

New Results on Fault Tolerant Geometric Spanners*

Tamás Lukovszki

Heinz Nixdorf Institute,
University of Paderborn, D-33102 Paderborn, Germany
tamas@hni.uni-paderborn.de

Abstract. We investigate the problem of constructing spanners for a given set of points that are tolerant for edge/vertex faults. Let $S \subset \mathbb{R}^d$ be a set of n points and let k be an integer number. A k -edge/vertex fault tolerant spanner for S has the property that after the deletion of k arbitrary edges/vertices each pair of points in the remaining graph is still connected by a short path.

Recently it was shown that for each set S of n points there exists a k -edge/vertex fault tolerant spanner with $O(k^2n)$ edges which can be constructed in $O(n \log n + k^2n)$ time. Furthermore, it was shown that for each set S of n points there exists a k -edge/vertex fault tolerant spanner whose degree is bounded by $O(c^{k+1})$ for some constant c .

Our first contribution is a construction of a k -vertex fault tolerant spanner with $O(kn)$ edges which is a tight bound. The computation takes $O(n \log^{d-1} n + kn \log \log n)$ time. Then we show that the same k -vertex fault tolerant spanner is also k -edge fault tolerant. Thereafter, we construct a k -vertex fault tolerant spanner with $O(k^2n)$ edges whose degree is bounded by $O(k^2)$. Finally, we give a more natural but stronger definition of k -edge fault tolerance which not necessarily can be satisfied if one allows only simple edges between the points of S . We investigate the question whether Steiner points help. We answer this question affirmatively and prove $\Theta(kn)$ bounds on the number of Steiner points and on the number of edges in such spanners.

1 Introduction

Geometric spanners have many applications in various areas of the computer science. They have been studied intensively in recent years. Let S be a set of n points in \mathbb{R}^d , where d is an integer constant. Let $G = (S, E)$ be a graph whose edges are straight line segments between the points of S . For two points $p, q \in \mathbb{R}^d$, let $dist_2(p, q)$ be the Euclidean distance between p and q . The length $length(e)$ of an edge $e = (a, b) \in E$ is defined as $dist_2(a, b)$. For a path P in G the length $length(P)$ is defined as the sum of the length of the edges of P . A path between two points $p, q \in S$ is called a pq -path. Let $t > 1$ be a real number. The graph G is a t -spanner for S if for each pair of points $p, q \in S$ there is a pq -path in G such that the length of the path is at most t times the Euclidean distance $dist_2(p, q)$ between p and q . We call such a path a t -spanner path and t is called the *stretch factor* of the spanner. If G is a directed graph and G contains a

* Partially supported by EU ESPRIT Long Term Research Project 20244 (ALCOM-IT) and DFG Graduiertenkolleg "Parallele Rechnernetzwerke in der Produktionstechnik" Me872/4-1.

directed t -spanner path between each pair of points then G is called a directed t -spanner. In order to distinguish the edges of a directed from an undirected graph we use $\langle a, b \rangle$ to denote an edge between the vertices a and b in a directed and (a, b) in an undirected graph.

Spanners were introduced by Chew [6]. They have applications in motion planning [7], they were used for approximating the minimum spanning tree [17], to solve a special searching problem which appears in walkthrough systems [8], and to a polynomial time approximation scheme for the traveling salesman and related problems [13].

The problem of constructing a t -spanner for a real constant $t > 1$, that has $O(n)$ edges, has been investigated by many researchers. Keil [10] gave a solution for this problem introducing the θ -graph¹, which was generalized by Ruppert and Seidel [14] and Arya et al. [2] to any fixed dimension d . These authors gave also an $O(n \log^{d-1} n)$ time algorithm to construct the θ -graph. Chen et al. [5] proved that the problem of constructing any t -spanner for $t > 1$ takes $\Omega(n \log n)$ time in the algebraic computation tree model [3]. Callahan and Kosaraju [4], Salowe [15] and Vaidya [16] gave optimal $O(n \log n)$ time algorithm for constructing t -spanners. Several interesting quantities related to spanners were studied by Arya et al. [1]. They gave constructions for bounded degree spanners, spanners with low weight, spanners with low diameter, and for spanners having more than one of these properties. The weight $w(G)$ of a graph G is the sum of the length of its edges.

Fault tolerant spanners were introduced by Levcopoulos et al. [12]. For the formal definition we need the following notions. For a set $S \subset \mathbb{R}^d$ of n points let K_S denote the complete Euclidean graph with vertex set S . If $G = (S, E)$ is a graph and $E' \subseteq E$ then $G \setminus E'$ denotes the graph $G' = (S, E \setminus E')$. Similarly, if $S' \subseteq S$ then the graph $G \setminus S'$ is the graph with vertex set $S \setminus S'$ and edge set $\{(p, q) \in E : p, q \in S \setminus S'\}$. Let $t > 1$ be a real number and k be an integer, $1 \leq k \leq n - 2$.

- A graph $G = (S, E)$ is called a k -edge fault tolerant t -spanner for S , or (k, t) -EFTS, if for each $E' \subseteq E$, $|E'| \leq k$, and for each pair p, q of points of S , the graph $G \setminus E'$ contains a pq -path whose length is at most t times the length of a shortest pq -path in the graph $K_S \setminus E'$.
- Similarly, $G = (S, E)$ is called a k -vertex fault tolerant t -spanner for S , or (k, t) -VFTS, if for each subset $S' \subseteq S$, $|S'| \leq k$, the graph $G \setminus S'$ is a t -spanner for $S \setminus S'$.

