A Spanner for the Day After

Kevin Buchin

Department of Mathematics and Computing Science, TU Eindhoven, The Netherlands
k.a.buchin@tue.nl

Sariel Har-Peled

Department of Computer Science, University of Illinois at Urbana-Champaign, USA
sariel@illinois.edu

Dániel Oláh

Department of Mathematics and Computing Science, TU Eindhoven, The Netherlands
d.olah@tue.nl

Abstract

We show how to construct $(1+\varepsilon)$-spanner over a set P of n points in $\mathbb{R}^d$ that is resilient to a catastrophic failure of nodes. Specifically, for prescribed parameters $\vartheta, \varepsilon \in (0,1)$, the computed spanner G has $\mathcal{O}\big(\varepsilon^{-7d} \log^7 \varepsilon^{-1} \cdot \vartheta^{-6} n \log n (\log\log n)^6\big)$ edges. Furthermore, for *any* k, and *any* deleted set $B \subseteq P$ of k points, the residual graph $G \setminus B$ is $(1+\varepsilon)$-spanner for all the points of P except for $(1+\vartheta)k$ of them. No previous constructions, beyond the trivial clique with $\mathcal{O}(n^2)$ edges, were known such that only a tiny additional fraction (i.e., ϑ) lose their distance preserving connectivity.

Our construction works by first solving the exact problem in one dimension, and then showing a surprisingly simple and elegant construction in higher dimensions, that uses the one dimensional construction in a black box fashion.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Geometric spanners, vertex failures, robustness

Digital Object Identifier 10.4230/LIPIcs.SoCG.2019.19

Related Version A full version of the paper is available at `http://arxiv.org/abs/1811.06898`.

Funding *Sariel Har-Peled*: Work on this paper was partially supported by a NSF AF awards CCF-1421231.
Dániel Oláh: Supported by the Netherlands Organisation for Scientific Research (NWO) through Gravitation-grant NETWORKS-024.002.003.

1 Introduction

Spanners. A Euclidean graph is a graph whose vertices are points in $\mathbb{R}^d$ and the edges are weighted by the Euclidean distance between their endpoints. Let $G = (P, E)$ be a Euclidean graph and $p, q \in P$ be two vertices of G. For a parameter $t \geq 1$, a path between p and q in G is a ***t*-path** if the length of the path is at most $t \|p - q\|$, where $\|p - q\|$ is the Euclidean distance between p and q. The graph G is a ***t*-spanner** of P if there is a t-path between any pair of points $p, q \in P$. Throughout the paper, n denotes the cardinality of the point set P, unless stated otherwise. We denote the length of the shortest path between $p, q \in P$ in the graph G by $\mathsf{d}(p, q)$.

Spanners have been studied extensively. The main goal in spanner constructions is to have small *size*, that is, to use as few edges as possible. Other desirable properties are low degrees [2, 10, 18], low weight [6, 12], low diameter [3, 4] or to be resistant against failures. The book by Narasimhan and Smid [17] gives a comprehensive overview of spanners.

Robustness. In this paper, our goal is to construct spanners that are robust according to the notion introduced by Bose *et al.* [7]. Intuitively, a spanner is robust if the deletion of k vertices only harms a few other vertices. Formally, a graph G is an $f(k)$-robust t-spanner, for some positive monotone function f, if for any set B of k vertices deleted in the graph, the remaining graph $G \setminus B$ is still a t-spanner for at least $n - f(k)$ of the vertices. Note, that the graph $G \setminus B$ has $n - k$ vertices – namely, there are at most $\mathcal{L}(k) = f(k) - k$ additional vertices that no longer have good connectivity to the remaining graph. The quantity $\mathcal{L}(k)$ is the ***loss***. We are interested in minimizing the loss.

The natural question is how many edges are needed to achieve a certain robustness (since the clique has the desired property). That is, for a given parameter t and function f, what is the minimal size that is needed to obtain an $f(k)$-robust t-spanner on any set of n points.

A priori it is not clear that such a sparse graph should exist (for t a constant) for a point set in $\mathbb{R}^d$, since the robustness property looks quite strong. Surprisingly, Bose *et al.* [7] showed that one can construct a $\mathcal{O}(k^2)$-robust $\mathcal{O}(1)$-spanner with $\mathcal{O}(n \log n)$ edges. Bose *et al.* [7] proved various other bounds in the same vein on the size for one-dimensional and higher-dimensional point sets. Their most closely related result is that for the one-dimensional point set $P = \{1, 2, \ldots, n\}$ and for any $t \geq 1$ at least $\Omega(n \log n)$ edges are needed to construct an $\mathcal{O}(k)$-robust t-spanner.

An open problem left by Bose *et al.* [7] is the construction of $\mathcal{O}(k)$-robust spanners – they only provide the easy upper bound of $\mathcal{O}(n^2)$ for this case.

ϑ-reliable spanners. We are interested in building spanners where the loss is only fractional. Specifically, given a parameter ϑ, we consider the function $f(k) = (1 + \vartheta)k$. The loss in this case is $\mathcal{L}(k) = f(k) - k = \vartheta k$. A $(1 + \vartheta)k$-robust t-spanner is ***ϑ-reliable t-spanner***.

Exact reliable spanners. If the input point set is in one dimension, then one can easily construct a 1-spanner for the points, which means that the exact distances between points on the line are preserved by the spanner. This of course can be done easily by connecting the points from left to right. It becomes significantly more challenging to construct such an exact spanner that is reliable.

Fault tolerant spanners. Robustness is not the only definition that captures the resistance of a spanner network against vertex failures. A closely related notion is fault tolerance [13, 14, 15]. A graph $G = (P, E)$ is an *r-fault tolerant t-spanner* if for any set B of failed vertices with $|B| \leq r$, the graph $G \setminus B$ is still a t-spanner. The disadvantage of r-fault tolerance is that each vertex must have degree at least $r + 1$, otherwise the vertex can be isolated by deleting its neighbors. Therefore, the graph has size at least $\Omega(rn)$. There are constructions that show $\mathcal{O}(rn)$ edges are enough to build r-fault tolerant spanners. However, depending on the chosen value r the size can be too large.

In particular, fault tolerant spanners cannot have a near-linear number of edges, and still withstand a widespread failure of nodes. Specifically, if a fault tolerant spanner has m edges, then it can withstand a failure of at most $2m/n$ vertices. In sharp contrast, ϑ-reliable spanners can withstand a widespread failure. Indeed, a ϑ-reliable spanner can withstand a failure of close to $n/(1 + \vartheta)$ of its vertices, and still have some vertices that are connected by short paths in the remaining graph.

1.1 Our results

In this paper, we investigate how to construct reliable spanners with very small loss – that is ϑ-reliable spanners. To the best of our knowledge nothing was known on this case before this work. For omitted proofs we refer the reader to the full version of the paper [8].

(a) **Expanders are reliable.** Intuitively, a constant degree expander is a robust/reliable graph under a weaker notion of robustness – that is, connectivity. As such, for a parameter $\vartheta > 0$, we show that constant degree expanders are indeed ϑ-reliable in the sense that all except a small fraction of the points stay connected. Formally, one can build a graph G with $\mathcal{O}(\vartheta^{-3}n)$ edges, such that for any failure set B of k vertices, the graph $G \setminus B$ has a connected component of size at least $n-(1+\vartheta)k$. We emphasize, however, that distances are not being preserved in this case. See Lemma 6 for the result.

(b) **Exact $\mathcal{O}(1)$-reliable spanner in one dimension.** Inspired by the reliability of constant degree expanders, we show how to construct an $\mathcal{O}(1)$-reliable exact spanner on any one-dimensional set of n points with $\mathcal{O}(n \log n)$ edges.[1] The idea of the construction is to build a binary tree over the points, and to build bipartite expanders between certain subsets of nodes in the same layer. One can think of this construction as building different layers of expanders for different resolutions. The construction is described in Section 3.2. See Theorem 12 for the result.

