
OPTIMAL $O(k)$ -ROBUST SPANNERS IN ONE DIMENSION

TBD

October 29, 2018

ABSTRACT. For any $\epsilon > 0$ and any $n \in \mathbb{N}$, we construct a graph $G = (V, E)$ with vertex set $V = \{0, \dots, n-1\}$ having $O(\epsilon^{-3} n \log n)$ edges, and such that, for any $F \subseteq V$, $G - F$ contains a monotone path of length at least $n - (1 + \epsilon)|F|$. For any constant $\epsilon > 0$, the number of edges in G matches a lower-bound of Bose et al. (2013) and improves the previous construction of Buchin, Hulshof and Oláh (2018) which has $O(n^{1+1/\ell})$ edges and guarantees a monotone path of length $n - c^\ell |F|$, for some constant $c \geq 3$ and any $\ell \in \mathbb{N}$.

1 Introduction

A geometric graph $G = (V, E)$ with vertex set $V \subset \mathbb{R}^d$ is a (geometric) t -spanner of a subset $X \subset V$ if, for every pair of vertices $v, w \in X$, $v \neq w$,

$$\frac{\|vw\|_G}{\|vw\|} \leq t, \quad (1)$$

where $\|vw\|$ denotes the Euclidean distance between v and w and $\|vw\|_G$ denote the Euclidean length of the shortest path between v and w in G , where we use the convention that $\|vw\|_G = \infty$ if v and w are in different components of G . Most of the research on spanners focuses on *sparse* spanners, where the number of edges in G is linear, or close to linear, in $|V|$. In addition to having natural applications to transportation networks, sparse t -spanners have found numerous applications in approximation algorithms and geometric data structures.

For any non-decreasing function $f: \mathbb{N} \rightarrow \mathbb{N}$, Bose et al. say that a geometric graph G is an $f(k)$ -robust t -spanner if, for every set $F \subset V(G)$, there exists a set $F^* \supseteq F$ such that $|F^*| \leq f(|F|)$ and the graph $G - F$ is a t -spanner of $V(G) \setminus F^*$. In networking applications, this definition captures the idea that the number of nodes harmed by a set faulty nodes should be bounded by a function of the number of faulty nodes.

Under this definition, the most robust spanner one could hope for would be a k -robust spanner, but it is straightforward to argue that, even for one dimensional point sets, the complete graph is the only k -robust spanner.¹ The complete graph is not sparse, and is therefore not suitable for many applications. A natural second-best option is a ck -robust spanner with a near-linear number of edges, for some $c > 1$.

Bose et al. [4, Theorem 3] show that for any constants $c, t > 1$ and any n -element set $V \subset \mathbb{R}$, any ck -robust t -spanner $G = (V, E)$ must have $\Omega(n \log n)$ edges. Their work does not give

¹Proof: Consider any pair of vertices $v, w \in V$ that are not adjacent in G and let $F = V \setminus \{v, w\}$. Then $\|vw\|_{G-F} = \infty$ so $G - F$ is not a t -spanner of $V \setminus F = V \setminus F^*$ for any $t < \infty$.




Figure 1: The binary tree T on $\{0, \dots, n-1\}$.

any construction matching this lower bound, leaving a gap in our understanding of robust spanners, even in one dimension. Buchin, Hulshof and Oláh [5] made the first progress in filling this gap by constructing, for any $\ell \in \mathbb{N}$, a $k2^{O(\ell)}$ -robust spanner 1-spanner having $O(n^{1+1/\ell})$ edges.² In the current paper, we prove the following result:

Theorem 1. *For every n -element set $V \subset \mathbb{R}$ and every $\epsilon > 0$, there exist a $(1 + \epsilon)k$ -robust 1-spanner $G = (V, E)$ with $O(\epsilon^{-3}n \log n)$ edges.*

The lower-bound discussed above implies that the number of edges in Theorem 1 is optimal, except possibly for the dependence on ϵ .

