We first show a lower bound on the dilation of any Steiner tree for S. The lower bound for this case uses the set S of n points $p_1, p_2, \ldots, p_n$ spaced equally on the unit circle, as shown in Fig. 1(a), and is based on similar arguments as in Ebbers-Baumann et al. [7] (Lemma 3).

Theorem 1. *For any $n > 1$, there is a set S of n points in convex position such that any Steiner tree on S has dilation at least $\frac{1}{\sin(\pi/n)} > \frac{n}{\pi}$.*

Proof. Consider the set S described above and illustrated in Fig. 1(a). Let o be the center of the circle, and let T be a Steiner tree on S. First, we assume that o does not lie on an edge of the tree.

Let x and y be any two points and let γ and γ' be two paths from x to y avoiding o. We call γ and γ' (*homotopy*) *equivalent* if γ can be deformed continuously into γ' without moving its endpoints or ever passing through the point o, that is, if γ and γ' belong to the same homotopy class in the punctured plane $\mathbb{R}^2 \setminus \{o\}$.

[2] A set of points is *in convex position* if they all lie on the boundary of their convex hull.

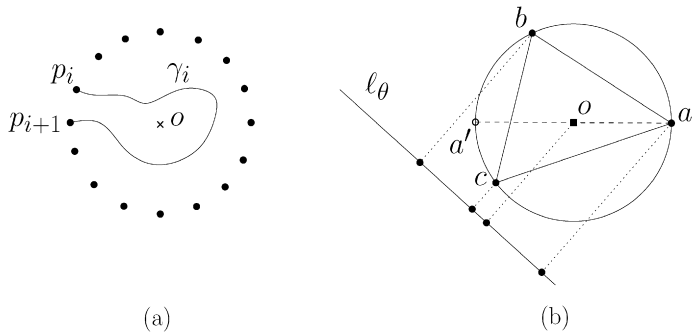

Fig. 1. (a) The homotopy class of the path γ_i. (b) Illustrating the proof of Lemma 1.

Let γ_i be the unique path in T from p_i to p_{i+1} (where $p_{n+1} := p_1$). We argue that there must be at least one index i for which γ_i is not equivalent to the straight segment $p_i p_{i+1}$, as illustrated in Fig. 1(a).

We argue by contradiction. Let Γ be the closed loop formed as the concatenation of $\gamma_1, \ldots, \gamma_n$, and let Γ' be the closed loop formed as the concatenation of the straight segments $p_i p_{i+1}$, for $i = 1 \ldots n$. If γ_i is equivalent to $p_i p_{i+1}$, for all i, then Γ and Γ' are equivalent. We now observe that, since Γ' is a simple closed loop surrounding o, it cannot be contracted to a point in the punctured plane (formally, it has winding number 1 around o). On the other hand, Γ is contained in the tree $T \not\ni o$ (viewed as a formal union[3] of its edges) and hence must be contractible in $\mathbb{R}^2 \setminus \{o\}$. Hence Γ and Γ' cannot be equivalent, a contradiction.

Consider now a path γ_i not equivalent to the segment $p_i p_{i+1}$. Then γ_i must "go around" o, and so its length is at least 2. The distance between p_i and p_{i+1}, on the other hand, is $2\sin(\pi/n)$, implying the theorem.

Now consider the case where o lies on an edge of T. Then some γ_i must pass through o, implying that again its length is at least 2, and the theorem follows as well. □

2.2. *The case $k = 0$*

If we require the tree to be a spanning tree without Steiner points, then the path γ_i in the above proof must not only "go around" o, but must do so using points p_j on the circle only. We can use this to improve the constant in Theorem 1, as follows. Let p_i and p_{i+1} be a pair of consecutive points such that the path γ_i is not equivalent to the segment $p_i p_{i+1}$. Consider the loop formed by γ_i and $p_{i+1} p_i$. It consists of straight segments visiting some of the points of S. Let C be the convex hull of this loop. The point o does not lie outside C (otherwise, the loop would be contractible in the punctured plane) and so there exist three vertices v_1, v_2, and v_3 of C such that $o \in \triangle v_1 v_2 v_3$. Since the loop visits each of these three vertices once, its length is at least the perimeter of $\triangle v_1 v_2 v_3$, which is at least 4 by Lemma 1 below. Therefore, we have proven

Corollary 1. *For any $n > 1$,*

$$\delta_C(n, 0) \geqslant \frac{4 - 2\sin(\pi/n)}{2\sin(\pi/n)} \geqslant \frac{2n}{\pi} - 1.$$

Lemma 1. *Any triangle inscribed in a unit circle and containing the circle center has perimeter at least* 4.

Proof. Let o be the circle center, and let a, b, and c be three points at distance one from o such that o is contained in the triangle $\triangle abc$. We will prove that the perimeter $p(\triangle abc)$ is at least 4.

We need the following definition: For a compact convex set C in the plane and $0 \leqslant \theta < \pi$, let $w(C, \theta)$ denote the *width of C in direction θ*. More precisely, if ℓ_θ is the line through the origin with normal vector $(\cos\theta, \sin\theta)$, then $w(C, \theta)$ is the length of the orthogonal projection of C to ℓ_θ. The Cauchy–Crofton formula [6] allows us to express the perimeter $p(C)$ of a compact convex set C in the plane as $p(C) = \int_0^\pi w(C, \theta)\,\mathrm{d}\theta$.

[3] Notice that T may not be properly embedded in the plane, that is, the edges of T may cross or overlap. However, viewed not as a subset of the plane, but rather as an abstract simplicial complex, T is certainly simply connected and Γ is a closed curve contained in it and thus contractible, *within* T, to a point. Therefore it is also contractible in the punctured plane, as claimed.