(c) **Exact ϑ-reliable spanner in one dimension.** One can get added redundancy by systematically shifting the layers. Done carefully, this results in a ϑ-reliable exact spanner. The construction is described in Section 3.3. See Theorem 13 for the result.

(d) **ϑ-reliable $(1+\varepsilon)$-spanners in higher dimensions.** We next show a *surprisingly simple and elegant* construction of ϑ-reliable spanners in two and higher dimensions, using a recent result of Chan *et al.* [11], which show that one needs to maintain only a "few" linear orders. This immediately reduces the d dimensional problem to maintaining a reliable spanner for each of this orderings, which is the problem we already solved. See Section 4 for details.

(e) **ϑ-reliable $(1+\varepsilon)$)-spanner in $\mathbb{R}^d$ with bounded spread.** Since both general constructions in $\mathbb{R}^d$ have some additional polylog factors that seems unnecessary, we present a better construction for the bounded spread case. Specifically, for points with spread Φ in $\mathbb{R}^d$, and for any $\varepsilon > 0$, we construct a ϑ-reliable $(1+\varepsilon)$-spanner with $\mathcal{O}\big(\varepsilon^{-d}\vartheta^{-2} n \log \Phi\big)$ edges. The basic idea is to construct a well-separated pair decomposition (WSPD) directly on the quadtree of the point set, and convert every pair in the WSPD into a reliable graph using a bipartite expander. The union of these graphs is the required reliable spanner. See Section 5 and Lemma 21 for details.

Shadow. Underlying our construction is the notion of identifying the points that loose connectivity when the failure set is removed. Intuitively, a point is in the shadow if it is surrounded by failed points. We believe that this concept is of independent interest – see Section 3.1 for details and relevant results in one dimension and the full version [8] for an additional result in higher dimensions.

Independently, Bose *et al.* [5] also obtained an upper bound on the size of reliable spanners in $\mathbb{R}^d$. Their construction has $\mathcal{O}(n \log^2 n \log\log n)$ edges, which is close to our bound of $\mathcal{O}(n \log n (\log\log n)^6)$ edges.

[1] This also improves an earlier preliminary construction by (some of) the authors `arXiv:1803.08719`.

2 Preliminaries

2.1 Problem definition and notations

Let $[n]$ denote the set $\{1, 2, \ldots, n\}$ and let $[i : j] = \{i, i+1, \ldots, j\}$.

▶ **Definition 1** (Robust spanner). *Let $G = (P, E)$ be a t-spanner for some $t \geq 1$ and let $f \colon \mathbb{N} \to \mathbb{R}_+$, and two point sets $P_1, P_2 \subseteq P$. The graph G is an **$f(k)$-robust t-spanner for $P_1 \oplus P_2$** if for any set of (failed) vertices $B \subseteq P$ there exists a set $B^+ \supseteq B$ with $|B^+| \leq f(|B|)$ such that the subgraph*

$$G \setminus B = G_{P \setminus B} = \Big(P \setminus B, \big\{uv \in E(G) \bigm| u, v \in P \setminus B\big\}\Big)$$

*induced by $P \setminus B$ is a t-spanner for $(P_1 \setminus B^+) \oplus (P_2 \setminus B^+)$. That is, $G \setminus B$ has a t-path between all pairs of points $p \in P_1 \setminus B^+$ and $q \in P_2 \setminus B^+$. If $P_1 = P_2 = P$, then G is a **$f(k)$-robust t-spanner**.*

*The vertices of $B^+ \setminus B$ are the vertices **harmed** by B, and the quantity $\mathcal{L}(k) = f(k) - k \geq |B^+| - |B|$ is the **loss**.*

▶ **Definition 2.** *For a parameter $\vartheta > 0$, a graph G that is $(1+\vartheta)k$-robust t-spanner is a **ϑ-reliable t-spanner**.*

▶ **Definition 3.** *For a number $x > 0$, let $\mathrm{pow}_2(x) = 2^{\lceil \log x \rceil}$ be the smallest number that is a power of 2 and is at least as large as x.*

2.2 Expander construction

For a set X of vertices in a graph $G = (V, E)$, let $\Gamma(X) = \big\{v \in V \bigm| uv \in E \text{ for a } u \in X\big\}$ be the ***neighbors*** of X in G. The following lemma, which is a standard expander construction (see e.g. [16, Section 5.3]), provides the main building block of our one-dimensional construction.

▶ **Lemma 4.** *Let L, R be two disjoint sets, with a total of n elements, and let $\xi \in (0, 1)$ be a parameter. One can build a bipartite graph $G = (L \cup R, E)$ with $\mathcal{O}(n/\xi^2)$ edges, such that*

(i) *for any subset $X \subseteq L$, with $|X| \geq \xi|L|$, we have that $|\Gamma(X)| > (1-\xi)|R|$, and*

(ii) *for any subset $Y \subseteq R$, with $|Y| \geq \xi|R|$, we have that $|\Gamma(Y)| > (1-\xi)|L|$.*

2.3 Expanders are reliable

Let P be a set with n elements, and let $\vartheta \in (0, 1)$ be a parameter. We next build a constant degree expander graph on P and show that it is ϑ-reliable. The following two lemmas are not surprising if one is familiar with expanders and their properties.

▶ **Lemma 5.** *Let n be a positive integer number, let $\alpha > 1$ be an integer constant, and let $\beta \in (0, 1)$ be some constant. One can build a graph $G = ([n], E)$, such that for all sets $X \subset V$, we have that $|\Gamma(X)| \geq \min\big((1-\beta)n, \alpha\,|X|\big)$. The graph G has $\mathcal{O}((\alpha/\beta)n)$ edges.*

▶ **Lemma 6.** *Let n and $\vartheta \in (0, 1/2)$ be parameters. One can build a graph $G = ([n], E)$ with $\mathcal{O}(\vartheta^{-3}n)$ edges, such that for any set $B \subseteq [n]$, we have that $G \setminus B$ has a connected component of size at least $n - (1+\vartheta)\,|B|$. That is, the graph G is ϑ-reliable.*

3 Building reliable spanners in one dimension

3.1 Bounding the size of the shadow

Our purpose is to build a reliable 1-spanner in one dimension. Intuitively, a point in $[n]$ is in trouble, if many of its close by neighbors belong to the failure set B. Such an element is in the shadow of B, defined formally next.

▶ **Definition 7.** *Consider an arbitrary set $B \subseteq [n]$ and a parameter $\alpha \in (0,1)$. A number i is in the* **left α-shadow** *of B, if and only if there exists an integer $j \geq i$, such that $\left|[i:j] \cap B\right| \geq \alpha \left|[i:j]\right|$. Similarly, i is in the* **right α-shadow** *of B, if and only if there exists an integer i, such that $h \leq i$ and $|[h:i] \cap B| \geq \alpha\,|[h:i]|$. The left and right α-shadow of B is denoted by $\mathcal{S}_{\rightarrow}(B)$ and $\mathcal{S}_{\leftarrow}(B)$, respectively. The combined shadow is denoted by $\mathcal{S}(\alpha, B) = \mathcal{S}_{\rightarrow}(B) \cup \mathcal{S}_{\leftarrow}(B)$.*

▶ **Lemma 8.** *Fix a set $B \subseteq [n]$ and let $\alpha \in (0,1)$ be a parameter. Then, we have that $|\mathcal{S}_{\rightarrow}(B)| \leq (1 + \lceil 1/\alpha \rceil)\,|B|$. In particular, the size of $\mathcal{S}(\alpha, B)$ is at most $2(1 + \lceil 1/\alpha \rceil)\,|B|$.*

Lemma 8 is somewhat restrictive because the shadow is at least twice larger than the failure set B. Intuitively, as $\alpha \rightarrow 1$, the shadow should converge to B. The following lemma, which is a variant of Lemma 8 quantify this.