In any graph G whose vertex set is a set of real numbers, a path $x_1, \dots, x_m$ in G is *monotone* if $x_1 < \dots < x_m$. The statement “ G is a 1-spanner” is equivalent to the statement “there is a monotone path that includes all vertices of G .” Therefore, to prove Theorem 1 it is sufficient to show that we can construct a graph $G = (V, E)$ with $V = \{0, \dots, n-1\}$, having $O(\epsilon^{-3}n \log n)$ and such that, for any $F \subseteq V$, $G - F$ contains a monotone path of size at least $n - (1 + \epsilon)|F|$. From this point on, we will focus on this problem.

2 The Construction

Without loss of generality, assume n is a power of 2 and refer to Figure 1. Construct a complete ordered binary tree T of height $\log_2 n$ whose nodes are contiguous sets of integers. The leaves of T , in left-to-right order, are the singleton sets $\{0\}, \dots, \{n-1\}$ and each internal node of T is the set obtained by taking the union of its two children. Stated another way, $T = T(0, \log_2 n)$, where $T(i, b)$ has a single node $\{i\}$ if $b = 0$. Otherwise $T(i, b)$ has a root node $\{i, i+1, \dots, i+2^b-1\}$ whose two children are the subtrees $T(i, b-1)$ and $T(i+2^{b-1}, b-1)$.

Note that the nodes of T can be partitioned in levels $L_0, \dots, L_{\log_2 n}$, where each $L_i = \{u_{i,0}, \dots, u_{i,2^i-1}\}$ and

$$u_{i,k} = \{kn/2^i, \dots, (k+1)n/2^i - 1\}.$$

For each consecutive pair $(u_{i,k}, u_{i,k+1})$ of nodes in L_i we say that $u_{i,k}$ and $u_{i,k+1}$ are *buddies* and we call $u_{i,k+1}$ the *right buddy* of $u_{i,k}$. For any $k \in \{0, \dots, 2^i-1\}$, we use the notation $r(u_{i,k})$ to denote the right-buddy, $u_{i,k+1}$, of $u_{i,k}$. We also use an iterated for version of this notation: $r^{(0)}(u) = u$ and, for any $j \in \mathbb{N}$, $r^{(j)}(u) = r(r^{(j-1)}(u))$.

²As they describe it, the construction of Buchin et al. [5] is a $6^\ell k$ -robust. However, using expanders, the base 6 can be reduced to $3 + \delta$ for any $\delta > 0$ while only increasing the number of edges by a factor of $(1/\delta)^2$. This is done by replacing their matchings with the exploder graphs described in Theorem 2.

2.1 Expanders

Our construction G will consist of the union of many *exploder graphs* whose properties are given by the following definition:

Definition 1. A bipartite graph $B = (X, Y, E)$ with $|X| = |Y| = n$ is called an *order- n ϵ -exploder* for (X, Y) if, for every $X' \subseteq X$, $|X'| \geq \epsilon n/4$, $|N_B(X')| \geq (1 - \epsilon/2)n$.

Exploders are, of course, closely related to expanders [6]. However, where work on expanders typically focuses on ensuring that small subsets X' expand to a subset Y' that is a constant factor larger than X' . With exploders, we only care about large subsets X' but we wish to ensure that these subsets expand to a subset Y' that contains a very large fraction of Y .

The proof of the following theorem is standard (see, e.g., [6, Lemma 1.9]). We include it only for completeness, because it is short and it establishes the dependence between the number of edges and ϵ .

Theorem 2. For any $\epsilon > 0$ and any $n \in \mathbb{N}$, there exists an order- n ϵ -exploder having $O(n/\epsilon^2)$ edges.

