

A spanner for the day after

Citation for published version (APA):

Buchin, K. A., Har-Peled, S., & Olah, D. (2018). A spanner for the day after. Cornell University.

Document status and date:

Published: 01/01/2018

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

A Spanner for the Day After

Kevin Buchin*

Sariel Har-Peled†

Dániel Oláh‡

January 28, 2019

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 $\mathcal{O}(\varepsilon^{-c} \vartheta^{-6} n \log n (\log \log n)^6)$ edges, where $c = \mathcal{O}(d)$. 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.

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 $d(p, q)$.

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 [AdBC⁺08, CC10, Smi06], low weight [BCF⁺10, GLN02], low diameter [AMS94, AMS99] or to be resistant against failures. The book by Narasimhan and Smid [NS07] 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.* [BDMS13]. 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 all but $n - f(k)$ 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.

*Department of Mathematics and Computing Science, TU Eindhoven, P.O. Box 513, 5600 MB Eindhoven, The Netherlands.

†Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@illinois.edu; <http://sarielhp.org/>. Work on this paper was partially supported by a NSF AF award CCF-1421231.

‡Department of Mathematics and Computing Science, TU Eindhoven, P.O. Box 513, 5600 MB Eindhoven, The Netherlands. Supported by the Netherlands Organisation for Scientific Research (NWO) through Gravitation-grant NETWORKS-024.002.003.

	robustness	# edges [BDMS13]	
$d = 1$	$\mathcal{O}(k \log k)$	$\mathcal{O}(n \log n)$	Theorem 1
	$\mathcal{O}(f(k)f^*(k))$	$\mathcal{O}(nf^*(n))$	Theorem 2
	$\mathcal{O}(f(k))$	$\mathcal{O}(nf^*(n))$	$f(k) \in k2^{\Omega(\sqrt{\log k})}$
	$\mathcal{O}(k)$	$\Omega(n \log n)$	Theorem 3
	$f(k)$	$\Omega(nf^*(n))$	Theorem 4
	$\mathcal{O}(k \log k)$ $\mathcal{O}(kc^{\sqrt{\log k}})$ $\mathcal{O}(k^c)$	$\Omega(n \log n / \log \log n)$ $\Omega(n\sqrt{\log n})$ $\Omega(n \log \log n)$	Corollary 2
$d > 1$	$\mathcal{O}(kf(k))$	$\mathcal{O}(nf^*(n))$	Theorem 5
	$\mathcal{O}(k^2)$	$\mathcal{O}(n \log n)$	Corollary 3

Table 1.1: Some of the results of Bose *et al.* [BDMS13]. Let t, c be constants larger than one. All results are for graphs that are t -spanners. In the above f^* is how many times you have to apply f to itself till it reaches the parameter (as such, for $f(k) = 2k$, we have $f^*(k) = \Theta(\log k)$).

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.* [BDMS13] 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.* [BDMS13] proved various other bounds in the same vein on the size for one-dimensional and higher-dimensional point sets – see Table 1.1 for a summary of their relevant results. Their most closely related result is that for the one-dimensional point set $P = \{1, 2, \dots, n\}$ 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.* [BDMS13] is the construction of $\mathcal{O}(k)$ -robust spanners – they only provide the easy upper bound of $\mathcal{O}(n^2)$ for this case. In this paper, we present several constructions for this case with optimal or near-optimal size. These results even hold for a stronger requirement on the spanners, which we call ϑ -reliable.

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 [LNS98, LNS02, Luk99]. 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.

Expanders are reliable. Intuitively, a constant degree expander is a robust/reliable graph under a weaker notion of robustness – that is, connectivity. To this end, we show that, for a parameter $\vartheta > 0$, that constant degree expanders are indeed ϑ -reliable in the sense that all except a small fraction of the vertices 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 2.6](#) for the statement and [Appendix B.3](#) for the proof. This is essentially already known, and we provide the proof (and statement) for the sake of completeness.

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.

- (A) **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 3.6](#) for the result.
- (B) **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 3.12](#) for the result.
- (C) **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.* [[CHJ18](#)], which shows 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. By applying a recursive scheme, using the same idea, we obtain the desired spanner of size $\mathcal{O}(n \log n (\log \log n)^6)$. See [Section 4](#) for details.
- (D) **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 $\mathcal{O}(\varepsilon^{-d} \vartheta^{-2} n \log \Phi)$ 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 5.10](#) 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 [Appendix A](#) for an additional result in higher dimensions.

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

1.2 The competition

An earlier version of this paper was posted to the arxiv on November 16, 2018, with a somewhat weaker construction (with a bound similar to [Lemma 4.2](#) but using a semi-separated pair decomposition instead of [Theorem 4.1](#)). Shortly after (December 24, 2018), and independently, Bose *et al.* [[BCDM18](#)] posted a slightly better construction with $\mathcal{O}(n \log^2 n \log \log n)$ edges (for ε and ϑ constants). They also use expanders in combination with different tools (such as fair-split trees and WSPDs). Of course, the current paper present an even better construction with only $\mathcal{O}(n \log n (\log \log n)^6)$ edges.

2 Preliminaries

2.1 Problem definition and notations

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

Definition 2.1 (Robust spanner). Let $G = (P, E)$ be a t -spanner for some $t \geq 1$ and let $f: \mathbb{N} \rightarrow \mathbb{R}_+$, 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

$$G \setminus B = G_{P \setminus B} = \left(P \setminus B, \{uv \in E(G) \mid u, v \in P \setminus B\} \right)$$

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

Definition 2.3. For a number $x > 0$, let $\text{pow}_2(x) = 2^{\lceil \log x \rceil}$ be the smallest number that is a power of 2 and is at least as large as x .

2.2 Expanders

2.2.1 Basic construction

For a set X of vertices in a graph $G = (V, E)$, let $\Gamma(X) = \{v \in V \mid uv \in E \text{ for a } u \in X\}$ be the **neighbors** of X in G . The following lemma, which is a standard expander construction, provides the main building block of our one-dimensional construction. Since this is pretty standard construction in the expanders literature, we provide the proof in [Appendix B.1](#).

Lemma 2.4. *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 $|\Gamma(Y)| > (1 - \xi)|L|$.*

2.2.2 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, as such, we delegate the proofs to the appendix.