▶ **Lemma 9.** *Fix a set $B \subseteq [n]$, let $\alpha \in (2/3, 1)$ be a parameter, and let $\mathcal{S}(\alpha, B)$ be the set of elements in the α-shadow of B. We have that $|\mathcal{S}(\alpha, B)| \leq |B| / (2\alpha - 1)$.*

Proof. Let $c = 1 - 1/\alpha < 0$. For $i = 1, \ldots, n$, let $x_i = c$ if $i \in B$, and $x_i = 1$ otherwise. For any interval I of length Δ, with $\tau\Delta$ elements in B, such that $x(I) = \sum_{i \in I} x_i \leq 0$, we have that

$$\begin{aligned} x(I) \leq 0 &\iff (1-\tau)\Delta + c\tau\Delta \leq 0 \iff 1 - \tau \leq -\tau c \iff 1/\tau \leq 1 - c \\ &\iff 1/\tau \leq 1 - (1 - 1/\alpha) \iff 1/\tau \leq 1/\alpha \iff \tau \geq \alpha. \end{aligned}$$

An element $j \in [n]$ is in the left α-shadow of B if and only if there exists an integer j', such that $|[j:j'] \cap B| \geq \alpha\,|[j:j']|$ and, by the above, $x\big([j:j']\big) \leq 0$. Namely, an integer j in the left α-shadow of B corresponds to some prefix sum of the x_is that starts at j and add up to some non-positive sum. From this point on, we work with the sequence of numbers $x_1, \ldots, x_n$, using the above summation criterion to detect the elements in the left α-shadow.

For a location $j \in [n]$ that is in the left α-shadow, let $W_j = [j:j']$ be the *witness interval* for j – this is the shortest interval that has a non-positive sum that starts at j. Let $I = W_k = [k:k']$ be the shortest witness interval, for any number in $\mathcal{S}(\alpha, B) \setminus B$. For any $j \in [k+1:k']$, we have $x\big([k:j-1]\big) + x([j:k']) = x([k:k']) \leq 0$. Thus, if $x_j = 1$, this implies that either j or k have shorter witness intervals than I, which is a contradiction to the choice of k. We conclude that $x_j < 0$ for all $j \in [k+1:k']$, that is, $[k+1:k'] \subseteq B$.

Letting $\ell = |I| = k' - k + 1$, we have that $(\ell - 1)/\ell \geq \alpha \iff \ell - 1 \geq \alpha\ell \iff \ell \geq 1/(1-\alpha) \iff \ell \geq \lceil 1/(1-\alpha) \rceil \geq 3$, as $\alpha \geq 2/3$. In particular, by the minimality of I, it follows that $\ell = \lceil 1/(1-\alpha) \rceil$.

Let $J = [k:k'-1] \subset I$. We have that $x(J) > 0$. For any $j \in \mathcal{S}(\alpha, B) \setminus B$, such that $j \neq k$, consider the witness interval W_j. If $j > k$, then $j > k'$, as all the elements of I, except k, are in B. If $j < k$ and $j' \in J$, then $\tau = x([k:j']) > 0$, which implies that $x\big([j:k-1]\big) = x\big(W_j\big) - \tau < 0$, but this is a contradiction to the definition of W_j. Namely, all the witness intervals either avoids J, or contain it in their interior. Given a witness interval W_j, such that $J \subset W_j$, we have $x(W_j \setminus J) = x(W_j) - x(J) < x(W_j) \leq 0$, since $x(J) > 0$.

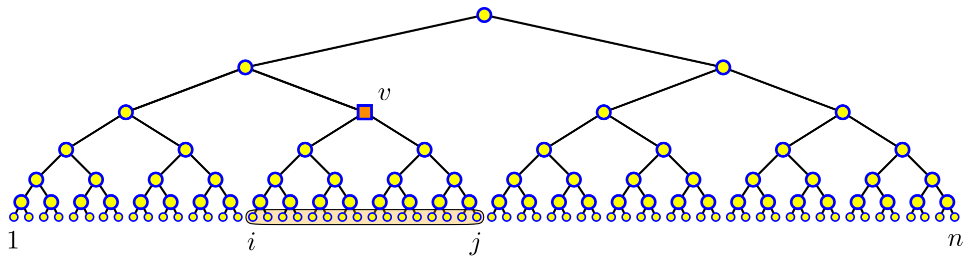

Figure 1 The binary tree built over $[n]$. The block of node v is the interval $[i : j]$.

So consider the new sequence of numbers $x_{[n]\setminus J} = x_1, \ldots, x_{k-1}, x_{k'}, \ldots x_n$ resulting from removing the elements that corresponds to J from the sequence. Reclassify which elements are in the left shadow in the new sequence. By the above, any element that was in the shadow before, is going to be in the new shadow. As such, one can charge the element k, that is in the left shadow (but not in B), to all the other elements of J (that are all in B). Applying this charging scheme inductively, charges all the elements in the left shadow (that are not in B) to elements in B. We conclude that the number of elements in the left shadow of B, that are not in B is bounded by

$$\frac{|B|}{|J|-1} = \frac{|B|}{\ell-2} = \frac{|B|}{\lceil 1/(1-\alpha)\rceil - 2} \leq \frac{1-\alpha}{1-2(1-\alpha)}\,|B| = \frac{1-\alpha}{2\alpha-1}\,|B|\,.$$

The above argument can be applied symmetrically to the right shadow. We conclude that

$$|\mathcal{S}(\alpha, B)| \leq |B| + 2\frac{1-\alpha}{2\alpha-1}\,|B| = \frac{2\alpha-1+2-2\alpha}{2\alpha-1}\,|B| = \frac{|B|}{2\alpha-1}. \qquad ◀$$

3.2 Construction of $\mathcal{O}(1)$-reliable exact spanners in one dimension

3.2.1 Constructing the graph H

Assume n is a power of two, and consider building the natural full binary tree T with the numbers of $[n]$ as the leaves. Every node v of T corresponds to an interval of numbers of the form $[i : j]$ its canonical interval, which we refer to as the block of v, see Figure 1. Let $\mathcal{I}$ be the resulting set of all blocks. In each level one can sort the blocks of the tree from left to right. Two adjacent blocks of the same level are neighbors. For a block $I \in \mathcal{I}$, let $\mathrm{next}(I)$ and $\mathrm{prev}(I)$ be the blocks (in the same level) directly to the right and left of I, respectively.

We build the graph of Lemma 4 with $\xi = 1/16$ for any two neighboring blocks in $\mathcal{I}$. Let H be the resulting graph when taking the union over all the sets of edges generated by the above.

3.2.2 Analysis

In the following we show that the resulting graph H is an $\mathcal{O}(k)$-robust 1-spanner and has $\mathcal{O}(n \log n)$ edges. We start by verifying the size of the graph.

▶ **Lemma 10.** *The graph H has $\mathcal{O}(n \log n)$ edges.*

Proof. Let $h = \log n$ be the depth of the tree T. In each level $i = 1, 2, \ldots, h$ of T there are 2^{h-i} nodes and the blocks of these nodes have size 2^i. The number of pairs of adjacent blocks in level i is $2^{h-i} - 1$ and each pair contributes $\mathcal{O}(2^i)$ edges. Therefore, each level of T contributes $\mathcal{O}(n)$ edges. We get $\mathcal{O}(n \log n)$ for the overall size by summing up for all levels. ◀

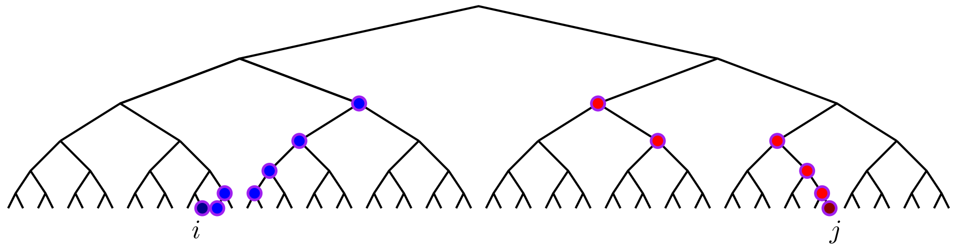

(a) The canonical path in the tree.