Proof. Consider the random bipartite graph $B = (X, Y, E)$ in which each vertex in X chooses d random elements of Y as neighbours, with replacement. Fix a subset $X' \subset X$ of size $\epsilon n/4$ and a subset $Y' \subseteq Y$ of size $(1 - \epsilon/2)n$. The probability that all elements of X' choose all their neighbours in Y' is $(1 - \epsilon/2)^{d\epsilon n/4}$. Applying the union bound over all choices of X' and Y' we get that the probability that B is not an ϵ -exploder is at most

$$\begin{aligned}
\binom{n}{\epsilon n/4} \binom{n}{(1 - \epsilon/2)n} (1 - \epsilon/2)^{d\epsilon n/4} &\leq \left(\frac{en}{\epsilon n/4} \right)^{\epsilon n/4} \left(\frac{en}{(1 - \epsilon/2)n} \right)^{(1 - \epsilon/2)n} (1 - \epsilon/2)^{d\epsilon n/4} \\
&= \left(\frac{e^{1 - \epsilon/4}}{(\epsilon/4)^{\epsilon/4} (1 - \epsilon/2)^{1 - \epsilon/2}} \right)^n (1 - \epsilon/2)^{d\epsilon n/4} \\
&\leq \left(2e^{1 - \epsilon/4} \right)^n (1 - \epsilon/2)^{d\epsilon n/4} \\
&= \left(2e^{1 - \epsilon/4} (1 - \epsilon/2)^{d\epsilon/4} \right)^n \\
&\approx \left(2e^{1 - \epsilon/4 - \epsilon^2 d/8} \right)^n \\
&< 1
\end{aligned}$$

for $d > 16/\epsilon^2$. Therefore, we conclude that there exists an order- n ϵ -exploder with at most $16n/\epsilon^2$ edges. $\square$

Our graph G will contain many ϵ -exploders for pairs (u, w) where $u, w \in L_i$ are nodes at the same level, i , of T . Given a set $F \subseteq V$, we say that a node u of T is *sparse* with respect to F if $|u \cap F| \leq (1 - \epsilon)|u|$ and u is *dense* (with respect to F) otherwise. In other words, if u is sparse then $|u \setminus F| \geq \epsilon|u|$.

The following observation is a consequence of the definitions so far:




Figure 2: An example with $\epsilon = 1/4$. Elements in F are denoted by $\times$, dense nodes are shaded blue, sparse nodes whose right buddy is dense are shaded in pink. $F^+ = \{4, 5, 6, 7, 8, 9, 10, 11, 18, 19, 20, 22, 23, 24, 25, 29, 30\}$ includes every value in every shaded node.

Observation 1. Let n , G , F , T , and $L_0, \dots, L_{\log_2 n}$ be as defined above and assume the following:

1. $x \in \{1, \dots, n\}$ is a vertex and $u, w \in L_i$ is a pair of nodes of T ;
2. $\{x\} < u < w$;
3. G contains an ϵ -exploder, B_{uw} , for (u, w) ; and
4. $u' \subseteq u$, $|u'| \geq \epsilon|u|/4$, and, for every $x' \in u'$, $G - F$ contains a monotone path from x to x' .

Then there exists a subset $w' \subseteq w$, $|w'| \geq (1 - \epsilon/2)|u| - |F \cap w|$, and, for every $y \in w'$, $G - F$ contains a monotone path from x to y .

Proof. By the definition of ϵ -exploder, $N_{B_{uw}}(u')$ is a subset of w of size at least $(1 - \epsilon/2)|w|$. Therefore $w' = N_{B_{uw}}(u') \setminus F$ has size $|w'| \geq (1 - \epsilon/2)|w| - |F \cap w|$. Every vertex y in w' has a neighbour $x' \in u'$. Appending the edge $x'y$ to the monotone path from x to x' gives a monotone path from x to y . $\square$

2.2 A $(2 + \epsilon)k$ -Robust Construction

We now describe a first construction of G that has $O(\epsilon^{-2}n \log n)$ edges and guarantees that, for any $F \subseteq V$, $G - F$ has a monotone path of length at least $n - (2 + O(\epsilon))|F|$. Changing notation slightly, this means that, for any $\delta > 0$, G is a $(2 + \delta)k$ -robust 1-spanner with $O(\delta^{-1}n \log n)$ edges.