Lemma 2.5. (Proof in [Appendix B.2](#).) *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 $|\Gamma(X)| \geq \min((1 - \beta)n, \alpha|X|)$. The graph G has $\mathcal{O}((\alpha/\beta)n)$ edges.*

Lemma 2.6. (Proof in [Appendix B.3](#).) *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 3.1. 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 $|[i : j] \cap B| \geq \alpha |[i : j]|$. Similarly, i is in the **right α -shadow** of B , if and only if there exists an integer $h \leq i$, such that $|[h : i] \cap B| \geq \alpha |[h : i]|$. 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 3.2. *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 $2(1 + \lceil 1/\alpha \rceil)|B|$.*

Proof: Let x_i , for $i = 1, \dots, n$, be a sequence of numbers, where $x_i = c = -\lceil 1/\alpha \rceil$ if $i \in B$, and $x_i = 1$ otherwise. If j is in the left α -shadow of B , then there exists an integer j' , such that $\beta = |[j : j'] \cap B| \geq \alpha |[j : j']|$. Setting $\Delta = |[j : j']| = j' - j + 1$, we have that

$$\sum_{i=j}^{j'} x_i = (\Delta - \beta) + c\beta = \Delta + (c - 1)\beta \leq \Delta + (c - 1)\alpha\Delta = \Delta(1 - \lceil 1/\alpha \rceil \alpha - \alpha) \leq -\alpha\Delta < 0.$$

Namely, an integer $j \in \mathcal{S}_{\rightarrow}(B)$ corresponds to some prefix sum of the x_i s that starts at location j and adds up to some negative sum. In order to bound the number of such locations, consider the minimal location i that has $x_i = c$. Mark the $1 + \lceil 1/\alpha \rceil$ consecutive locations ending at i (including i itself) as potentially being in the shadow, and delete them from the sequence. In this way the sum of the x_i s for the locations we delete is zero. If $i < 1 + \lceil 1/\alpha \rceil$ then we are naturally marking fewer locations as being in the shadow.

Clearly, a location that had a negative prefix sum in the original sequence also has a negative prefix sum starting at this location in the new sequence. Every such operation deletes $1 + \lceil 1/\alpha \rceil$ elements from the sequence, and reduces the number of elements in B by one. We conclude that the number of elements that start a negative prefix sum is at most $(1 + \lceil 1/\alpha \rceil)|B|$. Therefore, $|\mathcal{S}_{\rightarrow}(B)| \leq (1 + \lceil 1/\alpha \rceil)|B|$ holds.

The above argument applied symmetrically also bounds the number of elements in the right α -shadow of B , and adding these two quantities implies that $|\mathcal{S}(\alpha, B)| \leq 2(1 + \lceil 1/\alpha \rceil)|B|$. \blacksquare

Lemma 3.2 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 3.2** quantify this.

Lemma 3.3. *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 $i = 1, \dots, n$, 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, $x([j : j']) \leq 0$. Namely, an integer j in the left α -shadow of B corresponds to some prefix sum of the x_i s 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, \dots, 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 $x([k : j - 1]) + x([j : k']) = x([k : k']) \leq 0$. 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 $x([j : k - 1]) = x(W_j) - \tau < 0$, 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 $x(J) > 0$.

So consider the new sequence of numbers $x_{[n] \setminus J} = x_1, \dots, x_{k-1}, x_{k'}, \dots, x_n$ 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

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

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 3.1](#). Let \mathcal{I} be the resulting set of all blocks. One can sort the blocks of the tree, that are of nodes in the same level, from left to right. Two adjacent blocks of the same level are neighbors. For a block $I \in \mathcal{I}$, let $\text{next}(I)$ and $\text{prev}(I)$ be the blocks (in the same level) directly to the right and left of I , respectively.

We build the graph of [Lemma 2.4](#) 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

For the sake of simplicity, assume for the time being that n is a power of 2.

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

Lemma 3.4. *The graph H has $\mathcal{O}(n \log n)$ edges.*




Figure 3.1: The binary tree built over $[n]$. The block of node v is the interval $[i : j]$.

Proof: Let $h = \log n$ be the depth of the tree T . In each level $i = 1, 2, \dots, h$ 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. ■

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

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

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

be the resulting walk on the blocks where we removed repeated blocks. Figure 3.2 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 3.5. *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, $I_j \subseteq I_{j+1}$. It is thus easy to verify that either $i \in I_t$, or $i \in \text{prev}(I_t)$, or $i \in \text{prev}(\text{prev}(I_t))$. 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. ■



(a) The canonical path in the tree.



(b) The canonical path on the blocks.

Figure 3.2: 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.

Theorem 3.6. *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 3.4. 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 3.2 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 3.5.

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 $\pi(i, j): I_0 \rightarrow I_1 \rightarrow \dots \rightarrow I_\ell$. 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 $I_t \rightarrow I_{t+1}$. There are two cases:

(i) $I_{t+1} = \text{next}(I_t)$. Then, the graph H includes the expander graph on I_t, I_{t+1} described in Lemma 2.4. 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 $I_t \cap \mathcal{S}$. 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. \blacksquare




Figure 3.3: 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.

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 $\{1, 2, 3, \dots, 2^h\}$. Since H is a 1-spanner, the 1-paths between any pair of vertices of $[n]$ does not use any vertices from $\{n+1, \dots, 2^h\}$. 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.

3.3 Construction of ϑ -reliable exact spanners in one dimension

3.3.1 The construction

Here, we show how to extend [Theorem 3.6](#), 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 3.6](#), 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

$$N = \text{pow}_2(c/\vartheta^2) \quad \text{and} \quad \xi = \frac{1}{32N}, \quad (3.1)$$

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 $i = i_0, \dots, h-1$, and $j = 1, \dots, N$, 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 $[1 - 2^i : 1 - 2^i/N]$, starting at its left endpoint. For $k = 0, \dots, n/2^i$, let

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

be the shifted interval of length 2^i that starts at $\Delta(i, j)$ and is shifted k blocks to the right, see [Figure 3.3](#). The set of all intervals of interest is

$$\mathcal{I} = \left\{ \mathcal{I}(i, j, k) \left| \begin{array}{l} i = i_0, \dots, \log n \\ j = 1, \dots, N \\ k = 0, \dots, n/2^i \end{array} \right. \right\}. \quad (3.2)$$

Constructing the graph H_ϑ . Let $G_E(i, j, k)$ denote the expander graph of [Lemma 2.4](#), constructed over $\mathcal{I}(i, j, k)$ and $\mathcal{I}(i, j, k+1)$, with the value of the parameter ξ as specified in [Eq. \(3.1\)](#). We define H_ϑ to be the union of all the graphs 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 $[\text{pow}_2(n)]$. In any case, the last step is to delete vertices from H_ϑ that are outside the range of interest $[n]$.