Levcopoulos et al. [12] presented an algorithm with running time $O(n \log n + k^2 n)$ which constructs a (k, t) -EFTS/VFTS with $O(k^2 n)$ edges for any real constant $t > 1$. The constants hidden in the O -notation are $(\frac{d}{t-1})^{O(d)}$ if $t \searrow 1$. They also showed that $\Omega(kn)$ is a lower bound for the number of edges in such spanners. This follows from the obvious fact that each k -edge/vertex fault tolerant spanner must be k -edge/vertex connected. Furthermore, they gave another algorithm with running time $O(n \log n + c^{k+1} n)$, for some constant c , which constructs a (k, t) -VFTS whose degree is bounded by $O(c^{k+1})$ and whose weight is bounded by $O(c^{k+1} w(MST))$.

¹ Yao [17] and Clarkson [7] used a similar construction to solve other problems.

1.1 New results

We consider directed and undirected fault tolerant spanners. Our first contribution is a construction of a (k, t) -VFST with $O(kn)$ edges in $O(n \log^{d-1} n + kn \log \log n)$ time. Then we show that the same k -vertex fault tolerant spanner is also a k -edge fault tolerant spanner. Our bounds for the number of edges in fault tolerant spanners are optimal up to a constant factor and they improve the previous $O(k^2 n)$ bounds significantly. Furthermore, we construct a k -vertex fault tolerant spanner with $O(k^2 n)$ edges whose degree is bounded by $O(k^2)$ which also improves the previous $O(c^{k+1})$ bound.

Then we study Steinerized fault tolerant spanners that are motivated by the following. In the definition of (k, t) -EFTS we only require that after deletion of k arbitrary edges E' in the remaining graph each pair of points p, q is still connected by a path whose length is at most t times the length of the shortest pq -path in $K_S \setminus E'$. Such a path can be arbitrarily long, much longer than $\text{dist}_2(p, q)$. To see this consider the following example. Let $r > 1$ be an arbitrarily large real number. Let $p, q \in S$ be two points such that $\text{dist}_2(p, q) = 1$ and let the remaining $n - 2$ points of S be placed on the ellipsoid $\{x \in \mathbb{R}^d : \text{dist}_2(p, x) + \text{dist}_2(q, x) = r \cdot t\}$. Clearly, each t -spanner G for S contains the edge between p and q , because each path which contains any third point $s \in S \setminus \{p, q\}$ has a length at least $r \cdot t$. Therefore, if the edge $(p, q) \in E'$ then the graph $G \setminus E'$ can not be a t -spanner for S . However, $G \setminus E'$ can contain a path satisfying the definition of the k -edge fault tolerance. In some applications one would need a stronger property. After deletion of k edges a pq -path would be desirable whose length is at most t times $\text{dist}_2(p, q)$. In order to solve this problem we extend the original point set S by Steiner points. Then we investigate the question how many Steiner points and how many edges do we need to satisfy the following natural but stronger condition of edge fault tolerance. Let $t > 1$ be a real number and $k \in \mathbb{N}$.

- The graph $G = (V, E)$ with $S \subseteq V$ is called a *k -edge fault tolerant Steiner t -spanner* for S , or (k, t) -EFTSS, if for each $E' \subset E$, $|E'| \leq k$ and for each two points $p, q \in S$, there is a pq -path P in $G \setminus E'$ such that $\text{length}(P) \leq t \cdot \text{dist}_2(p, q)$.
- Similarly, $G = (V, E)$ with $S \subseteq V$ is a *k -vertex fault tolerant Steiner t -spanner* for S , or (k, t) -VFSTSS, if for each $V' \subset V$, $|V'| \leq k$ and for each two points $p, q \in S \setminus V'$, there is a pq -path P in $G \setminus V'$ such that $\text{length}(P) \leq t \cdot \text{dist}_2(p, q)$.

To our knowledge, fault tolerant Steiner spanners have not been investigated before. First we show that for each set S of n points, $t > 1$ real constant, and $k \in \mathbb{N}$, a (k, t) -EFTSS/VFSTSS for S can be constructed which contains $O(kn)$ edges and $O(kn)$ Steiner points. Then we show that there is a set S of n points in \mathbb{R}^d , $d \geq 1$, such that for each $t > 1$ and $k \in \mathbb{N}$, each (k, t) -EFTSS for S contains $\Omega(kn)$ edges and $\Omega(kn)$ Steiner points. In this paper we assume that the dimension d is a constant.

2 A k -vertex fault tolerant t -spanner with $O(kn)$ edges

The construction of a k -vertex fault tolerant t -spanner with $O(kn)$ edges is based on a generalization of the θ -graph [17, 7, 10, 11, 14, 12]. First we introduce the notion of *i th order θ -graph* of the point set S , for $1 \leq i \leq n - 1$. Then we prove that for appropriate θ , the $(k + 1)$ th order θ -graph is a (k, t) -VFST for the given set of points.

2.1 The i th order θ -graph

For the formal description we need the notion of simplicial cones. We assume that the points of \mathbb{R}^d are represented by coordinate vectors. Let p_0, p_1, \dots, p_d be points in \mathbb{R}^d such that the vectors $(p_i - p_0)$, $1 \leq i \leq d$, are linearly independent. Then the set $\{p_0 + \sum_{i=1}^d \lambda_i (p_i - p_0) : \lambda_i \geq 0 \text{ for all } i\}$ is called a *simplicial cone* and p_0 is called the *apex* of the cone (see, e.g., in [9]). Let θ be a fixed angle $0 < \theta \leq \pi$ and C be a collection of simplicial cones such that

1. each cone $c \in C$ has its apex at the origin,
2. $\bigcup_{c \in C} c = \mathbb{R}^d$,
3. for each cone $c \in C$ there is a fixed halfline l_c having the endpoint at the origin such that for each halfline l , which has the endpoint at the origin and is contained in c , the angle between l_c and l is at most $\theta/2$.

We call such a collection C of simplicial cones a θ -*frame*². Yao [17] and Ruppert and Seidel [14] showed methods how a θ -frame C of $(\frac{d}{\theta})^{O(d)}$ cones can be constructed. Assuming that the dimension d and the angle θ are constant we obtain a constant number of cones. In the following, the number of cones in C is denoted by $|C|$.