The graph G contains, for each pair (u, w) of buddies in T , an exploder graph B_{uw} for the pair (u, w) . The graph B_{uw} has $O(\epsilon^{-2}|u|)$ edges. Therefore, for each level L_i , the buddy pairs in L_i contribute $O(\epsilon^{-2}n)$ edges to G . Summing over all levels we see that G has $O(\epsilon^{-2}n \log n)$ edges.

Next we define a set $F^+ \supseteq F$ as follows:

$$F^+ = \cup \{u \in V(T) : u \text{ is dense or } r(u) \text{ is dense}\}.$$

Refer to Figure 2. First we show that $F^+ \supseteq F$ is not overly large:

Claim 1. $F^+ \supseteq F$ and, for any $\epsilon < 1/2$, $|F^+| \leq (2 + 4\epsilon)|F|$.

Proof. Consider the set M of *maximal* dense nodes containing exactly those nodes u of T that are dense and such that no strict ancestor of u is dense. Observe that every $x \in F$ appears in exactly one node of M . This immediately implies that $F^+ \supseteq F$.

To upper bound $|F^+|$, observe that $|F^+| \leq \sum_{u \in M} 2|u|$ and that

$$|F| = \sum_{u \in M} |F \cap u| \geq \sum_{u \in M} (1 - \epsilon)|u| ,$$

so

$$\frac{|F^+|}{|F|} \leq \frac{2}{1 - \epsilon} = 2 \sum_{r=0}^{\infty} \epsilon^r = 2(1 + \epsilon + \epsilon \sum_{r=1}^{\infty} \epsilon^r) \leq 2(1 + 2\epsilon) . \quad \square$$

Claim 2. For any $F \subset V(G)$, $G - F$ contains a monotone path of size at least $n - (2 + O(\epsilon))|F|$.

Proof. For a sparse node u of T with $x < \min(u)$ we say that a subset $u' \subset u$ is a *witness set* for (x, u) if $|u'| \geq (1 - \epsilon/2)|u| - |F \cap u|$ and, for every $x' \in u$, $G - F$ contains a monotone path from x to x' .

We will construct the path in $G - F$ in a sequence of phases. At the beginning of phase p we begin at a vertex $x_p \notin F^+$ and the phase will end when we are able to construct a monotone path to some vertex $y_p > x_p$, $y_p \notin F^+$. For the first phase, we begin with x_1 as the first vertex not in F^+ , i.e., $x_1 = \min(V(G) \setminus F^+)$. For each subsequent phase p , the starting vertex $x_p = y_{p-1}$ found in the previous phase.

The p th phase consists of a sequence of rounds. Refer to Figure 3. At the beginning of the i th round we have a triple (a_i, u_i, u'_i) where a_i is a node of T containing x_p , $u_i = r(a_i)$, and u'_i is a witness set for (x_p, u_i) . To begin phase p we set $a_0 = \{x_p\}$, $u_0 = r(a_0) = \{x_p + 1\}$, and $u'_0 = u_0$. Since $x_p \notin F^+$, neither a_0 nor u_0 is dense. Since a_0 and u_0 are buddies, G contains the exploder $B_{a_0 u_0}$ which contains the single edge with endpoints x_p and $x_p + 1$. Therefore, $u'_0 = \{x_p + 1\}$ is indeed a witness set for (x_p, u_0) .

Now consider a general round that begins with (a_i, u_i, u'_i) . Since $x_p \in a_i$ and $x_p \notin F^+$, each of a_i and u_i are sparse with respect to F . There are three cases to consider. These cases are not mutually exclusive, but when more than one applies, we consider the first applicable case.