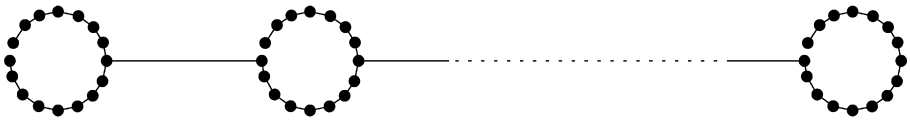

Fig. 2. Illustrating the point set S constructed in the proof of Theorem 2.

We apply this formula to $\triangle abc$, and consider its projection on the line ℓ_θ (for a given θ). We choose a coordinate system where ℓ_θ is horizontal. By mirroring and renaming the points a, b, c, we can assume that b is the leftmost point, that a is the rightmost point, that o lies below the edge ab, and that c does not lie to the right of o. This immediately implies that $w(oc, \theta) \leqslant w(ob, \theta)$. Consider now the point a' obtained by mirroring a at o. Since c must lie below the line aa' and not left of b, we have $w(oc, \theta) \leqslant w(oa', \theta) = w(oa, \theta)$. This implies $3w(oc, \theta) \leqslant w(oa, \theta) + w(ob, \theta) + w(oc, \theta)$, and therefore

$$\begin{aligned} w(\triangle abc, \theta) &= w(oa, \theta) + w(ob, \theta) \\ &= w(oa, \theta) + w(ob, \theta) + w(oc, \theta) - w(oc, \theta) \\ &\geqslant \frac{2}{3}\big(w(oa, \theta) + w(ob, \theta) + w(oc, \theta)\big). \end{aligned}$$

Integrating θ from 0 to π and applying the Cauchy–Crofton formula gives $p(\triangle abc) \geqslant \frac{2}{3}(p(oa) + p(ob) + p(oc))$. Since oa, ob, and oc are segments of length 1, each has perimeter 2, and so we have $p(\triangle abc) \geqslant \frac{2}{3} \cdot 6 = 4$. □

2.3. The general case

We now turn to the general case and consider graphs with $n - 1 + k$ edges for $0 < k < n$.

Theorem 2. *For any n and any k with $0 < k < n$,*

$$\delta(n, k) \geqslant \frac{2}{\pi} \cdot \left\lfloor \frac{n}{k+1} \right\rfloor - 1.$$

Proof. Our example S consists of $k + 1$ copies of the set used in Theorem 1. More precisely, we choose sets S_i, for $1 \leqslant i \leqslant k + 1$, each consisting of at least $\lfloor n/(k + 1) \rfloor$ points. We place the points in S_i equally spaced on a unit-radius circle with center at $(2ni, 0)$, as in Fig. 2. The set S is the union of $S_1, \ldots, S_{k+1}$; we choose the cardinalities of the sets S_i such that S contains n points.

Let G be a graph with vertex set S and $n - 1 + k$ edges. We call an edge of G *short* if its endpoints lie in the same set S_i, and *long* otherwise. Since G is connected, there are at least k long edges, and therefore at most $n - 1$ short edges. Since $\sum |S_i| = n$, this implies that there is a set S_i such that the number of short edges with both endpoints in S_i is at most $|S_i| - 1$. Let G' be the induced subgraph of S_i. If G' is not connected, then S_i contains two points connected in G using a long edge, and therefore with dilation at least n. If G' is connected, then it is a tree, and so the argument of Corollary 1 implies that its dilation is at least

$$\frac{2 - \sin(\pi/\lfloor n/(k+1) \rfloor)}{\sin(\pi/\lfloor n/(k+1) \rfloor)} \geqslant \frac{2}{\pi} \cdot \left\lfloor \frac{n}{k+1} \right\rfloor - 1.$$

This implies the claimed lower bound on the dilation of G. □

2.4. Points in convex position

The point set of Theorem 1 is in convex position, but works as a lower bound only for $k = 0$. In fact, by adding a single edge (the case $k = 1$) one can reduce the dilation to a constant. Now consider n points that lie on the boundary of a planar convex figure with aspect ratio at most ρ, that is, with the ratio of diameter to width at most ρ. It is not difficult to see that connecting the points along the boundary—hence, using n edges—leads to a graph with dilation $\Theta(\rho)$. However, the following theorem shows that for large aspect ratio, one cannot do much better than in the general case.

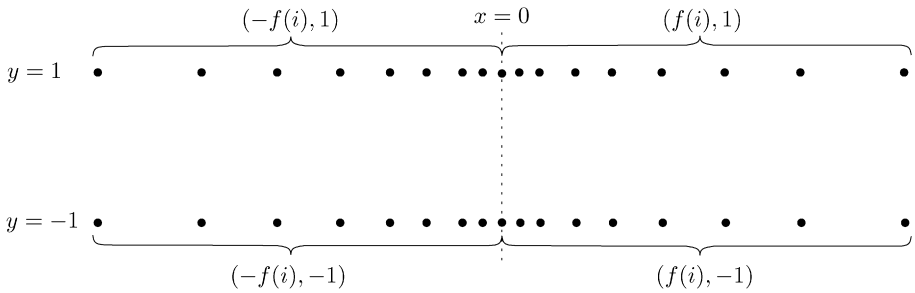

Fig. 3. Illustrating the lower bound construction used in the proof of Theorem 3 with $n = 34$.