Let $0 < \theta \leq \pi$ be an angle and C be a corresponding θ -frame. For a simplicial cone $c \in C$ and for a point $p \in \mathbb{R}^d$, let $c(p)$ be the translated cone $\{x + p : x \in c\}$ and let $l_c(p)$ be the translated cone axis $\{x + p : x \in l_c\}$. For $c \in C$ and $p, q \in \mathbb{R}^d$ such that $q \in c(p)$, let $dist_c(p, q)$ denote the Euclidean distance between p and the orthogonal projection of q to $l_c(p)$.

Now we define the *i th order θ -graph* $G_{\theta,i}(S)$ for a set S of n points in \mathbb{R}^d and for an integer $1 \leq i \leq n - 1$ as follows. For each point $p \in S$ and each cone $c \in C$, let $S_{c(p)} := c(p) \cap S \setminus \{p\}$, i.e., $S_{c(p)}$ is the set of points of $S \setminus \{p\}$ that are contained in the cone $c(p)$. For any integer i , $1 \leq i \leq n - 1$, let $N_{i,c}(p) \subseteq S_{c(p)}$ be the set of the $\min(i, |S_{c(p)}|)$ -nearest neighbors of p in the cone $c(p)$ w.r.t. the distance $dist_c$, i.e., for each $q \in N_{i,c}(p)$ and $q' \in S_{c(p)} \setminus N_{i,c}(p)$ holds that $dist_c(p, q) \leq dist_c(p, q')$. Let $G_{\theta,i}(S)$ be the directed graph with vertex set S such that for each point $p \in S$ and each cone $c \in C$ there is a directed edge $\langle p, q \rangle$ to each point $q \in N_{i,c}(p)$.

2.2 The vertex fault tolerant spanner property

In [14] it is proved that for $0 < \theta < \pi/3$, the graph $G_{\theta,1}(S)$ is a spanner for S with stretch factor $t \leq \frac{1}{1-2\sin(\theta/2)}$. The proof is based on the following lemma which will be also crucial to show the fault tolerant spanner property of $G_{\theta,i}(S)$ for $i > 1$.

Lemma 1. [14] *Let $0 < \theta < \pi/3$. Let $p \in \mathbb{R}^d$ be a point and $c \in C$ be a cone. Furthermore, let q and r be two points in $c(p)$ such that $dist_c(p, r) \leq dist_c(p, q)$. Then $dist_2(r, q) \leq dist_2(p, q) - (1 - 2\sin(\theta/2)) dist_2(p, r)$.*

Theorem 1. *Let $S \subset \mathbb{R}^d$ be a set of n points. Let $0 < \theta < \pi/3$ and $1 \leq k \leq n - 2$ be an integer number. Then the graph $G_{\theta,k+1}(S)$ is a directed $(k, \frac{1}{1-2\sin(\theta/2)})$ -VFTS for S . $G_{\theta,k+1}(S)$ contains $O(|C|kn)$ edges and it can be constructed in $O(|C|(n \log^{d-1} n + kn \log \log n))$ time.*

² The notion of θ -frame was introduced by Yao [17]. We use a slightly modified θ -frame definition which is suggested by Ruppert and Seidel [14].

Proof. Let $S' \subset S$ be a set of at most k points. We show that for each two points $p, q \in S \setminus S'$ there is a (directed) pq -path P in $G_{\theta, k+1}(S) \setminus S'$ such that the length of P is at most $\frac{1}{1-2\sin(\theta/2)} \text{dist}_2(p, q)$. The proof is similar to the proof of Ruppert and Seidel [14]. Consider the path constructed in the following way. Let $p_0 := p, i := 0$ and let P contain the single point p_0 . If the edge $\langle p_i, q \rangle$ is present in the graph $G_{\theta, k+1}(S) \setminus S'$ then add the vertex q to P and stop. Otherwise, let $c(p_i)$ be the cone which contains q . Choose an arbitrary point $p_{i+1} \in N_{k+1, c}(p_i)$ as the next vertex of the path P and repeat the procedure with p_{i+1} .

Consider the i th iteration of the above algorithm. If $\langle p_i, q \rangle \in G_{\theta, k+1}(S)$ then the algorithm terminates. Otherwise, if $\langle p_i, q \rangle \notin G_{\theta, k+1}(S)$ then by definition the cone $c(p_i)$ contains at least $k+1$ points that are not further from p_i than q w.r.t. the distance dist_c . Hence, in the graph $G_{\theta, k+1}(S)$ the point p_i has $k+1$ neighbors in $c(p_i)$ and, therefore, in the graph $G_{\theta, k+1}(S) \setminus S'$ it has at least one neighbor in $c(p_i)$. Consequently, the algorithm is well defined in each step. Furthermore, Lemma 1 implies that $\text{dist}_2(p_{i+1}, q) < \text{dist}_2(p_i, q)$ and hence, each point is contained in P at most once. Therefore, the algorithm terminates and finds a pq -path P in $G_{\theta, k+1}(S) \setminus S'$. The bound on the length of P follows by applying Lemma 1 iteratively in the same way as in [14]: Let p_0, \dots, p_m be the vertices on P , $p_0 = p$ and $p_m = q$. Then

$$\sum_{0 \leq i < m} \text{dist}_2(p_{i+1}, q) \leq \sum_{0 \leq i < m} \left(\text{dist}_2(p_i, q) - (1 - 2\sin(\theta/2)) \text{dist}_2(p_i, p_{i+1}) \right).$$

