

Sparse Fault-Tolerant Spanners for Doubling Metrics with Bounded Hop-Diameter or Degree

T.-H. Hubert Chan, Mingfei Li, and Li Ning

The University of Hong Kong

Abstract. We study *fault-tolerant* spanners in doubling metrics. A subgraph H for a metric space X is called a k -vertex-fault-tolerant t -spanner ((k, t) -VFTS or simply k -VFTS), if for any subset $S \subseteq X$ with $|S| \leq k$, it holds that $d_{H \setminus S}(x, y) \leq t \cdot d(x, y)$, for any pair of $x, y \in X \setminus S$.

For any doubling metric, we give a basic construction of k -VFTS with stretch arbitrarily close to 1 that has optimal $O(kn)$ edges. In addition, we also consider bounded hop-diameter, which is studied in the context of fault-tolerance for the first time even for Euclidean spanners. We provide a construction of k -VFTS with bounded hop-diameter: for $m \geq 2n$, we can reduce the hop-diameter of the above k -VFTS to $O(\alpha(m, n))$ by adding $O(km)$ edges, where α is a functional inverse of the Ackermann's function.

Finally, we construct a fault-tolerant single-sink spanner with bounded maximum degree, and use it to reduce the maximum degree of our basic k -VFTS. As a result, we get a k -VFTS with $O(k^2n)$ edges and maximum degree $O(k^2)$.

1 Introduction

A metric space (X, d) can be represented by a complete graph $G = (X, E)$, where the edge weight $w(e)$ on an edge $e = \{x, y\}$ is $d(x, y)$. A t -spanner of X , is a weighted subgraph $H = (X, E')$ of G that preserves all pairwise distance within a factor of t , i.e., $d_H(x, y) \leq t \cdot d(x, y)$ for all $x, y \in X$. Here, $d_H(x, y)$ denotes the shortest-path distance between x and y in H . The factor t is called the *stretch* of H . A path between x and y in H with length at most $t \cdot d(x, y)$ is called a *t -spanner path*. Spanners have been studied extensively since the mid-eighties (see [2, 8, 1, 16, 12, 3, 4, 10] and the references therein; also refer to [15] for an excellent survey).

Spanners are important structures, as they enable approximation of a metric space in a much more economical form. One natural requirement is that spanners should be sparse, ideally with the number of edges being linear in the number of points in the metric space. In addition, for some applications, it might also be required that a spanner should have small maximum degree, or a small *hop-diameter*, i.e., every pair of points x and y should be connected by a t -spanner path with small number of edges.

In many applications of spanners, we want our spanner to be robust to failures, meaning that even when some of the points in the spanner fail, the remaining part is still a t -spanner. Formally, given $1 \leq k \leq n - 2$, a spanner H of X is

called a k -vertex-fault-tolerant t -spanner ((k, t) -VFST or simply k -VFST if the stretch t is clear from context), if for any subset $S \subseteq X$ with $|S| \leq k$, $H \setminus S$ is a t -spanner for $X \setminus S$.

The notion of fault-tolerant spanners was introduced by Levcopoulos et al. [13] in the context of *Euclidean spanners* (the special case when X is a finite subset of low dimensional Euclidean space and $d(x, y) = \|x - y\|_2$). They presented an algorithm that constructs Euclidean $(k, 1 + \epsilon)$ -VFST with $O(k^2 n)$ edges. This result was later improved by Lukovszki in [14]. They provided two constructions of $(k, 1 + \epsilon)$ -VFST: one with optimal $O(kn)$ edges, and the other with $O(k^2 n)$ edges and maximum degree $O(k^2)$. There has also been research on the trade-off between maximum degree and weight in fault-tolerant Euclidean spanners [13, 7], and fault-tolerant spanners for general graphs [6, 9].

In this paper, we study fault-tolerant spanners for *doubling metrics*. The *doubling dimension* of a metric space (X, d) , denoted by $\dim(X)$ (or \dim when the context is clear), is the smallest value ρ such that every ball in X can be covered by 2^ρ balls of half the radius [11]. A metric space is called *doubling*, if its doubling dimension is bounded by some constant. Doubling dimension is a generalization of Euclidean dimension to arbitrary metric spaces, as the space \mathbb{R}^T equipped with ℓ_p -norm has doubling dimension $\Theta(T)$ [11]. Spanners for doubling metrics have been studied in [12, 3, 4, 10, 16].

1.1 Our Results and Techniques

Basic spanner with small number of edges. Our first result is a construction of $(k, 1 + \epsilon)$ -VFST for doubling metrics with $O(kn)$ edges. Note that the size is optimal up to a constant factor [14].

Theorem 1 ($(k, 1 + \epsilon)$ -VFST with $O(kn)$ Edges). *Let (X, d) be a doubling metric with n points and let $0 < \epsilon < \frac{1}{2}$ be a constant. Given $1 \leq k \leq n - 2$, there exists a $(k, 1 + \epsilon)$ -VFST of X with $\epsilon^{-O(\dim)} \cdot kn$ edges.*

Our technique of the basic k -VFST construction is an extension of that in [3]. Specifically, we give a k -fault-tolerant version of the hierarchical nets and net-trees used in [3], which guarantee that even under the failure of at most k points, for any functioning point x and level i , there exists a net-tree in which there is a path with length at most $O(2^i)$ between x and some level- i net point. We also add cross edges between net points at the same level that are reasonably close. Then, we show that under the failure of at most k points, a $(1 + \epsilon)$ -spanner path between x and y can be formed by first climbing from x to some net point x' at an appropriate level i , then going along a cross edge $\{x', y'\}$, and finally going from y' down to y .

The upper bound on the number of edges in the spanner is established in a way similar to [3] by carefully assigning a direction to each edge, and then showing that the out-degree of each point is bounded by $O(k)$. However, we simplify and improve the analysis because in our case cross edges can be added between arbitrarily close points.

