

Improved algorithms for constructing fault-tolerant spanners*

Christos Levcopoulos[†] Giri Narasimhan[‡] Michiel Smid[§]

Abstract

Let S be a set of n points in a metric space, and k a positive integer. Algorithms are given that construct k -fault-tolerant spanners for S . If in such a spanner at most k vertices and/or edges are removed, then each pair of points in the remaining graph is still connected by a “short” path. First, an algorithm is given that transforms an arbitrary spanner into a k -fault-tolerant spanner. For the Euclidean metric in \mathbb{R}^d , this leads to an $O(n \log n + c^k n)$ -time algorithm that constructs a k -fault-tolerant spanner of degree $O(c^k)$, whose total edge length is $O(c^k)$ times the weight of a minimum spanning tree of S , for some constant c . For constant values of k , this result is optimal. In the second part of the paper, algorithms are presented for the Euclidean metric in \mathbb{R}^d . These algorithms construct (i) in $O(n \log n + k^2 n)$ time, a k -fault-tolerant spanner with $O(k^2 n)$ edges, and (ii) in $O(kn \log n)$ time, such a spanner with $O(kn \log n)$ edges.

1 Introduction

Spanners have applications in the design of networks. Consider a set S of n points in a metric space. A network on S can be modeled as an undirected

*Portions of this work appear, in preliminary form, in [12]. The present paper improves some of the results presented in [12].

[†]Department of Computer Science, Lund University, Box 118, S-22100 Lund, Sweden. E-mail: Christos.Levcopoulos@dna.lth.se.

[‡]Department of Mathematical Sciences, The University of Memphis, Memphis TN 38152. E-mail: giri@msci.memphis.edu. Research supported by NSF Grant CCR-940-9752, and a grant by Cadence Design Systems.

[§]Department of Computer Science, University of Magdeburg, D-39106 Magdeburg, Germany. E-mail: michiel@isg.cs.uni-magdeburg.de.

graph G with vertex set S and with edges $e = (a, b)$ of length $|e|$ that is defined as the distance $|ab|$ between its two endpoints a and b . Let p and q be two points of S , and let P be a pq -path in G , i.e., a path in G between p and q . The length $|P|$ of P is defined as the sum of the lengths of the edges of P .

Let $t > 1$ be a real number. We say that G is a t -spanner for S , if for each pair of points $p, q \in S$, there exists a pq -path in G of length at most t times the distance between p and q . If S is a set of points in \mathbb{R}^d for some constant d , and the metric is the Euclidean metric, then we call the graph G a *Euclidean t -spanner*.

The problem of constructing spanners has been investigated by many researchers. For general metric spaces, Althöfer *et al.* [1], and Chandra *et al.* [7] showed that a natural greedy algorithm computes, for any constant $t > 1$, a t -spanner with $O(n^{1+2/(t-1)})$ edges, in $O(n^{3+4/(t-1)})$ time.

For the Euclidean case in \mathbb{R}^2 , Keil and Gutwin [11] showed that for any constant $t > 1$, a t -spanner for S having $O(n)$ edges can be constructed in $O(n \log n)$ time. Salowe [15], Vaidya [17] and Callahan and Kosaraju [5] showed the same result for any fixed dimension d .

In this paper, we show that it is possible to incorporate *fault-tolerance* into such networks. Fault tolerance is intimately related to the graph-theoretic concept of *connectivity*. The edge (vertex) connectivity of a graph is the minimum number of edges (vertices) that need to be removed in order to disconnect it. Fault-tolerant networks are usually designed by making them highly connected.

We construct networks that are more than just resilient to edge or vertex faults. Our networks have the property that after removing at most k vertices and/or edges, the remaining graph still contains “short” paths between each pair of points. Before we can define this formally, we have to introduce the following notation.

If S is a set of points, then K_S denotes the complete graph on S . Let $G = (S, E)$ be a graph, E' a subset of E , and S' a subset of S . We denote by $G \setminus S'$ the graph with vertex set $S \setminus S'$, and edge set the set of all edges of E that have both endpoints in $S \setminus S'$. Similarly, $G \setminus E'$ denotes the graph $(S, E \setminus E')$. Finally, $G \setminus (S', E')$ denotes the graph with vertex set $S \setminus S'$, and edge set the set of all edges of $E \setminus E'$ that have both endpoints in $S \setminus S'$.

Definition 1 *Let S be a set of n points in a metric space, $t > 1$ a real number, k a positive integer, and $G = (S, E)$ an undirected graph.*

1. G is called a k -vertex fault-tolerant t -spanner for S , or (k, t) -VF T S, if for each subset S' of S having size at most k , the graph $G \setminus S'$ is a t -spanner for the points of $S \setminus S'$.
2. G is called a k -edge fault-tolerant t -spanner for S , or (k, t) -EF T S, if the following holds for each subset E' of E having size at most k :
 - For each pair p and q of points in S , the graph $G \setminus E'$ contains a pq -path having length at most t times the length of a shortest pq -path in the graph $K_S \setminus E'$.
3. G is called a k -fault-tolerant t -spanner for S , or (k, t) -F T S, if the following holds for each subset S' of S and each subset E' of E such that $|S'| + |E'| \leq k$:
 - For each pair p and q of points in $S \setminus S'$, the graph $G \setminus (S', E')$ contains a pq -path having length at most t times the length of a shortest pq -path in the graph $K_S \setminus (S', E')$.