Theorem 3. *For any n and any k with $0 \leqslant k < n$,*

$$\delta_{\mathrm{C}}(n,k) = \Omega\left(\frac{n}{(k+1)\log n}\right).$$

Proof. We present the proof for $k = 0$, but the construction can be generalized to hold for any $k > 0$ by using the same idea as in the proof of Theorem 2: placing $k+1$ copies of the construction along a horizontal line, with sufficient space between consecutive copies.

For the case when $k = 0$, set $m = \lfloor n/4 - 1/2 \rfloor$ and let $o := (0,0)$. Consider the function $f(i) = (1 + \frac{\ln m}{m})^i - 1$. Our construction consists of the $4m + 2 \leqslant n$ points S with coordinates $(f(i), 1)$, $(f(i), -1)$, $(-f(i), 1)$, and $(-f(i), -1)$, for $i \in \{0, \ldots, m\}$, as shown in Fig. 3. The points of S lie on the boundary of a rectangle, so S is in convex position. (A slight perturbation of S would even give a set in strictly convex position, that is, a set where every point is extreme.) Consider a minimum-dilation tree T of S. As in the proof of Lemma 1, we can now argue that there must exist two *consecutive* points p and q in S such that the path in T connecting them is not (homotopy) equivalent to the straight-line segment between them in the punctured plane $\mathbb{R}^2 \backslash \{o\}$. By symmetry, we can assume that p and q lie in the half-plane $x \geqslant 0$, and that p lies above the x-axis.

The first case is $p = (f(i), 1)$ and $q = (f(i+1), 1)$, where $0 \leqslant i < m$. The Euclidean distance between them is $f(i+1) - f(i) = \frac{\ln n}{n}(1 + f(i))$, and the length of the shortest path in T is at least $\sqrt{2}(1 + f(i))$. The two bounds imply that the dilation of T is at least $\sqrt{2}\frac{n}{\ln n}$ in this case.

The second case is $p = (f(m), 1)$ and $q = (f(m), -1)$. Then the Euclidean distance between them is 2, and the length of the shortest path in T between them is at least $2f(m)$, thus T has dilation $f(m)$ in this case. It remains to bound $f(m)$, which can be done by using the inequality $(1 + t/m)^m \geqslant e^t(1 - t^2/m)$, which holds for $|t| \leqslant m$ [19, Proposition B.3]. We obtain

$$f(m) = \left(1 + \frac{\ln m}{m}\right)^m - 1 \geqslant m\left(1 - \frac{\ln^2 m}{m}\right) - 1 = \Omega(m),$$

which concludes the proof of the theorem. □

3. A constructive upper bound

In this section we show an upper bound on the dilation achievable with k extra edges. We make use of the fact (observed by Eppstein [9]) that a minimum spanning tree has dilation at most $n - 1$. We will use the following lemma.

Lemma 2. *Let S be a set of n points in $\mathbb{R}^D$, and let T be a minimum spanning tree of S. Then T has dilation at most $n - 1$. Hence $\Delta(S, 0) \leqslant \Delta(T) \leqslant n - 1$ and, as this holds for all S with $|S| = n$, we have $\delta(n, 0) \leqslant n - 1$.*

Proof. Let $p, q \in S$ and consider the path γ connecting p and q in T. Since T is a minimum spanning tree, any edge in γ has length at most $d(p, q)$. Since γ consists of at most $n - 1$ edges, the dilation of γ is at most $n - 1$. □

3.1. The planar case, $D = 2$

Algorithm SPARSESPANNER builds a spanner with at most $n - 1 + k$ edges.

Algorithm 1. SPARSESPANNER(S, k).
Input A set S of n points in the plane and an integer $k \in \{0, \ldots, 2n - 5\}$.
Output A graph $G = (S, E)$ with dilation $O(n/(k + 1))$ and at most $n - 1 + k$ edges.

1: Compute a Delaunay triangulation of S.
2: Compute a minimum spanning tree T of S.
3: **if** $k = 0$ **then**
4: **return** T.
5: Let $m \leftarrow \lfloor (k + 5)/2 \rfloor$.
6: Compute m disjoint subtrees of T, each containing $O(n/m)$ points, by removing $m - 1$ edges.
7: $E \leftarrow \emptyset$.
8: **for** each subtree T' **do**
9: add the edges of T' to E.
10: **for** each pair of subtrees T' and T'' **do**
11: **if** there is a Delaunay edge (p, q) with $p \in T'$, $q \in T''$ **then**
12: add the shortest such edge (p, q) to E.
13: **return** $G = (S, E)$.

We first prove the correctness of the algorithm.

Lemma 3. *Algorithm* SPARSESPANNER *returns a graph G with at most $n - 1 + k$ edges and dilation bounded by* $O(n/(k + 1))$.

Proof. Lemma 2 shows that our algorithm is correct if $k = 0$, so from now on we assume that $k \geqslant 1$, and thus $m \geqslant 3$. The output graph G is a subset of a Delaunay triangulation of S and is therefore a planar graph, see Fig. 4(a). Consider now the graph G' obtained from G by contracting each subtree T' created in step 6 to a single node. Since G is planar and since two subtrees are connected with at most one edge, G' is a planar graph with $m \geqslant 3$ vertices, without loops or multiple edges, and so it has at most $3m - 6$ edges. The total number of edges in the output graph is therefore at most

$$\big(n - 1 - (m - 1)\big) + (3m - 6) = n + 2m - 6 \leqslant n + k - 1.$$

Next we prove the dilation bound. Consider two points $x, y \in S$. Let $x = x_0, x_1, \ldots, x_j = y$ be a shortest path from x to y in the Delaunay graph. The dilation of this path is bounded by $2\pi/(3\cos(\pi/6)) = O(1)$ [14]. We claim that each edge $x_i x_{i+1}$ can be replaced by a path in G with dilation $O(n/k)$. The concatenation of these paths yields a path from x to y with dilation $O(n/k)$, proving the lemma.