Spanners with small hop-diameters. We also consider the hop-diameter, which is studied in the context of fault-tolerance for the first time even for Euclidean spanners. The k -vertex-fault-tolerant hop-diameter is defined as follows.

Definition 1 (k -Vertex-Fault-Tolerant Hop-Diameter). *Let H be a (k, t) -VFTS for the metric space (X, d) . The k -vertex-fault-tolerant hop-diameter (or simply hop-diameter) of H is at most D , if for any set of points $S \subseteq X$ with $|S| \leq k$, there exists a t -spanner path in $H \setminus S$ with at most D edges (hops) between every pair of $x, y \in X$.*

We show that by adding a few extra edges to our basic k -VFTS, we can significantly reduce its hop-diameter.

Theorem 2 ($(k, 1 + \epsilon)$ -VFTS with Small Hop-Diameter). *Let $m \geq 2n$. We can add $O(km)$ extra edges to the spanner in Theorem 1 to get a $(k, 1 + \epsilon)$ -VFTS with hop-diameter at most $O(\alpha(m, n))$, where α is the functional inverse of Ackermann's function.*

The technique of reducing the hop-diameter is similar to that in [4]. Let H be the spanner in Theorem 1. Recall that when the points in S fail, the $(1 + \epsilon)$ -spanner path in $H \setminus S$ between any two points of x and y is the concatenation of a path P_1 in some net-tree T_1 , the cross edge $\{x', y'\}$ and a path P_2 in some net-tree T_2 . We add edges to net-trees to shortcut the paths P_1 and P_2 , and hence obtain a spanner with small hop-diameter.

Spanners with bounded degree. We also give a construction of (k, t) -VFTS with bounded maximum degree. This is achieved with a sacrifice of increasing the number of edges. The result matches the state-of-the-art result of the bounded-degree Euclidean $(k, 1 + \epsilon)$ -VFTS in [14].

Theorem 3 ($(k, 1 + \epsilon)$ -VFTS with Bounded Degree). *Let (X, d) be a doubling metric with n points and let $0 < \epsilon < 1$ be a constant. Given $1 \leq k \leq n - 2$, there exists a $(k, 1 + \epsilon)$ -VFTS with $\epsilon^{-O(\dim)} \cdot k^2 n$ edges and maximum degree $\epsilon^{-O(\dim)} \cdot k^2$.*

In [3], it is shown how to reduce the maximum degree of net-tree based spanners for doubling metrics. This is achieved by replacing some cross edges with inter-level edges. As a result, the end points of a replaced cross edge $\{u, v\}$ are connected by a path $\{u, w_1, w_2, \dots, w_i, v\}$, with approximately the same length. However, in the context of fault-tolerance, some of the w_j 's might fail, and it is unclear how to make this procedure resilient to failures.

However, we note that the degree-reduction techniques similar to those in [1, 14] can be applied. Recall that the edges of the k -VFTS in Theorem 1 can be directed such that the out-degree of each point is bounded by $O(k)$. To reduce the in-degrees, for each point x , we replace the star consisting of x and edges going into x with a k -vertex-fault-tolerant single-sink spanner that approximately preserves the distances to x , and has maximum degree $O(k)$.

We show how to construct k -vertex-fault-tolerant single-sink spanners with bounded degree for doubling metrics. The construction for those in Euclidean

space is based on Θ -graphs [1, 14]. We provide a novel technique called *ring-partition*, which can be seen as a replacement for the Θ -graph in doubling metrics. Given a specific root point from X , the metric space X is partitioned into rings centered at the root with geometrically increasing radii, and each ring is further partitioned into small clusters. From each cluster, we select some *portals*, which are connected to the root by short paths. In addition, points in each cluster are connected to the portals with short paths as well. As a result, every point's distance to the root is approximately preserved.

1.2 Preliminaries

For any positive integer m , we denote $[m] := \{1, 2, \dots, m\}$.

Throughout this paper, let (X, d) be a metric space with n points, $1 \leq k \leq n-2$ be an integer representing the number of faults allowed, and let $0 < \epsilon < \frac{1}{2}$ be a constant. Without loss of generality, we also assume that the minimum inter-point distance of X is strictly greater than 1. We denote $\Delta := \max_{x, y \in X} d(x, y)$ as the *diameter* of X .

Suppose $r > 0$. The ball of radius r centered at x is $B(x, r) := \{y \in X : d(x, y) \leq r\}$. We say that a cluster $C \subseteq X$ has radius at most r , if there exists $x \in C$ such that $C \subseteq B(x, r)$. Let $r_2 > r_1 > 0$. The ring of inner radius r_1 and outer radius r_2 centered at x is $R(x, r_1, r_2) := B(x, r_2) \setminus B(x, r_1)$.

A set $Y \subseteq X$ is an *r-cover* for X if for any point $x \in X$ there is a point $y \in Y$ such that $d(x, y) \leq r$. A set Y is an *r-packing* if for any pair of distinct points $y, y' \in Y$, it holds that $d(y, y') > r$. We say that a set $Y \subseteq X$ is an *r-net* for X if Y is both an *r-cover* for X and an *r-packing*. Note that if X is finite, an *r-net* can be constructed greedily.

By recursively applying the definition of doubling dimension, we can get the following key proposition [11].

Proposition 1 (Nets Have Small Size). *Let $R \geq 2r > 0$ and let $Y \subseteq X$ be an *r-packing* contained in a ball of radius R . Then, $|Y| \leq (\frac{R}{r})^{2\dim}$.*

2 Basic Construction of Sparse Fault-Tolerant Spanners

In this section, we extend the $O(n)$ -edge spanner construction in [4] and build a k -VFST with $O(kn)$ edges. We construct $k+1$ sequences of hierarchical nets and assign each sequence with a distinct "color". Then based on the hierarchy of each color, we extract a net-tree. We show some properties similar to those in [3], and in addition we show that the fault-tolerance property can be established.