Note that in a vertex and/or edge fault-tolerant t -spanner, our definition insists that between every pair of points there is a path whose length is at most t times the best possible path under the circumstances, i.e., the shortest path in the graph $K_S \setminus S'$, $K_S \setminus E'$, or $K_S \setminus (S', E')$.

In the definition of a (k, t) -VF T S, we could have required that for each pair p and q of points in $S \setminus S'$, the graph $G \setminus S'$ contains a pq -path having length at most t times the length of a shortest pq -path P in the graph $K_S \setminus S'$. Since $K_S \setminus S'$ is the complete graph on the point set $S \setminus S'$, this shortest path P , however, consists of the single edge (p, q) . Hence, $G \setminus S'$ is indeed a t -spanner for $S \setminus S'$.

To our knowledge, the concept of fault-tolerant spanners has not been investigated before. The only related work that we are aware of is by Ueno *et al.* [16], who showed that if an unweighted graph has a sufficiently high connectivity, then it contains a spanner having a linear number of edges.

1.1 Our results

Unless stated otherwise, all algorithms in this paper belong to the algebraic computation tree model. (See Ben-Or [3] and Preparata and Shamos [14].)

It is clear that any (k, t) -F T S is also a (k, t) -VF T S and a (k, t) -EF T S. In Section 2, we will prove the converse. That is, we show that any (k, t) -VF T S

is in fact a (k, t) -FTS and, hence, in particular, a (k, t) -EFTS. As a result, it suffices to show how to construct spanners that are resilient to vertex faults.

In Section 3, we give a simple construction that transforms any t -spanner G_0 into a (k, t) -VFTS G . If the degree of each vertex in G_0 is bounded by D , then each vertex of G has degree $O(D^{k+1})$. Moreover, in this case the total edge length of G is proportional to $k \cdot D^k$ times that of G_0 . The running time of the algorithm that transforms G_0 into G is bounded by $O(D^{k+1}n)$.

For the Euclidean metric in \mathbb{R}^d , Arya *et al.* [2] show how to compute in $O(n \log n)$ time, a t -spanner G_0 , in which each point has a degree that is bounded by a constant D , that only depends on d and t . Combining this with our transformation of Section 3 and our result of Section 2, gives an algorithm that constructs in $O(n \log n + D^{k+1}n)$ time, a Euclidean (k, t) -FTS in which each point has degree $O(D^{k+1})$. If k is a constant, then this result is optimal. The optimality of the running time follows from Chen *et al.* [8], who proved that computing any Euclidean t -spanner takes $\Omega(n \log n)$ time in the algebraic computation tree model.

Arya *et al.* [2] also claim an $O(n \log n)$ -time algorithm that computes a t -spanner in which the degree of each point is bounded by a constant D , and whose total edge length is bounded by a constant times the weight of a minimum spanning tree of the points. Unfortunately, this result seems to be incorrect. Recently, Gudmundsson *et al.* [10] achieved this result in the *real RAM* model. (See Preparata and Shamos [14].) Hence, if we combine this result with the transformation of Section 3, then we get a real RAM algorithm that constructs in $O(n \log n + D^{k+1}n)$ time, a Euclidean (k, t) -FTS of degree $O(D^{k+1})$, whose total edge length is proportional to $O(k \cdot D^k)$ times the weight of a minimum spanning tree of the points.

In Section 4, we show that a Euclidean (k, t) -FTS having $O(k^2n)$ edges can be constructed in $O(n \log n + k^2n)$ time, and that such a spanner having $O(kn \log n)$ edges can be constructed in $O(kn \log n)$ time, where the constant factors only depend on t and the dimension d . Our construction is based on the *well-separated pair decomposition* of Callahan and Kosaraju [6]. They show in [5] that a Euclidean t -spanner with $O(n)$ edges can be obtained from this decomposition. We extend this result to fault-tolerant spanners.

2 It suffices to construct vertex fault-tolerant spanners

In this section, we prove the following theorem.

Theorem 1 *Let S be a set of n points in a metric space, k a positive integer, $t > 1$ a real constant, and $G = (S, E)$ an undirected graph. Then G is a (k, t) -VF TS for S if and only if it is a (k, t) -F TS for S .*

It is clear that a (k, t) -F TS is also a (k, t) -VF TS . To prove the converse, assume that G is a (k, t) -VF TS for S . Let S' be a subset of S of size k' , and let E' be a subset of E of size k'' , such that $k' + k'' \leq k$. We may assume without loss of generality that no edge of E' is incident to any point of S' ; otherwise, we can decrease k'' accordingly.

Let p and q be two distinct points of $S \setminus S'$. We have to show that the graph $G \setminus (S', E')$ contains a pq -path of length at most t times the length of a shortest pq -path in $K_S \setminus (S', E')$. This follows from the following two lemmas.