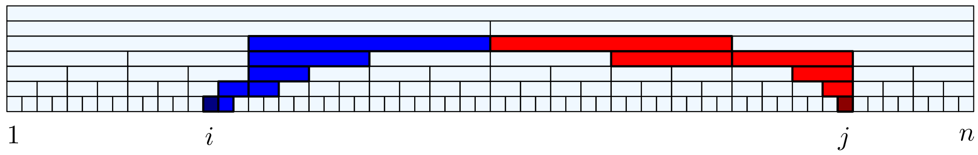

(b) The canonical path on the blocks.

Figure 2 The canonical path between the vertices i and j in two different representations. The blue nodes and blocks correspond to the ascent part and the red nodes and blocks correspond to the descent part of the walk.

There is a natural path between two leaves in the tree T, described above, going through their lowest common ancestor. However, for our purposes we need something somewhat different – intuitively because we only want to move forward in the 1-path.

Given two numbers i and j, where $i < j$, consider the two blocks $I, J \in \mathcal{I}$ that correspond to the two numbers at the bottom level. Set $I_0 = I$, and $J_0 = J$. We now describe a canonical walk from I to J, where initially $\ell = 0$. During the walk we have two active blocks I_ℓ and J_ℓ, that are both in the same level. For any block $I \in \mathcal{I}$ we denote its parent by $p(I)$. At every iteration we bring the two active blocks closer to each other by moving up in the tree.

Specifically, repeatedly do the following:

(a) If I_ℓ and J_ℓ are neighbors then the walk is done.

(b) If I_ℓ is the right child of $p(I_\ell)$, then set $I_{\ell+1} = \mathrm{next}(I_\ell)$ and $J_{\ell+1} = J_\ell$, and continue to the next iteration.

(c) If J_ℓ is the left child of $p(J_\ell)$, then set $I_{\ell+1} = I_\ell$ and $J_{\ell+1} = \mathrm{prev}(J_\ell)$, and continue to the next iteration.

(d) Otherwise – the algorithm ascends. It sets $I_{\ell+1} = p(I_\ell)$, and $I_{\ell+1} = p(J_\ell)$, and it continues to the next iteration.

It is easy to verify that this walk is well defined, and let

$$\pi(i,j) \equiv \underbrace{I_0 \to I_1 \to \cdots \to I_\ell}_{\text{ASCENT}} \to \underbrace{J_\ell \to \cdots \to J_0}_{\text{DESCENT}}$$

be the resulting walk on the blocks where we removed repeated blocks. Figure 2 illustrates the path of blocks between two vertices i and j.

In the following, consider a fixed set $B \subseteq [n]$ of faulty nodes. A block $I \in \mathcal{I}$ is ***α-contaminated***, for some $\alpha \in (0,1)$, if $|I \cap B| \geq \alpha\,|I|$.

▶ **Lemma 11.** *Consider two nodes $i, j \in [n]$, with $i < j$, and let $\pi(i,j)$ be the canonical path between i and j. If any block of $\pi = \pi(i,j)$ is α-contaminated, then i or j are in the $\alpha/3$-shadow of B.*

Proof. Assume the contamination happens in the left half of the path, i.e., at some block I_t, during the ascent from i to the connecting block to the descent path into j. By construction, there could be only one block before I_t on the path of the same level, and all previous blocks are smaller, and there are at most two blocks at each level. Furthermore, for two consecutive I_j, I_{j+1} that are blocks of different levels, $I_j \subseteq I_{j+1}$. Thus we have that either $i \in I_t$, or $i \in \mathrm{prev}(I_t)$, or $i \in \mathrm{prev}(\mathrm{prev}(I_t))$, since there are at most $|I_t| + |I_t|/2 + \cdots + 2 + 1 = 2|I_t| - 1$ vertices that are contained in the path before the block I_t. Notice that if $i \in I_t$, then it is the leftmost point of I_t.

In particular, let r be the maximum number in I_t, and observe $|[i:r] \cap B| \geq \alpha |I_t| \geq (\alpha/3)|[i:r]|$. Thus, the number i is the $\alpha/3$-shadow, as claimed.

The other case, when the contamination happens in the right part during the descent, is handled symmetrically. ◀

▶ **Theorem 12.** *The graph H constructed above on the set $[n]$ is an $\mathcal{O}(1)$-reliable exact spanner and has $\mathcal{O}(n \log n)$ edges.*

Proof. The size is proved in Lemma 10. Let $\alpha = 1/32$. Let B^+ be the set of vertices that are in the $\alpha/3$-shadow of B, that is, $B^+ = \mathcal{S}(\alpha/3, B)$. By Lemma 8 we have that $|B^+| \leq 2(1 + \lceil 3/\alpha \rceil)|B| \leq 200|B|$.

Consider any two vertices $i, j \in [n] \setminus B^+$. Let $\pi(i,j)$ be the canonical path between i and j. None of the blocks in this path are α-contaminated, by Lemma 11.

Let $\mathcal{S}$ be the set of all vertices that have a 1-path from i to them. Consider the ascent part of the path $\pi(i,j)\colon I_0 \to I_1 \to \cdots \to I_\ell$. The claim is that for every block I_t in this path, we have that at least $\frac{3}{4}$ of the vertices have 1-paths from i (i.e., $|I_t \cap \mathcal{S}| \geq \frac{3}{4}|I_t|$).

This claim is proven by induction. The claim trivially holds for I_0. Now, consider two consecutive blocks $I_t \to I_{t+1}$. There are two cases:

(i) $I_{t+1} = \mathrm{next}(I_t)$. Then, the graph H includes the expander graph on I_t, I_{t+1} described in Lemma 4. At least $\frac{3}{4}|I_t|$ vertices of I_t are in $\mathcal{S}$. As such, at least $\frac{15}{16}|I_{t+1}|$ vertices of I_{t+1} are reachable from the vertices of I_t. Since I_{t+1} is not α-contaminated, at most an α-fraction of vertices of I_{t+1} are in B, and it follows that $|I_{t+1} \cap \mathcal{S}| \geq (\frac{15}{16} - \alpha)|I_{t+1}| \geq \frac{3}{4}|I_{t+1}|$, as claimed.

(ii) I_{t+1} is the parent of I_t. In this case, I_t is the left child of I_{t+1}. Let I'_t be the right child of I_{t+1}. Since I_{t+1} is not α-contaminated, we have that $|I_{t+1} \cap B| \leq \alpha |I_{t+1}|$. As such,

$$|I'_t \cap B| \leq |I_{t+1} \cap B| \leq 2\alpha |I'_t|$$

Now, by the expander construction on (I_t, I'_t), and arguing as above, we have

$$|I'_t \cap \mathcal{S}| \geq \left(\frac{15}{16} - 2\alpha\right)|I'_t| \geq \frac{3}{4}|I'_t|,$$

which implies that $|I_{t+1} \cap \mathcal{S}| \geq \frac{3}{4}|I_{t+1}|$.