Fault-Tolerant Hierarchical Nets. We color each point in X with one of $k+1$ colors and let X_c be the set of points with color c . For each color $c \in [k+1]$, we build a sequence of hierarchical nets of $\ell := \lceil \log_2 \Delta \rceil$ levels, $X_c = N_0^c \supseteq N_1^c \supseteq \dots \supseteq N_\ell^c$. We denote by $N_i := \cup_{c \in [k+1]} N_i^c$ the set of all level- i net points. Let $r_i := 2^i$ be the *distance scale* of level i . Fault-tolerant hierarchical nets should satisfy the following properties:

1. *Packing.* For each $0 \leq i \leq \ell$ and $c \in [k+1]$, N_i^c is an r_i -packing;
2. *Covering.* For any $1 \leq i \leq \ell$, if $x \in X$ is not a net point in N_i , then for each color $c \in [k+1]$, there exists a net point $y_c \in N_i^c$ such that $d(y_c, x) \leq r_i$.

Construction. The hierarchical nets can be constructed in a top-down approach. Initially, each N_ℓ^c consists of a distinct point in X . Note that $k \leq n-2$ and hence the initialization is well defined. Also, the single point in N_ℓ^c is colored with c and points not included in any cluster N_ℓ^c stay uncolored.

Suppose all nets on level $i+1$ have been built and we construct the level- i nets as follows. For c from 1 to $k+1$, let U_c be the set of uncolored points when we start to build N_i^c , i.e., after finishing the construction of $N_i^1, N_i^2, \dots, N_i^{c-1}$. We initialize $N_i^c := N_{i+1}^c$, and extend N_{i+1}^c to get N_i^c by greedily adding points in U_c to N_i^c such that the resulting N_i^c is an r_i -net for U_c ; we color the points in $N_i^c \cap U_c$ with color c .

Note that the packing property and the covering property follow directly from the net construction.

Fault-Tolerant Net-Trees. For each color $c \in [k+1]$, we define a *net-tree* T_c , which spans all nodes in X except nodes in the highest level N_ℓ with colors different from c . The construction is given in Algorithm 1. It follows from the construction that all internal nodes of T_c have color c , and all points excluding $N_\ell \cup X_c$ are leaves of T_c .

```

1 Initialize  $T_c$  to be the only point in  $N_\ell^c$ , which is the root of  $T_c$ ;
2 for  $i = \ell - 1$  to 0 do
3   for each point  $x \in N_i \setminus N_{i+1}$  do
4     let  $y_c \in N_{i+1}^c$  be a point such that  $d(y_c, x) \leq r_{i+1}$  (such a point exists by
5       the covering property of fault-tolerant hierarchical nets);
6       add  $x$  to  $T_c$  and set  $y_c$  as its parent in  $T_c$  by adding the edge  $\{y_c, x\}$  to
7        $T_c$ ;
8   end
9 end
10 return  $T_c$ ;

```

Algorithm 1: Construction of net-tree T_c for color c

For any $c \in [k+1]$, a path $P = \{x_0, x_1, \dots, x_i\}$ is called a c -path, if all edges on P are contained in T_c . Note that for a c -path $P = \{x_0, x_1, \dots, x_i\}$, any point x_j with $0 < j < i$ has degree at least 2 in T_c and hence is an internal node. Thus, its color must be c . The length of P is defined as $\text{length}(P) := \sum_{j=1}^i d(x_{j-1}, x_j)$. The following lemma shows that for any point $x \notin N_\ell$, any color $c \in [k+1]$ and any $0 \leq i \leq \ell$, there is a c -path from x to some node in N_i with length at most $2r_i$.

Lemma 1 (Climbing Path in Net-tree). *Let T_c be a net-tree obtained above and let $x \in X \setminus N_\ell$ be a non-root point. For any $0 \leq i \leq \ell$, there exists a c -path P_i*

starting from x and ending at some net point $x_i \in N_i$, such that $\text{length}(P_i) \leq 2 \cdot r_i$. In addition, x_i also has color c if $x_i \neq x$.

Proof. Let $i^*(x)$ be the largest i such that $x \in N_i$. For $0 \leq i \leq i^*(x)$, let $P_i = \{x\}$ and $x_i = x$. The conclusion holds trivially for P_i .

Now suppose $i^*(x) < i \leq \ell$ and we use induction on i . The base case is for $i = i^*(x) + 1$. We let x_i be x 's parent in T_c and let $P_i := \{x, x_i\}$. Note that P_i is a c -path. By the construction of T_c , we know that $x_i \in N_i^c$ and $d(x, x_i) \leq r_i$. Hence, x_i has color c and $\text{length}(P_i) = d(x, x_i) \leq r_i \leq 2 \cdot r_i$.

Suppose $i > i^*(x) + 1$ and there exists an x_{i-1} with color c , and a c -path $P_{i-1} = \{x, \dots, x_{i-1}\}$ such that $\text{length}(P_{i-1}) \leq 2 \cdot r_{i-1}$. If $x_{i-1} \in N_i$, let $x_i = x_{i-1}$ and $P_i = P_{i-1}$. In this case, we have $\text{length}(P_i) = \text{length}(P_{i-1}) \leq 2 \cdot r_{i-1} \leq 2 \cdot r_i$. Other properties of P_i follow directly from the properties of P_{i-1} , and hence P_i is a c -path.