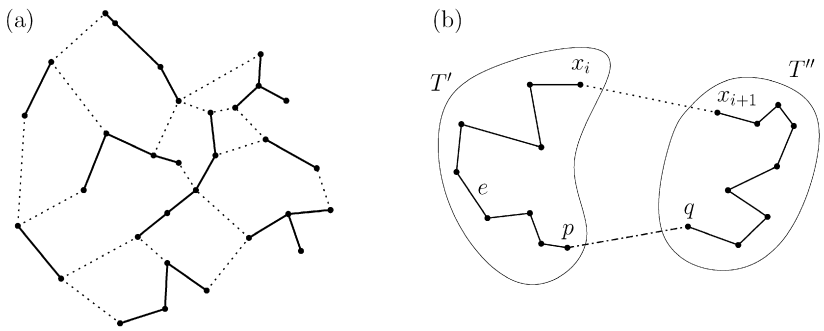

Fig. 4. (a) The solid edges connect vertices within a subtree, while the dotted edges connect different subtrees. (b) Illustrating the proof of Lemma 3.

If x_i and x_{i+1} fall into the same subtree T', then Lemma 2 implies a dilation of $O(n/m) = O(n/k)$.

It remains to consider the case $x_i \in T'$, $x_{i+1} \in T''$, where T' and T'' are distinct subtrees of T, see Fig. 4(b). Let (p, q) be the edge with $p \in T'$, $q \in T''$ inserted in step 12—such an edge exists, since there is at least one Delaunay edge between T' and T'', namely (x_i, x_{i+1}). Let μ' be the path from x_i to p in T', and let μ'' be the path from q to x_{i+1} in T''. Both paths have $O(n/k)$ edges. By construction, we have $d(p, q) \leqslant d(x_i, x_{i+1})$. We claim that every edge e on μ' and μ'' has length at most $d(x_i, x_{i+1})$. Indeed, if e has length larger than $d(x_i, x_{i+1})$, then the original tree T (of which T' and T'' are subtrees) cannot be a minimum spanning tree of S: we can remove e from T and insert either (x_i, x_{i+1}) or (p, q) to obtain a better spanning tree. Therefore the concatenation of μ', the edge (p, q), and μ'' is a path of $O(n/k)$ edges, each of length at most $d(x_i, x_{i+1})$, and so the dilation of this path is $O(n/k)$. □

Theorem 4. *Given a set S of n points in the plane and an integer $0 \leqslant k \leqslant 2n - 5$, a graph G with vertex set S, $n - 1 + k$ edges, and dilation $O(n/(k + 1))$ can be constructed in $O(n \log n)$ time.*

Proof. We use algorithm SPARSESPANNER. Its correctness has been proven in Lemma 3. Steps 1 and 2 can be implemented in $O(n \log n)$ time [9]. Step 6 can be implemented in linear time as follows. Orient T by choosing any node of degree at most five as root node. Since the degree of a minimum spanning tree is at most six [9,27], every node in T has at most five children. We will partition T into disjoint subtrees by removing edges, such that each subtree has size at least n/m and at most $5(n/m) - 4$. The only exception is the top subtree—that is, the subtree containing the root r of T—which has size at most $5(n/m) - 4$ but could be smaller than n/m. To obtain this partition, we first partition the trees rooted at the at most five children of r recursively. All subtrees we get in this manner will have sizes between n/m and $5(n/m) - 4$, except possibly for the subtrees containing the children of r which could be too small. The subtrees that are too small are grouped together with r into one new subtree of size at most $5(n/m) - 4$. Since all but one subtree have size at least n/m, the number of subtrees is at most m. To obtain exactly m subtrees, we can simply remove some more edges arbitrarily. To implement steps 10–12, we scan the edges of the Delaunay triangulation, and keep the shortest edge connecting each pair of subtrees. This can be done in $O(n \log n)$ time. □

3.2. The higher-dimensional case, $D \geqslant 2$

The algorithm of Section 3.1 uses the fact that the Delaunay triangulation has constant dilation, a linear number of edges, and can be computed efficiently. These properties do not generalize to dimensions $D \geqslant 3$, and so we have to turn to another bounded-degree spanner of our point set S. Since the minimum spanning tree is unlikely to be computable in near-linear time in dimension $D \geqslant 3$, we will use a minimum spanning tree *of this spanner* instead of a minimum spanning tree of the point set itself. These two ideas will allow us to generalize our result to any constant dimension $D \geqslant 2$. Moreover, we show that this can be achieved by a graph with degree at most five.

3.2.1. Properties of the minimum spanning tree of a spanner

Let $t \geqslant 1$ be a real number, let G be an arbitrary t-spanner for S, and let T be a minimum spanning tree of G. In the following three lemmas, we prove that T has "approximately" the same properties as an exact minimum spanning tree of the point set S.

Lemma 4. *Every edge on the path in T between two points p and q in S has length at most $t \cdot d(p, q)$.*

Proof. Let (x, y) be an arbitrary edge on the path in T from p to q. For a contradiction, assume that $d(x, y) > t \cdot d(p, q)$.