The symmetric claim for the descent part of the path is handled in a similar fashion, therefore, at least $\frac{3}{4}$ of the points in J_ℓ can reach j with a 1-path. Using these and the expander construction between I_ℓ and J_ℓ, we conclude that there is a 1-path from i to j in $H \setminus B$, as claimed. ◀

Note that it is easy to generalize the construction for arbitrary n. Let h be an integer such that $2^{h-1} < n < 2^h$ and build the graph H on $\{1, 2, 3, \ldots, 2^h\}$. Since H is a 1-spanner, the 1-paths between any pair of vertices of $[n]$ does not use any vertices from $\{n+1, \ldots, 2^h\}$. Therefore, we can simply delete the part of H that is beyond n to obtain an $\mathcal{O}(1)$-reliable 1-spanner on $[n]$. Since we defined B^+ to be the shadow of B, the $\mathcal{O}(1)$-reliability is inherited automatically.

We also note that no effort was made to optimize the constants in the above construction.

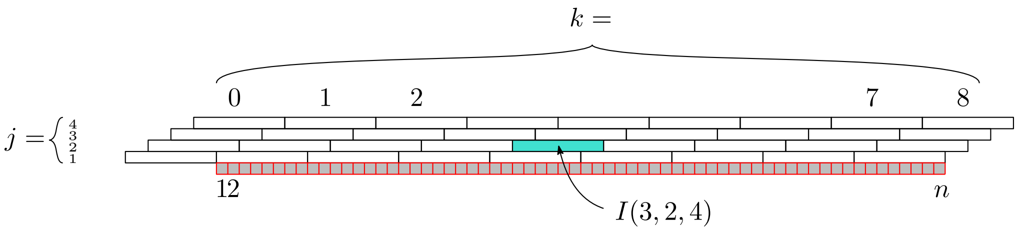

Figure 3 The shifted intervals $I(i, \cdot, \cdot)$ for $i = 3$ with $N = 4$ and $n = 64$. Each interval has length $2^i = 8$, there are $N = 4$ different shifts and there are $\frac{n}{2^i} + 1 = 9$ blocks per each shift.

3.3 Construction of ϑ-reliable exact spanners in one dimension

Here, we show how to extend Theorem 12, to build a one dimensional graph, such that for any fixed $\vartheta > 0$ and any set B of k deleted vertices, at most $(1 + \vartheta)k$ vertices are no longer connected (by a 1-path) after the removal of B. The basic idea is to retrace the construction of Theorem 12, and extend it to this more challenging case. The main new ingredient is a shifting scheme.

Let $[n]$ be the ground set, and assume that n is a power of two, and let $h = \log n$. Let

$$N = \mathrm{pow}_2\left(c/\vartheta^2\right) \qquad \text{and} \qquad \xi = \frac{1}{32N}, \tag{1}$$

where c is a sufficiently large constant ($c \geq 512$). We first connect any $i \in [n]$, to all the vertices that are in distance at most $3N$ from it, by adding an edge between the two vertices. Let G_0 be the resulting graph.

Let $i_0 = \log N$. For $i = i_0, \ldots, h - 1$, and $j = 1, \ldots, N$, let

$$\Delta(i, j) = 1 + (j - 1)2^i/N - 2^i$$

be the *shift* corresponding to i and j. For a fixed i, the $\Delta(i, j)$s are N equally spaced numbers in the block $\left[1 - 2^i : 1 - 2^i/N\right]$, starting at its left endpoint. For $k = 0, \ldots, n/2^i$, let

$$\mathtt{I}(i, j, k) = \left[\Delta(i, j) + k2^i : \Delta(i, j) + (k + 1)2^i\right]$$

be the shifted interval of length 2^i that starts at $\Delta(i, j)$ and is shifted k blocks to the right, see Figure 3. The set of all intervals of interest is

$$\mathcal{I} = \left\{ \mathtt{I}(i, j, k) \;\middle|\; \begin{array}{c} i = i_0, \ldots, \log n \\ j = 1, \ldots, N \\ k = 0, \ldots, n/2^i \end{array} \right\}.$$

Constructing the graph H_ϑ. Let $\mathrm{G_E}(i, j, k)$ denote the expander graph of Lemma 4, constructed over $\mathtt{I}(i, j, k)$ and $\mathtt{I}(i, j, k + 1)$, with the value of the parameter ξ as specified in Eq. (1). We define H_ϑ to be the union of all the graphs $\mathrm{G_E}$ over all choices of i, j, k, and also including the graph G_0 (described above). In the case that n is not a power of two, do the construction on $[\mathrm{pow}_2(n)]$. In any case, the last step is to delete vertices from H_ϑ that are outside the range of interest $[n]$.

▶ **Theorem 13.** *For parameters n and $\vartheta > 0$, the graph H_ϑ constructed over $[n]$, is a ϑ-reliable exact spanner. Furthermore, H_ϑ has $\mathcal{O}(\vartheta^{-6} n \log n)$ edges.*

4 Building a reliable spanner in $\mathbb{R}^d$

In the following, we assume that $P \subseteq [0,1)^d$ – this can be done by an appropriate scaling and translation of space. We use a novel result of Chan *et al.* [11], called locality-sensitive orderings. These orderings can be thought as an alternative to quadtrees and related structures. For an ordering σ of $[0,1)^d$, and two points $p, q \in [0,1)^d$, such that $p \prec q$, let $(p,q)_\sigma = \{z \in [0,1)^d \mid p \prec z \prec q\}$ be the set of points between p and q in the order σ.

▶ **Theorem 14** ([11]). *For $\varsigma \in (0,1)$ fixed, there is a set $\Pi^+(\varsigma)$ of $M(\varsigma) = \mathcal{O}(\varsigma^{-d} \log \varsigma^{-1})$ orderings of $[0,1)^d$, such that for any two (distinct) points $p, q \in [0,1)^d$, with $\ell = \|p - q\|$, there is an ordering $\sigma \in \Pi^+$, and a point $z \in [0,1)^d$, such that*

(i) $p \prec_\sigma q$,
(ii) $(p,z)_\sigma \subseteq \mathrm{ball}(p, \varsigma\ell)$,
(iii) $(z,q)_\sigma \subseteq \mathrm{ball}(q, \varsigma\ell)$, *and*
(iv) $z \in \mathrm{ball}(p, \varsigma\ell)$ *or* $z \in \mathrm{ball}(q, \varsigma\ell)$.

Furthermore, given such an ordering σ, and two points p, q, one can compute their ordering, according to σ, using $\mathcal{O}(d \log \varsigma^{-1})$ arithmetic and bitwise-logical operations.

4.1 Construction

Given a set P of n points in $[0,1)^d$, and parameters $\varepsilon, \vartheta \in (0,1)$, let $\varsigma = \varepsilon/c$,

$$M = M(\varsigma) = \mathcal{O}(\varsigma^{-d} \log \varsigma^{-1}) = \mathcal{O}\big(\varepsilon^{-d} \log \varepsilon^{-1}\big),$$

and c be some sufficiently large constant. Next, let $\vartheta' = \vartheta/(3N \cdot M)$ where $N = \lceil \log\log n \rceil + 1$, and let $\Pi^+ = \Pi^+(\varsigma)$ be the set of orderings of Theorem 14. For each ordering $\sigma \in \Pi^+$, compute the ϑ'-reliable exact spanner G_σ of P, see Theorem 13, according to σ. Let G be the resulting graph by taking the union of G_σ for all $\sigma \in \Pi^+$.

4.2 Analysis

▶ **Theorem 15.** *The graph G, constructed above, is a ϑ-reliable $(1+\varepsilon)$-spanner and has size*

$$\mathcal{O}\left(\varepsilon^{-7d} \log^7 \frac{1}{\varepsilon} \cdot \vartheta^{-6} n \log n (\log\log n)^6\right).$$

Proof. First, we show the bound on the size. There are M different orderings for which we build the graph of Theorem 13. Each of these graphs has $\mathcal{O}\big((\vartheta')^{-6} n \log n\big)$ edges. Therefore, the size of G is

$$M \cdot \mathcal{O}\big((\vartheta')^{-6} n \log n\big) = \mathcal{O}\left(M \left(\frac{3NM}{\vartheta}\right)^6 n \log n\right) = \mathcal{O}\left(\varepsilon^{-7d} \log^7 \frac{1}{\varepsilon} \cdot \vartheta^{-6} n \log n (\log\log n)^6\right).$$

Next, we identify the set of harmed vertices B^+ given a set of failed vertices $B \subseteq P$. First, let B_1 be the union of all the $(1 - \vartheta'/4)$-shadows resulting from B in G_σ, for all $\sigma \in \Pi^+$. Then, for $i = 2, \ldots, N$, we define B_i in a recursive manner to be the union of all the $(1 - \vartheta'/4)$-shadows resulting from B_{i-1} in G_σ, for all $\sigma \in \Pi^+$. We set $B^+ = B_N$.