Otherwise, $x_{i-1} \notin N_i$ and we know that x_{i-1} is not the root of T_c and let x_i be x_{i-1} 's parent in T_c . By our construction of net-trees, $x_i \in N_i^c$ and hence has color c ; in addition, $d(x_{i-1}, x_i) \leq r_i$. We let $P_i = P_{i-1} \oplus x_i$, which is formed by appending x_i to the end of P_{i-1} . Note that the edge $\{x_{i-1}, x_i\} \in T_c$ and hence $P_i \subseteq T_c$ is a c -path. Also, $\text{length}(P_i) = \text{length}(P_{i-1}) + d(x_{i-1}, x_i) \leq 2 \cdot r_{i-1} + r_i = 2 \cdot r_i$. \square

Fault-Tolerant Spanners. We have added inter-level edges in the net-trees. Now we add edges connecting net points at the same level to achieve small stretch. Define $\gamma := 4 + \frac{32}{\epsilon}$. For any $0 \leq i \leq \ell$, we call $\{x, y\}$ a *cross edge* at level i , iff x and y are both in N_i , and $d(x, y) \leq \gamma \cdot r_i$. (An edge can be a cross edge at more than one level.) We construct a spanner H by taking the union of all edges in the net-trees of all colors, and all cross edges at all levels, and claim that H is a k -VFTS with stretch at most $1 + \epsilon$.

Lemma 2 (Fault-Tolerant Stretch). *Let $S \subseteq X$ be any set with $|S| \leq k$. For any $x, y \in X \setminus S$, $d_{H \setminus S}(x, y) \leq (1 + \epsilon) \cdot d(x, y)$.*

Proof. Fix $x \neq y \in X \setminus S$ and suppose $r_i < d(x, y) \leq r_{i+1}$. Let q be some integer such that $\frac{8}{2^q} \leq \epsilon < \frac{16}{2^q}$, say $q := \lceil \log \frac{8}{\epsilon} \rceil$.

If $i \leq q - 1$, then $d(x, y) < r_{i+1} \leq 2^q < \frac{16}{\epsilon} < \gamma \cdot r_0$. Hence, $\{x, y\}$ is a cross edge at level 0 and $d_{H \setminus S}(x, y) = d(x, y)$.

Now suppose $i \geq q$ and let $j := i - q \geq 0$. Since $|S| \leq k$, there exists some $c \in [k + 1]$ such that $X_c \cap S = \emptyset$. We first show that there exist x' and y' such that $d_{H \setminus S}(x, x') \leq 2 \cdot r_j$ and $d_{H \setminus S}(y, y') \leq 2 \cdot r_j$. Note that either the node x is in N_j , or $i^*(x) < j$ and hence $x \notin N_\ell$. In the latter case, by Lemma 1, there exists an $x_j \in N_j$ with color c and a c -path P connecting x and x_j . We let $x' = x$ and $x' = x_j$ in respective cases. If $x' = x$, then $d_{H \setminus S}(x, x') = d(x, x') = 0 \leq 2 \cdot r_j$. Otherwise, since P is a c -path and x' has color c , no point on P is contained in S and hence $P \subseteq H \setminus S$. Therefore, we have $d_{H \setminus S}(x, x') \leq 2 \cdot r_j$. Similarly, we can choose some $y' \in N_j \setminus S$ such that $d_{H \setminus S}(y, y') \leq 2 \cdot r_j$.

Note that $d(x, x') \leq d_{H \setminus S}(x, x') \leq 2r_j$ and $d(y, y') \leq d_{H \setminus S}(y, y') \leq 2r_j$. Hence, we have $d(x', y') \leq d(x', x) + d(x, y) + d(y, y') \leq 2r_j + r_{i+1} + 2r_j <$

$(4 + \frac{32}{\epsilon}) \cdot r_j = \gamma \cdot r_j$. It follows that $\{x', y'\}$ is a cross edge at level j and thus $d_{H \setminus S}(x', y') = d(x', y')$.

Note that $d(x', y') \leq d(x', x) + d(x, y) + d(y, y') \leq 4r_j + d(x, y)$, and $d_{H \setminus S}(x, y) \leq d_{H \setminus S}(x, x') + d_{H \setminus S}(x', y') + d_{H \setminus S}(y', y)$. Hence we conclude that $d_{H \setminus S}(x, y) \leq 8 \cdot r_j + d(x, y) = \frac{8}{2^j} \cdot r_i + d(x, y) \leq (1 + \epsilon) \cdot d(x, y)$. \square

Remark. Note that in the proof, there exists a color c such that no node in S has color c . Also, the spanner path in $H \setminus S$ is the concatenation of a path $P_1 = \{x, \dots, x'\}$, a cross edge $\{x', y'\}$, and a path $P_2 = \{y', \dots, y\}$, such that each of P_1 and P_2 is either a c -path or a trivial path with only one point. This property is useful in our later construction of k -VFTS with small hop-diameter.

Bounding the Number of Edges. We show that the number of edges in H is $O(kn)$. We actually show a stronger result: we can direct the edges in H in a way such that every point has out-degree $O(k)$.

Lemma 3 (Bounding the Number of Edges). *Let H be the $(k, 1 + \epsilon)$ -VFTS we construct above. Then, the number of edges in H is $\epsilon^{-O(\dim)} \cdot kn$. Moreover, the edges of H can be directed such that the out-degree of each point is bounded by $\epsilon^{-O(\dim)} \cdot k$.*

Proof. Note that the edges of H come from two sources: the net-trees and cross edges; we bound them separately.

In any net-tree T_c , we direct the edge $\{x, p(x)\}$ from x to $p(x)$, where $p(x)$ is the parent of x in T_c . Note that every point has out-degree at most 1 in each tree, and hence the out-degree due to the net-tree edges is bounded by $k + 1$.

Now we bound the out-degree due to the cross edges. Recall that $i^*(x)$ is the maximum i such that $x \in N_i$. Given an edge $\{x, y\}$, we direct it from x to y if $i^*(x) < i^*(y)$, and direct it arbitrarily if $i^*(x) = i^*(y)$.