1. $u'_i \setminus F^+$ is non-empty. In this case we take $y_p = \min(u'_i \setminus F^+)$. We have found a monotone path from x_p to $y_p \notin F^+$ so we complete the phase and begin the next phase starting from $x_{p+1} = y_p$.
2. u_i contains $n - 1$ (see round $i = 4$ in Figure 3). In this case, $r(u_i)$ does not exist. This is the final phase. We take y_p to be any element of u'_i . Since $y_p \in F^+$ we can not begin another phase and we terminate the entire procedure, having constructed a path that terminates at any node y_p .
3. a_i and u_i are siblings with a common parent a_{i+1} (see rounds $i = 1, 2$ in Figure 3). In this case, we set $u_{i+1} = r(a_{i+1})$. Notice that $u'_i \subseteq u_i \subset a_{i+1}$ and that $|u'_i| \geq \epsilon|u_i|/2 = \epsilon|a_{i+1}|/4$. Therefore, by Observation 1, $u'_{i+1} = N_{B_{a_{i+1} u_{i+1}}}(u') \setminus F$ is a witness set for (x_p, u_{i+1}) .




Figure 3: The nodes a_i , u_i , and w_i visited during phase p . The node a_0 is the leaf containing x_p .

4. Otherwise, the parent a_{i+1} of a_i and the parent u_{i+1} of u_i are distinct (see rounds $i = 0, 3$ in Figure 3). In this case u_i is the left child of u_{i+1} . Let $w_i = r(u_i)$ be the right child of u_{i+1} . Then G contains the exploder $B_{u_i w_i}$, so by Observation 1 $w'_i = N_{B_{uw}}(u') \setminus F$ is a witness set for (x_p, w) .

Now, for every $y \in u'_{i+1} = u'_i \cup w'_i$, $G - F$ contains a monotone path from x to y and

$$\begin{aligned} |u'_{i+1}| &\geq |u'_i| + |w'_i| \\ &\geq (1 - \epsilon/2)|u_i| - |F \cap u_i| + (1 - \epsilon/2)|w_i| - |F \cap w_i| \\ &= (1 - \epsilon/2)|u_{i+1}| - |F \cap u_{i+1}|. \end{aligned}$$

Therefore u'_{i+1} is a witness set for (x_p, u_{i+1}) .

This completes the description of phase p that constructs a monotone path in $G - P$ that begins at x_p and ends at y_p . We will now argue that the vast majority of vertices in $\{x_p + 1, \dots, y_p - 1\}$ are in the set F^+ . More specifically, we will show that $|F^+ \cap \{x_p + 1, \dots, y_p - 1\}| \geq (1 - O(\epsilon))(y_p - x_p - 1)$.

Consider the sequence $u_0, \dots, u_k$. This sequence begins at $u_0 = \{x_p + 1\}$ and each set in the sequence is twice as large as its predecessor. If, for some node u_i in this sequence $u_i \supset u_{i-1}$, then u_{i-1} is the left child of u_i in T and w_{i-1} is the right child of u_i . We create a new sequence $v_0, \dots, v_k$ where

$$v_i = \begin{cases} w_{i-1} & \text{if } i > 0 \text{ and } u_{i-1} \subset u_i \\ u_i & \text{otherwise} \end{cases}$$

See Figure 4. The nodes in the modified sequence $v_0, \dots, v_k$ form a partition of $\{\min(v_0), \dots, \max(v_k)\}$ and each node v_i is at most four times the size of v_{i-1} .




Figure 4: The sequence $v_1, \dots, v_4$ derived from the sequence $u_1, \dots, u_4$ in Figure 3.

In Case 3, above, the phase will terminate in round $i + 1$ if $|u_{i+1} \setminus F^+| > \epsilon|u_{i+1}|/2$. In Case 4, the phase will terminate in round $i + 1$ if $|w_i \setminus F^+| > \epsilon|w_i|/2$. This implies that, for each $i \in \{1, \dots, k-1\}$, $|v_i \setminus F^+| \leq \epsilon|v_i|/2$, otherwise the phase would have terminated in round $i < k$. Stated another way,

$$|v_i| \leq \frac{|v_i \cap F^+|}{1 - \epsilon/2} \leq (1 + \epsilon)|v_i \cap F^+| .$$