3.3.2 Analysis of H_ϑ

Lemma 3.7. *The graph H_ϑ has $\mathcal{O}(\vartheta^{-6}n \log n)$ edges.*

Proof: There are $\log n$ resolutions. For every resolution there are $N = \mathcal{O}(1/\vartheta^2)$ different shifts. For every shift, the number of edges created is $\mathcal{O}(n\xi^{-2}) = \mathcal{O}(n/\vartheta^4)$, by [Lemma 2.4](#). Thus, H_ϑ has $\mathcal{O}(\vartheta^{-6}n \log n)$ edges. ■

In the following, let $\llbracket s, \ell \rrbracket = [s : s + \ell - 1]$ be the set of consecutive integers starting at s containing ℓ numbers.

Definition 3.8. For two vertices $x, y \in [n]$, y is a **descendant** of x (and x is an **ancestor** of y) in G , if $x < y$ and there is a 1-path between x and y in G . For a set $B \subseteq [n]$, and a vertex s , let $\mathcal{D} = \mathcal{D}(G, s, B)$ be the set of all descendants of s in $G \setminus B$. Similarly, for a vertex t , let $\mathcal{A} = \mathcal{A}(G, t, B)$ be the set of ancestors of t in $G \setminus B$.

For an interval $I \subseteq [n]$, the set $I \cap \mathcal{D}$ is the set of all nodes in I that are descendants of s in the graph $G \setminus B$. In a symmetric fashion, the set of ancestors in I that can reach a node t is denoted by $I \cap \mathcal{A}$.

Lemma 3.9. *Let $B \subseteq [n]$ be the set of deleted locations, $\alpha = 1 - \vartheta/4$ and s be a location in $[n]$ that is not in the α -shadow of B . Let $h \geq N$ be an integer number, and let $c \geq 512$ be the constant from the construction. Let $\mathcal{D} = \mathcal{D}(G, s, B)$, and assume that $|\llbracket s, h \rrbracket \cap \mathcal{D}| \geq (\vartheta/32)h$. Then, for some number H , $8h/\vartheta \leq H \leq c'h/\vartheta$, we have $|\llbracket s, H \rrbracket \cap \mathcal{D}| \geq (\vartheta/c')H$, where $c' = c/8$.*

Proof: The idea is to choose the right resolution in the construction of H_ϑ . As a first step, let

$$\Delta = \text{pow}_2(\vartheta h/64) \quad \implies \quad \vartheta h/64 \leq \Delta \leq \vartheta h/32$$

be the desired shift. We pick the resolution i such that the shift used $2^i/N$ is equal to Δ (i.e., $\Delta = 2^i/N$). This implies that $i = \log(N\Delta)$. There is a choice of j and k , such that the right endpoint of $L = I(i, j, k)$ lies in the interval $\llbracket s + h, \Delta \rrbracket$. Notice that $\llbracket s, h \rrbracket \subseteq L$, since

$$h + \Delta \leq (1 + 64/\vartheta)\Delta = \left(1 + \frac{64}{\vartheta}\right) \frac{2^i}{N} \leq \left(1 + \frac{64}{\vartheta}\right) \frac{\vartheta^2}{c} 2^i \leq 2^i$$

holds. Let $R = I(i, j, k + 1)$ and $H = \text{right}(R) - s + 1$, where $\text{right}(R)$ is the right endpoint of the interval R , see [Figure 3.4](#). Observe that $\vartheta h/64 \leq \Delta \leq \vartheta h/32$ and

$$H \geq 2^i = N\Delta \geq \frac{c}{\vartheta^2} \cdot \frac{\vartheta h}{64} = \frac{c}{64} \cdot \frac{h}{\vartheta} \geq \frac{8h}{\vartheta},$$

since $c \geq 512$. Similarly,

$$H \leq 2 \cdot 2^i = 2N\Delta \leq 2 \cdot \frac{2c}{\vartheta^2} \cdot \frac{\vartheta h}{32} = \frac{c}{8} \cdot \frac{h}{\vartheta}.$$

Let $U = \llbracket s, h \rrbracket \cap \mathcal{D}$. By assumption, $|U| \geq (\vartheta/32)h$. Since the interval L is of length 2^i , we have

$$\frac{|L \cap \mathcal{D}|}{|L|} \geq \frac{|U|}{2^i} \geq \frac{(\vartheta/32)h}{N\Delta} \geq \frac{(\vartheta/32)h}{N(\vartheta/32)h} = \frac{1}{N} \geq \xi.$$




Figure 3.4: The intervals L and R and their relation to s, h, Δ and H .

Since s is not in the α -shadow of B , it follows that the interval $\llbracket s, H \rrbracket$ contains at least $(\vartheta/4)H$ elements that are not in B . Let τ be the fraction of elements of \mathbb{R} that are not in B . We have that

$$\begin{aligned} \tau = \frac{|\mathbb{R} \setminus B|}{|\mathbb{R}|} &\geq \frac{(\vartheta/4)H - h - \Delta}{2^i} \geq \frac{(\vartheta/4)(2^i + h) - (h + \Delta)}{2^i} \\ &\geq \frac{(\vartheta/4)(2^i + (32/\vartheta N)2^i) - (64/\vartheta + 1)2^i/N}{2^i} \\ &= \frac{\vartheta}{4} + \frac{8}{N} - \left(1 + \frac{64}{\vartheta}\right) \frac{1}{N} \geq \frac{\vartheta}{4} - \frac{64}{\vartheta N} \geq \frac{\vartheta}{4} - \frac{64\vartheta}{c} \geq \frac{\vartheta}{8}. \end{aligned}$$