Rearranging the sum we get

$$\begin{aligned} \sum_{0 \leq i < m} \text{dist}_2(p_i, p_{i+1}) &\leq \frac{1}{1-2\sin(\theta/2)} \sum_{0 \leq i < m} \left(\text{dist}_2(p_i, q) - \text{dist}_2(p_{i+1}, q) \right) \\ &= \frac{1}{1-2\sin(\theta/2)} \text{dist}_2(p_0, q). \end{aligned}$$

Hence, the graph $G_{\theta, k+1}(S)$ is a $(k, \frac{1}{1-2\sin(\theta/2)})$ -VFTS for S . Clearly, it contains $O(|C|kn)$ edges, where $|C| = (d/\theta)^{O(d)}$. It can be constructed in $O(|C|(n \log^{d-1} n + kn \log \log n))$ time using the algorithm of Levcopoulos et al. [12]. They compute for each point $p \in S$ and each cone $c \in C$ the set $N_{k, c}(p)$ in order to determine so-called strong approximated neighbors. \square

Corollary 1. *Let S be a set of n points in \mathbb{R}^d , $t > 1$ a real constant, and k an integer, $1 \leq k \leq n - 2$. Then there is a (k, t) -VFTS for S with $O(kn)$ edges. Such a spanner can be constructed in $O(n \log^{d-1} n + kn \log \log n)$ time.*

Proof. We set θ such that $t \geq \frac{1}{1-2\sin(\theta/2)}$ and $0 < \theta < \pi/3$ and construct $G_{\theta, k+1}(S)$. If $t \searrow 1$ then the constant factors hidden in the O -calculus are $(\frac{d}{t-1})^{O(d)}$.

3 k -edge fault tolerant t -spanners

Levcopoulos et al. [12] claimed that any (k, t) -VFTS is also a (k, t) -EFTS. We give our own proof of this fact. The proof is simple and holds also for directed spanners.

Theorem 2. *Let S be a set of n points in \mathbb{R}^d , $t > 1$ a real constant, and k an integer, $1 \leq k \leq n - 2$. Then every (directed) (k, t) -VFTS for S is also a (directed) (k, t) -EFTS for S .*

Proof. Let $G = (S, E)$ be a (directed) (k, t) -VFTS for S . Let $E' \subset E$ be a set of at most k edges. Consider two arbitrary points $p, q \in S$. Let P^* be the shortest (directed) pq -path in $K_S \setminus E'$. Such a path exists, since the set of pq -paths in $K_S \setminus E'$ is not empty. It contains, for example, at least one of the $n - 2$ paths in K_S of two edges $P_s = p, s, q$, for $s \in S \setminus \{p, q\}$, or the immediate path $P_0 = p, q$, because at least one of them is distinct from E' .

We have to show that there is a (directed) pq -path P in $G \setminus E'$ such that the length of P is at most t times the length of P^* . The edges e in P^* that are contained in G will also be contained in P . Consider an edge (u, v) ($\langle u, v \rangle$ in the directed case) in P^* which is not contained in G . We show that this edge can be substituted by a uv -path P_{uv} in $G \setminus E'$ such that $\text{length}(P_{uv}) \leq t \cdot \text{dist}_2(u, v)$: For each edge $e' \in E'$ (for each $e' \in E' \setminus \{\langle v, u \rangle\}$ in the directed case) we fix one of its endpoints $p_{e'}$ such that $p_{e'} \in S \setminus \{u, v\}$. Let $S'_{uv} := \{p_{e'} : e' \in E'\}$ ($S'_{uv} := \{p_{e'} : e' \in E' \setminus \{\langle v, u \rangle\}\}$ in the directed case). Note that $|S'_{uv}| \leq |E'| \leq k$. Since G is a (directed) (k, t) -VFTS for S , there is a (directed) uv -path P_{uv} in $G \setminus S'_{uv}$ such that P_{uv} does not contain any edge of E' and $\text{length}(P_{uv}) \leq t \cdot \text{dist}_2(u, v)$. The desired pq -path P is composed of the edges of $P^* \cap G$ and the uv -paths for the edges $(u, v) \in P^* \setminus G$ ($\langle u, v \rangle \in P^* \setminus G$ in the directed case). Clearly, $\text{length}(P) \leq t \cdot \text{length}(P^*)$. \square

This, together with Corollary 1, leads to

Corollary 2. *Let S be a set of n points in \mathbb{R}^d , $t > 1$ a real constant, and k an integer, $1 \leq k \leq n - 2$. Then there is a (directed) (k, t) -EFTS for S with $O(kn)$ edges. Such a spanner can be constructed in $O(n \log^{d-1} n + kn \log \log n)$ time.*

The proof of Theorem 2 implies also the following for directed graphs.

Theorem 3. *Let S be a set of n points in \mathbb{R}^d , $t > 1$ a real constant, and k an integer, $1 \leq k \leq n - 2$. Let $G = (V, E)$ be a directed (k, t) -VFTS for S . Let $E' \subset E$ be a set of at most k edges and let $E'' := \{\langle v, u \rangle : (u, v) \in E'\}$. Then for each two points $p, q \in S$ the graph $G \setminus (E' \cup E'')$ contains a pq -path P such that the length of P is at most t times the length of the shortest pq -path in $K_S \setminus (E' \cup E'')$.*

4 A k -vertex fault tolerant t -spanners with degree $O(k^2)$

We now turn to the problem of constructing fault tolerant spanners with bounded degree. We proceed similar to the method in [1] which constructs a spanner with constant degree. However, we must take much more care, because of the fault tolerant property and the goal of keeping the number of edges small. We have shown that for any real constant $t > 1$ we can construct a directed (k, t) -VFTS/EFTS for S whose outdegree is $O(k)$. In this section we give a method to construct a (k, t) -VFTS whose degree is $O(k^2)$ from a directed $(k, t^{1/3})$ -VFTS whose outdegree is $O(k)$.

