

ROBUST GEOMETRIC SPANNERS*

PROSENJIT BOSE[†], VIDA DUJMOVIĆ[‡], PAT MORIN[†], AND MICHEL SMID[†]

Abstract. Highly connected and yet sparse graphs (such as expanders or graphs of high treewidth) are fundamental, widely applicable, and extensively studied combinatorial objects. We initiate the study of such highly connected graphs that are, in addition, geometric spanners. We define a property of spanners called robustness. Informally, when one removes a few vertices from a robust spanner, this harms only a small number of other vertices. We show that robust spanners must have a superlinear number of edges, even in one dimension. On the positive side, we give constructions, for any dimension, of robust spanners with a near-linear number of edges.

Key words. spanners, stretch-factor, spanning-ratio, treewidth, connectivity, expansion

AMS subject classifications. 68M10, 05C10, 65D18

DOI. 10.1137/120874473

1. Introduction. The cost of building a network, such as a computer network or a network of roads, is closely related to the number of edges in the underlying graph that models this network. This gives rise to the requirement that this graph be sparse. However, sparseness typically has to be counterbalanced with other desirable graph (that is, network design) properties such as reliability and efficiency.

The classical notion of graph connectivity provides some guarantee of reliability. In particular, an r -connected graph remains connected as long as fewer than r vertices are removed. However, these graphs are not sparse for large values of r ; an r -connected graph with n vertices has at least $rn/2$ edges.

For many applications, disconnecting a small number of nodes from the network is an inconvenience for the nodes that are disconnected but has little effect on the rest of the network. In contrast, disconnecting a large part (say, a constant fraction) of the network from the rest is catastrophic. For example, it may be tolerable that the failure of one network component cuts off internet access for the residents of a small village. However, the failure of a single component that eliminates all communications between North America and Europe would be disastrous.

This global notion of connectivity is captured in graph theory by expanders and graphs of high treewidth, each of which can have a linear number of edges. These two properties of graphs have an enormous number of applications and have been the subject of intensive research for decades. See, for example, the book by Klops [30] or the surveys by Bodlaender [11, 12] on treewidth and the survey by Hoory, Linial, and Wigderson [27] on expanders.

In this paper, we consider how to combine this global notion of connectivity with another desirable property of geometric graphs: low spanning ratio (a.k.a., low stretch factor or low dilation), the property of approximately preserving Euclidean distances

*Received by the editors April 23, 2012; accepted for publication (in revised form) June 10, 2013; published electronically August 15, 2013. A preliminary version of this paper appears in *Proceedings of the 29th ACM Symposium on Computational Geometry*, 2013. This work was partly funded by NSERC and CFI.

<http://www.siam.org/journals/sicomp/42-4/87447.html>

[†]School of Computer Science, Carleton University, Ottawa, Ontario K1S 5B6, Canada (jit@scs.carleton.ca, morin@scs.carleton.ca, michiel@scs.carleton.ca).

[‡]School of Mathematics and Statistics and Department of Systems and Computer Engineering, Carleton University, Ottawa, Ontario K1S 5B6, Canada (vida@scs.mcgill.edu).

between vertices. In particular, given a set of n points in \mathbb{R}^d , we study the problem of constructing a graph on these points, where the weights of the edges are given by the Euclidean distance between their endpoints. We wish to construct a graph such that

1. the graph is sparse: the graph has $o(n^2)$ edges;
2. the graph is a spanner: (weighted) shortest paths in the graph do not exceed the Euclidean distance between their endpoints by more than a constant factor; and
3. the graph has high global connectivity: removing a small number of vertices leaves a graph in which a set of vertices of size $n - o(n)$ are all in the same component and all vertices in this set have spanning paths between them.

This is the first paper to consider combining low spanning ratio with high global connectivity. This is somewhat surprising, since many variations on sparse geometric spanners have been studied, including spanners of low degree [6, 19, 36], spanners of low weight [14, 24, 26], spanners of low diameter [8, 9], planar spanners [5, 21, 23, 29], spanners of low chromatic number [13], fault-tolerant spanners [2, 22, 31, 32], low-power spanners [4, 34, 37], kinetic spanners [1, 3], angle-constrained spanners [20], and combinations of these [7, 10, 15, 16, 17, 18]. The closest related work is that on fault-tolerant spanners [2, 22, 31, 32], but r -fault-tolerance is analogous to the traditional definition of r -connectivity in graph theory and suffers the same shortcoming: every r -fault-tolerant spanner has $\Omega(rn)$ edges.

In the next few subsections, we formally define robust spanners and discuss, at a more rigorous level, the relationship between robust-spanners, fault-tolerant spanners, and expanders. From this point onward, all graphs we discuss have vertices that are points in \mathbb{R}^d ; n refers to the number points/vertices; all distances between pairs of points are Euclidean distances; and any shortest path in a graph refers to the shortest (Euclidean) path that uses only edges of the graph.

1.1. Robustness. Let $V \subset \mathbb{R}^d$ be a set of n points in \mathbb{R}^d . An undirected graph $G = (V, E)$ is a (geometric) t -spanner of $V^- \subseteq V$ if for every pair $x, y \in V^-$,

$$\frac{\|xy\|_G}{\|xy\|} \leq t ,$$

where $\|xy\|$ denotes the Euclidean distance between x and y and $\|xy\|_G$ denotes the length of the Euclidean shortest path from x to y that uses only edges in G . Here we use the convention that $\|xy\|_G = \infty$ if there is no path, in G , from x to y . We say simply that G is a t -spanner if it is a t -spanner of V (i.e., $V^- = V$). We point out that, although t is always at least 1, it need not be an integer.

Geometric t -spanners have been studied extensively and have applications in robotics, graph theory, data structures, wireless networks, and network design. A book [33] and handbook chapter [25] provide extensive discussions of geometric t -spanners and their applications.

For a graph $G = (V, E)$ and a subset $S \subseteq V$ of G 's vertices, we denote by $G \setminus S$ the subgraph of G induced by $V \setminus S$. A graph G is an $f(k)$ -robust t -spanner of V if for every subset $S \subseteq V$, there exists a superset $S^+ \supseteq S$, $|S^+| \leq f(|S|)$, such that $G \setminus S$ is a t -spanner of $V \setminus S^+$.

An example is shown in Figure 1.1 which suggests that the $\sqrt{n} \times \sqrt{n}$ grid graph is an $O(k^2)$ -robust 3-spanner. The set S^+ in this example is obtained by choosing "disjoint" squares that cover the vertices of S and adding to S^+ any vertices contained in these squares. A short path between any two vertices in $V \setminus S^+$ is obtained by




FIG. 1.1. From the set S (whose elements are denoted by \bullet) we find a superset S^+ (whose elements are denoted by \times and \bullet) so that $G \setminus S$ is a 3-spanner of $G \setminus S^+$ (whose vertices are denoted by \circ).

starting with some shortest path in G between these two vertices and then routing around any of the square holes encountered by this path. (A proof that the grid graph is indeed an $O(k^2)$ -robust 3-spanner is sketched in section 4.)