Lemma 1 *Assume that (p, q) is an edge of $K_S \setminus (S', E')$. Then $G \setminus (S', E')$ contains a pq -path of length at most t times the distance between p and q .*

Proof. Let S'' be any set of at most k'' vertices of $S \setminus \{p, q\}$, that is obtained by taking for each edge of E' an arbitrary endpoint that is not equal to p or q . Since (p, q) is not an edge of E' , this is possible. (For example, if (a, b) and (b, c) are edges of E' , then S'' can contain the endpoints a and b ; or a and c ; or b and c ; or only b .) Define $G' := G \setminus (S' \cup S'')$. Note that

$$|S' \cup S''| = |S'| + |S''| \leq k' + k'' \leq k.$$

Since G is a (k, t) -VF TS for S , the graph G' is a t -spanner for $S \setminus (S' \cup S'')$. Since p and q are vertices of G' , this graph contains a pq -path P of length at most $t|pq|$. This path neither contains vertices of S' , nor edges of E' . That is, P is a pq -path in $G \setminus (S', E')$. ■

Lemma 2 *The graph $G \setminus (S', E')$ contains a pq -path of length at most t times the length of a shortest pq -path in $K_S \setminus (S', E')$.*

Proof. Let $P = (p_0 = p, p_1, p_2, \dots, p_l = q)$ be a shortest pq -path in $K_S \setminus (S', E')$. Then for each i , $0 \leq i < l$, (p_i, p_{i+1}) is an edge of $K_S \setminus (S', E')$. Hence by Lemma 1, the graph $G \setminus (S', E')$ contains a path Q_i between p_i and p_{i+1} having length at most $t|p_i p_{i+1}|$. Let Q be the concatenation of Q_0, Q_1, \dots, Q_{l-1} . Then, Q is a pq -path in $G \setminus (S', E')$, having length

$$\sum_{i=0}^{l-1} |Q_i| \leq \sum_{i=0}^{l-1} t|p_i p_{i+1}| = t|P|.$$

This proves the lemma. ■

3 Fault-tolerant spanners in general metric spaces

In this section, we give a simple transformation that turns any spanner G_0 into a fault-tolerant spanner G . If the degree of G_0 is bounded by D , then the degree of G is proportional to D^{k+1} . Moreover, in this case, the transformation increases the total edge length by at most a factor proportional to $k \cdot D^k$.

Let S be a set of n points in a metric space, $t > 1$ a real number, and k a positive integer. Let G_0 be an arbitrary t -spanner for S . For each vertex $p \in S$, let $N(p)$ be the set of all vertices of $S \setminus \{p\}$ that are connected to p , in G_0 , by a path consisting of at most $k + 1$ edges. Define $E_p := \{(p, q) : q \in N(p)\}$. The transformed graph G has the points of S as its vertices, and it has edge set $E := \bigcup_{p \in S} E_p$. Note that G_0 is a subgraph of G .

Lemma 3 *The graph G is a (k, t) -FTS for S .*

Proof. By Theorem 1, it suffices to show that G is a (k, t) -VFTS for S .

Let S' be a subset of S having size at most k , and let p and q be two distinct points of $S \setminus S'$. We will show that the graph $G \setminus S'$ contains a pq -path of length at most t times the distance between p and q .

Since G_0 is a t -spanner for S , there is a pq -path

$$P = (q_0 = p, q_1, q_2, \dots, q_l = q)$$

in G_0 of length at most $t|pq|$. We will construct a pq -path Q in $G \setminus S'$ that is a subpath of P . Then, the triangle inequality implies that the length of Q is at most that of P . This will prove the lemma.

First assume that $l \leq k + 1$. Then, $q \in N(p)$ and, hence, (p, q) is an edge of G . Since p and q are both vertices of $S \setminus S'$, (p, q) is an edge of $G \setminus S'$, and we can take for Q the path consisting of this single edge.

Assume that $k + 2 \leq l$. The following algorithm constructs the path $Q = (p_0, p_1, \dots)$ incrementally.

Step 1: Define $p_0 := p$, $i := 0$, and $j := 0$. Go to Step 2.

Step 2: At this moment, $Q = (p_0, p_1, \dots, p_i)$ is a path in $G \setminus S'$, j is the index such that $p_i = q_j$, and $j + k + 2 \leq l$. (In particular, $p_i \neq q$, and $q_j \in S \setminus S'$.)

If there is an index m , $j + 1 \leq m \leq j + k + 1$, such that (i) $m + k + 2 \leq l$ and (ii) q_m is a vertex of $S \setminus S'$, then go to Step 3. Otherwise, go to Step 4.

Step 3: Since q_j and q_m are both vertices of $S \setminus S'$, and $q_m \in N(q_j)$, we know that (q_j, q_m) is an edge of $G \setminus S'$. Therefore, we define $p_{i+1} := q_m$, set $i := i + 1$ and $j := m$, and go to Step 2.

Step 4: We know that $p_i = q_j$ and $j + k + 2 \leq l$. Moreover, for all m , $j + 1 \leq m \leq j + k + 1$, such that q_m is a vertex of $S \setminus S'$, we have $m + k + 1 \geq l$.