Let $U' \subseteq \mathbb{R}$ be the set of all nodes that are connected by an edge of H_ϑ to U . Note, that all the nodes of U' are descendants of s . The graph $G_E(i, j, k)$ guarantees that $|U'| \geq (1 - \xi)|\mathbb{R}|$, where $G_E(i, j, k)$ is the expander graph built over \mathbb{L} and \mathbb{R} . We have that

$$\begin{aligned} |\llbracket s, H \rrbracket \cap \mathcal{D}| &\geq |(\mathbb{R} \setminus B) \cap U'| = |\mathbb{R} \setminus B| - |(\mathbb{R} \setminus B) \cap \overline{U'}| \\ &\geq |\mathbb{R} \setminus B| - |\mathbb{R} \cap \overline{U'}| \geq |\mathbb{R} \setminus B| - \xi 2^i = (\tau - \xi)2^i. \end{aligned}$$

Since $\xi \leq \vartheta/16$, we have $\frac{|\llbracket s, H \rrbracket \cap \mathcal{D}|}{H} \geq \frac{(\tau - \xi)2^i}{2 \cdot 2^i} = \frac{\tau - \xi}{2} \geq \frac{\vartheta/8 - \vartheta/16}{2} = \frac{\vartheta}{32}$. ■

Remark 3.10. One can state a symmetric version of [Lemma 3.9](#) about the number of ancestors that can reach a target node t .

Lemma 3.11. *Let $B \subseteq [n]$ be the set of faulty vertices, and let $\mathcal{S}(\alpha, B)$ be its α -shadow with $\alpha = 1 - \vartheta/4$. Let s, t be two vertices in $[n] \setminus \mathcal{S}(\alpha, B)$, such that $s < t$. Then, there is a 1-path between s and t in $H_\vartheta \setminus B$. Further, this path between s and t uses at most $2 \log n$ edges.*

Proof: If $|s - t| \leq 3N$, then the two vertices are connected by an edge in H_ϑ by construction, and the claim holds.

Let \mathbb{L} and \mathbb{R} be two adjacent consecutive blocks of the same size in \mathcal{I} (see [Eq. \(3.2\)](#)), such that $s \in \mathbb{L}$ and $t \in \mathbb{R}$, and these are the smallest blocks for which this property holds. If there are several pairs of intervals of the same size that have the desired property, we pick the pair such that $\min(\text{right}(\mathbb{L}) - s, t - \text{left}(\mathbb{R}))$ is maximized (i.e., the common boundary between the two intervals is as close to the middle $(s + t)/2$ as possible). Let $2^i = |\mathbb{R}| = |\mathbb{L}|$. It is easy to verify that $2^i/2 \leq |s : t| \leq 2 \cdot 2^i$. Indeed, the lower bound holds by the minimality of \mathbb{L} and \mathbb{R} . Otherwise, the right half of \mathbb{L} and the left half of \mathbb{R} would also be a valid choice and would have smaller size. The upper bound follows from the fact that $|\mathbb{L}| + |\mathbb{R}| = 2 \cdot 2^i$.

Set $L_0 = \llbracket s, N \rrbracket$ and $R_0 = [t - N + 1 : t]$. Since s and t are not in the α -shadow, we have that $|L_0 \setminus B| \geq (\vartheta/4)|L_0|$ and $|R_0 \setminus B| \geq (\vartheta/4)|R_0|$. For $i > 0$, in the i th iteration, let L_i be the interval starting at s of length $\Theta(|L_{i-1}|/\vartheta)$ such that at least $\vartheta/32$ fraction of its elements are descendants of s that are not in B . The existence of such an interval is guaranteed by [Lemma 3.9](#). Similarly, we expand the right interval R_{i-1} in a symmetric way.

Let j be the first iteration such that $L_{j+1} \not\subseteq \mathbb{L}$. By the choice of \mathbb{L} and \mathbb{R} and by [Lemma 3.9](#), we have

$$\frac{2^i}{4} - \frac{2^i}{N} \leq |L_{j+1}| \leq \frac{c}{8\vartheta} |L_j|.$$

This implies that

$$\frac{|\mathbb{L} \cap \mathcal{D}|}{|\mathbb{L}|} \geq \frac{|L_j \cap \mathcal{D}|}{|L_j|} \geq \frac{(\vartheta/32)|L_j|}{2^i} \geq \frac{\vartheta}{32} \cdot \frac{8\vartheta}{c} \left(\frac{1}{4} - \frac{1}{N}\right) \geq \frac{\vartheta}{32} \cdot \frac{8\vartheta}{c} \cdot \frac{1}{8} \geq \frac{1}{32N} = \xi.$$

Applying the same argumentation, using [Lemma 3.9](#) for the reachable ancestors, we have that

$$|\mathbb{R} \cap \mathcal{A}|/|\mathbb{R}| \geq \xi$$

(i.e., there are at least $\xi |R|$ elements in R that have a 1-path to t in $H_\vartheta \setminus B$). The graph H_ϑ contains an expander $G_E(i, j, k)$ built over L and R . By the pigeonhole principle and the properties of the expander between L and R , there is an edge between a vertex of $L \cap \mathcal{D}$ and a vertex of $R \cap \mathcal{A}$. That is, there is a 1-path between s and t in $H_\vartheta \setminus B$, as desired.

By [Lemma 3.9](#) we have $8|L_i| \leq (8/\vartheta)|L_i| \leq |L_{i+1}|$ for $i = 0, \dots, j$. Therefore, the number of iterations we do to expand L_0 is less than $\log n$. The same is true for R_0 . Thus, the number of edges that we used for the 1-path is bounded by $2 \log n$. \blacksquare

Theorem 3.12. *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.*

Proof: The bound on the number of edges is from [Lemma 3.7](#).

Next, fix the set B . Define the set B^+ to be the $(1 - \vartheta/4)$ -shadow of B . By [Lemma 3.3](#) we have that $|B^+| \leq |B| / (2(1 - \vartheta/4) - 1) = |B| / (1 - \vartheta/2) \leq (1 + \vartheta) |B|$.

A 1-path in $H_\vartheta \setminus B$ between any two vertices in $[n] \setminus B^+$ exists by [Lemma 3.11](#). \blacksquare

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

4.1 A first construction

In the following, we assume that $P \subseteq [0, 1]^d$ – this can be done by an appropriate scaling and translation of space. 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 . We need the following minor variant of a result of Chan *et al.* [[CHJ18](#)].