In order to show this construction we need the notion of k -vertex fault tolerant single sink spanner. This is a generalization of *single sink spanners* introduced in [1]. Let V be

a set of m points in \mathbb{R}^d , $v \in V$, $\hat{t} > 1$ a real constant, and k an integer, $1 \leq k \leq m - 2$. A directed graph $G = (V, E)$ is a k -vertex fault tolerant v -single sink \hat{t} -spanner, or (k, \hat{t}, v) -VFTssS for V if for each $u \in V \setminus \{v\}$ and each $V' \subseteq V \setminus \{v, u\}$, $|V'| \leq k$, there is an \hat{t} -spanner path in $G \setminus V'$ from u to v .

Now let V be a set of m points in \mathbb{R}^d , $v \in V$ a fixed point, $1 \leq i \leq m - 1$ an integer, θ an angle, $0 < \theta < \pi/3$, and C a θ -frame. We define a directed graph $\hat{G}_{v, \theta, i}(V) = (V, E)$ whose edges are directed straight line segments between points of V as follows. First we partition the set V in clusters such that each cluster contains at most i points. Then we build a tree-like structure based on these clusters. For the clustering we use the cones of C . Now we describe this procedure more precisely.

First we create a cluster $cl(\{v\})$ containing the unique point v . For each cluster that we create, we choose a point as the *representative* of the cluster. The representative of $cl(\{v\})$ is v . The clustering of the set $V \setminus \{v\}$ is recursive. The recursion stops if $V \setminus \{v\}$ is the empty set. Otherwise, we do the following. For each cone $c \in C$ let $V_{c(v)}$ be the set of points of $V \setminus \{v\}$ contained in c . If a point is contained in more than one cone then assign the point only to one of them. If one cone, say c , contains more than $m/2$ points, then partition the points of $V_{c(v)}$ arbitrarily into two sets $V_{c(v)}^1$ and $V_{c(v)}^2$ both having at most $m/2$ points. For each nonempty set $V_{c(v)}$, $c \in C$ (or in the case if $V_{c(v)}$ had to be partitioned, for each $V_{c(v)}^1$ and $V_{c(v)}^2$), let $N_{i,c}(v) \subseteq V_{c(v)}$ be the set of the $\min(i, |V_{c(v)}|)$ -nearest neighbors of v in $V_{c(v)}$ w.r.t. the distance $dist_c$. The points contained in the same $N_{i,c}(v)$ define a new cluster $cl(N_{i,c}(v))$. Note that in this way we obtain at most $|C| + 1$ new clusters. We say that these clusters are the *children* of $cl(\{v\})$ and $cl(\{v\})$ is the *parent* of these clusters. For each new cluster $cl(N_{i,c}(v))$ we choose a representative $u_c \in N_{i,c}(v)$ such that $dist_c(v, u_c) = \max\{dist_c(v, u) : u \in N_{i,c}(v)\}$. Then, for each set $V_{c(v)}$, $c \in C$ (and $V_{c(v)}^1$, $V_{c(v)}^2$ if exist), we recursively cluster $V_{c(v)} \setminus N_{i,c}(v)$ using the cones around u_c .

After the clustering is done, for each cluster $cl \neq cl(\{v\})$ we add an edge in $\hat{G}_{v, \theta, i}(V)$ from each point $u \in cl$ to each point w of the parent cluster of cl . Figure 1 shows an example for $\hat{G}_{v, \theta, i}(V)$. The dotted lines represent the boundaries of the cones at the representatives of the clusters.




Fig. 1. The directed graph $\hat{G}_{v, \theta, 3}(V)$ for a point set V in \mathbb{R}^2 .

Lemma 2. Let V be a set of m points in \mathbb{R}^d , $v \in V$ a fixed point and $1 \leq k \leq m - 2$ an integer number. Let $0 < \theta < \pi/3$ be an angle and C be a θ -frame. Then the graph $\hat{G}_{v, \theta, k+1}(V)$ is a $(k, (\frac{1}{1-2\sin(\theta/2)})^2, v)$ -VFTssS for V . Its degree is bounded by $O(|C|k)$ and it can be computed in $O(|C|(m \log m + km))$ time.

Proof. For each point $u \in V$ let $cl(u)$ denote the cluster containing it. The outdegree of each point $u \in V \setminus \{v\}$ in $\hat{G}_{v,\theta,k+1}(V)$ is bounded by $k+1$, because each point u has only edges to the points contained in the parent cluster of $cl(u)$ and the number of points in each cluster is bounded by $k+1$. (Each internal cluster – i.e., a cluster which is different from $cl(v)$ and has at least one child – contains exactly $k+1$ points). Since each cluster has at most $|C|+1$ children, the indegree of the points is bounded by $(|C|+1)(k+1)$. The bound for the construction time follows from the fact that the recursion has depth $O(\log m)$.

Now we prove the fault tolerant single sink spanner property. Consider an arbitrary point $u \in V \setminus \{v\}$. Let $P_0 := u_0, \dots, u_l$, $u_0 = u$ and $u_l = v$, be the unique path from u to v in $\hat{G}_{v,\theta,k+1}(V)$ such that each internal vertex u_i , $1 \leq i < l$, is the representative of a cluster. Note that $l = O(\log m)$. The length of P_0 is at most $\frac{1}{1-2\sin(\theta/2)} \text{dist}_2(u, v)$. If the edge $\langle u, v \rangle \in \hat{G}_{v,\theta,k+1}(V)$, this claim holds trivially, otherwise, it follows by applying Lemma 1 iteratively for the triples u_{i+1}, u_i, u_{i-1} , $i = 1, \dots, l-1$, in the same way as in the proof of Theorem 1.