We claim that there is an index m , $j + 1 \leq m \leq j + k + 1$, such that (q_j, q_m) and (q_m, q) are both edges of $G \setminus S'$.

Assume this claim is true. Then we define $p_{i+1} := q_m$ and $p_{i+2} := q$, and the construction of the pq -path Q is complete.

It remains to prove the claim. Since S' has size at most k , there is an index m , $j + 1 \leq m \leq j + k + 1$, such that $q_m \in S \setminus S'$. Hence, $q_m \in N(q_j)$ and (q_j, q_m) is an edge of $G \setminus S'$. Our assumption implies that $m + k + 1 \geq l$. Therefore, $q = q_l \in N(q_m)$ and (q_m, q) is an edge of G . Since q_m and q are both contained in $S \setminus S'$, edge (q_m, q) is contained in $G \setminus S'$. This proves the claim.

Why does this algorithm terminate? Each time Step 3 is executed, path Q is extended by a new point. Therefore, at some moment, Step 4 must be executed. At that moment, Q reaches q , and the algorithm terminates. ■

Lemma 4 *Assume that each point of S has degree at most D in G_0 . Then*

1. *each point of S has degree at most $2 \cdot D^{k+1}$ in G , and*
2. *the total edge length of G is at most $8(k + 1) \cdot D^k$ times that of G_0 .*

Proof. Let $p \in S$. Then

$$|N(p)| \leq D + D^2 + D^3 + \dots + D^{k+1} \leq 2 \cdot D^{k+1}.$$

Since $q \in N(p)$ if and only if $p \in N(q)$, it follows that each point has degree at most $2 \cdot D^{k+1}$ in G .

To bound the total edge length of G , we use the following charging scheme. Let (p, q) be any edge of G , and consider any pq -path $P = (p_0 = p, p_1, p_2, \dots, p_l = q)$ in G_0 containing $l \leq k + 1$ edges. (Note that P exists.) We charge the length $|pq|$ of edge (p, q) to the edges of P , in such a way that no edge (p_i, p_{i+1}) , $0 \leq i < l$, is charged by more than $|p_i p_{i+1}|$. Since $|pq| \leq |P|$, this is possible. We do this for all edges of G .

For each edge e of G_0 , let n_e be the number of times this edge is charged. Then the total edge length of G is at most equal to $\sum_{e \in G_0} n_e \cdot |e|$. We will show that $n_e \leq 8(k + 1) \cdot D^k$. This will imply that the total edge length of G is at most $8(k + 1) \cdot D^k \cdot \sum_{e \in G_0} |e|$, which is equal to $8(k + 1) \cdot D^k$ times the total edge length of G_0 .

Let e be an edge of G_0 , and let it have endpoints a and b . Every time e is charged, there are two points p and q , such that there is a pq -path in G_0 containing at most $k + 1$ edges, e being one of them. Assume w.l.o.g. that a occurs before b on this path. Let i be the number of edges on the subpath from p to a . Then $0 \leq i \leq k$. If j denotes the number of edges on the subpath from b to q , then $0 \leq j \leq k - i$.

If we fix i and j , then the number of possibilities for p is at most

$$D + D^2 + D^3 + \dots + D^i \leq 2 \cdot D^i,$$

and the number of possibilities for q is at most

$$D + D^2 + D^3 + \dots + D^j \leq 2 \cdot D^j.$$

It follows that

$$\begin{aligned} n_e &\leq \sum_{i=0}^k 2 \cdot D^i \sum_{j=0}^{k-i} 2 \cdot D^j \\ &= 4 \sum_{i=0}^k D^i (1 + D + D^2 + \dots + D^{k-i}) \\ &\leq 8 \sum_{i=0}^k D^i \cdot D^{k-i} \\ &= 8(k + 1)D^k. \end{aligned}$$

We now apply these results to the Euclidean case. ■

Theorem 2 *Let S be a set of n points in \mathbb{R}^d , k a positive integer, and $t > 1$ a real constant. There exists a Euclidean (k, t) -FTS for S , in which each point has degree less than or equal to α^{dk+d} , for some constant α that only depends on t . This (k, t) -FTS can be computed in $O(n \log n + \alpha^{dk+d}n)$ time. If $t \downarrow 1$, then $\alpha \sim c/(t-1)$ for some constant c .*

Proof. In [2], it is shown that in $O(n \log n + \beta_{dt}n)$ time, a Euclidean t -spanner G_0 can be constructed whose degree D is bounded by β_{dt} . The value of β_{dt} only depends on d and t , and if $t \downarrow 1$, then $\beta_{dt} \sim (c'/(t-1))^d$ for some constant c' .