By the recursion and Lemma 9 we have that

$$\begin{aligned}
|B_i| &\le \left(\frac{|B_{i-1}|}{2(1 - \frac{\vartheta'}{4}) - 1} - |B_{i-1}|\right) M + |B_{i-1}| = \frac{|B_{i-1}| - (1 - \frac{\vartheta'}{2}) |B_{i-1}|}{1 - \frac{\vartheta'}{2}} M + |B_{i-1}| \\
&= \frac{\vartheta' |B_{i-1}|}{2 - \vartheta'} M + |B_{i-1}| \le (1 + \vartheta' M) |B_{i-1}| = \left(1 + \frac{\vartheta}{3N}\right) |B_{i-1}|.
\end{aligned}$$

Therefore, we obtain

$$\left|B^{+}\right| = |B_N| \le \left(1 + \frac{\vartheta}{3N}\right)^N |B| \le \exp\left(N\frac{\vartheta}{3N}\right)|B| \le (1+\vartheta)\,|B|\,,$$

using $1 + x \le e^x \le 1 + 3x$, for $x \in [0,1]$.

Now we show that there is a $(1+\varepsilon)$-path $\hat{\pi}$ between any pair of vertices $p, q \in P \setminus B^{+} \equiv P \setminus B_N$ such that the path $\hat{\pi}$ does not use any vertices of B. By Theorem 13, the graph $G_\sigma \setminus B_{N-1} \subseteq G \setminus B_{N-1}$ contains a monotone path connecting p to q, according to σ. Observe that there is a unique edge (p', q') on this path that "jumps" from the locality of p to the locality of q. Formally, we have the following:

(a) $\|p' - q'\| \le \|p - q\| + 2\varsigma \|p - q\| = (1 + 2\varepsilon/c)\,\|p - q\|\,.$

(b) $\|p - p'\| \le 2\varsigma \|p - q\| = 2(\varepsilon/c)\,\|p - q\|$ and similarly $\|q - q'\| \le 2(\varepsilon/c)\,\|p - q\|$.

(c) $p', q' \in P \setminus B_{N-1}$.

We fix the edge (p', q') to be used in the computed path $\hat{\pi}$ connecting p to q. We still need to build the two parts of the path $\hat{\pi}$ between p, p' and q, q'.

This procedure reduced the problem of computing a reliable path between two points $p, q \in P \setminus B_N$, into computing two such paths between two points of $P \setminus B_{N-1}$ (i.e., p, p' and q, q'). The benefit here is that $\|p - p'\|, \|q - q'\| \ll \|p - q\|$. We now repeat this refinement process $N - 1$ times.

To this end, let Q_i be the set of active pairs that needs to be connected in the ith level of the recursion. Thus, we have $Q_0 = \{(p,q)\}$, $Q_1 = \{(p,p'), (q,q')\}$, and so on. We repeat this $N - 1$ times. In the ith level there are $|Q_i| = 2^i$ active pairs. Let $(x, y) \in Q_i$ be such a pair. Then, there is an edge (x', y') in the graph $G \setminus B_{N-(i+1)}$, such that we have the following:

(a) $\|x' - y'\| \le \|x - y\|\,(1 + 2\varepsilon/c) \le (2\varepsilon/c)^i(1 + 2\varepsilon/c)\,\|p - q\|\,.$

(b) $\|x - x'\| \le 2(\varepsilon/c)\,\|x - y\| \le (2\varepsilon/c)^{i+1}\,\|p - q\|$ and $\|y - y'\| \le (2\varepsilon/c)^{i+1}\,\|p - q\|$.

(c) $x', y' \in P \setminus B_{N-(i+1)}$.

The edge (x', y') is added to the path $\hat{\pi}$. After $N - 1$ iterations the set of active pairs is Q_{N-1} and for each pair $(x, y) \in Q_{N-1}$ we have that $x, y \in P \setminus B_1$. For each of these pairs $(x, y) \in Q_{N-1}$ we apply Theorem 14 and Theorem 13 to obtain a path of length at most $\|x - y\|\, 2\log n$ between x and y (and this subpath of course does not contain any vertex in B). We add all these subpaths to connect the active pairs in the path $\hat{\pi}$, which completes $\hat{\pi}$ into a path.

Now, we bound the length of path $\hat{\pi}$. Since, for all $(x, y) \in Q_{N-1}$, we have $\|x - y\| \le \|p - q\| \cdot (2\varepsilon/c)^{N-1}$ and $|Q_{N-1}| = 2^{N-1}$, the total length of the subpaths calculated, in the last step, is

$$\begin{aligned}
\sum_{(x,y)\in Q_{N-1}} \operatorname{length}\bigl(\hat{\pi}[x,y]\bigr) &\le 2^{N-1}\,\|p - q\| \cdot \left(\frac{2\varepsilon}{c}\right)^{N-1} 2\log n \\
&\le \|p - q\| \cdot \left(\frac{4\varepsilon}{c}\right)^{\log\log n} 2\log n \le \|p - q\| \cdot \varepsilon^{\log\log n}\left(\frac{4}{c}\right)^{\log\log n} 2\log n \\
&\le \|p - q\| \cdot \frac{\varepsilon}{4} \cdot \frac{1}{\log n} \cdot 2\log n \le \frac{\varepsilon}{2}\,\|p - q\|\,,
\end{aligned}$$

for large enough n and $c \ge 8$. The total length of the long edges added to $\hat{\pi}$ in the recursion, is bounded by

$$\sum_{i=0}^{N-2} 2^i \|p-q\| \left(\frac{2\varepsilon}{c}\right)^i \left(1+\frac{2\varepsilon}{c}\right) \leq \|p-q\| \left(1+\frac{2\varepsilon}{c}\right) \sum_{i=0}^{\infty} \left(\frac{4\varepsilon}{c}\right)^i$$
$$= \|p-q\| \left(1+\frac{2\varepsilon}{c}\right) \frac{1}{1-4\varepsilon/c} = \|p-q\| \left(1+\frac{6\varepsilon}{c-4\varepsilon}\right) \leq \left(1+\frac{\varepsilon}{2}\right) \|p-q\|,$$

which holds for $c \geq 16$. Therefore, the computed path $\hat{\pi}$ between p and q is a $(1+\varepsilon)$-path in $G \setminus B$, which concludes the proof of the theorem. ◀

5 Construction for points with bounded spread in $\mathbb{R}^d$

The input is again a set $P \subset \mathbb{R}^d$ of n points, and parameters $\vartheta \in (0, 1/2)$ and $\varepsilon \in (0, 1)$. The goal is to build a ϑ-reliable $(1+\varepsilon)$-spanner on P that has optimal size under some condition on P. The condition is that the spread $\Phi(P)$ is bounded by a polynomial of n. The construction is based on well-separated pair decompositions (WSPD), which was introduced by Callahan and Kosaraju [9]. For preliminaries see the full version [8].