Now let $V' \subset V \setminus \{u, v\}$ be a set of at most k points. We show that there is a uv -path P in $\hat{G}_{v,\theta,k+1}(V) \setminus V'$ such that $\text{length}(P) \leq \frac{1}{1-2\sin(\theta/2)} \text{length}(P_0)$. This will imply the desired stretch factor $(\frac{1}{1-2\sin(\theta/2)})^2$. Let P be the path constructed as follows. Let $v_0 := u$, $i := 0$ and let P contain the single point v_0 . If $v_i = v$ then stop. Otherwise, let v_{i+1} be an arbitrary point with $\langle v_i, v_{i+1} \rangle \in \hat{G}_{v,\theta,k+1}(V) \setminus V'$. Add the vertex v_{i+1} to P and repeat the procedure with v_{i+1} .




Fig. 2. The paths $P_0 := u_0, \dots, u_l$ and $P := v_0, \dots, v_l$. The dotted lines show the cone boundaries.

The above algorithm is well defined in each step. To see this, consider the i th iteration. If the cluster $cl(v)$ is the parent of $cl(v_i)$ then the algorithm chooses v as v_{i+1} and terminates. Otherwise, the parent of $cl(v_i)$ contains $k+1$ points and, hence, at least one point disjoint from V' . The algorithm chooses such a point as v_{i+1} . Clearly, the algorithm terminates after $l = O(\log m)$ steps and constructs a uv -path $P = v_0, \dots, v_l$ (Figure 2) with

$$\begin{aligned}
 \text{length}(P) &= \sum_{0 \leq i < l} \text{dist}_2(v_i, v_{i+1}) \\
 &\leq \sum_{0 \leq i < l} \left(\text{dist}_2(v_i, u_{i+1}) + \text{dist}_2(u_{i+1}, v_{i+1}) \right) \quad (1) \\
 &= \sum_{0 \leq i < l} \left(\text{dist}_2(v_i, u_{i+1}) + \text{dist}_2(u_i, v_i) \right) + \underbrace{\text{dist}_2(u_l, v_l)}_{=0} - \underbrace{\text{dist}_2(u_0, v_0)}_{=0} \\
 &\leq \sum_{0 \leq i < l} \frac{1}{1-2\sin(\theta/2)} \text{dist}_2(u_i, u_{i+1}) \quad (2) \\
 &= \frac{1}{1-2\sin(\theta/2)} \text{length}(P_0).
 \end{aligned}$$

(1) holds because of the triangle inequality and (2) follows by applying Lemma 1 for the triples u_{i+1}, v_i, u_i , $i = 0, \dots, l-1$. Hence, the claimed stretch factor of $\hat{G}_{v, \theta, k+1}(V)$ follows. \square

Theorem 4. *Let S be a set of n points in \mathbb{R}^d , $t > 1$ a real constant, and k an integer, $1 \leq k \leq n-2$. Then there is a (k, t) -VFTS G for S whose degree is bounded by $O(k^2)$. The total number of edges in G is $O(k^2 n)$ and G can be constructed in $O(n \log^{d-1} n + kn \log n + k^2 n)$ time.*

Proof. Let G_0 be a directed $(k, t^{1/3})$ -VFTS for S whose outdegree is $O(k)$, for example, let G_0 be the $(k+1)$ th order θ -graph $G_{\theta, k+1}(S)$ with $t^{1/3} \geq \frac{1}{1-2\sin(\theta/2)}$. For each point $p \in S$ let $N_{in}(p) := \{q \in S : \langle q, p \rangle \in G_0\}$. Let G be the directed graph with vertex set S which is created such that for each $p \in S$ we construct the graph $\hat{G}_{p, \theta, k+1}(N_{in}(p) \cup \{p\})$ and we add the edges of $\hat{G}_{p, \theta, k+1}(N_{in}(p) \cup \{p\})$ to G .

We can bound the degree of G as follows. For each $q \in S$, the graph G contains the edges of $\hat{G}_{q, \theta, k+1}(N_{in}(q) \cup \{q\})$. In this VFTSS each vertex p has an in- and outdegree $O(k)$. Now for each $p \in S$, we have to count the graphs $\hat{G}_{q, \theta, k+1}(N_{in}(q) \cup \{q\})$, $q \in S$, that contain p . Clearly, the number of such graphs is equal to one plus the outdegree of p in G_0 , which is $O(k)$. Therefore, the degree of each $p \in S$ in G is $O(k^2)$.

Now we show that G is a (k, t) -VFTS for S . Let $S' \subset S$, $|S'| \leq k$. Consider two arbitrary points $p, q \in S \setminus S'$. Since G_0 is a $(k, t^{1/3})$ -VFTS for S , there is an $t^{1/3}$ -spanner path P_0 in $G_0 \setminus S'$ between p and q . Furthermore, for each edge $\langle u, v \rangle \in P_0$, there is a $t^{2/3}$ -spanner path P_{uv} in $G \setminus S'$, because G contains all edges of the graph $\hat{G}_{v, \theta, k+1}(N_{in}(v) \cup \{v\})$ which is, by Lemma 2, a $(k, t^{2/3}, v)$ -VFTSS for $N_{in}(v) \cup \{v\}$. Therefore, the path $P := \cup_{\langle u, v \rangle \in P_0} P_{uv}$ is contained in $G \setminus S'$ and P is an t -spanner path between p and q . \square

5 Fault tolerant spanners with Steiner points

In this section we show a very simple method which constructs for an arbitrary set S of n points in \mathbb{R}^d , $t > 1$, and $k \in \mathbb{N}$, a (k, t) -EFTSS and (k, t) -VFTSS for S with $O(kn)$ edges and kn Steiner points. Then we prove the surprising fact that these upper bounds on the number of edges and on the number Steiner points in a (k, t) -EFTSS are optimal up to constant factors.