Let G be the graph obtained by applying our transformation to G_0 . By Lemma 3, G is a Euclidean (k, t) -FTS. The bound on the degree of G follows from Lemma 4. The definition of G immediately leads to an algorithm for constructing it from G_0 , in time $O(\sum_{p \in S} |N(p)|) = O(D^{k+1}n)$. ■

If the t -spanner G_0 has low total edge length and low degree, then our transformation results in a (k, t) -FTS, and Lemma 4 gives bounds on its total edge length and its degree. Das and Narasimhan [9] showed how to compute such a spanner G_0 in $O(n \log^2 n)$ time. The running time was recently improved by Gudmundsson *et al.* [10] to $O(n \log^2 n / \log \log n)$ in the algebraic computation tree model, and to $O(n \log n)$ in the real RAM model. Hence, we have the following result.

Theorem 3 *Let S be a set of n points in \mathbb{R}^d , k a positive integer, and $t > 1$ a real constant. There exists a Euclidean (k, t) -FTS for S , in which the degree of each point is at most c^k , and whose total edge length is at most $O(c^k)$ times the weight of a minimum spanning tree of S . Here, c is a constant that depends on t and d .*

1. *In the algebraic computation tree model, this (k, t) -FTS can be computed in $O(n \log^2 n / \log \log n + c^k n)$ time.*
2. *In the real RAM model, this (k, t) -FTS can be computed in $O(n \log n + c^k n)$ time.*

4 Euclidean fault-tolerant spanners with a polynomial number of edges

The number of edges in the Euclidean (k, t) -FTS of Theorems 2 and 3 is exponential in k . In this section, we give an algorithm for constructing a (k, t) -FTS that uses only a polynomial number of edges. Unfortunately, we are not able to prove non-trivial bounds on the degree and total edge length of this spanner. Before we give our construction, we recall some facts about well-separated pairs.

4.1 Well-separated pairs

Our construction of fault-tolerant spanners is based on the notion of *well-separated pairs*, which is due to Callahan and Kosaraju [6].

Definition 2 *Let $s > 0$ be a real number, and let A and B be two finite sets of points in \mathbb{R}^d . We say that A and B are well-separated w.r.t. s , if there are two disjoint d -dimensional balls C_A and C_B , having the same radius, such that (i) C_A contains all points of A , (ii) C_B contains all points of B , and (iii) the distance between C_A and C_B is at least equal to s times the radius of C_A .*

See Figure 1 for an illustration. In this paper, s will always be a constant, called the *separation constant*.

Definition 3 ([6]) *Let S be a set of n points in \mathbb{R}^d , and $s > 0$ a real number. A well-separated pair decomposition (WSPD) for S (w.r.t. s) is a sequence of pairs of non-empty subsets of S ,*

$$\{A_1, B_1\}, \{A_2, B_2\}, \dots, \{A_m, B_m\},$$

such that

1. $A_i \cap B_i = \emptyset$, for all $i = 1, 2, \dots, m$,
2. for any two distinct points p and q of S , there is exactly one pair $\{A_i, B_i\}$ in the sequence, such that
 - (a) $p \in A_i$ and $q \in B_i$, or




Figure 1: Two planar point sets A and B that are well-separated w.r.t. s . Both circles have radius ρ ; their distance is at least $s\rho$.

(b) $p \in B_i$ and $q \in A_i$,

3. A_i and B_i are well-separated w.r.t. s , for all $i = 1, 2, \dots, m$.

The integer m is called the size of the WSPD.

Theorem 4 ([4, 6]) Let S be a set of n points in \mathbb{R}^d , and $s > 0$ a separation constant.

1. In $O(n \log n + \alpha_{ds}n)$ time, we can compute a WSPD for S of size less than or equal to $\alpha_{ds}n$.
2. In $O(\alpha_{ds}n \log n)$ time, we can compute a WSPD for S of size $O(\alpha_{ds}n \log n)$, in which each pair $\{A_i, B_i\}$ contains at least one singleton set.

The constants in the Big-Oh bounds do not depend on s . Moreover, for a large separation constant s , the value of α_{ds} is proportional to $((c+1)s)^d$ for some constant c .




Figure 2: *Illustrating the construction of the fault-tolerant spanner.*

4.2 Definition of the graph G

Let S be a set of n points in \mathbb{R}^d , $t > 1$ a real constant, and k a positive integer. Consider an arbitrary WSPD

$$\{A_1, B_1\}, \{A_2, B_2\}, \dots, \{A_m, B_m\}$$

for S , with separation constant $s = 4(t + 1)/(t - 1)$.

We will define a graph G based on the well-separated pair decomposition, and show that it is a fault-tolerant spanner.

Our graph G has the points of S as its vertices. Below, we define for each i , $1 \leq i \leq m$, a set E_i of edges. The edge set E of G is then defined as $E := \bigcup_{i=1}^m E_i$.

Let $1 \leq i \leq m$, and consider the well-separated pair $\{A_i, B_i\}$. We assume without loss of generality that $|A_i| \geq |B_i|$. To define E_i , we distinguish three cases.

Case 1: $|B_i| \geq k + 1$.

Choose $k + 1$ arbitrary, but pairwise distinct points $a_j \in A_i$, $1 \leq j \leq k + 1$, and $k + 1$ arbitrary, but pairwise distinct points $b_j \in B_i$, $1 \leq j \leq k + 1$. The edge set E_i consists of the $k + 1$ edges (a_j, b_j) , $1 \leq j \leq k + 1$. (See Figure 2.)