Theorem 4.1 ([\[CHJ18\]](#)). *For $\varsigma \in (0, 1)$, 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) $(p, z)_\sigma \subseteq \text{ball}(p, \varsigma \ell)$,
- (iii) $(z, q)_\sigma \subseteq \text{ball}(q, \varsigma \ell)$, and
- (iv) $z \in \text{ball}(p, \varsigma \ell)$ or $z \in \text{ball}(q, \varsigma \ell)$.

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.

First, we give a very simple construction and analysis, for building reliable spanners, using the theorem above and our one-dimensional construction. We present it to convey the basic principle of this technique. Then, by tuning the parameters, we repeat the construction to obtain a reliable spanner of size $\mathcal{O}(n \log n (\log \log n)^6)$. This construction has a more elaborate analysis, where we use the same idea, but in an iterative manner.

4.1.1 Construction in detail

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

$$M = M(\varsigma) = \mathcal{O}(\varsigma^{-d} \log \varsigma^{-1}) = \mathcal{O}\left(\varepsilon^{-d} \log^d n \log \frac{\log n}{\varepsilon}\right),$$

and c be some sufficiently large constant. Next, let $\vartheta' = \vartheta / M$, and let $\Pi^+ = \Pi^+(\varsigma)$ be the set of orderings of [Theorem 4.1](#). For each ordering $\sigma \in \Pi^+$, compute the ϑ' -reliable exact spanner G_σ of P , see [Theorem 3.12](#), according to σ . Let G be the resulting graph by taking the union of G_σ for all $\sigma \in \Pi^+$.

4.1.2 Analysis

Lemma 4.2. *The graph G , constructed above, is a ϑ -reliable $(1 + \varepsilon)$ -spanner and has size*

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

Proof: Given a (failure) set $B \subseteq P$, let B^+ be the union of all the harmed sets resulting from B in G_σ , for all $\sigma \in \Pi^+$. We have that $|B^+| \leq (1 + M \cdot \vartheta') |B| = (1 + \vartheta) |B|$.

Consider any two points $p, q \in P \setminus B^+$. By **Theorem 4.1**, for $\ell = \|p - q\|$, there exists an ordering $\sigma \in \Pi^+$, and a point $z \in [0, 1)^d$, such that $(p, z)_\sigma \subseteq \text{ball}(p, \varsigma \ell)$ and $(z, q)_\sigma \subseteq \text{ball}(q, \varsigma \ell)$ (and z is in one of these balls).

By **Theorem 3.12**, the graph $G_\sigma \setminus B \subseteq G \setminus B$ contains a monotone path π , according to σ , with $h = \mathcal{O}(\log n)$ hops, connecting p to q . Let $p = p_1, \dots, p_{h+1} = q$ be this path. Observe that there is a unique index i , such that $z \in (p_i, p_{i+1})$. We have the following:

$$(A) \quad \forall j \neq i \quad \|p_j - p_{j+1}\| \leq 2\varsigma \ell.$$

$$(B) \quad \|p_i - p_{i+1}\| \leq \ell + 2\varsigma \ell.$$

As such, the total length of π is $\sum_{j=1}^h \|p_j - p_{j+1}\| = (1 + 2\varsigma h)\ell \leq (1 + \varepsilon)\ell$, as desired, if c is sufficiently large. Namely, G is the desired reliable spanner.

The number of edges of G is

$$M \cdot \mathcal{O}((\vartheta')^{-6} n \log n) = \mathcal{O}(M(M/\vartheta)^6 n \log n) = \mathcal{O}\left(\varepsilon^{-7d} \vartheta^{-6} n \log^{7d} n \log^7 \frac{\log n}{\varepsilon}\right). \quad \blacksquare$$

4.2 An improved construction

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

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

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 4.1**. For each ordering $\sigma \in \Pi^+$, compute the ϑ' -reliable exact spanner G_σ of P , see **Theorem 3.12**, according to σ . Let G be the resulting graph by taking the union of G_σ for all $\sigma \in \Pi^+$.

Theorem 4.3. *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 3.12**. Each of these graphs has $\mathcal{O}((\vartheta')^{-6} n \log n)$ edges. Therefore, the size of G is

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

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 $i = 2, \dots, N$, 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 3.3](#) we have that

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

Therefore, we obtain

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

using $1+x \leq e^x \leq 1+3x$, 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 3.12](#), 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) $\|p' - q'\| \leq \|p - q\| + 2\varsigma \|p - q\| = (1 + 2\varepsilon/c) \|p - q\|$.
- (B) $\|p - p'\| \leq 2\varsigma \|p - q\| = 2(\varepsilon/c) \|p - q\|$ and similarly $\|q - q'\| \leq 2(\varepsilon/c) \|p - q\|$.
- (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 i th 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 i th 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) $\|x' - y'\| \leq \|x - y\| (1 + 2\varepsilon/c) \leq (2\varepsilon/c)^i (1 + 2\varepsilon/c) \|p - q\|$.
- (B) $\|x - x'\| \leq 2(\varepsilon/c) \|x - y\| \leq (2\varepsilon/c)^{i+1} \|p - q\|$ and $\|y - y'\| \leq (2\varepsilon/c)^{i+1} \|p - q\|$.
- (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 4.1](#) and [Theorem 3.12](#) 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 $\|x - y\| \leq \|p - q\| \cdot (2\varepsilon/c)^{N-1}$ and $|Q_{N-1}| = 2^{N-1}$, the total length of the subpaths calculated, in the last step, is

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

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

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

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. \blacksquare

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 the condition that the spread $\Phi(P)$ is bounded by a polynomial of n .

5.1 Preliminaries

Definition 5.1. For a point set $P \subseteq \mathbb{R}^d$, let $\text{diam}(P) = \max_{p, q \in P} \|p - q\|$ denote the **diameter** of P . Let $\text{cp}(P) = \min_{p, q \in P, p \neq q} \|p - q\|$ denote the **closest pair** distance in P . Furthermore, let $\Phi(P) = \text{diam}(P)/\text{cp}(P)$ be the **spread** of P .