Let T' be the graph obtained by removing the edge (x, y) from T. It consists of two components, one containing p, the other containing q. Now, since G is a t-spanner, it contains a path γ between p and q of length at most $t \cdot d(p, q)$. The union of T' with the path γ is a spanning graph of G, and its weight is less than the weight of T, a contradiction. □

Lemma 5. *T is a $t(n - 1)$-spanner for S.*

Proof. Let p and q be distinct points of S, and let γ be the path in T between p and q. By Lemma 4, each edge of γ has length at most $t \cdot d(p, q)$. Since γ contains at most $n - 1$ edges, it follows that the length of γ is at most $t(n - 1) \cdot d(p, q)$. □

Lemma 6. *Let m be an integer with $1 \leqslant m \leqslant n-1$, and let T' and T'' be two vertex-disjoint subtrees of T, each consisting of at most m vertices. Let p be a vertex of T', let q be a vertex of T'', and let γ be the path in T between p and q. If x is a vertex of T' that is on the subpath of γ within T', and y is a vertex of T'' that is on the subpath of γ within T'', then*

$$d(x,y) \leqslant \big(2t(m-1)+1\big) \cdot d(p,q).$$

Proof. Let γ' be the subpath of γ between p and x. By Lemma 4, each edge of $\gamma' \subset \gamma$ has length at most $t \cdot d(p,q)$. Since γ' contains at most $m-1$ edges, it follows that this path has length at most $t(m-1) \cdot d(p,q)$. On the other hand, since γ' is a path between p and x, its length is at least $d(p,x)$. Thus, we have $d(p,x) \leqslant t(m-1) \cdot d(p,q)$. A symmetric argument gives $d(q,y) \leqslant t(m-1) \cdot d(p,q)$. Therefore, we have

$$\begin{aligned} d(x,y) &\leqslant d(x,p) + d(p,q) + d(q,y) \\ &\leqslant t(m-1) \cdot d(p,q) + d(p,q) + t(m-1) \cdot d(p,q), \end{aligned}$$

completing the proof of the lemma. □

3.2.2. A graph with at most $n-1+k$ edges and dilation $\mathrm{O}(n/(k+1))$

Fix a constant $t > 1$. Let c be a constant such that for any set X of $2m$ points, the t-spanner of Das and Heffernan [5] has at most cm edges. Let G be a t-spanner for S whose degree is bounded by a constant. Since then the minimum spanning tree T of G has bounded degree as well, we can use the algorithm described in the proof of Theorem 4 (step 6) to remove $m = \lfloor k/(c-1) \rfloor$ edges from T and obtain vertex-disjoint subtrees $T_0, T_1, \ldots, T_m$, each containing $\mathrm{O}(n/(k+1))$ vertices. The vertex sets of these subtrees form a partition of S. Let X be the set of endpoints of the m edges that are removed from T. The size of X is at most $2m$.

We define G' to be the graph with vertex set S that is the union of

1. the trees $T_0, T_1, \ldots, T_m$, and
2. Das and Heffernan's t-spanner G'' for the set X (G'' is empty if $m = 0$).

We first observe that the number of edges of G' is bounded from above by $n-1-m+cm \leqslant n-1+k$.

Lemma 7. *The graph G' has dilation* $\mathrm{O}(n/(k+1))$.

Proof. The statement follows from Lemma 5 if $m = 0$. Let $m > 0$, and let p and q be two distinct points of S. Let i and j be the indices such that p is a vertex of the subtree T_i and q is a vertex of the subtree T_j.

First assume that $i = j$. Let γ be the path in T_i between p and q. Then, γ is a path in G'. By Lemma 4, each edge on γ has length at most $t \cdot d(p,q)$. Since T_i contains $\mathrm{O}(n/k)$ vertices, the number of edges on γ is $\mathrm{O}(n/k)$. Therefore, since t is a constant, the length of γ is $\mathrm{O}(n/k) \cdot d(p,q)$.

Now assume that $i \neq j$. Let γ be the path in T between p and q. Let (x,x') be the edge of γ for which x is a vertex of T_i, but x' is not a vertex of T_i. Similarly, let (y,y') be the edge of γ for which y is a vertex of T_j, but y' is not a vertex of T_j. Then, both (x,x') and (y,y') are edges of T that have been removed when the subtrees were constructed. Hence, x and y are both contained in X and, therefore, are vertices of G''. Let γ_i be the path in T_i between p and x, let γ_{xy} be a shortest path in G'' between x and y, and let γ_j be the path in T_j between y and q. The concatenation γ' of γ_i, γ_{xy}, and γ_j is a path in G' between p and q.

Since both γ_i and γ_j are subpaths of γ, it follows from Lemma 4 that each edge on γ_i and γ_j has length at most $t \cdot d(p,q)$. Since T_i and T_j contain $\mathrm{O}(n/k)$ vertices, it follows that the sum of the lengths of γ_i and γ_j is $\mathrm{O}(n/k) \cdot d(p,q)$. The length of γ_{xy} is at most $t \cdot d(x,y)$ which, by Lemma 6, is also $\mathrm{O}(n/k) \cdot d(p,q)$. Thus, the length of γ' is $\mathrm{O}(n/k) \cdot d(p,q)$. □

Now let G be the t-spanner of Das and Heffernan [5]. This spanner can be computed in $\mathrm{O}(n \log n)$ time, and each vertex has degree at most three. Given G, its minimum spanning tree T and the subtrees $T_0, T_1, \ldots, T_m$ can be computed in $\mathrm{O}(n \log n)$ time. Finally, G'' can be computed in $\mathrm{O}(m \log m) = \mathrm{O}(n \log n)$ time, and each vertex has

degree at most three. Thus G' has dilation $\mathrm{O}(n/(k+1))$, it contains at most $n-1+k$ edges, and it can be computed in $\mathrm{O}(n \log n)$ time.