Case 2: $|B_i| \leq k$ and $|A_i| \geq k + 1$.

Choose $k + 1$ arbitrary, but pairwise distinct points $a_j \in A_i$, $1 \leq j \leq k + 1$. Let $B_i = \{b_1, b_2, \dots, b_x\}$, where $x = |B_i| \leq k$. The edge set E_i consists of the $x(k + 1)$ edges (a_j, b_l) , $1 \leq j \leq k + 1$, $1 \leq l \leq x$. Hence, E_i has size at most $k(k + 1)$.

Case 3: $|A_i| \leq k$.

In this case, the set E_i is defined as the edge set of the complete bipartite Euclidean graph on the points of $A_i \cup B_i$. Note that E_i has size $|A_i| \cdot |B_i| \leq k^2$.

This concludes the definition of our graph G . Note that E , the edge set of G , has size $O(k^2m)$.

4.3 The graph G is a (k, t) -FTS for S

We now prove that the above construction does have the requisite properties. By Theorem 1, it suffices to show that G is a (k, t) -VFTS. Let S' be an arbitrary subset of S of size at most k , and let p and q be two points of $S \setminus S'$. We will prove that the graph $G \setminus S'$ contains a pq -path having length at most t times the Euclidean distance between p and q . The proof is by induction on the rank of the distance $|pq|$ in the sorted sequence of distances in $S \setminus S'$.

If $p = q$, then the claim clearly holds. So assume that $p \neq q$. Moreover, assume that for any pair $a, b \in S \setminus S'$ with $|ab| < |pq|$, the graph $G \setminus S'$ contains an ab -path of length at most $t|ab|$.

Let i , $1 \leq i \leq m$, be the index such that (i) $p \in A_i$ and $q \in B_i$, or (ii) $p \in B_i$ and $q \in A_i$. According to Definition 3, i exists and is, in fact, unique. We assume without loss of generality that (i) holds, and that $|A_i| \geq |B_i|$.

Since the sets A_i and B_i are well-separated, there are two balls C_{A_i} and C_{B_i} having the same radius, say ρ , that contain the sets A_i and B_i , respectively, and that have distance at least $s\rho$. We distinguish three cases.

Case 1: $|B_i| \geq k + 1$.

Consider the $k + 1$ points $a_j \in A_i$, $1 \leq j \leq k + 1$, and the $k + 1$ points $b_j \in B_i$, $1 \leq j \leq k + 1$, that were chosen in the construction of G .

Lemma 5 *There is an index j , $1 \leq j \leq k + 1$, such that the graph $G \setminus S'$ contains*

1. the edge (a_j, b_j) ,
2. a path P between p and a_j of length at most $2t\rho$, and
3. a path Q between q and b_j of length at most $2t\rho$.

Proof. Since S' has size at most k , there is an index j , $1 \leq j \leq k + 1$, such that a_j and b_j are both contained in $S \setminus S'$. Let j be an arbitrary index having this property. Then (a_j, b_j) is an edge of $G \setminus S'$.

If $p = a_j$, then we take for P the empty path, having length zero. So assume that $p \neq a_j$. Since $|pq| \geq s\rho$, $|pa_j| \leq 2\rho$, and $s > 2$, we must have $|pa_j| < |pq|$. Therefore, by the induction hypothesis, the graph $G \setminus S'$ contains a path P between p and a_j having length at most $t|pa_j|$. Clearly, P has length at most $2t\rho$.

In exactly the same way, it can be shown that $G \setminus S'$ contains a qb_j -path Q of length at most $2t\rho$. \blacksquare

We can now complete the proof for Case 1. Consider the index j , and the paths P and Q , of Lemma 5. Let R be the pq -path in $G \setminus S'$ obtained by concatenating path P , edge (a_j, b_j) , and path Q . We will show that $|R| \leq t|pq|$.

First note that $|R| \leq 4t\rho + |a_j b_j|$. The triangle inequality implies that $|a_j b_j| \leq |a_j p| + |pq| + |qb_j|$. Furthermore, $|a_j p| \leq 2\rho$ and $|qb_j| \leq 2\rho$. Therefore,

$$|R| \leq (4t + 4)\rho + |pq|.$$

Since $|pq| \geq s\rho$ and $s = 4(t + 1)/(t - 1)$, it follows that $|R| \leq t|pq|$. This completes the proof for Case 1.

Case 2: $|B_i| \leq k$ and $|A_i| \geq k + 1$.

Consider the $k + 1$ points $a_j \in A_i$, $1 \leq j \leq k + 1$, that were chosen in the construction of G . Let b_j , $1 \leq j \leq x = |B_i|$, be the points of B_i . Note that q is one of the b_j 's. Also, in G , point q is connected to each point a_j , $1 \leq j \leq k + 1$.