Let $\hat{v}_k = \{z \in v_k : z < y_p\}$. Observe that $\hat{v}_k$ contains at most $\epsilon|v_k|/2$ vertices that are not in F^+ , otherwise $v'_k \cap \hat{v}_k$ is non-empty, contradicting the minimality of y_p . This means

$$|\hat{v}_k| \leq |F^+ \cap \hat{v}_k| + \epsilon|v_k|/2 .$$

Now we have

$$\begin{aligned} y_p - x_p - 1 &= \sum_{i=0}^{k-1} |v_i| + |\hat{v}_k| \\ &\leq (1 + \epsilon) \sum_{i=0}^{k-1} |v_i \cap F^+| + |\hat{v}_k| \\ &\leq (1 + \epsilon) \sum_{i=0}^{k-1} |v_i \cap F^+| + |\hat{v}_k \cap F^+| + \epsilon|v_k|/2 \\ &\leq (1 + \epsilon) \sum_{i=0}^{k-1} |v_i \cap F^+| + |\hat{v}_k \cap F^+| + 2\epsilon|v_{k-1}| \\ &\leq (1 + \epsilon) \sum_{i=0}^{k-1} |v_i \cap F^+| + |\hat{v}_k \cap F^+| + (2\epsilon + 2\epsilon^2)|v_{k-1} \cap F^+| \\ &\leq (1 + 3\epsilon + \epsilon^2) \sum_{i=0}^{k-1} |v_i \cap F^+| + |\hat{v}_k \cap F^+| \\ &< (1 + 4\epsilon) \sum_{i=0}^{k-1} |v_i \cap F^+| + |\hat{v}_k \cap F^+| \\ &= (1 + 4\epsilon)|\{z \in F^+ : x_p < z < y_p\}| \end{aligned}$$

Suppose the procedure terminates after $s - 1$ phases, so that $x_1 < \dots < x_{s-1}$ are the vertices that begin each phase. For convenience we define $x_s = y_{s-1}$, $x_0 = -1$ and $x_{s+1} = n$.

For each $p \in \{0, \dots, s\}$, let $F_p^+ = \{z \in F^+ : x_p < z < x_{p+1}\}$. Since none of $x_0, \dots, x_{s+1}$ are in F^+ , $F_0^+, \dots, F_s^+$ is a partition of F^+ .

For each $p \in \{1, \dots, s\}$, the argument above shows that $(1 + 4\epsilon)|F_p^+| \geq x_{p+1} - x_p - 1$. Furthermore, the minimality of x_1 ensures that $(1 + 4\epsilon)|F_0^+| \geq |F_0^+| = x_1 - x_0 - 1$. Therefore,

$$\begin{aligned} (1 + 4\epsilon)|F^+| &= (1 + 4\epsilon) \sum_{p=0}^s |F_p^+| \\ &\geq \sum_{p=0}^s (x_{p+1} - x_p - 1) \\ &= n - s. \end{aligned}$$

so

$$\begin{aligned} s &\geq n - (1 + 4\epsilon)|F^+| \\ &\geq n - (1 + 4\epsilon)(2 + 4\epsilon)|F| \\ &= n - (2 + O(\epsilon))|F|. \end{aligned}$$

Since s is a lower bound on the size of the monotone path in $G - S$ constructed by this procedure, this completes the proof. $\square$

3 Achieving $1 + \epsilon$

Next, we show how to generalize our construction so that, for any $F \subset \{1, \dots, n\}$, $G - F$ contains a monotone path of size $n - (1 + O(\epsilon))|F|$. To achieve this, we add additional exploder graphs. For each node u of T , and each $i \in \{1, \dots, q\}$ such that $r^{(i)}(u)$ exists, G contains the edges of an exploder graph for the pair $(u, r^{(i)}(u))$. Observe that this construction, with $q = 1$, corresponds exactly to the graph G described in the previous section. In this generalized construction, the number of edges is $O(q\epsilon^{-2}n \log n)$ edges.

Given a set $F \subseteq V$, we define F^+ as follows: The vertices in each node u of T are all included in F^+ if

1. u is dense with respect to F ; or
2. $r^{(i)}(u)$ exists and is dense with respect to F for all $i \in \{1, \dots, q\}$.