Definition 5.2. Let $s > 0$ be a real number and let B and C be sets of points in \mathbb{R}^d . The sets B and C are **s -separated** if $\text{d}(B, C) \geq s \cdot \max(\text{diam}(B), \text{diam}(C))$, where $\text{d}(B, C) = \min_{p \in B, q \in C} \|p - q\|$.

Definition 5.3. Let P be a set of n points in the plane and let $s > 0$ be a real number. An **s -well-separated pair decomposition** (**s -WSPD**) of P is a collection $\{(B_i, C_i)\}_{i=1}^m$ of pairs of subsets of P such that

- B_i and C_i are s -separated for all $i = 1, 2, \dots, m$
- for any $p, q \in P$ ($p \neq q$) there exists a unique pair (B_i, C_i) such that $p \in B_i$ and $q \in C_i$ (or $q \in B_i$ and $p \in C_i$).

The well-separated pairs decomposition was introduced by Callahan and Kosaraju [CK95]. The **size** of a WSPD is the number of pairs m , and the **weight** of a pair decomposition \mathcal{W} is defined as $\omega(\mathcal{W}) = \sum_{i=1}^m (|B_i| + |C_i|)$.

There are several ways to compute an s -WSPD. In this paper we use a quadtree-based approach, which has important properties that we can exploit. More precisely, we use the following result of Abam and Har-Peled [AH12, Lemma 2.8] for computing a WSPD.

Lemma 5.4. *Let P be a set of n points in \mathbb{R}^d , with spread $\Phi = \Phi(P)$, and let $\varepsilon > 0$ be a parameter. Then, one can compute an ε^{-1} -WSPD for P of total weight $\mathcal{O}(n\varepsilon^{-d} \log \Phi)$. Furthermore, any point of P participates in at most $\mathcal{O}(\varepsilon^{-d} \log \Phi)$ pairs.*

5.2 The construction of G_Φ

First, compute a quadtree T for the point set P . For any node $v \in T$, let \square_v denote the **cell** (i.e. square or cube, depending on the dimension) represented by v . Let $P_v = \square_v \cap P$ be the point set stored in the subtree of v . Compute a $(6/\varepsilon)$ -WSPD \mathcal{W} over T for P using Lemma 5.4. The pairs in \mathcal{W} can be represented by pairs of nodes $\{u, v\}$ of the quadtree T . Note that the algorithm of Lemma 5.4 uses the diameters and distances of the cells of the quadtree, that is, for a pair $\{u, v\} \in \mathcal{W}$, we have

$$(6/\varepsilon) \cdot \max(\text{diam}(\square_u), \text{diam}(\square_v)) \leq \text{d}(\square_u, \square_v).$$

For any pair $\{u, v\} \in \mathcal{W}$, we build the bipartite expander of [Lemma 2.4](#) 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 2.4](#) 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.

5.3 Analysis

Lemma 5.5. *The graph G_Φ has $\mathcal{O}(\xi^{-2}\varepsilon^{-d}n \log \Phi(P))$ edges.*

Proof: By [Lemma 5.4](#), every point participates in $\mathcal{O}(\varepsilon^{-d} \log \Phi(P))$ WSPD pairs. By [Lemma 2.4](#) 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 $\mathcal{O}(\xi^{-2}\varepsilon^{-d}n \log \Phi(P))$. \blacksquare

Definition 5.6. 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 5.7. For a node x in T , let $n(x) = |P_x|$, and $b(x) = |P_x \cap B|$.

Lemma 5.8. *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 $H' = \{u_1, \dots, u_m\}$ be the set of all maximal nodes in H , and observe that $\bigcup_{u \in H'} P_u = \bigcup_{v \in H} P_v = \mathcal{S}(\gamma, B)$. For any two nodes $x, y \in H'$, we have $P_x \cap P_y = \emptyset$. 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 $1/\gamma = 1/(1 - \vartheta/2) \leq 1 + \vartheta$. \blacksquare

Lemma 5.9. *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, $R_u(p) = \{q \in P_u \setminus B \mid d_H(p, q) \leq 2 \cdot \text{diam}(\square_u)\}$. Then, we have $|R_u(p)| \geq 3\xi |P_u|$.*

Proof: Let $u_1, u_2, \dots, u_j = u$ 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 $p q_1, q_1 q_2, \dots, q_{m-1} q_m$, 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 $i = 1, \dots, j$, 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, \dots, u_{j-1} . Let v_1, \dots, 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

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




Figure 5.1: 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 $G_\Phi \setminus B$.

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

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

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

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

Lemma 5.10. *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 $\mathcal{O}(\varepsilon^{-d} \vartheta^{-2} n \log \Phi(P))$ edges, where $\Phi(P)$ is the spread of P .*

Proof: Let $\xi = \vartheta/8$ and $\gamma = 1 - \vartheta/2$. The bound on the number of edges follows by [Lemma 5.5](#).

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 5.8](#). 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 5.1](#). 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 $2 \cdot \text{diam}(\square_u)$ (resp. $2 \cdot \text{diam}(\square_v)$). By [Lemma 5.9](#), $|R_u(p)| \geq 3\xi n(u) \geq \xi n(u)$ 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 2.4](#), 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

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

\blacksquare

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 $\mathcal{O}(n \log n (\log \log n)^6)$ edges, which is nearly optimal.

It is still an open question whether ϑ -reliable spanners can be constructed with $\mathcal{O}(n \log n)$ edges for general point sets. Another natural open question is how to construct reliable spanners that are required to be subgraphs of a given graph.

References

- [AdBC⁺08] 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.
- [AH12] 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.
- [AMS94] 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.
- [AMS99] 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.
- [BCDM18] P. Bose, P. Carmi, V. Dujmovic, and P. Morin. Near-optimal $O(k)$ -robust geometric spanners. *CoRR*, abs/1812.09913, 2018.
- [BCF⁺10] 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.
- [BDMS13] P. Bose, V. Dujmović, P. Morin, and M. Smid. Robust geometric spanners. *SIAM Journal on Computing*, 42(4):1720–1736, 2013.
- [CC10] P. Carmi and L. Chaitman. Stable roommates and geometric spanners. In *Proc. 22nd Canad. Conf. Comput. Geom. (CCCG)*, pages 31–34, 2010.
- [CHJ18] T. M. Chan, S. Har-Peled, and M. Jones. On locality-sensitive orderings and their applications. *CoRR*, abs/1809.11147, 2018.
- [CK95] 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.
- [GLN02] J. Gudmundsson, C. Levkopoulos, and G. Narasimhan. Fast greedy algorithms for constructing sparse geometric spanners. *SIAM J. Comput.*, 31(5):1479–1500, May 2002.
- [LNS98] C. Levkopoulos, 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.
- [LNS02] C. Levkopoulos, G. Narasimhan, and M. Smid. Improved algorithms for constructing fault-tolerant spanners. *Algorithmica*, 32(1):144–156, 2002.
- [Luk99] 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.
- [MR95] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, USA, 1995.
- [NS07] G. Narasimhan and M. Smid. *Geometric spanner networks*. Cambridge University Press, New York, NY, USA, 2007.
- [Smi06] 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.