We analyze the degree of G': Consider any vertex p of G'. If $p \notin X$, then the degree of p in G' is equal to the degree of p in T, which is at most three. Assume that $p \in X$. The graph G'' contains at most three edges that are incident to p. Similarly, the tree T contains at most three edges that are incident to p, but, since $p \in X$, at least one of these three edges is not an edge of G'. Therefore, the degree of p in G' is at most five. Thus, each vertex of G' has degree at most five.

Thus, we have proved the following result:

Theorem 5. *Given a set S of n points in $\mathbb{R}^D$, for $D \geqslant 2$, and an integer k with $0 \leqslant k < n$, a graph G with vertex set S, $n-1+k$ edges, degree at most five, and dilation $\mathrm{O}(n/(k+1))$ can be constructed in $\mathrm{O}(n \log n)$ time.*

4. Bounded spread

In this section we consider the case when the set of input points has bounded spread. The *spread* of a set of n points S, denoted $s(S)$, is the ratio between the longest and shortest pairwise distance in S. In $\mathbb{R}^D$ we have $s(S) = \Omega(n^{1/D})$.

We define the function

$$\delta(n, s, k) := \sup\bigl\{\Delta(S, k) \mid S \subset \mathbb{R}^D, |S| = n, s(S) \leqslant s\bigr\},$$

that is, the best dilation one can guarantee for any set S of n points in $\mathbb{R}^D$ with spread s if one is allowed to use $n-1+k$ edges.

Theorem 6. *For any n, any s, and any k with $0 \leqslant k < n$,*

$$\delta(n, s, k) = \mathrm{O}\bigl(s/(k+1)^{1/D}\bigr).$$

Proof. Assume without loss of generality that the smallest interpoint distance in S is 1.

The case $k \leqslant 2$ is nearly trivial: we pick an arbitrary point $u \in S$ and connect all other points to u. Since the shortest interpoint distance is 1, the dilation is at most $2s$.

For the case $k > 2$, let B be an axis-parallel cube with side length s containing S, and let $m := \lfloor (k-1)^{1/D} \rfloor$. We partition B into m^D small cubes with side length s/m. Let us call each small cube a *cell*. Let $X \subset S$ be a set of *representative points*, that is, one point of S taken from each non-empty cell. Let G' be a t-spanner for X with $2|X|$ edges and constant t [5]. We obtain our final graph G from G' by connecting each point $p \in S$ to the representative point of the cell containing p.

Since $|X| \leqslant m^D \leqslant k-1$, the total number of edges of G is at most $n - |X| + 2|X| \leqslant n-1+k$.

It remains to prove that the dilation of G is $\mathrm{O}(s/m)$ for every two points $p, q \in S$. If p and q lie in the same cell, then the bound follows immediately. If p and q lie in different cells, then let p' and q' be their representative points. The length of the shortest path between p and q in G is then

$$\begin{aligned}
d_G(p,q) &= d_G(p,p') + d_G(p',q') + d_G(q',q) && (1)\\
&\leqslant d(p,p') + t \cdot d(p',q') + d(q',q) && (2)\\
&\leqslant d(p,p') + t \cdot \bigl(d(p',p) + d(p,q) + d(q,q')\bigr) + d(q',q) && (3)\\
&= (1+t)\bigl(d(p,p') + d(q',q)\bigr) + t \cdot d(p,q) && (4)\\
&\leqslant (1+t)2\sqrt{D}\frac{s}{m} + t \cdot d(p,q) && (5)\\
&\leqslant \left((1+t)2\sqrt{D}\frac{s}{m} + t\right) \cdot d(p,q). && (6)
\end{aligned}$$

The last inequality uses $d(p,q) \geqslant 1$. □

To prove a corresponding lower bound, we need a lemma in the spirit of Theorem 1.

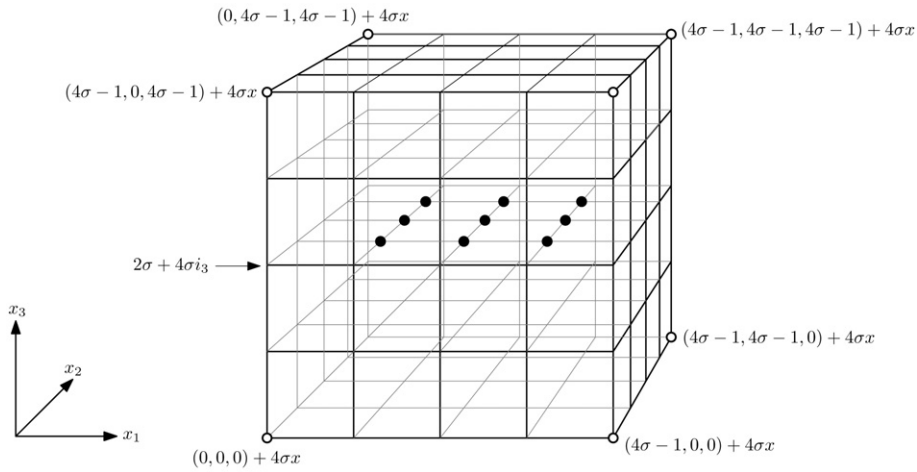

Fig. 5. The cube illustrates the cell of an integer vector x in 3 dimensions and the set of solid black points illustrates the square of the cell.