One can think of an $f(k)$ -robust t -spanner in terms of network reliability. If a network is an $f(k)$ -robust t -spanner, and k nodes of the network fail, then the network remains a t -spanner of $n - f(k)$ of its nodes. Intuitively, most of the network survives the removal of k nodes, provided that k is small enough that $f(k) \ll n$.

A slightly stronger version of robustness, which is achieved by some of our constructions, requires that $G \setminus S^+$ induces a spanner. Under this definition, the graph $G \setminus S^+$ must be a t -spanner of $V \setminus S^+$. For example, the grid graph in Figure 1.1 satisfies this stronger definition since the vertices inside the squares are not used in the short paths between vertices outside the squares. In some applications, this stronger definition may be preferable since the nodes in S^+ , which no longer gain the full benefits of the network G , are not required to help with the routing of messages between nodes of $V \setminus S^+$. (An open problem related to this stronger definition of robustness is discussed in section 4.)

1.2. Robustness versus fault-tolerance. Robustness is related to, but different from, r -fault tolerance. An r -fault-tolerant t -spanner, $G = (V, E)$, has the property that $G \setminus S$ is a t -spanner of $V \setminus S$ for any subset $S \subseteq V$ of size at most r . In our terminology, an r -fault tolerant spanner is $f(k)$ -robust with

$$f(k) = \begin{cases} k & \text{for } k \leq r, \\ \infty & \text{for } k > r. \end{cases}$$




FIG. 1.2. In an r -fault-tolerant spanner, removing $r + 1$ vertices may disconnect the graph in such a way that no component has size greater than $n/2$.

At a minimum, an r -fault-tolerant spanner must remain connected after the removal of any r vertices. This immediately implies that any r -fault-tolerant spanner with $n > r$ vertices has at least $(r + 1)n/2$ edges, since every vertex must have degree at least $r + 1$. Several constructions of r -fault-tolerant spanners with $O(rn)$ edges exist [22, 31, 32].

In contrast, surprisingly sparse $f(k)$ -robust t -spanners exist. For example, we show that for one-dimensional point sets, there exists $O(k \log k)$ -robust 1-spanners with $O(n \log n)$ edges; the removal of any set of $o(n/\log n)$ vertices leaves a subgraph of size $n - o(n)$ that is a 1-spanner. An r -fault-tolerant spanner with $r = n/\log n$ also has this property, but all such graphs have $\Omega(n^2/\log n)$ edges.

We suggest that in many applications where an r -fault-tolerant spanner is used, an $f(k)$ -robust spanner may be a better choice. For example, one might build an r -fault-tolerant spanner so that a network survives up to r faults, perhaps because more than r faults is viewed as unlikely. Using an $f(k)$ -robust spanner instead means that if $r' \leq r$ faults do occur, then an additional $f(r') - r'$ nodes suffer, but the remaining $n - f(r')$ nodes are unaffected. In one case, the network loses r' nodes, while in the other case $f(r')$ nodes are affected. For slow-growing functions f this may be perfectly acceptable.

The use of an $f(k)$ -robust spanner in place of an r -fault-tolerant spanner has the additional advantage that the maximum number of faults need not be known in advance. In the unlikely event that $r' > r$ faults occur, the network continues to remain usable. In particular, after $r' > r$ faults, the usable network has size at least $n - f(r')$. In contrast, even with $r' = r + 1$ faults, an r -fault-tolerant spanner may have no component of size larger than $n/2$; see Figure 1.2 for an example.

1.3. Robustness and magnification. A function h is called a *magnification* (or *vertex-expansion*) function [28, p. 390] for the graph $G = (V, E)$ if for all $S \subseteq V$,

$$|N(S)| \geq h(|S|) ,$$

where $N(S)$ denotes the set of vertices in $V \setminus S$ that are adjacent to vertices in S . Of particular interest are graphs that have a magnification function $h(x) = cx$ for fixed $c > 1$ and all $x \in \{1, \dots, \lfloor V/2 \rfloor\}$. Such graphs are called *vertex expanders* and have a long history and an enormous number of applications [27].

If G is $f(k)$ -robust, then there exists a magnification function, h , for G that satisfies $h(x) \geq k$ for all $x > f(k) - k$ and every $k \in \{1, \dots, \lfloor \max\{k' : f(k') \leq n/2\} \rfloor\}$; see Figure 1.3. This can be proven by contradiction: If $h(x)$ must be less than k for some $x > f(k) - k$, then there exists a set S'' of size $x > f(k) - k$ such that $|N(S'')| < k$. Taking $S = N(S'') \cup \{x_1, \dots, x_{k-|N(S'')|}\}$, where each x_i is chosen arbitrarily from $V \setminus N(S'')$, yields a set, S , of size k , such that $G \setminus S$ has no component of size greater than $n - x < n - f(k)$.




FIG. 1.3. If G does not have a magnification function, h , with $h(x) \geq k$ for all $x > f(k) - k$, then G is not $f(k)$ -robust.

If we think of $f(k)$ as a continuous increasing function (and hence invertible), then the above argument says that any $f(k)$ -robust spanner with n vertices has a magnification function $h(x)$ such that $h(x) \in \Omega(\min\{n - x, f^{-1}(x)\})$. This implies, for example, that the smallest separator in an $f(k)$ -robust spanner with n vertices has size $\Omega(f^{-1}(n/2))$.

Unfortunately, achieving $f(k)$ -robustness is considerably more difficult than just obtaining a magnification function of the preceding form; there exist vertex expanders with a linear number of edges [27], so they have magnification functions of the form $h(x) = cx$ with fixed $c > 1$. However, these graphs cannot be $f(k)$ -robust since, in Theorem 4, we show that $f(k)$ -robust spanners have a superlinear number of edges for any function $f(k)$.

1.4. Overview of results. In this paper, we prove upper and lower bounds on the size (number of edges) needed to achieve $f(k)$ -robustness. These bounds are expressed as a dependence on the function $f(k)$. In particular, the number of edges depends on the function $f^*(n)$, which is the maximum number of times one can iterate the function f on an initial input k_0 before exceeding n . As a concrete example, if $f(k) = 2k$, then $f^*(n) = \lceil \log_2 n \rceil$ (with the initial input $k_0 = 1$).

Our most general lower-bound, Theorem 4, states that for any constant, $t > 1$, there exists one-dimensional point sets of size n for which any $f(k)$ -robust t -spanner has size $\Omega(nf^*(n))$. For one-dimensional point sets, we can almost match this lower-bound: Theorem 2 states that any one-dimensional point set of size n has an $O(f(k)f^*(k))$ -robust 1-spanner of size $O(nf^*(n))$. Furthermore, if $f(k)$ is sufficiently fast-growing, this construction is $O(f(k))$ -robust and hence has optimal size. For point sets in dimension $d > 1$, our upper and lower bounds diverge by a factor of k . Theorem 5 shows that for any set of n points in \mathbb{R}^d and any fixed $t > 1$, there exists an $O(kf(k))$ -robust t -spanner of size $O(nf^*(n))$.