A Living in the shadow of a point set

Let P be a set of n points in \mathbb{R}^d , and let $F \subseteq P$ be an unknown subset of “bad” points. Let $\text{ball}(p, r)$ denote the close ball of radius r centered at p . For a parameter $\alpha \in (0, 1]$, a point $p \in P$ is in the α -*shadow* of F , if there exists a radius $r > 0$, such that

$$|\text{ball}(p, r) \cap F| \geq \alpha |\text{ball}(p, r) \cap P|.$$

Put differently, a ball is not in the α -shadow of F , if any ball centered at p contains at most α -fraction of points of P that belongs to F .

Lemma A.1. *Given a set P of n points in \mathbb{R}^d , and for any set $F \subseteq P$ of k points, and a parameter $\alpha \in (0, 1)$, there are at most $\mathcal{O}(k/\alpha)$ points of P that are in the shadow of F .*

Proof: Assume for the sake of simplicity of exposition that all the pairwise distances of points of P are unique. The idea is to mark points of P that are potentially in the shadow, and argue that this marks all relevant points. To this end, let \mathcal{C} be a set of interior disjoint cones in \mathbb{R}^d centered in the origin, that covers the space, such that the angular diameter of each cone in this set is $< \pi/3$. One can construct such a set of size $\mathcal{O}(1)$.

Let $\mathcal{W} = P$ be a working set. We collect points that are unsafe into a set U . For each point $q \in F$ do the following:

- (A) Add q to U , and remove it from \mathcal{W} .
- (B) For every cone $C \in \mathcal{C}$, let $S(q, C)$ be the set of $\lceil 1/\alpha \rceil$ closest points to q in $(\mathcal{W} \setminus F) \cap (q + C)$, where $q + C$ is the translation of the cone C , such that its apex is in q . Remove the points of $S(q, C)$ from \mathcal{W} , and add them to U .

In the end of the process, we have $F \subseteq U$, and $|U| = |F| (1 + \lceil 1/\alpha \rceil) = \mathcal{O}(k/\alpha)$.

We claim that all the points of $P \setminus U$ are not in the α -shadow of F . To this end, consider a point $p \in P \setminus U$, and assume, for the sake of contradiction, that it is in the α -shadow of F , witnessed by a ball $b = \text{ball}(p, r)$ – namely, $|b \cap F| \geq \alpha |b \cap P|$.

First, observe that by construction $p \notin F$. Next, consider a bad point $q \in F \cap b$. There is a cone $C \in \mathcal{C}$, such that $C + q$ contains p . Let C' be the cone with apex at q , with angular radius $\pi/3$, with axis of symmetry along the line spanning qp . By construction, $q + C \subseteq C'$.

Now, consider the set $X = S(q, C')$. The points of $X \subseteq U$. Namely, p is not one of the $\lceil 1/\alpha \rceil$ closest point to q in P . As such, the maximum distance of points of X from q is bounded by $\ell = \|q - p\|$. Observe that the set $Z = (q + C) \cap \text{int}(\text{ball}(q, \ell))$ is contained in $C' \cap \text{ball}(q, \ell)$. Furthermore, $X \subseteq Z$.

We charge $q \in F \cap b$ to the $\lceil 1/\alpha \rceil$ points of $X \subset (P \setminus F) \cap b \subseteq b$. Since the points of X are associated with q , it follows that they would be charged at most once. We repeat this charging till all points in $F \cap b$ are handled. Let $m = |F \cap b|$. This process found $\lceil 1/\alpha \rceil m$ points in $b \cap (P - p)$ that are not in F , and there are exactly m points of F in the ball b . We conclude that

$$|b \cap F| = m \leq \alpha \lceil 1/\alpha \rceil m < \alpha |b \cap P|,$$

which is a contradiction to the assumption that p is in the shadow of F . We conclude that all shadowed points are contained in U . ■

B Properties of expanders

Here we prove various properties of expanders that we need for our work. These results are standard in the expanders literature, and we provide the proofs for the sake of completeness.

B.1 Building expanders

Restatement of Lemma 2.4. *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 $|\Gamma(Y)| > (1 - \xi)|L|$.*

Proof: This is a variant of an expander graph. See [MR95, Section 5.3] for a similar construction.

Let $c = \lceil 3/\xi^2 \rceil$. For every vertex in L , pick randomly and uniformly (with repetition) $\ell = c \lceil n/|L| \rceil$ neighbors in R . Do the same for every vertex in R , picking $c \lceil n/|R| \rceil$ neighbors at random from L . Let G be the resulting graph, after removing redundant parallel edges. Clearly, the number of edges is as required.

As for the claimed properties, there are at most 2^n subsets of L of size at least ξn . Fix such a subset $X \subseteq L$, and fix a subset on the right, $Z \subseteq R$ of size $\leq (1 - \xi)|R|$ (there are at most 2^n such subsets). The probability that all the edges we picked for the vertices of X , stay inside Z , is at most

$$(1 - \xi)^{\ell|X|} \leq (1 - \xi)^{\ell\xi|L|} \leq (1 - \xi)^{c\xi n} \leq \exp\left(-\xi \cdot \frac{3}{\xi^2} \cdot \xi n\right) \leq \exp(-3n) \leq 1/8^n,$$

since $c \geq 4/\xi^2$ and $1 - \xi \leq \exp(-\xi)$. In particular, for a given X the probability that this happens for any subset Z is less than $2^n/8^n = 1/4^n$. Thus, with probability less than $2^n/4^n = 1/2^n$ there is an $X \subseteq L$ with $|\Gamma(X)| \leq (1 - \xi)n$. Using the same argument for $Y \subseteq R$ we get that the random graph does not have the desired properties with probability $2/2^n < 1$ (for $n > 1$). This implies that a graph with the desired properties exists. ■

B.2 Stronger expander

Restatement of Lemma 2.5. *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 $|\Gamma(X)| \geq \min((1 - \beta)n, \alpha|X|)$. The graph G has $\mathcal{O}((\alpha/\beta)n)$ edges.*

Proof: Let $c = 64 \lceil \alpha/\beta \rceil$. Let $V = [n] = \{1, \dots, n\}$. For each node $i \in [n]$, choose independently and uniformly, c neighbors in V (with repetition). Let G be the resulting graph after removing self loops.