Lemma 8. *Let Q be a two-dimensional square of side length 2σ in $\mathbb{R}^D$, let S be a set of points on the boundary of Q, and let T be a Steiner tree for S (with vertices in $\mathbb{R}^D$). Then there are two consecutive points $p, q \in S$ such that $d_T(p,q) \geqslant 2\sigma$.*

Proof. Let T' be the orthogonal projection of T onto the plane spanned by Q. Edges of T may cross or overlap in the projection, vertices may even coincide, but viewed as a formal union of its edges, T' is now a Steiner tree for S. The argument made in the proof of Theorem 1 therefore goes through; there are consecutive points p, q in S such that the path from p to q in T' goes around the center of Q, and therefore its length is at least 2σ. □

Theorem 7. *For any integer $r \geqslant 4$, any integer m with $1 \leqslant m \leqslant r/4$, and any integer n with $2rm^{D-1} \leqslant n \leqslant r^D$, there is a subset S of n points of the $r \times r \times \cdots \times r$ integer grid $\{0, 1, \ldots, r-1\}^D$ such that*

$$\delta(S, m^D - 1) \geqslant \frac{r}{2m} - 1.$$

Proof. Let $\sigma := \lfloor r/(4m) \rfloor$. We define m^D small subgrids as follows: For an integer vector $x = (i_1, i_2, \ldots, i_D) \in \{0, \ldots, m-1\}^D$, let the *cell* $C(x) := \{0, 1, \ldots, 4\sigma - 1\}^D + 4\sigma x$. For cell $C(x)$, we define its *square* $Q(x) := (\{\sigma, \ldots, 3\sigma\}^2 \times \{2\sigma\}^{D-2}) + 4\sigma x$, as illustrated in Fig. 5. Note that cells are D-dimensional and contain $4^D\sigma^D$ points each, squares are two-dimensional and contain $(2\sigma+1)^2$ points each. We are in fact only interested in the boundary points of each square, there are 8σ such boundary points in each square, and therefore $8m^D\sigma$ such points in total. Our set S contains all these square boundary points, and then an arbitrary subset of the remaining grid points $\{0, \ldots, r-1\}^D$.

It remains to prove that any network G with $n - 1 + (m^D - 1)$ edges on S has dilation at least 2σ. Let us call an edge e of G *short* if both endpoints of e are in the same cell, and *long* otherwise. Let T be any spanning tree of G. We color the $m^D - 1$ edges of G that are not in T red. Since there are at most $m^D - 1$ short red edges in G and m^D cells, there must be a cell $C(x)$ without a short red edge. We apply Lemma 8 to the boundary points S' of the square $Q(x)$ inside $C(x)$. The tree T is a Steiner tree for S', and so there are consecutive points $p, q \in S'$ such that $d_T(p,q) \geqslant 2\sigma$. If $d_G(p,q) = d_T(p,q)$, then this implies $\Delta(G) \geqslant 2\sigma$. Otherwise, the shortest path from p to q in G must use a red edge. Since $C(x)$ contains no short red edge, the path must therefore pass through a point outside $C(x)$, which implies that it has length at least 2σ, and so $\Delta(G) \geqslant 2\sigma$. □

Combining Theorems 6 and 7, we obtain an asymptotically sharp bound for the dilation.

Corollary 2. *For any n, any k with $0 \leqslant k < n$, and any s with $s = \Omega(n^{1/D})$ and $s = \mathrm{O}(n/(k+1)^{1-1/D})$, we have*

$$\delta(n, s, k) = \Theta\big(s/(k+1)^{1/D}\big).$$

Proof. The upper bound is Theorem 6. For the lower bound, set $r := \lceil s/\sqrt{D} \rceil$ and $m := \lceil (k+1)^{1/D} \rceil$, apply Theorem 7 and observe that the resulting grid set has spread at most s. □

The "most regular" point set one might imagine is the regular grid. It turns out that the spread-based lower bound remains true even for this set.

Corollary 3. *Let $\mathcal{G}$ be a set of $n = r^D$ points forming a $r \times r \times \cdots \times r$-grid, and let $0 \leqslant k < n$. Then*

$$\Delta(\mathcal{G}, k) = \Theta\big(r/(k+1)^{1/D}\big).$$

Proof. The upper bound follows from the fact that $\mathcal{G}$ has spread $s = \Theta(r)$ (better constants can be obtained by a direct construction). The lower bound follows from Theorem 7, by observing that if you set n to r^D, then the construction results in $S = \mathcal{G}$. □

5. Open problems

We have shown that for any n-point set S in $\mathbb{R}^D$ and any parameter $0 \leqslant k < n$, there is a graph G with vertex set S, $n-1+k$ edges, degree at most five, and dilation $\mathrm{O}(n/(k+1))$. We also proved a lower bound of $\Omega(n/(k+1))$ on the maximum dilation of such a graph. An interesting open problem is whether the degree can be reduced to four, or even three.

The constant in our lower bound for $\delta(n/k)$ is $2/\pi$, but we have only proven asymptotically matching upper bounds. Even for the case $k = 0$, the upper bound is $n-1$ while the lower bound is only $2n/\pi - 1$, and it would be interesting to establish the right constant. For $k > 0$ the discrepancy between upper and lower bounds is even larger.

Minimum-dilation graphs are not well understood yet. As mentioned in the introduction, it is NP-hard to decide whether there is a t-spanner of a point set with at most $n-1+k$ edges, even in the case when the spanner is a tree [4]. However, if the spanner is restricted to be a star then the minimum dilation graph can be computed in polynomial time [10]. Are there other restrictions on the spanner or the point set that allow for efficient algorithms?