As a concrete example, we can consider a function $f(k) \in O(k^2)$. Removing any set S of vertices from a n vertex $O(k^2)$ -robust t -spanner leaves a set of at least $n - O(|S|^2)$ vertices which continue to have t -spanning paths between them. Our results show that in one dimension, $O(k^2)$ -robust spanners can be constructed that have $O(n \log \log n)$ edges, and this is optimal. In two and higher dimensions, $O(k^2)$ -robust spanners can be constructed that have $O(n \log n)$ edges.

The remainder of the paper is organized as follows: section 2 gives results for one-dimensional point sets, section 3 gives results for d -dimensional point sets, and section 4 summarizes and concludes with directions for further research.




FIG. 2.1. The graph G' .




FIG. 2.2. Constructing the set S^+ (whose elements are denoted by \times and \bullet) from the set S (whose elements are denoted by \bullet).

2. One-dimensional point sets. In this section, we consider constructions of robust t -spanners for one-dimensional point sets. Throughout this section $V = \{x_1, \dots, x_n\}$ is a set of real numbers with $x_1 < x_2 < \dots < x_n$. We begin by giving a construction of an $O(k \log k)$ -robust 1-spanner having $O(n \log n)$ edges. This construction contains most of the ideas needed for the construction of $O(f(k))$ -robust 1-spanners for more general f .

2.1. An $O(k \log k)$ -robust spanner with $O(n \log n)$ edges. We now consider the following graph, $G_{2\times} = (V, E)$, which is closely related to the hypercube. The edge set, E , of $G_{2\times}$ consists of

$$E = \{x_i x_{i+2^j} : j \in \{0, \dots, \lfloor \log n \rfloor\}, i \in \{1, \dots, n - 2^j\}\} .$$

Notice that $G_{2\times}$ is a 1-spanner since it contains every edge of the form $x_i x_{i+1}$ for $i \in \{1, \dots, n - 1\}$. Furthermore, $G_{2\times}$ has size $O(n \log n)$ since every vertex has degree at most $2\lfloor \log n \rfloor + 2$. We now prove an upper-bound on the robustness of $G_{2\times}$ by using the probabilistic method.

THEOREM 1. *Let $V \subset \mathbb{R}$ be any set of n real numbers. Then there exists an $O(k \log k)$ -robust 1-spanner of V of size $O(n \log n)$.*

Proof. Let S be any nonempty subset of V and let $k = |S|$. Select a random integer $r \in \{0, 1, 2, 3, \dots, 2^{\lfloor \log n \rfloor} - 1\}$ and consider the subgraph, G' , of $G_{2\times}$ consisting only of the edges of the form $x_i x_{i+2^j}$, where $i - r \equiv 0 \pmod{2^j}$. One can think of the edges of G' as a set $O(\log n)$ monotone paths that all contain x_r ; one of these paths contains every vertex in V , another contains every second vertex, yet another contains every fourth vertex, and so on. (For readers with a background in data structures, G' looks a lot like a perfect skiplist in which x_r appears at the top level; see Figure 2.1.)

For a vertex $x_i \in S$, let $j = j(i)$ be the largest integer such that $i - r$ is a multiple of 2^j . Then we say that x_i kills the vertices $x_{i-2^{j+1}}, \dots, x_{i+2^{j+1}}$ in G' ; see Figure 2.2. When this happens, the cost of x_i is $c(x_i) = 2^{j+1} - 1$, which is the number of vertices killed by x_i . Observe that unless $i < 2^j$ or $i > n - 2^j$, G' contains the edge $x_{i-2^j} x_{i+2^j}$ that “jumps over” all the vertices killed by x_i . Therefore, if we define S^+ to be the set of all vertices killed by vertices in S , then $G' \setminus S$ (and hence also $G_{2\times} \setminus S$) is a 1-spanner of $V \setminus S^+$; it contains a path that visits all vertices of $V \setminus S^+$ in order.

We say that a vertex $x \in S$ is cheap if $c(x) < 4k$ and expensive otherwise. We call our choice of r a failure if

1. \mathcal{A} : some vertex of S is expensive; or
2. \mathcal{B} : the total cost of all cheap vertices exceeds $4k \log k + 12k$.

We declare our choice of r a *success* if neither \mathcal{A} nor \mathcal{B} holds. Observe that in the case of a success, we obtain a set S^+ , $|S^+| \in O(k \log k)$, such that $G_{2^x} \setminus S$ is a 1-spanner of $V \setminus S^+$. Therefore, all that remains is to show that the probability of success is greater than 0.

We first note that the probability that any particular $x_i \in S$ is expensive is at most $1/4k$. This is because x_i is expensive if and only if $(i - r) \equiv 0 \pmod{2^{\lceil \log(4k) \rceil}}$. The probability of selecting r with this property is only $1/2^{\lceil \log(4k) \rceil} \leq 1/4k$. Therefore, by the union bound,

$$\Pr\{\mathcal{A}\} \leq k/4k = 1/4 .$$

To upper-bound the total expected cost of cheap vertices, we note that if $x_i \in S$ kills $2^{j+1} - 1$ vertices, then $i - r \equiv 0 \pmod{2^j}$. The probability that this happens is $1/2^j$. Letting S^c denote the set of cheap vertices in S , the total expected cost of all cheap vertices is at most

$$\begin{aligned} \mathbb{E} \left[\sum_{x \in S^c} c(x) \right] &\leq \mathbb{E} \left[\sum_{x \in S} \min\{2^{\lceil \log 4k \rceil}, c(x)\} \right] \\ &\leq k \sum_{j=0}^{\lceil \log(4k) \rceil} (2^{j+1} - 1)/2^j \\ &\leq k \sum_{j=0}^{\lceil \log(4k) \rceil} 2 \\ &\leq 2k \log k + 6k . \end{aligned}$$

Therefore, by Markov's inequality, $\Pr\{\mathcal{B}\} \leq 1/2$. By the union bound,

$$\Pr\{\mathcal{A} \text{ or } \mathcal{B}\} \leq 1/4 + 1/2 < 1 . \quad \square$$

2.2. A general construction. Let $k_0 \geq 1$ be a constant and let $f : \mathbb{R} \rightarrow \mathbb{R}$ be any function that is convex, increasing over the interval $[k_0, \infty)$, and such that $f(k_0 + 1) - f(k_0) > 1$. Let $f^i(k)$ be the function f iterated i times on the initial value k , i.e.,

$$f^i(k) = \underbrace{f(f(\dots f(k) \dots))}_i .$$

We use the convention that $f^0(x) = k_0$ for all x . We define the *iterated f -inverse function*

$$f^*(n) = \max\{i : f^i(k_0) \leq n\} .$$

Notice that for any $k > k_0$, there exists i such that

$$f^i(k_0) < k \leq f(f^i(k_0)) .$$

In particular, the sequence $k_0, f(k_0), f^2(k_0), \dots$, contains a value $f^{i+1}(k_0)$ such that

$$k \leq f^{i+1}(k_0) < f(k) .$$




FIG. 2.3. The vertices killed by x_i .