Fix $x \in X$. We bound the number of edges coming out of x and going into some point with a fixed color c . Let $i := i^*(x)$. For any directed edge (x, y) such that y has color c , we know that $i^*(y) \geq i^*(x) = i$ and hence $y \in N_i^c$. Note that the existence of a cross edge $\{x, y\}$ implies that $d(x, y) \leq \gamma \cdot r_i$. Also note that N_i^c is an r_i -packing. Then, by Proposition 1, the number of such edges is $\gamma^{O(\dim)} = (4 + \frac{32}{\epsilon})^{O(\dim)}$. Since there are at most $k + 1$ colors, the number of cross edges coming out of x is bounded by $(k + 1) \cdot (4 + \frac{32}{\epsilon})^{O(\dim)} = \epsilon^{-O(\dim)} \cdot k$.

The upper bound on the number of edges follows directly from the analysis of out-degree. \square

3 Achieving Small Hop-diameter

In this section, we show that a technique similar to that in [4] can be used to reduce the hop-diameter of our basic k -VFTS.

Let T be a tree metric with n nodes. It is shown in [5, 4] that for $m \geq 2n$, we can add m edges to T to obtain a spanner R , such that for the unique tree path P between x and y in T , there is a path P' in R that connects x and y via at most $O(\alpha(m, n))$ nodes on P (in the same order), where $\alpha(\cdot, \cdot)$ is defined below. By the triangle inequality, $\text{length}(P') \leq \text{length}(P)$.

Definition 2 (Ackermann's Function [17]). Let $A(i, j)$ be a function defined for integers $i, j \geq 0$ as the following.

$$\begin{aligned} A(0, j) &= 2j && \text{for } j \geq 0 \\ A(i, 0) &= 0, A(i, 1) = 2 && \text{for } i \geq 1 \\ A(i, j) &= A(i-1, A(i, j-1)) && \text{for } i \geq 1, j \geq 2 \end{aligned}$$

Define the function $\alpha(m, n) := \min\{i \geq 1, A(i, 4 \lceil \frac{m}{n} \rceil) > \log_2 n\}$.

Adding Edges to Reduce Hop-diameter. Let H be the $(k, 1 + \epsilon)$ -VFTS constructed in Section 2. Now we show how to add edges to H to reduce the hop-diameter. For each net-tree T_c , we use the technique in [5] to add m edges to T_c to get a spanner R_c such that between any two points in T_c , there is a path between them in R_c with $O(\alpha(m, n))$ hops which preserves their original path distance in T_c . Let H' denote the spanner constructed by taking the union of all edges in R_c 's for all colors, and all cross edges at all levels. Hence, H' has $k(m + \epsilon^{-O(\text{dim})} \cdot n)$ edges. We prove in the following lemma that H' has small hop-diameter.

Lemma 4 (Bounded Hop-Diameter). For $m \geq 2n$, the spanner H' constructed above has $k(m + \epsilon^{-O(\text{dim})} \cdot n)$ edges. Let $S \subseteq X$ be a set with $|S| \leq k$. For any pair $x, y \in X \setminus S$, there exists a path between x and y in $H' \setminus S$ with $O(\alpha(m, n))$ hops, and the path has length at most $(1 + \epsilon) \cdot d(x, y)$.

Proof. In Lemma 2, we have proved that there is a color c such that no point in S has color c , and the $(1 + \epsilon)$ -spanner path in $H \setminus S$ connecting x and y is a concatenation of a path P_1 , a cross edge and a path P_2 , where each of P_1 and P_2 is either a c -path or a trivial path consisting of only one point. Note that if P_1 is not a trivial path, it can be substituted by a path P'_1 in $R_c \setminus S \subseteq H' \setminus S$ consisting of $O(\alpha(m, n))$ hops. Similarly, P_2 can also be substituted by a path P'_2 with $O(\alpha(m, n))$ hops. The new spanner path connecting x and y in $H' \setminus S$ after the substitution has length at most $(1 + \epsilon) \cdot d(x, y)$, as $\text{length}(P'_1) \leq \text{length}(P_1)$ and $\text{length}(P'_2) \leq \text{length}(P_2)$. \square

4 Achieving Bounded Degree

4.1 Fault-tolerant Single-sink Spanners

Our technique of reducing degrees in fault-tolerant spanners is based on single-sink spanners. Given a point $v \in X$, a spanner H for X is a k -vertex-fault-tolerant v -single-sink t -spanner $((k, t, v)$ -VFTssS), if for any subset $S \subseteq X \setminus \{v\}$ with $|S| \leq k$, and any point $x \in X \setminus S$, it holds that $d_{H \setminus S}(v, x) \leq t \cdot d(v, x)$. Here, t is called the *root-stretch* of H . In this section, we show a construction of a $(k, 1 + \epsilon, v)$ -VFTssS with maximum degree $O(k)$. Throughout this section, we assume a point $v \in X$ is given. Without loss of generality, we assume $0 < \epsilon < \frac{1}{9}$ is a constant and build a $(k, 1 + 9\epsilon, v)$ -VFTssS. Our construction is based on a technique called *ring-partition*.

Ring-Partition. Let $\ell = \lceil \log_{\frac{1}{\epsilon}} \Delta \rceil$ and $r_i = \frac{1}{\epsilon^i}$ with $i \in [\ell]$. For convenience, let $r_0 = 1$ (recall that we assume inter-point distances are larger than 1). Consider the rings, denoted by R_1, \dots, R_ℓ , where $R_i := R(v, r_{i-1}, r_i)$. For convenience, let $R_0 := \{v\}$. The rings are pairwise disjoint and their union covers X . For each $i \in [\ell]$, we build an ϵr_{i-1} -net N_i for R_i . By Proposition 1, N_i contains at most $(\frac{r_i}{\epsilon r_{i-1}})^{2\dim} = \epsilon^{-4\dim}$ points. We denote this upper bound by $\Gamma := \lceil \epsilon^{-4\dim} \rceil$ and then we have $|N_i| \leq \Gamma$. Let $N := \cup_{i>0} N_i$ be the set of net points. Then, for each net point $y \in N_i$, we construct a *net cluster* C_y , such that a point $x \in X$ is in C_y iff x is in R_i , and among all points in N_i , y is the closest one to x (breaking ties arbitrarily). For each $y \in N$, we arbitrarily choose $k+1$ *portals* $Q_y \subseteq C_y$ (if $|C_y| < k+1$, we let $Q_y = C_y$). Let $Q_i := \cup_{y \in N_i} Q_y$ be the portals in R_i , and $Q := \cup_{y \in N} Q_y$ be the set of all portals. Note that $|Q_i| \leq (k+1) \cdot |N_i| \leq \Gamma \cdot (k+1)$.