For any maximal node u that is not dense but satisfies Condition 2 above, the nodes $r^{(1)}(u), \dots, r^{(q)}(u)$ all exist and are dense. It follows that

$$\begin{aligned} |F^+| &\leq \left(\frac{1 + 1/q}{1 - \epsilon} \right) |F| \\ &\leq (1 + 1/q)(1 + 2\epsilon)|F| && \text{(for any } \epsilon \leq 1/2) \\ &\leq (1 + 4\epsilon)|F| && \text{(for } q = \lceil 1/\epsilon \rceil). \end{aligned}$$

Claim 3. For any $F \subseteq V(G)$, $G - F$ contains a monotone path of size $n - (1 + O(\epsilon))|F|$.

Proof. This proof is similar to the proof of Claim 2 so we only focus on the differences.

The path is constructed in phases and each phase is made up of several rounds. Phase p begins at a vertex x_p and ends when we find a monotone path from x_p to some vertex $y_p \notin F^+$. During each round, we consider the unique node a_i of T that has size 2^i and that contains x_p . For Round 0 we begin at the unique leaf $\{x_p\}$ of T that contains x_p .

In addition to a_i we will consider the nodes $v_{i,1}, \dots, v_{i,d_i}$ where $v_{i,j} = r^{(j)}(a_i)$ and d_i is the smallest integer such that $r^{(d_i)}(a_i)$ is sparse. If no such integer exists, the procedure for finding the monotone path has completed and finishes at x_p . To remain consistent with notation in the proof of Claim 2, we define $u_i = v_{i,d_i}$.

For each round i , we will maintain the following *witness invariant*: Each node $v_{i,j}$, $j \in \{1, \dots, d_i\}$ has a subset $v'_{i,j} \subseteq v_{i,j}$ that is a witness set for $(x_p, v_{i,j})$.

To proceed from round i to round $i + 1$ we consider the first of following cases that is applicable:

1. $u'_i \setminus F^+$ is non-empty. In this case, we take $y_p = \min(u'_i \setminus F^+)$. We have found a monotone path from x_p to y_p and we terminate the phase. The next phase begins at $x_{p+1} = y_p$.
2. Each of $r^{(1)}(a_{i+1}), \dots, r^{(q)}(a_{i+1})$ is non-existent or dense. In this case, the procedure terminates. This is the final phase and the path we have constructed ends at x_p .
3. u_i is the right child of u_{i+1} . This case is easy. For each node $v \in \{v_{i+1,1}, \dots, v_{i+1,d_{i+1}}\}$, the two children of v are contained in $\{v_{i,1}, \dots, v_{i,d_i}\}$. For each such node v , the witness invariant ensures that v 's two children each have witness sets. The union of these two witness sets is a witness set for v . Thus we are able proceed from round i to round $i + 1$ while maintaining the witness invariant.
4. u_i is the left child of u_{i+1} . Let $w_i = r(u_i)$ be the right child of u_{i+1} . Then G contains the exploder graph B_{u_i, w_i} . Since u_i is sparse and has a witness set u'_i , $N_{B_{u_i, w_i}}(u'_i) \subset w_i$ is a witness set for w_i . Now we proceed as in the previous case since, for each node $v \in \{v_{i+1,1}, \dots, v_{i+1,d_{i+1}}\}$, the two children of v are contained in $\{v_{i,1}, \dots, v_{i,d_i}\} \cup \{w_i\}$.
5. Otherwise, u_i is not a child of u_{i+1} . In this case, select $j \in \{1, \dots, d_{i+1} - 1\}$ so that $v_{i+1,j}$ is the parent of u_i . By using one of the two preceding arguments, we conclude that each of $v_{i,1}, \dots, v_{i,j}$ has a witness set. All that remains is to show the existence of a witness set for each $v \in \{v_{i,j+1}, \dots, v_{i+1,d_{i+1}}\}$. To see why this is the case, observe that the witness set u'_i is a subset of $v_{i,j}$ of size at least $\epsilon|u_i|/2 \geq \epsilon|v_{i,j}|/4$. Now, G contains the exploder graph $B_{v_{i,j}, v}$. Therefore, by Observation 1, $N_{B_{v_{i,j}, v}}(u'_i)$ is a witness set for v . Thus, we are able to maintain the witness invariant when proceeding from round i to round $i + 1$.