We define the event $\mathcal{E}_j = \{\exists X \subseteq V \text{ s.t. } |X| = j \text{ and } |\Gamma(X)| < \min(\alpha|X|, (1 - \beta)n)\} \subseteq \{\exists X \subseteq V \text{ s.t. } |X| = j \text{ and } |\Gamma(X)| < \alpha|X|\}$ for $j = 1, \dots, n$. For all subsets of size $s < n/(4\alpha)$, we have

$$\begin{aligned} \mathbb{P}[\mathcal{E}_s] &\leq \binom{n}{s} \binom{n}{\alpha s} \left(\frac{\alpha s}{n}\right)^{cs} \leq \left(\frac{en}{s}\right)^s \left(\frac{en}{\alpha s}\right)^{\alpha s} \left(\frac{\alpha s}{n}\right)^{cs} \leq \left[\frac{en}{s} \left(\frac{en}{\alpha s}\right)^\alpha \left(\frac{\alpha s}{n}\right)^c\right]^s \\ &= \left[\exp(1 + \alpha) \left(\frac{s}{n}\right)^{-1-\alpha+c} \alpha^{c-\alpha}\right]^s \leq \left[\exp(1 + \alpha) \left(\frac{\alpha}{4\alpha}\right)^{-1-\alpha+c} \alpha\right]^s < \frac{1}{4^s}, \end{aligned}$$

if $c > 4\alpha + 1$.

For subsets of size $s \geq n/(4\alpha)$, let $t = \min((1 - \beta)n, \alpha s)$. Then, we have

$$\mathbb{P}[\mathcal{E}_s] \leq \binom{n}{s} \binom{n}{t} \left(\frac{t}{n}\right)^{cs} \leq 4^n (1 - \beta)^{cs} \leq 4^n \exp\left(-\beta c \frac{n}{4\alpha}\right) \leq 4^n \exp(-4n) \leq \exp(-2n),$$

if $c > 16\alpha/\beta$. Therefore, we have that $\sum_{s=1}^n \mathbb{P}[\mathcal{E}_s] < 1/2$, which establish that with probability $\geq 1/2$, the generated graph has the desired properties. ■

B.3 Expanders are reliable

Restatement of Lemma 2.6. *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.*

Proof: Let $\alpha = \lceil 100/\vartheta \rceil$ and $\beta = \vartheta/\alpha$, and let G be the graph of [Lemma 2.5](#). Consider any failure set $B \subseteq [n]$, and let $k = |B|$. If $k \geq n/(1 + \vartheta)$, then the claim trivially holds. As such, in the following $k < n/(1 + \vartheta)$.

Let C_1, \dots, C_t be the connected components of $G \setminus B$. Let $n_i = |C_i|$, for all i , and assume that $n_1 \geq n_2 \geq \dots \geq n_t$. If $n_1 \geq n - (1 + \vartheta)k$, then we are done. Assume, that we have $n_1 < n - (1 + \vartheta)k$. This implies that

$$\sum_{i=2}^t n_i = n - n_1 - k > n - (n - (1 + \vartheta)k) - k = \vartheta k.$$

Let $\nu \geq 2$ be the maximal index, such that $\sum_{i=\nu}^t n_i > (\vartheta/4)k$, and let $X = \bigcup_{i=\nu}^t C_i$. Observe, that we have $\Gamma(X) \subseteq X \cup B$, which implies $|\Gamma(X)| \leq |X| + |B|$.

If $|\Gamma(X)| \geq \alpha |X| \geq \alpha(\vartheta/4)k \geq 25k$, since $\alpha \geq 100/\vartheta$. We conclude that $|\Gamma(X)| > \max(3k, 3|X|) \geq |X| + |B|$. But this implies that some vertex in X has a neighbor outside $B \cup X$, which is a contradiction.

Otherwise, by [Lemma 2.5](#), if $|\Gamma(X)| < \alpha |X|$, then it must be that $|\Gamma(X)| \geq (1 - \beta)n$. This in turn implies that

$$|X| \geq |\Gamma(X)| - |B| = (1 - \beta)n - k.$$

There are two possibilities:

(i) If $\vartheta k \geq 4\beta n$ (i.e., k is “large”) then

$$|X| \geq n - \beta n - k \geq n - (1 + \vartheta/4)k.$$

This implies that $n_1 \leq (\vartheta/4)k$. This in turn implies that $(\vartheta/4)k < |X| = \sum_{i=\nu}^t n_i \leq (\vartheta/2)k$, by the maximality of ν , and since $(\vartheta/4)k \geq n_1 \geq n_2 \geq \dots$. This implies that

$$(\vartheta/2)k \geq |X| \geq n - (1 + \vartheta/4)k \iff (1 + (3/4)\vartheta)k \geq n \iff k \geq \frac{n}{1 + (3/4)\vartheta},$$

which is impossible, as $k < n/(1 + \vartheta)$.

(ii) It must be that $\vartheta k < 4\beta n$ – namely, $k < 4(\beta/\vartheta)n \leq n/25$. By the construction of X , we have that $\nu \geq 2$,

$$|X| - n_\nu < (\vartheta/4)k \leq n/100 \quad \text{and} \quad |X| \geq (1 - \beta)n - k \geq (11/12)n.$$

Namely, $n_\nu \geq (3/4)n$, but this is of course impossible, since $n_1 \geq n_\nu$, and $n \geq n_1 + n_\nu \geq 2(3/4)n > n$, a contradiction. ■