We construct a $(k, 1 + 9\epsilon, v)$ -VFTssS in two stages. First, we add edges to connect portals to the root and obtain a $(k, 1 + \epsilon, v)$ -VFTssS for $Q \cup \{v\}$. Then, we add edges to connect the points in each cluster with their portals by short paths.

Connecting Portals to the Root. Assign each point $q \in Q$ a unique identifier $\text{id}(q) \in [|Q|]$, such that for any $q \in R_i$ and $q' \in R_j$ with $i < j$, it holds that $\text{id}(q) < \text{id}(q')$; also let $\text{id}(v) = 0$. In other words, points closer to the root v have smaller identifiers. We divide the points in Q into groups of size $k+1$. Specifically, let $A_j := \{q \in Q \mid (j-1) \cdot (k+1) + 1 \leq \text{id}(q) \leq j \cdot (k+1)\}$ for $j \geq 1$. The edges to connect portals with the root are added as follows.

- For $q \in Q$ with $1 \leq \text{id}(q) \leq (2\Gamma + 1) \cdot (k+1)$, add an edge $\{q, v\}$. Let E_0 denote the set of such edges.
- For $j > 2\Gamma + 1$, we add an edge between every point in A_j and every point in $A_{j-2\Gamma-1}$, and let E_j denote the set of such edges, i.e. $E_j := \{\{x, y\} \mid x \in A_j \text{ and } y \in A_{j-2\Gamma-1}\}$.

Define $\tilde{E} := E_0 \cup (\cup_{j>2\Gamma+1} E_j)$ and $\tilde{H} := (Q \cup \{v\}, \tilde{E})$. Note that the degree of v in \tilde{H} is at most $(2\Gamma + 1) \cdot (k+1) = \epsilon^{-O(\dim)} \cdot k$, and the degree of any point $q \in Q$ is at most $2(k+1) = O(k)$.

Lemma 5 (Edges Connect Portals at least Two Levels Apart). *Let $q \in Q$, and let $\{q, x\} \in \tilde{E}$ be an edge with $\text{id}(x) < \text{id}(q)$. Then, either $x = v$ or there exist $i > 2$ and $0 < i' \leq i - 2$ such that $q \in Q_i$ and $x \in Q_{i'}$.*

Proof. We first consider $\{q, x\}$ with $q \in Q_1 \cup Q_2$. Since $|Q_1| \leq \Gamma \cdot (k+1)$ and $|Q_2| \leq \Gamma \cdot (k+1)$, it holds that $\text{id}(q) \leq 2\Gamma \cdot (k+1) \leq (2\Gamma + 1) \cdot (k+1)$. Hence, q is connected to the root and $x = v$.

Now consider the case that $q \in Q_i$ for some $i > 2$. If $x = v$, then we are done. Now suppose $x \neq v$. Then, we know that $q \in A_j$ for some $j > 2\Gamma + 1$ and $x \in A_{j-2\Gamma-1}$. Note that there are exactly $2\Gamma \cdot (k+1)$ points in $A_{j-2\Gamma} \cup \dots \cup A_{j-1}$, and hence $\text{id}(x) < \text{id}(q) - 2\Gamma \cdot (k+1)$. On the other hand, since there are at most $2\Gamma \cdot (k+1)$ points in $Q_i \cup Q_{i-1}$, x cannot be in $Q_i \cup Q_{i-1}$. Hence, $x \in Q_{i'}$ for some $i' \leq i - 2$. \square

Lemma 6. *Let $S \subseteq X \setminus \{v\}$ be a set of at most k points. Then, for any point $q \in Q \setminus S$, $d_{\tilde{H} \setminus S}(v, q) \leq (1 + 3\epsilon) \cdot d(v, q)$.*

Proof. We use induction on $\text{id}(q)$. For $q \in Q \setminus S$ with $1 \leq \text{id}(q) \leq (2\Gamma+1) \cdot (k+1)$, we know that q is connected to v in \tilde{H} and hence $d_{\tilde{H}}(v, q) = d(v, q)$.

Now suppose $\text{id}(q) > (2\Gamma+1) \cdot (k+1)$ and for any point $q' \in Q \setminus S$ with $\text{id}(q') < \text{id}(q)$, it holds that $d_{\tilde{H} \setminus S}(v, q') \leq (1+3\epsilon)d(v, q')$. Let $j > 2\Gamma+1$ be such that $q \in A_j$. From the construction of \tilde{H} , we know that for any point $p \in A_{j-2\Gamma-1}$, there exists an edge $\{p, q\} \in \tilde{H}$. Since $A_{j-2\Gamma-1}$ contains $k+1$ points, there exists a functioning point $p^* \notin S$ in $A_{j-2\Gamma-1}$. Note that $\text{id}(p^*) < \text{id}(q)$. Hence, we have

$$d_{\tilde{H} \setminus S}(v, p^*) \leq (1+3\epsilon)d(v, p^*) \quad (1)$$

by induction hypothesis.