This completes the description of how to construct the path from x_p to $y_p = x_{p+1}$ in phase p of the procedure. The remainder of this proof proceeds almost exactly as the proof

of Claim 2. In particular, the same arguments used in the proof of Claim 2 show that there are nodes $v_1, \dots, v_k$ and a partial node $\hat{v}_k = \{z \in v_k : z < y_p\}$ such that (i) $v_1, \dots, v_{r-1}, \hat{v}_k$ is a partition of $\{x_p + 1, \dots, y_p - 1\}$; (ii) for each $i \in \{1, \dots, r-1\}$, $|v_{i+1}| \leq 4|v_i|$ and $|v_i \cap F^+| \geq (1 - O(\epsilon))|v_i|$; and (iii) $|\hat{v}_k \setminus F^+| \leq \epsilon|v_k|/2$. This establishes that $y_p - x_p - 1 \leq (1 - O(\epsilon))|\{z \in F^+ : x_p < z < y_p\}|$. At this point (after taking special care of the first and final phase), the same double-counting argument establishes that the number, s , of phases is at least $n - (1 + O(\epsilon))|F^+| = n - (1 + O(\epsilon))|F|$. Since the monotone path we construct contains all the vertices $x_1, \dots, x_s$, this completes the proof. $\square$

4 Remarks

Claim 3 and the fact that the graph G in Section ?? has $O(q\epsilon^{-2}n \log n) = O(\epsilon^{-3}n \log n)$ completes the proof of Theorem 1.

Except possibly for the dependence on ϵ , this work settles the question of $f(k)$ -robust spanners in 1-dimension. The obvious extension of this work is to 2-dimensions. There, larger gaps remain between the upper and lower bounds. Indeed, for point sets in $\mathbb{R}^2$, no non-trivial constructions of $f(k)$ -robust t -spanners exist for any constant $t > 1$ and any $f(k) \in o(k^2)$. We conjecture that, for every constant $t > 1$, every n -point set $V \subset \mathbb{R}^2$ has an $O(k)$ -robust t -spanner with $O(n \log n)$ edges. More specifically, we conjecture that if one considers a semi-separated pair decomposition $\{(A_i, B_i) : i \in \{1, \dots, r\}\}$ [2, 1, 3] with an appropriate separation parameter s (that depends on t) and constructs a graph G that contains an ϵ -exploder for each pair (A_i, B_i) then the resulting graph is an $O(k)$ -robust t -spanner.

References

- [1] Mohammad Ali Abam, Paz Carmi, Mohammad Farshi, and Michiel H. M. Smid. On the power of the semi-separated pair decomposition. *Comput. Geom.*, 46(6):631–639, 2013.
- [2] Mohammad Ali Abam, Mark de Berg, Mohammad Farshi, and Joachim Gudmundsson. Region-fault tolerant geometric spanners. *Discrete & Computational Geometry*, 41(4):556–582, 2009.
- [3] Mohammad Ali Abam, Mark de Berg, Mohammad Farshi, Joachim Gudmundsson, and Michiel H. M. Smid. Geometric spanners for weighted point sets. *Algorithmica*, 61(1):207–225, 2011.
- [4] Prosenjit Bose, Vida Dujmović, Pat Morin, and Michiel H. M. Smid. Robust geometric spanners. *SIAM J. Comput.*, 42(4):1720–1736, 2013.
- [5] Kevin Buchin, Tim Hulshof, and Dániel Oláh. $\mathcal{O}(k)$ -robust spanners in one dimension. *CoRR*, abs/1803.08719, 2018.
- [6] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43:439–561, 2006.