Another important property is that, since $f(k)$ is increasing and convex and $f(k_0 + 1) - f(k_0) > 1$, the function $f(x)/x$ is nondecreasing for $x \geq k_0$: For every $\delta \geq 0$ and every $x \geq k_0$, $f(x + \delta)/(x + \delta) \geq f(x)/x$.

For a positive number x , we define $\lceil x \rceil = 2^{\lceil \log x \rceil}$ as the smallest power of 2 greater than or equal to x . From the function, f , we define the graph $G_f = (V, E_f)$ to have the edge set

$$E_f = \{x_i x_{i+1} : i \in \{1, \dots, n - 1\}\} \cup \{x_i x_{i+\lceil f^j(k_0) \rceil} : j \in \{0, \dots, f^*(n)\}, i \in \{1, \dots, n - \lceil f^j(k_0) \rceil\}\}.$$

The graph G_f clearly has $O(nf^*(n))$ edges. The following theorem shows that this graph is a robust spanner.

THEOREM 2. *Let f, f^*, k_0 , and G_f be defined as above. Then the graph G_f has $O(nf^*(n))$ edges and is*

1. *an $O(f(4k)f^*(k))$ -robust 1-spanner; and*
2. *an $O(f(4k))$ -robust 1-spanner if $f(k) \in k2^{\Omega(\sqrt{\log k})}$.*

Proof. The proof is very similar to the proof of Theorem 1. Let S be any nonempty subset of V and let $k = |S|$. Select a random integer r from the set $\{0, 1, 2, 3, \dots, \lceil f^{f^*(n)+1}(k_0) \rceil - 1\}$. We consider the subgraph G' of G_f that contains only the edges $x_i x_{i+\ell}$, where $i - r \equiv 0 \pmod{\ell}$. We say that an edge $x_i x_{i+\ell}$ has *span* ℓ .

For an integer i , let $j = j(i)$ be the smallest integer such that $i - r \not\equiv 0 \pmod{\lceil f^j(k_0) \rceil}$; see Figure 2.3. Informally, if G' has any edge that jumps over x_i , then it has an edge of span $\lceil f^j(k_0) \rceil$ that jumps over x_i . Then we say that x_i *kills* $x_{i-p+1}, \dots, x_{i+q-1}$, where

$$p = ((i - r) \bmod \lceil f^j(k_0) \rceil) \text{ and} \\ q = ((i - r) \bmod \lceil f^j(k_0) \rceil) .$$

As before, we define S^+ to be the set of all vertices killed by vertices in S . It is easy to verify, since all edges have spans that are powers of 2, that the graph $G' \setminus S^+$ (and hence also $G_f \setminus S$) contains a path that visits all the vertices of $V \setminus S^+$ in order. Therefore, $G_f \setminus S$ is a 1-spanner of $V \setminus S^+$.

What remains is to show that, with some positive probability, S^+ is sufficiently small to satisfy the appropriate condition, 1 or 2, of the theorem. Define $c(x_i)$ as the number of vertices killed by x_i . We say that x_i is *expensive* if $c(x_i) > f(4k)$ and *cheap* otherwise. If x_i is expensive, then $f^{j(i)-1}(k_0) \geq 4k$ and $i - r \equiv 0 \pmod{\lceil f^{j(i)-1}(k_0) \rceil}$. Therefore, the probability that x_i is expensive is at most $1/f^{j(i)-1}(k_0) \leq 1/4k$. Therefore, by the union bound, the probability that S contains some expensive vertex is at most $1/4$. All that remains is to bound the expected cost of all cheap vertices.

Letting S^c denote the set of cheap vertices in S , we obtain

$$\begin{aligned}
 \mathbb{E} \left[\sum_{x \in S^c} c(x) \right] &\leq k \sum_{j=0}^{f^*(4k)} \frac{\|f^{j+1}(k_0)\|}{\|f^j(k_0)\|} \\
 &\leq 2k \sum_{j=0}^{f^*(4k)} f^{j+1}(k_0)/f^j(k_0) \\
 &= 2k \sum_{j=0}^{f^*(4k)} f(f^j(k_0))/f^j(k_0) \\
 &\leq 2k \sum_{j=0}^{f^*(4k)} f(4k)/4k && \text{(since } f(x)/x \text{ is nondecreasing)} \\
 &\leq (1/2)(f(4k)(f^*(4k) + 1)) .
 \end{aligned}$$

Again, Markov’s inequality implies that the probability that the total cost of all cheap vertices exceeds $f(4k)(f^*(4k) + 1)$ is at most $1/2$. Therefore, the probability of finding a set S^+ of size at most $f(4k)(f^*(4k) + 1)$ is at least

$$1 - 1/2 - 1/4 > 0,$$

which proves the existence of such a set S^+ .

To prove the second part of the theorem, we proceed exactly the same way, except that the sequence $f^{j+1}(k_0)/f^j(k_0)$, $j = 0, 1, 2, \dots$, becomes geometric,¹ so it is dominated by its last term. This yields

$$\begin{aligned}
 \mathbb{E} \left[\sum_{x \in S^c} c(x) \right] &\leq k \sum_{j=0}^{f^*(4k)} \frac{\|f^{j+1}(k_0)\|}{\|f^j(k_0)\|} \\
 &\leq 2k \sum_{j=0}^{f^*(4k)} f^{j+1}(k_0)/f^j(k_0) \\
 &\leq 2ck \left(\frac{f(f^{f^*(4k)}(k_0))}{f^{f^*(4k)}(k_0)} \right) && \text{(for some } c, \text{ since the sum is geometric)} \\
 &\leq 2ck \left(\frac{f(4k)}{4k} \right) && \text{(since } f(x)/x \text{ is nondecreasing)} \\
 &\leq (c/2)f(4k) ,
 \end{aligned}$$

as required. \square

Applying Theorem 2 with different functions $f(k)$ yields the following results.

COROLLARY 1. *For any set V of n real numbers and any constant $\varepsilon > 0$, there exist $f(k)$ -robust 1-spanners $G = (V, E)$ with*

1. $f(k) \in O(k \log k)$ and $O(n \log n)$ edges;
2. $f(k) \in O(k(1 + \varepsilon)^{\sqrt{\log k}})$ and $O(n\sqrt{\log n})$ edges; and
3. $f(k) \in O(k^{1+\varepsilon})$ and $O(n \log \log n)$ edges.

¹This is most easily seen by taking $f(k) = k\delta^{2\sqrt{(\log k)/(\log \delta)+1}}$. Then it is straightforward to verify that $f^j(\delta) = \delta^{(j+1)^2}$, so that $f^{j+1}(\delta)/f^j(\delta) = \delta^{2j+3}$, so the sequence is exponentially increasing. Taking $\delta = 1 + \varepsilon$ for a sufficiently small $\varepsilon > 0$ allows us to lower-bound any function $f(k) \in k^{2\Omega(\sqrt{\log k})}$ this way.




FIG. 2.4. After removing S (denoted by \bullet), there are still two vertices $u, w \in V \setminus S^+$ such that $\|uw\| \leq 2ck$ but $\|uw\|_{G \setminus S} > 2ctk$.