5.1 The construction of G_Φ

First, compute a quadtree T for the point set P. For any node $v \in T$, let $\Box_v$ denote the ***cell*** (i.e. square or cube, depending on the dimension) represented by v. Let $P_v = \Box_v \cap P$ be the point set stored in the subtree of v. Compute a $(6/\varepsilon)$-WSPD $\mathcal{W}$ over T for P using the construction of Abam and Har-Peled [1, Lemma 2.8]. The pairs in $\mathcal{W}$ can be represented by pairs of nodes $\{u, v\}$ of the quadtree T. Note that the algorithm uses the diameters and distances of the cells of the quadtree, that is, for a pair $\{u, v\} \in \mathcal{W}$, we have

$$(6/\varepsilon) \cdot \max\bigl(\operatorname{diam}(\Box_u), \operatorname{diam}(\Box_v)\bigr) \leq \mathsf{d}(\Box_u, \Box_v).$$

For any pair $\{u, v\} \in \mathcal{W}$, we build the bipartite expander of Lemma 4 on the sets P_u and P_v such that the expander property holds with $\xi = \vartheta/8$. Furthermore, for every two node u and v that have the same parent in the quadtree T we add the edges of the bipartite expander of Lemma 4 between P_u and P_v. Let G_Φ be the resulting graph when taking the union over all the sets of edges generated by the above.

5.2 Analysis

▶ **Lemma 16.** *The graph G_Φ has $\mathcal{O}\bigl(\xi^{-2}\varepsilon^{-d} n \log \Phi(P)\bigr)$ edges.*

Proof. By Lemma 5.4 of [8], every point participates in $\mathcal{O}(\varepsilon^{-d} \log \Phi(P))$ WSPD pairs. By Lemma 4 the average degree in all the expanders is at most $\mathcal{O}(1/\xi^2)$, resulting in the given bound on the number of edges. There are also the additional pairs between a node in T and its parent, but since every point participates in only $\mathcal{O}(\log \Phi(P))$ such pairs, the number of edges is dominated by the expanders on the WSPD pairs. It follows that the number of edges in the resulting graph is $\mathcal{O}(\xi^{-2}\varepsilon^{-d} n \log \Phi(P))$. ◀

▶ **Definition 17.** *For a number $\gamma \in (0, 1)$, and failed set of vertices $B \subseteq P$, a node v of the quadtree T is in the* γ-shadow *if $|B \cap P_v| \geq \gamma |P_v|$. Naturally, if v is in the γ-shadow, then the points of P_v are also in the shadow. As such, the* γ**-shadow** *of B is the set of all the points in the shadow – formally, $\mathcal{S}(\gamma, B) = \bigcup_{v \in T:\ |B \cap P_v| \geq \gamma |P_v|} P_v$.*

Let $\gamma = 1 - \vartheta/2$. Note that $B \subseteq \mathcal{S}(\gamma, B)$, since every point of B is stored as a singleton in a leaf of T.

▶ **Definition 18.** *For a node x in T, let $n(x) = |P_x|$, and $b(x) = |P_x \cap B|$.*

▶ **Lemma 19.** *Let $\gamma = 1 - \vartheta/2$ and $B \subseteq P$ be fixed. Then, the size of the γ-shadow of B is at most $(1+\vartheta)\,|B|$.*

Proof. Let H be the set of nodes of T that are in the γ-shadow of B. A node $u \in H$ is *maximal* if none of its ancestors is in H. Let $H' = \{u_1, \ldots, u_m\}$ be the set of all maximal nodes in H, and observe that $\cup_{u \in H'} P_u = \cup_{v \in H} P_v = \mathcal{S}(\gamma, B)$. For any two nodes $x, y \in H'$, we have $P_x \cap P_y = \varnothing$. Therefore, we have

$$|B| = \sum_{u \in H'} b(u) \geq \sum_{u \in H'} \gamma n(u) = \gamma\,|\mathcal{S}(\gamma, B)|\,.$$

Dividing both sides by γ implies the claim, since $1/\gamma = 1/(1-\vartheta/2) \leq 1+\vartheta$. ◀

▶ **Lemma 20.** *Let $\gamma = 1 - \vartheta/2$. Fix a node $u \in T$ of the quadtree, the failure set $B \subseteq P$, its shadow $B^+ = \mathcal{S}(\gamma, P)$, and the residual graph $H = G_\Phi \setminus B$. For a point $p \in P_u \setminus B^+$, let $R_u(p)$ be the set of all reachable points in P_u with stretch two, formally, $R_u(p) = \left\{q \in P_u \setminus B \,\middle|\, \mathsf{d}_H(p,q) \leq 2 \cdot \operatorname{diam}(\square_u)\right\}$. Then, we have $|R_u(p)| \geq 3\xi\,|P_u|$.*

Proof. Let $u_1, u_2, \ldots, u_j = u$ be the sequence of nodes in the quadtree from the leaf u_1 that contains (only) p, to the node u. A *level* of a point $q \in P_u$, denoted by $\ell(q)$, is the first index i, such that $q \in P_{u_i}$. A *skipping path* in G_Φ, is a sequence of edges $pq_1, q_1q_2, \ldots q_{m-1}q_m$, such that $\ell(q_i) < \ell(q_{i+1})$, for all i.

Let Q_i be the set of all points in $P_{u_i} \setminus B$ that are reachable by a skipping path in H from p. We claim, for $i = 1, \ldots, j$, that

$$|Q_i| \geq (1-\xi)n(u_i) - b(u_i) \geq (1-\xi-\gamma)n(u_i) = (\vartheta/2 - \xi)n(u_i) = 3\xi n(u_i),$$

since $\xi = \vartheta/8$ and p is not in the γ-shadow. The claim clearly holds for u_1. So, assume inductively that the claim holds for $u_1, \ldots, u_{j-1}$. Let $v_1, \ldots, v_m$ be the children of u_j that have points stored in them (excluding u_{j-1}). There is an expander between $P_{u_{j-1}}$ and P_{v_i}, for all i, as a subgraph of G_Φ. It follows, by induction, that

$$\begin{aligned}|Q_j| &\geq (1-\xi)n(u_{j-1}) - b(u_{j-1}) + \sum_i \big((1-\xi)n(v_i) - b(v_i)\big)\\ &= (1-\xi)n(u_{j-1}) + \sum_i (1-\xi)n(v_i) - \Big(b(u_{j-1}) + \sum_i b(v_i)\Big) = (1-\xi)n(u_j) - b(u_j).\end{aligned}$$

Observe that a skipping path from p to $q \in P_{u_j}$ has length at most

$$\sum_{i=1}^{j} \operatorname{diam}(\square_{u_i}) \leq \operatorname{diam}\big(\square_{u_j}\big) \sum_{i=1}^{j} 2^{1-j} \leq 2 \cdot \operatorname{diam}\big(\square_{u_j}\big).$$

Thus, $Q_j \subseteq R_u(p)$, and the claim follows. ◀

Now we are ready to prove that G_Φ is a reliable spanner.

▶ **Lemma 21.** *For a set $P \subseteq \mathbb{R}^d$ of n points and parameters $\varepsilon \in (0,1)$ and $\vartheta \in (0,1/2)$, the graph G_Φ is a ϑ-reliable $(1+\varepsilon)$-spanner with $\mathcal{O}\big(\varepsilon^{-d}\vartheta^{-2} n \log \Phi(P)\big)$ edges, where $\Phi(P)$ is the spread of P.*

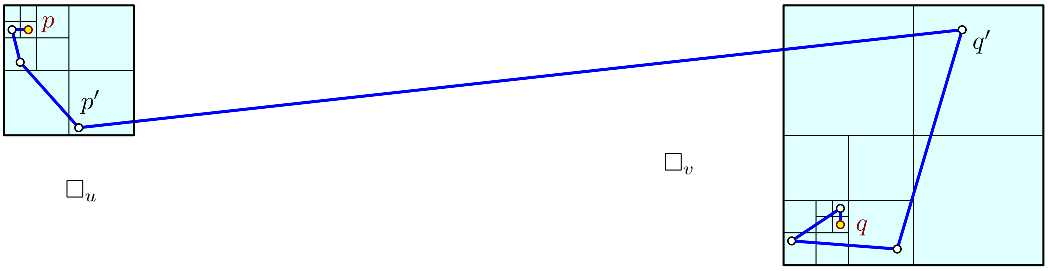

Figure 4 The pair $\{u, v\} \in \mathcal{W}$ that separates p and q. The blue path is a $(1+\varepsilon)$-path between p and q in the graph $G_\Phi \setminus B$.