Note that $\text{id}(q) > (2\Gamma+1) \cdot (k+1)$ implies $q \in Q_i$ for some $i > 2$. By Lemma 5, we have $p^* \in Q_{i'} \subseteq R_{i'}$ for some $0 < i' \leq i-2$. Hence,

$$d(v, p^*) \leq r_{i'} = \epsilon r_{i'+1} \leq \epsilon r_{i-1} < \epsilon d(v, q) \quad (2)$$

By the triangle inequality, $d_{\tilde{H} \setminus S}(v, q) \leq d_{\tilde{H} \setminus S}(v, p^*) + d(p^*, q) \leq d_{\tilde{H} \setminus S}(v, p^*) + d(p^*, v) + d(v, q)$, which by (1) and (2) is at most $(1+\epsilon(2+3\epsilon))d(v, q) \leq (1+3\epsilon)d(v, q)$, where the last inequality holds when $\epsilon \leq \frac{1}{3}$. \square

Connecting Points in Clusters to Portals. Fix $i \in [\ell]$ and a point $y \in N_i$. Recall that C_y denotes the net cluster centered at y , whose radius is at most $r := \epsilon \cdot r_{i-1}$, and Q_y is the set of portals for C_y . We call the portals in Q_y r -portals since they are portals for clusters with radius at most r . We define a procedure $\text{Add}(C_y, Q_y, r)$ which adds edges to connect points in C_y with portals in Q_y .

1. *Sub-clustering.* We return immediately if $C_y = Q_y$. Suppose $C_y \neq Q_y$. We build an $\frac{r}{2}$ -net \hat{N} for $C_y \setminus Q_y$. Recall that r is an upper bound on C_y 's radius. By Proposition 1, $|\hat{N}| \leq 4^{\dim}$. Then for each node $z \in \hat{N}$, we construct a cluster \hat{C}_z , such that a point $x \in C_y$ is in \hat{C}_z iff x is in $C_y \setminus Q_y$, and among all points of \hat{N} , z is the closest one to x (breaking ties arbitrarily).
2. *Connecting sub-portals.* For each sub-cluster \hat{C}_z , we arbitrarily select $k+1$ sub-portals \hat{Q}_z (called $\frac{r}{2}$ -portals) in \hat{C}_z (select all points if $|\hat{C}_z| < k+1$). Then, for each sub-portal in \hat{Q}_z and each portal in Q_y , we add an edge between them. Note that since $\hat{N} \leq 4^{\dim}$, and thus each portal $q \in Q_y$ is connected with at most $4^{\dim} \cdot (k+1) = 2^{O(\dim)} \cdot k$ sub-portals.
3. *Recursion.* For every $z \in \hat{N}$, recursively call $\text{Add}(\hat{C}_z, \hat{Q}_z, \frac{r}{2})$.

Let \hat{H}_y be the resulting spanner returned by $\text{Add}(C_y, Q_y, r)$. We have the following lemma.

Lemma 7. *Let $S \subseteq X$ be a set of at most k points. For any $x \in C_y \setminus (Q_y \cup S)$ and any r -portal $q \in Q_y \setminus S$, it holds that $d_{\hat{H}_y \setminus S}(x, q) \leq 2r$.*

Proof. Suppose that x is an $\frac{r}{2^i}$ -portal. Note that x is connected to $k+1$ distinct $\frac{r}{2^{i-1}}$ -portals, and at least one of them must be functioning (i.e., not in S). We let x_{i-1} be such a $\frac{r}{2^{i-1}}$ -portal. Using this argument, we can find a sequence of portals $x = x_i, x_{i-1}, x_{i-2}, \dots, x_0 = q$, such that for all $0 \leq j \leq i$, $x_j \notin S$ is an $\frac{r}{2^j}$ -portal. In addition, for $j \in [i]$, $\{x_j, x_{j-1}\} \in \widehat{H}_y$ and $d(x_j, x_{j-1}) \leq \frac{r}{2^{j-1}}$. Hence, $d_{\widehat{H}_y \setminus S}(x, q) \leq \sum_{j=1}^i \frac{r}{2^{j-1}} \leq 2r$. \square

Obtaining the $(k, 1 + 9\epsilon, v)$ -VFTssS. Our final $(k, 1 + 9\epsilon, v)$ -VFTssS, denoted by H_v , is the union of \widehat{H} and \widehat{H}_y 's for all $y \in N$. Note that the degree of v is bounded by $\epsilon^{-O(\dim)} \cdot k$ and the degree of any other point in H_v is bounded by $2^{O(\dim)} \cdot k$. It remains to show that H_v has root-stretch at most $1 + 9\epsilon$ under the failure of at most k points.

Lemma 8. *Let $S \subseteq X \setminus \{v\}$ be a set of at most k points. For any $x \in X \setminus S$, $d_{H_v}(v, x) \leq (1 + 9\epsilon) \cdot d(v, x)$.*

Proof. Suppose $x \neq v$. Otherwise the conclusion holds trivially. Let $y \in N$ be the net point covering x , i.e., $x \in C_y$, and let R_i be the ring that contains C_y .

If x is a portal for C_y , then by Lemma 6, we know that $d_{\widehat{H} \setminus S}(v, x) \leq (1 + 3\epsilon)d(v, x)$ and hence $d_{H_v \setminus S}(v, x) \leq d_{\widehat{H} \setminus S}(v, x) \leq (1 + 9\epsilon)d(v, x)$.