2.3. Lower bounds. In this section, we give lower-bounds on the number of edges in $f(k)$ -robust t -spanners. These lower-bounds hold already for a specific one-dimensional point set (the $1 \times n$ grid), and therefore they apply to all dimensions $d \geq 1$.

2.3.1. A lower bound for linear robustness. We begin by focusing on the hardest case, $f(k) \in O(k)$.

THEOREM 3. *Let $V = \{1, \dots, n\}$ and let $t \geq 1$ be a constant. Then any $O(k)$ -robust t -spanner of V has $\Omega(n \log n)$ edges.*

Proof. To simplify the following discussion, we will assume that $G = (V, E)$ is a ck -robust t -spanner. Note that we have gone from $O(k)$ -robust in the statement of the theorem to ck -robust in the proof. This does not cause a problem so long as we only consider values of k greater than some constant k_0 hidden in the O notation.

We claim that for every natural number k divisible by 4 and every $i \in \{ck + 1, \dots, n - ck - 1\}$, G has at least $k/2$ good edges, xy , such that $x < i - k/4 < i + k/4 < y$ and such that $y - x \leq 2ctk$.

To see why the preceding claim is true, consider the set S that contains $\{i - k/4, \dots, i + k/4\}$ as well as the left endpoint of each good edge (see Figure 2.4). The set S has size at most k , and in $G \setminus S$, the only edges xy with $x < i$ and $y > i$ have length greater than $2ctk$. Now consider any $S^+ \supseteq S$ with $|S^+| \leq ck$. Since $|S^+| \leq ck$ there is at least one element $u \in \{i - ck, \dots, i - 1\}$ that is not in S^+ and at least one element $w \in \{i + 1, \dots, i + ck\}$ that is not in S^+ . Now,

$$w - u \leq 2ck,$$

and in $G \setminus S$, every path from u to w uses an edge of length greater than $2ctk$. Therefore,

$$\frac{\|uw\|_{G \setminus S}}{\|uw\|} > \frac{2ctk}{2ck} = t .$$

This contradicts the assumption that G is ck -robust t -spanner, so we conclude that there are, indeed, at least $k/2$ good edges.

Applying the above argument to $i = ck + j2ctk$ for $j \in \{0, \dots, \lfloor (n - ck)/2ctk \rfloor\}$ implies that G contains $\Omega(n/tc) = \Omega(n)$ edges whose length is in the range $[k/2 + 1, 2ctk]$. Applying this argument for $k \in \{\lceil (4tck_0)^j \rceil : j \in \{0, \dots, \lfloor (\log n)/\log(2tck_0) \rfloor\}\}$ proves that for any constants $c, k_0, t > 1$, G has $\Omega(n \log n)$ edges. \square

2.3.2. A general lower bound. Using the iterated functions from section 2.2, we obtain a whole class of lower-bounds.

THEOREM 4. *Let k_0, f , and f^* be defined as in section 2.2, let $V = \{1, \dots, n\}$, and let $t \geq 1$ be a constant. Then any $f(k)$ -robust t -spanner of V has $\Omega(nf^*(n))$ edges.*

Proof. The proof is similar to the proof of Theorem 3. We need only consider $f(k) \in \omega(k)$ since otherwise we can apply Theorem 3. We group the edges of the graph G into $\Omega(f^*(n))$ classes and show that each class contains $\Omega(n)$ vertices.

In particular, using the same argument one can show that for any $i \in \{1, \dots, f^*(n/t)\}$, any $f(k)$ -robust t -spanner of V has $\Omega(n)$ edges whose lengths are in the range $[f^i(k_0)/2, 2tf^i(k_0)]$. Since $f(k)$ is superlinear, there exists a constant i_0 such that for any $i > i_0$, $f^{i+1}(k_0)/2 > 2tf^i(k_0)$. Thus, the number of edges in any $f(k)$ -robust t -spanner of V is at least

$$\Omega(n) \times (f^*(n/t) - i_0) = \Omega(nf^*(n)) \quad . \quad \square$$

COROLLARY 2. *Let $V = \{1, \dots, n\}$ and let $c > 1$ and $t > 1$ be constants. Then any $f(k)$ -robust t -spanner with*

1. $f(k) \in O(k \log k)$ has $\Omega(n \log n / \log \log n)$ edges;
2. $f(k) \in O(kc^{\sqrt{\log k}})$ has $\Omega(n\sqrt{\log n})$ edges; and
3. $f(k) \in O(k^c)$ has $\Omega(n \log \log n)$ edges.

Note that the lower bounds in parts 2 and 3 of this corollary match the corresponding upper-bounds while the lower-bound in part 1 is off by a factor of $\log \log n$.

Remark. The dependence of our lower bounds on the value of t is not given in the statements of Theorems 3 and 4 or in Corollary 2. However, it is readily extracted from their proofs. In Theorem 3, each value of k shows the existence of $\Omega(n/t)$ edges and there are $\Omega(\log_t n)$ values of k , so the lower-bound is $\Omega((n \log n)/(t \log t))$.

In Theorem 4, each value of k shows the existence of $\Omega(n/t)$ edges, but now the number of values of k is $f^*(n/t) - f^*(x_0)$, where x_0 is the minimum value such that $f(x_0) \geq 4tx_0$. (Informally, x_0 is where the slope of f exceeds $4t$.) Thus, in Theorem 4, the lower-bound is $\Omega((n/t)(f^*(n) - f^*(x_0)))$. It is fairly straightforward to apply this bound to the choices of f used in Corollary 2 or to other choices of f . For example, applying it to case 3 of Corollary 2 we get $x_0 = \Theta(t^{1/(c-1)})$ and the result that any $O(k^c)$ -robust t -spanner has $\Omega((n/t)(\log \log n - \log \log t - \log(1/(c-1))))$ edges.

3. Higher dimensions. In this section, we give a family of constructions for point sets $V \subset \mathbb{R}^d$, $d \geq 1$. These constructions make use of dumbbell tree spanners [33, Chapter 11]. In particular, they make use of binary dumbbell trees, first used by Arya et al. [7] in the construction of low-diameter spanners. A full description of the construction (and proof of existence) of binary dumbbell trees can be found in the notes by Smid [35].

A (binary) dumbbell tree spanner of V is defined by a set of $O(1)$ binary trees $\mathcal{T} = \{T_1, \dots, T_p\}$, each having n leaves. Each node, u , in each of these trees is associated with one element, $r(u) \in V$. For each $i \in \{1, \dots, p\}$ and each $x \in V$, T_i contains exactly one leaf, u , such that $r(u) = x$ and at most one internal node, w , such that $r(w) = x$. For any two points $x, y \in V$, there exists some tree, T_i , with two leaves, u and v , such that $r(u) = x$, $r(v) = y$ and the path, u, \dots, v in T_i defines a path $r(u), \dots, r(v)$ whose Euclidean length is at most $t' \|xy\|$, where $t' > 1$ is a parameter in the construction of the dumbbell tree. Thus, the graph $G_{\infty} = (V, E_{\infty})$ obtained by taking

$$E_{\infty} = \bigcup_{i=1}^p \{r(u)r(v) : uv \text{ is an edge of } T_i\}$$

is a t' -spanner of V .