Theorem 5. *Let $S \subset \mathbb{R}^d$ be a set of n points, $k \in \mathbb{N}$, and let $t > 1$ a real constant. Then there is a (k, t) -EFTSS and (k, t) -VFTSS G for S with kn Steiner points and $O(kn)$ edges.*

Proof. Assume that the Euclidean distance between the closest pair of S is one. Otherwise, we scale S accordingly. Let ε be a real number such that $0 < \varepsilon \leq (t-1)/3$. Let $t^* = t - 2\varepsilon$ and let $G^* = (S, E^*)$ be a t^* -spanner for S with $O(n)$ edges. G^* can be computed, for example, using the method described in [4] or in [14]. We construct from G^* a (k, t) -EFTSS/VFTSS G for S in the following way. Let $o \in \mathbb{R}^d$ be a fixed point and let $D := \{x \in \mathbb{R}^d : \text{dist}_2(o, x) = \varepsilon\}$ be the sphere with radius ε whose center is o . Let s^1, \dots, s^k be k distinct points on D . (In the case if $d = 1$, let s^1, \dots, s^k be k distinct points such that $0 < \text{dist}_2(o, p_i) \leq \varepsilon$, $1 \leq i \leq k$.) For each point $p \in S$ translate the sphere D

and the points s^1, \dots, s^k on D such that p becomes the center of the sphere. Let p^1, \dots, p^k denote the translated points around p . We construct the graph $G = (V, E)$ such that

$$\begin{aligned} V &:= \{p, p^1, \dots, p^k : p \in S\} \quad \text{and} \\ E &:= \{(p, q) : (p, q) \in E^*\} \cup \{(p, p^i) : p \in S, 1 \leq i \leq k\} \cup \\ &\quad \{(p^i, q^i) : (p, q) \in E^*, 1 \leq i \leq k\}. \end{aligned}$$




Fig. 3. Example for the graphs G^* and G for $k = 3, d = 2$.

Clearly, the graph G has kn Steiner points and $O(kn)$ edges. It is obvious that G is a k -EFTSS and k -VFTSS for S , because for each pair of points $p, q \in S$ and for each t^* -spanner path $P^* = p, p_1, \dots, p_{l-1}, q$ in G^* between p and q , there are $k + 1$ edge disjoint and up to the endpoints vertex disjoint pq -paths $P^0 = p, p_1, \dots, p_{l-1}, q$ and $P^i = p, p^i, p_1^i, \dots, p_{l-1}^i, q^i, q$, $1 \leq i \leq k$, in G whose length is at most

$$\text{length}(P^*) + 2\epsilon \leq t^* \cdot \text{dist}_2(p, q) + 2\epsilon \leq t \cdot \text{dist}_2(p, q).$$

Figure 3 shows an example. □

Now we prove a lower bound on the number of edges and Steiner points which shows that the above upper bound is optimal up to a constant factor.

Theorem 6. *For each $k \in \mathbb{N}$, $n \geq 2$, and $t > 1$, there exists a set $S \subset \mathbb{R}^d$ of n points such that each (k, t) -EFTSS for S contains at least $\Omega(kn)$ Steiner points and $\Omega(kn)$ edges.*

Proof. We give an example for a set S of n points in the plane for which we show that each (k, t) -EFTSS for S contains $\Omega(kn)$ Steiner points and $\Omega(kn)$ edges. For two points $p, q \in \mathbb{R}^d$ let

$$el(p, q) := \{x \in \mathbb{R}^d : \text{dist}_2(p, x) + \text{dist}_2(q, x) \leq t \cdot \text{dist}_2(p, q)\}.$$

If p, q are two points in S and G is a (k, t) -EFTSS for S then each t -spanner path between p and q must be contained entirely in $el(p, q)$. Clearly, a path which contains a point v outside $el(p, q)$ has a length at least $\text{dist}_2(p, v) + \text{dist}_2(q, v)$ which is greater than $t \cdot \text{dist}_2(p, q)$. For $p, q \in S$ let G_{pq} be the smallest subgraph of G which contains all t -spanner paths between p and q . Since G is a (k, t) -EFTSS, G_{pq} must be k -edge connected. Otherwise, we could separate p from q in G_{pq} by deletion of a set E' of k edges, and therefore, we would not have any t -spanner path in $G \setminus E'$. Since the graph G_{pq} is k -edge connected, Menger's Theorem implies that it contains at least $k + 1$ edge disjoint pq -paths. Hence, G_{pq} – and, therefore, $el(p, q)$ – contains at least k vertices different from p and q and at least $2k - 1$ edges of G .

Now we show how to place the points of S in order to get the desired lower bounds. We construct the set S of n points in the plane hierarchically bottom-up. For simplicity of the description we assume that n is a power of two. Let l be a horizontal line

and let o be any fixed point of l . We place the points of S on l . We put $p_1 \in S$ to o and $p_2 \in S$ right from p_1 such that $dist_2(p_1, p_2) = 1$. Let $\varepsilon > 0$ be a fixed real number. We translate $el(p_1, p_2)$ with the points p_1 and p_2 right on l , by a Euclidean distance $t + \varepsilon$. This translation guarantees that $el(p_1, p_2)$ and the translated ellipsoid are distinct. Let p_3 and p_4 denote the translated points p_1 and p_2 , respectively. In general, in the i th step, $1 \leq i < \log n$, we translate the ellipsoid $el(p_1, p_{2^i})$ with the points p_1, \dots, p_{2^i} right on l , by a Euclidean distance $t \cdot dist_2(p_1, p_{2^i}) + \varepsilon$. Denote the translated points by $p_{1+2^i}, \dots, p_{2^{i+1}}$ (Figure 4). Then the ellipsoids $el(p_1, p_{2^i})$ and $el(p_{2^i+1}, p_{2^{i+1}})$ are distinct. We say that the ellipsoid $el(p_1, p_{2^{i+1}})$ is the *parent* of $el(p_1, p_{2^i})$ and $el(p_{1+2^i}, p_{2^{i+1}})$. Furthermore, we call the two children of an ellipsoid *siblings* of one another. We denote by $parent(el(\cdot, \cdot))$ and $sib(el(\cdot, \cdot))$ the parent and the sibling of an ellipsoid $el(\cdot, \cdot)$, respectively.