Let j , $1 \leq j \leq k + 1$, be an index such that a_j is a vertex of $G \setminus S'$. Then (a_j, q) is an edge of $G \setminus S'$. It follows in exactly the same way as in the proof of Lemma 5, that $G \setminus S'$ contains a pa_j -path P of length at most $2t\rho$. Then, just as in Case 1, it can be shown that the path consisting of P , followed by edge (a_j, q) , is a pq -path in $G \setminus S'$ of length at most $t|pq|$.

Case 3: $|A_i| \leq k$.

In this case, G contains the complete bipartite Euclidean graph on $A_i \cup B_i$ as a subgraph. Since p and q are both contained in $S \setminus S'$, (p, q) is an edge of $G \setminus S'$. That is, $G \setminus S'$ contains a pq -path of length $|pq|$, which is at most $t|pq|$.

We have proved the following result.

Theorem 5 *Let S be a set of n points in \mathbb{R}^d , k a positive integer, and $t > 1$ a real constant. Let*

$$\{A_1, B_1\}, \{A_2, B_2\}, \dots, \{A_m, B_m\}$$

be an arbitrary WSPD for S , with separation constant $s = 4(t + 1)/(t - 1)$. The graph $G = (S, E)$ defined above is a (k, t) -FTS for S . This graph contains $O(k^2 m)$ edges.

4.4 Constructing the graph G

The algorithm for constructing the graph G follows immediately from the results of the previous sections. Given the set S , the positive integer k , and the real constant $t > 1$, we use the algorithm of [6] (see Theorem 4) to compute a WSPD for S of size $m = O(n)$, in $O(n \log n)$ time. For each pair $\{A_i, B_i\}$ in this WSPD, we construct the corresponding edge set E_i . If Case 1 applies, then we construct E_i in $O(k)$ time. If Case 2 or 3 applies, then we need $O(k^2)$ time to construct E_i .

Theorem 6 *Let S be a set of n points in \mathbb{R}^d , k a positive integer, and $t > 1$ a real constant.*

1. *There exists a (k, t) -FTS for S containing at most $\gamma_{dt} k^2 n$ edges. The value of γ_{dt} only depends on d and t , and if $t \downarrow 1$, then $\gamma_{dt} \sim (c/(t-1))^d$ for some constant c .*
2. *This (k, t) -FTS can be computed in $O(n \log n + \gamma_{dt} k^2 n)$ time.*

Proof. Let $s = 4(t + 1)/(t - 1)$. By Theorem 4, constructing the graph G takes time $O(n \log n + \alpha_{ds} k^2 n)$, where $\alpha_{ds} \sim ((c' + 1)s)^d$ for some constant c' . For $t \downarrow 1$, we have $s \sim 8/(t - 1)$, and $\alpha_{ds} \sim (8(c' + 1)/(t - 1))^d$. This graph has $O(\alpha_{ds} k^2 n)$ edges. By Theorem 1, G is a (k, t) -FTS for S . ■

Now consider a WSPD for S , in which each pair $\{A_i, B_i\}$ contains at least one singleton set. Theorem 4 states that such a WSPD of size $O(\alpha_{ds} n \log n)$ exists and can be computed in $O(\alpha_{ds} n \log n)$ time. If we construct the corresponding graph G , then Case 1 never occurs; in Case 2, only $k + 1$ edges are added; and in Case 3, at most k edges are added. Hence, the total number of edges of G is less than or equal to $k + 1$ times the size of the WSPD. Therefore, Theorems 4 and 6 lead to the following result.

Theorem 7 *Let S be a set of n points in \mathbb{R}^d , k a positive integer, and $t > 1$ a real constant.*

1. *There exists a (k, t) -FTS for S containing at most $\gamma_{dt}kn \log n$ edges. The value of γ_{dt} only depends on d and t , and if $t \downarrow 1$, then $\gamma_{dt} \sim (c/(t-1))^d$ for some constant c .*
2. *This (k, t) -FTS can be computed in $O(\gamma_{dt}kn \log n)$ time.*

5 Concluding remarks

We have presented efficient algorithms for constructing spanners that are resilient to k vertex and/or edge faults. In particular, Theorems 6 and 7 give constructions that use a polynomial (i.e., $O(k^2n)$ and $O(kn \log n)$, respectively) number of edges. On the other hand, the constructions of Theorems 2 and 3 use a number of edges that is exponential in k . In the latter constructions, however, upper bounds on the degree and/or total edge length can be guaranteed.

If k is a constant, then the best results are those of Theorems 2 and 3. Theorem 2 gives a Euclidean k -fault-tolerant t -spanner, in which the degree of each vertex is bounded by a constant. Moreover, this spanner can be constructed in $O(n \log n)$ time. Chen *et al.* [8] showed that constructing *any* t -spanner—that is not necessarily resilient to faults—takes $\Omega(n \log n)$ time in the algebraic computation tree model. Therefore, the result of Theorem 2 is optimal for constant values of k . Theorem 3 gives a Euclidean k -fault-tolerant t -spanner of bounded degree, whose total edge length is bounded by a constant times the weight of a minimum spanning tree of the points. In the algebraic computation tree model, this spanner can be constructed in $O(n \log^2 n / \log \log n)$ time, whereas in the real RAM model, the running time is bounded by $O(n \log n)$. We leave open the problem of obtaining the latter running time in the algebraic computation tree model.