FIG. 3.1. A dumbbell tree decomposed in components of size $O(k')$ by the removal of a set X of $O(n/k')$ vertices (each denoted by \circ).

The size (number of edges) of a dumbbell tree spanner is clearly $O(pn) = O(n)$. For a fixed dimension, d , as a function of t and as t approaches 1, the number of trees, p , is $O(\log(1/(t - 1))/(t - 1)^d)$. In particular, for $t = 1 + \varepsilon$, $p \in O(\log(1/\varepsilon)/\varepsilon^d)$.

In the following, we will often treat the nodes of each tree, T_i , in a dumbbell tree decomposition as if the nodes are elements of V . This will happen, for example, when we make statements like “the path in T_i from the leaf containing x to the leaf containing y has length at most $t'\|xy\|$.” We do this to avoid the cumbersome phraseology required to distinguish between a node $u \in T_i$ and the node $r(u) \in V$ associated with u . Hopefully the reader can tolerate this informality.

THEOREM 5. *Let k_0 , f , and f^* be defined as in section 2.2 and let $d \geq 1$ and $t > 1$ be constants. Let $V \subset \mathbb{R}^d$ be any set of n points in \mathbb{R}^d . Then, for any constant $t > 1$, there exists an $O(kf(k))$ -robust t -spanner of V with $O(nf^*(n))$ edges.*

Proof. Fix a value $k' > 1$ and recall that in any binary tree, T , with n nodes, there exists a vertex whose removal disconnects T into at most three components, each of size at most $n/2$. Repeatedly applying this fact to any component of size greater than k' yields a set of $O(n/k')$ vertices whose removal disconnects T into components each of size at most k' [33, Lemma 12.1.5]; see Figure 3.1.

Perform the above decomposition for each of the trees T_1, \dots, T_p defining a dumbbell tree t' -spanner, G_∞ , of V with $t' = \sqrt{t}$. This yields a set, X , of $O(n/k')$ vertices whose removal disconnects every dumbbell tree into components each of size at most k' . Using any of the k' -fault-tolerant spanner constructions cited in the introduction, we can construct a k' -fault-tolerant t' -spanner for X having $O(k'|X|) = O(n)$ edges. Let $G_{k'} = (V, E_{k'})$ denote the graph whose edge set contains all edges of the dumbbell spanner G_∞ and all edges of a k' -fault-tolerant spanner on X .

Suppose that we are now given a set $S \subseteq V$, $|S| = k \leq k'$. Any vertex $x \in S$ appears at most twice in each tree T_i . For each $i \in \{1, \dots, p\}$, we say that x kills all the vertices in any component of $T_i \setminus X$ that contains x . Furthermore, if x is an element of X , then x kills all the vertices in the (at most three) components of T_i that have a vertex adjacent to x . The total number of vertices killed by x is therefore $O(pk') = O(k')$; see Figure 3.2.

Let S^+ be the set of all vertices killed by all vertices in S . The size of S^+ is $O(kk')$. Consider some pair of vertices $x, y \in V \setminus S^+$. There exists a tree T_i such that the path in T_i from the leaf containing x to the leaf containing y has length at most $t'\|xy\|$. If x and y are in the same component of $T_i \setminus X$, then this path is also a path in $G_{k'} \setminus S$.




FIG. 3.2. The set S (whose elements are denoted by \bullet) kills $O(|S|k)$ vertices in each dumbbell tree.

If x and y are in different components of $T_i \setminus X$, then consider the path from the leaf containing x to the leaf containing y in T_i . Let x' denote the first node on this path that is in X and let y' denote the last node on this path that is in X . The graph $G_{k'} \setminus S$ contains a path from the leaf containing x , to x' , to y' , and then finally to y , where the path from x' to y' uses the k' -fault tolerant spanner; see Figure 3.2. Therefore,

$$\begin{aligned} \|xy\|_{G \setminus S} &\leq \|xx'\|_{T_i} + t' \|x'y'\| + \|y'y\|_{T_i} \\ &\leq t'(\|xx'\|_{T_i} + \|x'y'\| + \|y'y\|_{T_i}) \\ &\leq t'(\|xx'\|_{T_i} + \|x'y'\|_{T_i} + \|y'y\|_{T_i}) \\ &= t' \|xy\|_{T_i} \\ &\leq (t')^2 \|xy\| \\ &= t \|xy\| . \end{aligned}$$

Since this is true for every pair $x, y \in V \setminus S^+$, this means that $G_{k'} \setminus S$ is a t -spanner of $V \setminus S^+$.

We have just shown how to construct a graph G_k that has $O(n)$ edges and is $O(kk')$ robust provided that $|S| \leq k'$. To obtain a graph that is $kf(k)$ -robust for any value of k , we take the graph G containing the edges of each $G_{k'}$ for $k' \in \{f^i(k_0) : i \in \{0, \dots, f^*(n)\}\}$. The graph G has $O(nf^*(n))$ edges. For any set $S \in \binom{V}{k}$, we can apply the above argument on the subgraph $G_{k'}$ with $k \leq k' < f(k)$ to show that G is $O(kf(k))$ -robust. \square

Applying Theorem 5 with different functions $f(k)$ yields the following results.

COROLLARY 3. For any constants $d > 1$, $t > 1$, $\varepsilon > 0$ and any set V of n points in \mathbb{R}^d , there exist $f(k)$ -robust t -spanners $G = (V, E)$ with

1. $f(k) \in O(k^2)$ and $O(n \log n)$ edges;
2. $f(k) \in O(k^2(1 + \varepsilon)^{\sqrt{\log k}})$ and $O(n\sqrt{\log n})$ edges; and
3. $f(k) \in O(k^{2+\varepsilon})$ and $O(n \log \log n)$ edges.

Remark. Note that, like our lower bounds, Theorem 5 and Corollary 3 do not express the relationship between the number of edges and the spanning ratio, t . As before, this relationship is not hard to work out. The number, p , of spanning trees in the dumbbell tree spanner is $O(\log(1/\varepsilon)/\varepsilon^d)$, where $\varepsilon = 1 - \sqrt{t}$. Each k' -fault-tolerant spanner has $O(n\varepsilon^{d-1})$ edges [32], and we construct one of these for $f^*(n)$ different values of k' . Thus, the total number of edges in our constructions is $O(n(f^*(n)/\varepsilon^{d-1} + \log(1/\varepsilon)/\varepsilon^d))$.

3.1. Linear-size (kind of) robust spanners. The lower bound in Theorem 4 shows that linear-size $f(k)$ -robust t -spanners do not exist for any function $f(k)$. In this section, we show that there are linear-sized graphs that satisfy a weaker definition of robustness.

We say that a graph $G = (V, E)$ is $f(k, n)$ -hardy if for every subset $S \subseteq V$, there exists a superset $S^+ \supseteq S$, $|S^+| \leq f(|S|, |V|)$, such that $G \setminus S$ is a t -spanner of $V \setminus S^+$. Note that this definition is almost identical to that of robustness except that the size of S^+ may also depend on $|V|$. In particular, any $f(k)$ -robust t -spanner is also an $f'(k, n)$ -hardy t -spanner with $f'(k, n) = f(k)$.