Otherwise, $Q_y \neq C_y$ and hence $|Q_y| = k + 1$. Therefore, there must be some $q \in Q_y$ which is functioning. Let $r := \epsilon r_{i-1}$ be an upper bound on the radius of C_y . From the construction of the ring-partition, we know that $d(q, x) \leq r$ and $r \leq \epsilon \cdot d(v, x)$. By Lemma 7, it holds that $d_{\widehat{H}_y \setminus S}(q, x) \leq 2r$. Hence, $d_{H_v \setminus S}(v, x) \leq d_{H_v \setminus S}(v, q) + d_{H_v \setminus S}(q, x) \leq d_{\widehat{H} \setminus S}(v, q) + d_{\widehat{H}_y \setminus S}(q, x) \leq (1 + 3\epsilon)d(v, q) + 2r \leq (1 + 3\epsilon)(d(v, x) + d(x, q)) + 2r \leq (1 + 3\epsilon)d(v, x) + 6r \leq (1 + 9\epsilon)d(v, x)$. \square

4.2 $(k, 1 + \epsilon)$ -VFTS with Bounded Degree

Now we construct a $(k, 1 + \epsilon)$ -VFTS with bounded degree as follows. We first construct a basic $(k, 1 + \frac{\epsilon}{3})$ -VFTS for X with $O(kn)$ edges, and denote it by H_0 . Recall that the edges of H_0 can be directed such that the out-degree of each point in H_0 is $\epsilon^{-O(\dim)} \cdot k$. Denote an edge $\{x, y\}$ by (x, y) if it is directed from x to y in H_0 . For any point $x \in X$, we let $N_{in}(x) := \{y \in X \mid (y, x) \in H_0\}$, and build a $(k, 1 + \frac{\epsilon}{3}, x)$ -VFTssS H_x for $N_{in}(x) \cup \{x\}$. We take the spanner $H := \cup_{x \in X} H_x$, and show that H is a $(k, 1 + \epsilon)$ -VFTS with maximum degree $O(k^2)$.

Lemma 9. *For $0 < \epsilon \leq \frac{1}{2}$, H is a $(k, 1 + \epsilon)$ -VFTS, in which the degree of any point $x \in X$ is $\epsilon^{-O(\dim)} \cdot k^2$. Consequently, H has $\epsilon^{-O(\dim)} \cdot k^2 n$ edges.*

Proof. We first prove that H is a $(k, 1 + \epsilon)$ -VFTS. Let $S \subseteq X$ be a set of at most k points. Since H_0 is a $(k, 1 + \frac{\epsilon}{3})$ -VFTS, for any $x, y \in X \setminus S$, there exists a $(1 + \frac{\epsilon}{3})$ -spanner path P_0 in $H_0 \setminus S$ between x and y .

For each edge $\{u, v\} \in P_0$, suppose the edge is directed as (u, v) . Then, since H_v is a $(k, 1 + \frac{\epsilon}{3}, v)$ -VFTssS and u, v are both functioning, there is a $(1 + \frac{\epsilon}{3})$ -spanner path $P_{uv} \subseteq H_v \setminus S$ between u and v . Let P denote the concatenation of

P_{uv} 's for all edges $\{u, v\} \in P_0$. Then, P is contained in $H \setminus S$ and is a spanner path between x and y with stretch at most $(1 + \frac{\epsilon}{3})^2 \leq 1 + \epsilon$.

Next we bound the degree of an arbitrary point $x \in X$ in H . The edges incident to x in H are contained in H_x and H_y 's such that the edge $\{x, y\}$ is directed from x to y in H_0 . Note that the number of H_y 's involving x is bounded by the out-degree of x in H_0 , which is $\epsilon^{-O(\dim)} \cdot k$. Also recall that the degree of x in H_x is $\epsilon^{-O(\dim)} \cdot k$ and the degree of x in each H_y is $2^{O(\dim)} \cdot k$. Hence, we conclude that the degree of x in H is $\epsilon^{-O(\dim)} \cdot k^2$.

The upper bound on the number of edges follows directly from the degree analysis. \square

References

1. S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. H. M. Smid. Euclidean spanners: short, thin, and lanky. In *STOC*, pages 489–498, 1995.
2. P. B. Callahan and S. R. Kosaraju. Faster algorithms for some geometric graph problems in higher dimensions. In *SODA*, pages 291–300, 1993.
3. H. T.-H. Chan, A. Gupta, B. M. Maggs, and S. Zhou. On hierarchical routing in doubling metrics. In *SODA*, pages 762–771, 2005.
4. T.-H. H. Chan and A. Gupta. Small hop-diameter sparse spanners for doubling metrics. *Discrete & Computational Geometry*, 41(1):28–44, 2009.
5. B. Chazelle. Computing on a free tree via complexity-preserving mappings. *Algorithmica*, 2:337–361, 1987.
6. S. Chechik, M. Langberg, D. Peleg, and L. Roditty. Fault-tolerant spanners for general graphs. In *STOC*, pages 435–444, 2009.
7. A. Czumaj and H. Zhao. Fault-tolerant geometric spanners. *Discrete & Computational Geometry*, 32(2):207–230, 2004.
8. G. Das and G. Narasimhan. A fast algorithm for constructing sparse euclidean spanners. In *Symposium on Computational Geometry*, pages 132–139, 1994.
9. M. Dinitz and R. Krauthgamer. Fault-tolerant spanners: better and simpler. In *PODC*, pages 169–178, 2011.
10. L.-A. Gottlieb and L. Roditty. An optimal dynamic spanner for doubling metric spaces. In *ESA*, pages 478–489, 2008.
11. A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *FOCS*, pages 534–543, 2003.
12. S. Har-Peled and M. Mendel. Fast construction of nets in low dimensional metrics, and their applications. In *Symposium on Computational Geometry*, pages 150–158, 2005.
13. C. Levcopoulos, G. Narasimhan, and M. H. M. Smid. Efficient algorithms for constructing fault-tolerant geometric spanners. In *STOC*, pages 186–195, 1998.
14. T. Lukovszki. New results of fault tolerant geometric spanners. In *WADS*, pages 193–204, 1999.
15. G. Narasimhan and M. H. M. Smid. *Geometric spanner networks*. Cambridge University Press, 2007.
16. S. Solomon and M. Elkin. Balancing degree, diameter and weight in euclidean spanners. In *ESA (1)*, pages 48–59, 2010.
17. R. E. Tarjan. Efficiency of a good but not linear set union algorithm. *J. ACM*, 22(2):215–225, 1975.