Some other interesting problems remain to be solved. Any graph on n vertices that remains connected after removing at most k edges, must have $\Omega(kn)$ edges. The reason is that each vertex in such a graph must have degree at least $k + 1$. In a follow-up paper to [12], Lukovszki [13] showed that a Euclidean (k, t) -FTS with $O(kn)$ edges can be computed, in $O(n \log^{d-1} n + kn \log \log n)$ time. Can such a fault-tolerant spanner be constructed in time $O(n \log n + kn)$?

Lukovszki [13] also showed that a Euclidean (k, t) -FTS in which the degree of each point is bounded by $O(k^2)$, can be computed in $O(n \log^{d-1} n + kn \log n + k^2 n)$ time. Does a Euclidean (k, t) -FTS of degree $O(k)$ exist and, if so, can it be computed in $O(n \log n + kn)$ time?

Let k be an even integer, and consider a set A of $1 + k/2$ points that are all very close to the origin. Let B be a set of $n - 1 - k/2$ points that are all very close together, but at distance roughly one from the origin. Let G be any Euclidean (k, t) -FTS for the set $S := A \cup B$, where t is a constant close to one. Then, since G is a (k, t) -EFTS, every point of A has to be connected to at least $1 + k/2$ points of B . Hence, G contains $\Omega(k^2)$ edges having length roughly equal to one. On the other hand, a minimum spanning tree of S has weight roughly equal to one.

Is it possible to construct, for any set S of n points in \mathbb{R}^d , a Euclidean (k, t) -FTS, such that each vertex has degree $O(k)$, and the weight of this graph is $O(k^2)$ times the weight of a minimum spanning tree of S ? Can such a graph be constructed in $O(n \log n + kn)$ time?

References

- [1] I. Althöfer, G. Das, D. P. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete Comput. Geom.*, 9:81–100, 1993.
- [2] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid. Euclidean spanners: short, thin, and lanky. In *Proc. 27th Annu. ACM Sympos. Theory Comput.*, pages 489–498, 1995.
- [3] M. Ben-Or. Lower bounds for algebraic computation trees. In *Proc. 15th Annu. ACM Sympos. Theory Comput.*, pages 80–86, 1983.
- [4] P. B. Callahan. *Dealing with higher dimensions: the well-separated pair decomposition and its applications*. Ph.D. thesis, Dept. Comput. Sci., Johns Hopkins University, Baltimore, Maryland, 1995.
- [5] P. B. Callahan and S. R. Kosaraju. Faster algorithms for some geometric graph problems in higher dimensions. In *Proc. 4th ACM-SIAM Sympos. Discrete Algorithms*, pages 291–300, 1993.

- [6] P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k -nearest-neighbors and n -body potential fields. *J. ACM*, 42:67–90, 1995.
- [7] B. Chandra, G. Das, G. Narasimhan, and J. Soares. New sparseness results on graph spanners. *Internat. J. Comput. Geom. Appl.*, 5:125–144, 1995.
- [8] D. Z. Chen, G. Das, and M. Smid. Lower bounds for computing geometric spanners and approximate shortest paths. In *Proc. 8th Canad. Conf. Comput. Geom.*, pages 155–160, 1996.
- [9] G. Das and G. Narasimhan. A fast algorithm for constructing sparse Euclidean spanners. *Internat. J. Comput. Geom. Appl.*, 7:297–315, 1997.
- [10] J. Gudmundsson, C. Levcopoulos, and G. Narasimhan. Improved greedy algorithms for constructing sparse geometric spanners. Manuscript. A shortened version will appear in Proc. SWAT’00 (Bergen, Norway), Lecture Notes in Computer Science, Springer-Verlag, 2000.
- [11] J. M. Keil and C. A. Gutwin. Classes of graphs which approximate the complete Euclidean graph. *Discrete Comput. Geom.*, 7:13–28, 1992.
- [12] C. Levcopoulos, G. Narasimhan, and M. Smid. Efficient algorithms for constructing fault-tolerant geometric spanners. In *Proc. 30th Annu. ACM Sympos. Theory Comput.*, pages 186–195, 1998.
- [13] T. Lukovszki. New results on fault tolerant geometric spanners. In *Proceedings of the 6th Workshop on Algorithms and Data Structures*, volume 1663 of *Lecture Notes in Computer Science*, pages 193–204, Berlin, 1999. Springer-Verlag.
- [14] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, Berlin, 1988.
- [15] J. S. Salowe. Constructing multidimensional spanner graphs. *Internat. J. Comput. Geom. Appl.*, 1:99–107, 1991.
- [16] S. Ueno, M. Yamazaki, and Y. Kajitani. Graph spanners and connectivity. In *Proc. 3rd Annu. Internat. Sympos. Algorithms Comput.*, volume 650 of *Lecture Notes Comput. Sci.* Springer-Verlag, 1992.

- [17] P. M. Vaidya. A sparse graph almost as good as the complete graph on points in K dimensions. *Discrete Comput. Geom.*, 6:369–381, 1991.