THEOREM 6. *If $f(k, n)$ -hardy t -spanners with $O(n \cdot s(n))$ edges exist for all $V \subset \mathbb{R}^d$, then $O(f(k, n) \cdot s(n))$ -hardy t -spanners with $O(n)$ edges exist for all $V \subset \mathbb{R}^d$.*

Proof. Perform the same dumbbell tree decomposition used in the proof of Theorem 5 to obtain a set X of $O(n/s(n))$ nodes whose removal partitions each dumbbell tree into components of size at most $s(n)$. Construct an $f(k, n)$ -hardy t -spanner on the elements of X . The size of the resulting graph is

$$\begin{aligned} O(n) + O(|X| \cdot s(|X|)) &= O(n) + O\left(\frac{n}{s(n)} \cdot s\left(\frac{n}{s(n)}\right)\right) \\ &\leq O(n) + O\left(\frac{n}{s(n)} \cdot s(n)\right) \\ &= O(n) . \end{aligned}$$

The same argument used to prove Theorem 5 shows that the resulting construction is $O(f(k, n)s(n))$ -hardy. (Each vertex of X that belongs to S results in the loss of at most three components in each dumbbell tree, each of size at most $s(n)$.) \square

The following corollary is obtained by combining Theorem 6 with some of our upper-bound constructions.

COROLLARY 4. *For any constant $\epsilon > 0$, there exist linear size*

1. $O(k \log k \log n)$ -hardy 1-spanners of any $V \subset \mathbb{R}$;
2. $O(k^{1+\epsilon} \log \log n)$ -hardy 1-spanners of any $V \subset \mathbb{R}$;
3. $O(k^2 \log n)$ -hardy t -spanners of any $V \subset \mathbb{R}^d$; and
4. $O(k^{2+\epsilon} \log \log n)$ -hardy t -spanners of any $V \subset \mathbb{R}^d$.

Remark. One can use the same argument used to prove Theorems 3 and 4 to study the hardness/space tradeoff in hardy spanners. For example, one can show that any $f(k)g(n)$ -hardy t -spanner of the $1 \times n$ grid has $\Omega((nf^*(n))/g(n))$ edges. This implies, for example, that part 2 of Corollary 4 is tight; it is not possible to asymptotically reduce the dependence on k or n while keeping a linear number of edges (apply the tradeoff result with $f(k) \in O(k^{1+\epsilon})$, and $f^*(n) = g(n) \in \Theta(\log \log n)$).

4. Summary. We have introduced the notion of $f(k)$ -robust t -spanners and given upper and lower-bounds on the number of edges in such spanners. Our lower bounds show that for any f , $f(k)$ -robust spanners sometimes require a superlinear number of edges, even in one dimension. Our one-dimensional constructions nearly match this lower-bound except when the function f is nearly linear.

Open problem: Tighter bounds. We understand the situation less clearly in two and higher dimensions. The lower bounds show that $f(k)$ -robust t -spanners must have $\Omega(nf^*(n))$ edges, but we have only been able to obtain $O(kf(k))$ -robust t -spanners with $O(nf^*(n))$ edges. Closing this gap is the main open problem left by this work.

To gain some intuition about which is closer to the truth, the lower bound or the upper bound, one can study the $\sqrt{n} \times \sqrt{n}$ grid graph; see Figure 1.1. An argument

similar to the proof of Theorem 1, based on randomly shifting a quadtree, shows that this graph is an $O(k^2)$ -robust 3-spanner. Therefore, the vertices of the $\sqrt{n} \times \sqrt{n}$ grid admit a linear-size $O(k^2)$ -robust 3-spanner. In contrast, Theorem 4 shows that any $f(k)$ -robust t -spanner for the $1 \times n$ grid has superlinear size. This suggests that one dimension is the hardest case.

CONJECTURE 1. *If $f(k)$ -robust t -spanners with $s_f(n)$ edges exist for all one-dimensional point sets, then $O(f(k))$ -robust t -spanners with $s_f(n)$ edges exist for all point sets in \mathbb{R}^d .*

Open problem: Low weight. In many cases, the cost of building a network is more closely related to the total length (rather than number) of edges. In these cases, one attempts to construct a graph whose total edge length is close to that of the minimum spanning tree of V . The same lower-bound argument used in Theorem 4 shows that in general, $f(k)$ -robust spanners may require edges whose total length is $\Omega(f^*(n))$ times that of the minimum spanning tree. Is there a (nearly) matching upper bound?

Open problem: $O(k)$ -robust spanners. Another fundamental open problem has to do with the number of edges needed in an $O(k)$ -robust t -spanner. We have no upper-bound better than the trivial $O(n^2)$ and the only lower-bound is $\Omega(n \log n)$. This is true even if we restrict our attention to constructing a t -spanner for the one-dimensional point set $V = \{1, \dots, n\}$.

Open problem: Induced spanners. Finally, we observe that the one-dimensional constructions of $f(k)$ -robust spanners actually satisfy a property that is slightly stronger than $f(k)$ -robustness: For each of these, the graph $G \setminus S^+$ is a t -spanner. In other words, vertices not in $V \setminus S^+$ are not needed in the short paths between pairs of vertices in $V \setminus S^+$. Our d -dimensional constructions do not have this stronger property. It would be interesting to know if d -dimensional constructions having this stronger property exist.

Acknowledgment. The research in this paper was started at the workshop on *Models of Sparse Graphs and Network Algorithms (12w5004)*, hosted at the Banff International Research Station (BIRS), February 5–10, 2012. The authors are grateful to the organizers, Nicolas Broutin, Luc Devroye, and Gábor Lugosi, the other participants, and the staff at BIRS for providing a stimulating research environment.

REFERENCES

- [1] M. A. ABAM AND M. DE BERG, *Kinetic spanners in \mathbb{R}^d* , *Discrete Comput. Geom.*, 45 (2011), pp. 723–736.
- [2] M. A. ABAM, M. DE BERG, M. FARSHI, AND J. GUDMUNDSSON, *Region-fault tolerant geometric spanners*, *Discrete Comput. Geom.*, 41 (2009), pp. 556–582.
- [3] M. A. ABAM, M. DE BERG, AND J. GUDMUNDSSON, *A simple and efficient kinetic spanner*, *Comput. Geom.*, 43 (2010), pp. 251–256.
- [4] A. K. ABU-AFFASH, R. ASCHNER, P. CARMI, AND M. J. KATZ, *Minimum power energy spanners in wireless ad hoc networks*, *Wireless Networks*, 17 (2011), pp. 1251–1258.
- [5] S. R. ARIKATI, D. Z. CHEN, L. P. CHEW, G. DAS, M. H. M. SMID, AND C. D. ZAROLIAGIS, *Planar spanners and approximate shortest path queries among obstacles in the plane*, in *ESA, J. Díaz and M. J. Serna, eds.*, *Lecture Notes in Comput. Sci.* 1136, Springer, New York, 1996, pp. 514–528.
- [6] B. ARONOV, M. DE BERG, O. CHEONG, J. GUDMUNDSSON, H. J. HAVERKORT, M. H. M. SMID, AND A. VIGNERON, *Sparse geometric graphs with small dilation*, *Comput. Geom.*, 40 (2008), pp. 207–219.
- [7] S. ARYA, G. DAS, D. M. MOUNT, J. S. SALOWE, AND M. H. M. SMID, *Euclidean spanners: Short, thin, and lanky*, in *Proceedings of the 27th Annual ACM Symposium on Theory of Computing*, F. T. Leighton and A. Borodin, eds., ACM, New York, 1995, pp. 489–498.