Given that the general problem is NP-hard, it would also be interesting to look for algorithms that approximate the best possible dilation for a given point set S (instead of giving only a guarantee in terms of n and k, as we do). We are not aware of any result showing how to approximate the minimum-dilation spanning tree with approximation factor $\mathrm{o}(n)$. For general *unweighted* graphs, an $\mathrm{O}(\log n)$ approximation is possible [8]. Another result in this direction is by Knauer and Mulzer [16], who described an algorithm that computes a triangulation whose dilation is within a factor of $1 + \mathrm{O}(1/\sqrt{n})$ of the optimum.

References

[1] K. Alzoubi, X.-Y. Li, Y. Wang, P.-J. Wan, O. Frieder, Geometric spanners for wireless ad hoc networks, IEEE Transactions on Parallel and Distributed Systems 14 (2003) 408–421.
[2] P.B. Callahan, S.R. Kosaraju, A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields, Journal of the ACM 42 (1995) 67–90.
[3] B. Chandra, G. Das, G. Narasimhan, J. Soares, New sparseness results on graph spanners, International Journal of Computational Geometry and Applications 5 (1995) 124–144.
[4] O. Cheong, H. Haverkort, M. Lee, Computing a minimum-dilation spanning tree is NP-hard, in: Proc. of Computing: The Australasian Theory Symposium, in: CRPIT, vol. 65, 2007, pp. 15–24.
[5] G. Das, P. Heffernan, Constructing degree-3 spanners with other sparseness properties, International Journal of Foundations of Computer Science 7 (1996) 121–136.
[6] M. do Carmo, Differential Geometry of Curves and Surfaces, Prentice-Hall, 1976.
[7] A. Ebbers-Baumann, A. Grüne, R. Klein, The geometric dilation of finite point sets, Algorithmica 44 (2006) 137–149.
[8] Y. Emek, D. Peleg, Approximating minimum max-stretch spanning trees on unweighted graphs (full version), in: Proc. ACM–SIAM Symposium on Discrete Algorithms, 2004, pp. 261–270.
[9] D. Eppstein, Spanning trees and spanners, in: J.-R. Sack, J. Urrutia (Eds.), Handbook of Computational Geometry, Elsevier, 2000, pp. 425–461.
[10] D. Eppstein, K. Wortman, Minimum dilation stars, in: Proc. ACM Symposium on Computational Geometry, 2005, pp. 321–326.
[11] M. Farshi, J. Gudmundsson, Experimental study of geometric t-spanners, in: Proc. of the 13th Annual European Symposium on Algorithms, in: Lecture Notes in Computer Science, vol. 3669, Springer-Verlag, 2005, pp. 556–567.

[12] A.M. Farley, A. Proskurowski, D. Zappala, K.J. Windisch, Spanners and message distribution in networks, Discrete Applied Mathematics 137 (2004) 159–171.
[13] J. Gudmundsson, C. Knauer, Dilation and detour in geometric networks, in: T.F. Gonzalez (Ed.), Handbook on Approximation Algorithms and Metaheuristics, Chapman & Hall/CRC, Boca Raton, 2007.
[14] J.M. Keil, C.A. Gutwin, Classes of graphs which approximate the complete Euclidean graph, Discrete & Computational Geometry 7 (1992) 13–28.
[15] R. Klein, M. Kutz, Computing geometric minimum-dilation graphs is NP-hard, in: Proc. of Graph Drawing, in: Lecture Notes in Computer Science, vol. 4372, Springer-Verlag, 2006, pp. 196–207, 556–567.
[16] C. Knauer, W. Mulzer, An exclusion region for the minimum dilation triangulation, in: Proc. of the 21st European Workshop on Computational Geometry, 2005, pp. 33–36.
[17] C. Levcopoulos, 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.
[18] X.-Y. Li, Applications of computational geometry in wireless ad hoc networks, in: X.-Z. Cheng, X. Huang, D.-Z. Du (Eds.), Ad Hoc Wireless Networking, Kluwer, 2003.
[19] R. Motwani, P. Raghavan, Randomized Algorithms, Cambridge University Press, 1995.
[20] G. Narasimhan, M. Smid, Geometric Spanner Networks, Cambridge University Press, 2007.
[21] G. Navarro, R. Paredes, Practical construction of metric t-spanners, in: Proc. 5th Workshop Algorithm Engineering and Experimentation, SIAM Press, 2003, pp. 69–81.
[22] G. Navarro, R. Paredes, E. Chávez, t-spanners as a data structure for metric space searching, in: Proc. 9th International Symposium on String Processing and Information Retrieval, in: Lecture Notes in Computer Science, vol. 2476, Springer-Verlag, 2002, pp. 298–309.
[23] J.S. Salowe, Constructing multidimensional spanner graphs, International Journal of Computational Geometry and Applications 1 (1991) 99–107.
[24] M. Sigurd, M. Zachariasen, Construction of minimum-weight spanners, in: Proc. of the 12th European Symposium on Algorithms, in: Lecture Notes in Computer Science, vol. 3221, Springer-Verlag, 2004, pp. 797–808.
[25] M. Smid, Closest point problems in computational geometry, in: J.-R. Sack, J. Urrutia (Eds.), Handbook of Computational Geometry, Elsevier Science Publishers, Amsterdam, 2000, pp. 877–935.
[26] P.M. Vaidya, A sparse graph almost as good as the complete graph on points in K dimensions, Discrete & Computational Geometry 6 (1991) 369–381.
[27] A.C. Yao, On constructing minimum spanning trees in k-dimensional spaces and related problems, SIAM Journal on Computing 11 (1982) 721–736.