Fig. 4. Example for a set S of n points for which each (k, t) -EFTSS contains at least $n/2$ Steiner points and $(3k + 3)n/2 - (k + 1)$ edges.

Now we count the Steiner points and the edges in an arbitrary (k, t) -EFTSS G for the set S . Consider a pair of points $p_{2j-1}, p_{2j} \in S$, $1 \leq j \leq n/2$. For this pair, there are at least $k + 1$ edge disjoint paths in G contained entirely in $el(p_{2j-1}, p_{2j})$. Since, for $j \neq j'$ the ellipsoids $el(p_{2j-1}, p_{2j})$ and $el(p_{2j'-1}, p_{2j'})$ are disjoint, each $el(p_{2j-1}, p_{2j})$ contains in the interior at least k Steiner points and $2k + 1$ edges. Furthermore, p_{2j-1} and p_{2j} must be $(k + 1)$ -edge connected with the points of $sib(el(p_{2j-1}, p_{2j}))$. Therefore, we have at least $k + 1$ edges contained entirely in $parent(el(p_{2j-1}, p_{2j}))$ that have exactly one endpoint in $el(p_{2j-1}, p_{2j})$. We can repeat these arguments at each level of the hierarchy of the ellipsoids. Then we obtain that the number of edges in G is at least $(2k + 1)n/2 + (k + 1)(n/2 - 1) = (3k + 2)n/2 - (k + 1)$ and the number of Steiner points is at least $kn/2$.

In the case if n is not a power of two, we place 2^i points, where $2^{i-1} < n < 2^i$, in the same way as described above. Then we remove the points p_{n+1}, \dots, p_{2^i} . Using the above arguments we obtain that for this point set, each (k, t) -EFTSS contains at least $k \lfloor n/2 \rfloor$ Steiner points and $(3k + 2) \lfloor n/2 \rfloor - (k + 1)$ edges. This proves the claim of the theorem. \square

6 Conclusion and open problems

Some interesting problems remain to be solved. Is it possible to construct a (k, t) -VFTS whose degree is bounded by $O(k)$? Levcopoulos et al. [12] studied fault tolerant spanners with low weight. Let $w(MST)$ be the weight of the minimum spanning tree of S . In [12] it is proven that for each S a (k, t) -VFTS can be constructed whose weight is

$O(c^{k+1}w(MST))$ for some constant c . Can this upper bound be improved? In [12] it is also proven that $\Omega(k^2w(MST))$ is a lower bound on the weight. Is it possible to construct a (k,t) -VFTS with lower weight using Steiner points? Finally, we do not know any results for fault tolerant spanners with low diameter.

Acknowledgment: I would like to thank Artur Czumaj, Matthias Fischer, and Silvia Götz for their helpful comments and suggestions.

References

1. S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid. Euclidean spanners: Short, thin, and lanky. In *27th ACM Symposium on Theory of Computing (STOC'95)*, pages 489–498, 1995.
2. S. Arya, D. M. Mount, and M. Smid. Randomized and deterministic algorithms for geometric spanners of small diameter. In *35th IEEE Symposium on Foundations of Computer Science (FOCS'94)*, pages 703–712, 1994.
3. M. Ben-Or. Lower bounds for algebraic computation trees. In *15th ACM Symposium on Theory of Computing (STOC'83)*, pages 80–86, 1983.
4. P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k -nearest neighbors and n -body potential fields. *Journal of the ACM*, 42:67–90, 1995.
5. D. Z. Chen, G. Das, and M. Smid. Lower bounds for computing geometric spanners and approximate shortest paths. In *8th Canadian Conference on Computational Geometry (CCCG'96)*, pages 155–160, 1996.
6. L. P. Chew. There is a planar graph almost as good as the complete graph. In *2nd Annual ACM Symposium on Computational Geometry (SCG'86)*, pages 169–177, 1986.
7. K. L. Clarkson. Approximation algorithms for shortest path motion planning. In *19th ACM Symposium on Theory of Computing (STOC'87)*, pages 56–65, 1987.
8. M. Fischer, T. Lukovszki, and M. Ziegler. Geometric searching in walkthrough animations with weak spanners in real time. In *6th Annual European Symposium on Algorithms (ESA'98)*, pages 163–174, 1998.
9. J. G. Hocking and G. S. Young. *Topology*. Addison-Wesley, 1961.
10. J. M. Keil. Approximating the complete Euclidean graph. In *1st Scandinavian Workshop on Algorithm Theory (SWAT'88)*, pages 208–213, 1988.
11. J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete & Computational Geometry*, 7:13–28, 1992.
12. C. Levcopoulos, G. Narasimhan, and M. Smid. Efficient algorithms for constructing fault-tolerant geometric spanners. In *30th ACM Symposium on Theory of Computing (STOC'98)*, pages 186–195, 1998.
13. S. B. Rao and W. D. Smith. Improved approximation schemes for geometrical graphs via ‘spanners’ and ‘banyans’. In *30th ACM Symposium on Theory of Computing (STOC'98)*, pages 540–550, 1998.
14. J. Ruppert and R. Seidel. Approximating the d -dimensional complete Euclidean graph. In *3rd Canadian Conference on Computational Geometry (CCCG'91)*, pages 207–210, 1991.
15. J. S. Salowe. Constructing multidimensional spanner graphs. *International Journal of Computational Geometry & Applications*, 1:99–107, 1991.
16. P. M. Vaidya. A sparse graph almost as good as the complete graph on points in k dimensions. *Discrete & Computational Geometry*, 6:369–381, 1991.
17. A. C. Yao. On constructing minimum spanning trees in k -dimensional spaces and related problems. *SIAM Journal on Computing*, 11:721–736, 1982.