Proof. Let $\xi = \vartheta/8$ and $\gamma = 1 - \vartheta/2$. The bound on the number of edges follows by Lemma 16.

Let B be a set of faulty vertices of G_Φ, and let $H = G_\Phi \setminus B$ be the residual graph. We define B^+ to contain the vertices that are in the γ-shadow of B. Then, we have $B \subseteq B^+$ and $|B^+| \leq (1+\vartheta)\,|B|$ by Lemma 19. Finally, we need to show that there exists a $(1+\varepsilon)$-path between any $p, q \in P \setminus B^+$.

Let $\{u, v\} \in \mathcal{W}$ be the pair that separates p and q with $p \in P_u$ and $q \in P_v$, see Figure 4. Let $R_u(p)$ (resp. $R_v(q)$) be the set of points in P_u (resp. P_v) that are reachable in H from p (resp. q) with paths that have lengths at most $2 \cdot \operatorname{diam}(\square_u)$ (resp. $2 \cdot \operatorname{diam}(\square_v)$). By Lemma 20, $|R_u(p)| \geq 3\xi n(u) \geq \xi n(u)$ and $|R_v(q)| \geq 3\xi n(v)$.

Since there is a bipartite expander between P_u and P_v with parameter ξ, by Lemma 4, the neighborhood Y of $R_u(p)$ in P_v has size at least $(1-\xi)n(v)$. Observe that $|Y \cap R_v(q)| = |R_v(q) \setminus (P_v \setminus Y)| \geq |R_v(q)| - |P_v \setminus Y| \geq 3\xi n(v) - \xi n(v) > 0$. Therefore, there is a point $q' \in Y \cap R_v(q)$, and a point $p' \in R_u(p)$, such that $p'q' \in E(G_\Phi)$. We have that

$$\begin{aligned}
\mathsf{d}_H(p,q) &\leq \mathsf{d}_H(p,p') + \mathsf{d}_H(p',q') + \mathsf{d}_H(q',q) \leq 2 \cdot \operatorname{diam}(\square_u) + \|p' - q'\| + 2 \cdot \operatorname{diam}(\square_v) \\
&\leq 3 \cdot \operatorname{diam}(\square_u) + \mathsf{d}(\square_u, \square_v) + 3 \cdot \operatorname{diam}(\square_v) \leq \left(1 + 6 \cdot \frac{\varepsilon}{6}\right) \cdot \mathsf{d}(\square_u, \square_v) \\
&\leq (1+\varepsilon) \cdot \|p - q\| .
\end{aligned}$$ ◀

6 Conclusions

In this paper we have shown several constructions for ϑ-reliable spanners. Our results for constructing reliable exact spanners in one dimension have size $\mathcal{O}(n \log n)$, which is optimal. In higher dimensions we were able to show a simple construction of a ϑ-reliable spanner with optimal size for the case of bounded spread. For arbitrary point sets in $\mathbb{R}^d$ we obtained a construction with an extra $(\log\log n)^6$ factor in the size.

It seems clear that our construction for the unbounded case is suboptimal in terms of extra factors, and we leave improving it as an open problem for further research. Another natural open question is how to construct reliable spanners that are required to be subgraphs of a given graph.

References

1. M. A. Abam and S. Har-Peled. New Constructions of SSPDs and their Applications. *Computational Geometry: Theory and Applications*, 45(5–6):200–214, 2012. `doi:10.1016/j.comgeo.2011.12.003`.
2. B. Aronov, M. de Berg, O. Cheong, J. Gudmundsson, H. J. Haverkort, M. H. M. Smid, and A. Vigneron. Sparse geometric graphs with small dilation. *Computational Geometry: Theory and Applications*, 40(3):207–219, 2008. `doi:10.1016/j.comgeo.2007.07.004`.
3. S. Arya, D. M. Mount, and M. Smid. Randomized and deterministic algorithms for geometric spanners of small diameter. In *Proc. 35th Annu. IEEE Sympos. Found. Comput. Sci.* (FOCS), pages 703–712, 1994. `doi:10.1109/SFCS.1994.365722`.
4. S. Arya, D. M. Mount, and M. Smid. Dynamic algorithms for geometric spanners of small diameter: Randomized solutions. *Computational Geometry: Theory and Applications*, 13(2):91–107, 1999. `doi:10.1016/S0925-7721(99)00014-0`.
5. P. Bose, P. Carmi, V. Dujmovic, and P. Morin. Near-Optimal O(k)-Robust Geometric Spanners. *CoRR*, abs/1812.09913, 2018. `arXiv:1812.09913`.
6. P. Bose, P. Carmi, M. Farshi, A. Maheshwari, and M. Smid. Computing the Greedy Spanner in Near-Quadratic Time. *Algorithmica*, 58(3):711–729, November 2010. `doi:10.1007/s00453-009-9293-4`.
7. P. Bose, V. Dujmović, P. Morin, and M. Smid. Robust Geometric Spanners. *SIAM Journal on Computing*, 42(4):1720–1736, 2013. `doi:10.1137/120874473`.
8. K. Buchin, S. Har-Peled, and D. Oláh. A Spanner for the Day After. *CoRR*, abs/1811.06898, 2018. `arXiv:1811.06898`.
9. P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. *Journal of the Association for Computing Machinery*, 42(1):67–90, 1995. `doi:10.1145/200836.200853`.
10. P. Carmi and L. Chaitman. Stable Roommates and Geometric Spanners. In *Proc. 22nd Canad. Conf. Comput. Geom.* (CCCG), pages 31–34, 2010. URL: `http://cccg.ca/proceedings/2010/paper11.pdf`.
11. T. M. Chan, S. Har-Peled, and M. Jones. On Locality-Sensitive Orderings and their Applications. *CoRR*, abs/1809.11147, 2018. `arXiv:1809.11147`.
12. J. Gudmundsson, C. Levcopoulos, and G. Narasimhan. Fast Greedy Algorithms for Constructing Sparse Geometric Spanners. *SIAM J. Comput.*, 31(5):1479–1500, May 2002. `doi:10.1137/S0097539700382947`.
13. C. Levcopoulos, G. Narasimhan, and M. Smid. Efficient Algorithms for Constructing Fault-Tolerant Geometric Spanners. In *Proc. 30th Annu. ACM Sympos. Theory Comput.* (STOC), pages 186–195. ACM, 1998. `doi:10.1145/276698.276734`.
14. C. Levcopoulos, G. Narasimhan, and M. Smid. Improved Algorithms for Constructing Fault-Tolerant Spanners. *Algorithmica*, 32(1):144–156, 2002. `doi:10.1007/s00453-001-0075-x`.
15. T. Lukovszki. New Results of Fault Tolerant Geometric Spanners. In *Proc. 6th Workshop Algorithms Data Struct.* (WADS), volume 1663 of *LNCS*, pages 193–204. Springer, 1999. `doi:10.1007/3-540-48447-7_20`.
16. R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, USA, 1995.
17. G. Narasimhan and M. Smid. *Geometric spanner networks*. Cambridge University Press, New York, NY, USA, 2007.
18. M. Smid. Geometric Spanners with Few Edges and Degree Five. In *Proc. 12th Australasian Theo. Sym.* (CATS), volume 51 of *CRPIT*, pages 7–9, 2006. URL: `http://crpit.com/abstracts/CRPITV51Smid.html`.