- [8] S. ARYA, D. M. MOUNT, AND M. H. M. SMID, *Randomized and deterministic algorithms for geometric spanners of small diameter*, in Proceedings of the 35th Annual Symposium on Foundations of Computer Science, IEEE Computer Society, 1994, pp. 703–712.
- [9] S. ARYA, D. M. MOUNT, AND M. H. M. SMID, *Dynamic algorithms for geometric spanners of small diameter: Randomized solutions*, Comput. Geom., 13 (1999), pp. 91–107.
- [10] S. ARYA AND M. H. M. SMID, *Efficient construction of a bounded-degree spanner with low weight*, Algorithmica, 17 (1997), pp. 33–54.
- [11] H. L. BODLAENDER, *A partial k -arboretum of graphs with bounded treewidth*, Theoret. Comput. Sci., 209 (1998), pp. 1–45.
- [12] H. L. BODLAENDER, *Treewidth: Structure and algorithms*, in SIROCCO, G. Prencipe and S. Zaks, eds., Lecture Notes in Comput. Sci. 4474, Springer, New York, 2007, pp. 11–25.
- [13] P. BOSE, P. CARMİ, M. COUTURE, A. MAHESHWARI, M. H. M. SMID, AND N. ZEH, *Geometric spanners with small chromatic number*, Comput. Geom., 42 (2009), pp. 134–146.
- [14] P. BOSE, P. CARMİ, M. FARSHI, A. MAHESHWARI, AND M. H. M. SMID, *Computing the greedy spanner in near-quadratic time*, Algorithmica, 58 (2010), pp. 711–729.
- [15] P. BOSE, R. FAGERBERG, A. VAN RENNSSEN, AND S. VERDONSCHOT, *On plane constrained bounded-degree spanners*, in LATIN, D. Fernández-Baca, ed., Lecture Notes in Comput. Sci. 7256, Springer, New York, 2012, pp. 85–96.
- [16] P. BOSE, J. GUDMUNDSSON, AND M. H. M. SMID, *Constructing plane spanners of bounded degree and low weight*, Algorithmica, 42 (2005), pp. 249–264.
- [17] P. BOSE, M. H. M. SMID, AND D. XU, *Delaunay and diamond triangulations contain spanners of bounded degree*, Internat. J. Comput. Geom. Appl., 19 (2009), pp. 119–140.
- [18] P. CARMİ AND L. CHAITMAN, *Bounded degree planar geometric spanners*, arXiv:1003.4963, 2010.
- [19] P. CARMİ AND L. CHAITMAN, *Stable roommates and geometric spanners*, in Proceedings of the 22nd Annual Canadian Conference on Computational Geometry, 2010, pp. 31–34.
- [20] P. CARMİ AND M. H. M. SMID, *An optimal algorithm for computing angle-constrained spanners*, in ISAAC (1), O. Cheong, K.-Y. Chwa, and K. Park, eds., Lecture Notes in Comput. Sci. 6506, Springer, New York, 2010, pp. 316–327.
- [21] P. CHEW, *There are planar graphs almost as good as the complete graph*, J. Comput. System Sci., 39 (1989), pp. 205–219.
- [22] A. CZUMAJ AND H. ZHAO, *Fault-tolerant geometric spanners*, Discrete Comput. Geom., 32 (2004), pp. 207–230.
- [23] G. DAS AND D. JOSEPH, *Which triangulations approximate the complete graph?*, in Optimal Algorithms, H. Djidjev, ed., Lecture Notes in Comput. Sci. 401, Springer, New York, 1989, pp. 168–192.
- [24] G. DAS AND G. NARASIMHAN, *A fast algorithm for constructing sparse Euclidean spanners*, Internat. J. Comput. Geom. Appl., 7 (1997), pp. 297–315.
- [25] D. EPPSTEIN, *Spanning trees and spanners*, in Handbook of Computational Geometry, J.-R. Sack and J. Urrutia, eds., Elsevier, New York, 1999, pp. 425–461.
- [26] J. GUDMUNDSSON, C. LEVCOPOULOS, AND G. NARASIMHAN, *Fast greedy algorithms for constructing sparse geometric spanners*, SIAM J. Comput., 31 (2002), pp. 1479–1500.
- [27] S. HOORY, N. LINIAL, AND A. WIGDERSON, *Expander graphs and their applications*, Bull. Amer. Math. Soc., 43 (2006), pp. 439–561.
- [28] G. KALAI, *The diameter of graphs of convex polytopes and f -vector theory*, in Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift, DIMACS Ser. Discrete Math. Theoret. Comput. Sci. 4, AMS, Providence, RI, 1991, pp. 387–412.
- [29] J. M. KEIL AND C. A. GUTWIN, *The Delauney triangulation closely approximates the complete Euclidean graph*, in WADS, F. K. H. A. Dehne, J.-R. Sack, and N. Santoro, eds., Lecture Notes in Comput. Sci. 382, Springer, New York, 1989, pp. 47–56.
- [30] T. KLOKS, *Treewidth: Computations and Approximations*, Lecture Notes in Comput. Sci. 842, Springer-Verlag, New York, 1994.
- [31] C. LEVCOPOULOS, G. NARASIMHAN, AND M. H. M. SMID, *Improved algorithms for constructing fault-tolerant spanners*, Algorithmica, 32 (2002), pp. 144–156.
- [32] T. LUKOVSKI, *New results of fault tolerant geometric spanners*, in WADS, F. K. H. A. Dehne, A. Gupta, J.-R. Sack, and R. Tamassia, eds., Lecture Notes in Comput. Sci. 1663, Springer, New York, 1999, pp. 193–204.
- [33] G. NARASIMHAN AND M. H. M. SMID, *Geometric Spanner Networks*, Cambridge University Press, Cambridge, 2007.
- [34] M. SEGAL AND H. SHPUNGIN, *Improved multi-criteria spanners for ad-hoc networks under energy and distance metrics*, in INFOCOM, IEEE, 2010, pp. 6–10.

- [35] M. SMID, *Notes on Binary Dumbbell Trees*, Carleton University, Ottawa, Ontario, Canada, 2012.
- [36] M. H. M. SMID, *Geometric spanners with few edges and degree five*, in CATS, J. Gudmundsson and C. B. Jay, eds., CRPIT 51, Australian Computer Society, 2006, pp. 7–9.
- [37] Y. WANG AND X.-Y. LI, *Minimum power assignment in wireless ad hoc networks with spanner property*, J. Comb. Optim., 11 (2006), pp. 99–112.