



A Spanner for the Day After

Kevin Buchin¹ · Sariel Har-Peled² · Dániel Oláh¹

Received: 17 July 2019 / Revised: 22 May 2020 / Accepted: 19 June 2020 / Published online: 6 August 2020
© The Author(s) 2020

Abstract

We show how to construct a $(1 + \varepsilon)$ -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

$$\mathcal{O}(\varepsilon^{-O(d)} \vartheta^{-6} n (\log \log n)^6 \log n)$$

edges. Furthermore, for *any* k , and *any* deleted set $B \subseteq P$ of k points, the residual graph $G \setminus B$ is a $(1 + \varepsilon)$ -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 with this resilience property (i.e., only a tiny additional fraction of vertices, $\vartheta|B|$, 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.

Keywords Geometric spanners · Reliable spanners · Robust spanners · Vertex failures

Editor in Charge: Kenneth Clarkson

A preliminary version of the paper appeared in SoCG 2019 [10]. Sariel Har-Peled was partially supported by NSF AF awards CCF-1421231 and CCF-1907400. Dániel Oláh was supported by the Netherlands Organisation for Scientific Research (NWO) through Gravitation-grant NETWORKS-024.002.003.

Kevin Buchin
k.a.buchin@tue.nl

Sariel Har-Peled
sariel@illinois.edu
<http://sarielhp.org/>

Dániel Oláh
d.olah@tue.nl

¹ Department of Mathematics and Computing Science, TU Eindhoven, P.O. Box 513, 5600 MB Eindhoven, The Netherlands

² Department of Computer Science, University of Illinois, 201 N. Goodwin Avenue, Urbana, IL 61801, USA

Mathematics Subject Classification 65D18 · 68R01

1 Introduction

Spanners The vertices of a *Euclidean graph* 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 [4,13,24], low weight [8,16], low diameter [5,6], or to be resistant to failures. The book by Narasimhan and Smid [22] gives a comprehensive overview of geometric spanners.

Robustness Here, our goal is to construct spanners that are robust according to the notion introduced by Bose et al. [9]. 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. [9] 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. [9] proved various other bounds in the same vein on the size of one-dimensional and higher-dimensional point sets. 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. [9] 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, defined next.

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 a 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—indeed, simply connect the points from left to right. It becomes significantly more challenging to construct such an exact spanner that is reliable. In case of robust spanners, for a function f under some general conditions, Bose et al. [9] gave a construction of exact $f(k)$ -robust spanners. In particular, their result implies that one can construct $\mathcal{O}(k \log k)$ -robust 1-spanners of size $\mathcal{O}(n \log n)$ and, for any $\varepsilon > 0$, $\mathcal{O}(k^{1+\varepsilon})$ -robust 1-spanners of size $\mathcal{O}(n \log \log n)$, for one-dimensional point sets.

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 [18–20]. 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.

Region fault tolerant spanners In a surprising result, Abam et al. [1] showed that one can build a geometric spanner with near linear number of edges, so that if the deleted set are all the points belonging to a convex region (they also delete the edges intersecting this region), then the residual graph is still a spanner for the remaining points.

1.1 Our Results

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.

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 at most $\mathcal{O}(n \log n)$ 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

¹ This also improves an earlier preliminary construction by (some of) the authors [11].

as building different layers of expanders for different resolutions. The construction is described in Sect. 3.2. See Theorem 3.6 for the result.

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 with $\mathcal{O}(\vartheta^{-6}n \log n)$ edges. The construction is described in Sect. 3.3. See Theorem 3.11 for the result.

ϑ -Reliable $(1 + \varepsilon)$ -spanners in higher dimensions We next show a simple construction of ϑ -reliable spanners in \mathbb{R}^d , for d being constant, using a recent result of Chan et al. [14], 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 these orderings, which is the problem we already solved. By applying a recursive scheme, using the same idea, we obtain the desired spanner of size $\mathcal{O}(\varepsilon^{-O(d)}\vartheta^{-6}n \log n(\log \log n)^6)$. See Sect. 4 for details.

ϑ -Reliable $(1 + \varepsilon)$ -spanner with bounded spread Since the general construction in \mathbb{R}^d has some additional factors that seem unnecessary, we present an optimal construction for the case when the point set has bounded spread, that is, the ratio of the largest and the smallest distance of point pairs is bounded by a polynomial of n . Specifically, for points with spread Φ in \mathbb{R}^d , for d being constant, 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 Sect. 5 and Lemma 5.10 for details.

Shadow Underlying our construction is the notion of identifying the points that lose 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 Sect. 3.1 for details and relevant results in one dimension.

The competition Independently of this work, Bose et al. [7] 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, but worse, than our bound of $\mathcal{O}(n \log n(\log \log n)^6)$ edges. For both constructions, the fundamental building blocks are expander graphs. However, the construction in [7] does not use the one-dimensional construction directly, but uses well-separated pair decomposition, centroid decomposition, and ideas for maintaining orders. In particular, inspired by the preliminary work of Buchin et al. [11], in October 2018, Bose et al.² achieved independently essentially the same one-dimensional results by using expanders. Our own improved one-dimensional results were announced earlier in a public talk on June 12, 2018 (this was a YRF talk in SoCG 2018). To keep things in perspective, the results of Bose et al. [7] are equivalent to our results (up to a log factor), while being technically quite different (in two and higher dimensions). In particular, their one-dimensional construction has better dependency on the reliability parameter. The main advantage of our result, beyond the

² Available online here: <http://cglab.ca/~morin/publications/drafts/robust2/robust2-2018-10-29.pdf>.

aforementioned (all-important) log factor, is that it is (conceptually) simpler—the use of locality sensitive orderings [14] makes all the difference.

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*) Assume that we are given

- (i) a parameter $t \geq 1$,
- (ii) a graph $G = (P, E)$ that is a t -spanner,
- (iii) a function $f : \mathbb{N} \rightarrow \mathbb{R}_+$, and
- (iv) two point sets $P_1, P_2 \subseteq P$.

The graph G is an $f(k)$ -robust t -spanner for $P_1 \times P_2$ 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

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

induced by G on $P \setminus B$ is a t -spanner for $(P_1 \setminus B^+) \times (P_2 \setminus B^+)$. 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 a $(1 + \vartheta)k$ -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

For a set X of vertices in a graph $G = (V, E)$, let

$$\Gamma(X) = \{v \in V \mid uv \in E \text{ and } u \in X\}$$

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. We provide the proof only for the sake of completeness, as this is well known. The survey by Hoory et al. [17] gives a comprehensive study of expanders.

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 $|\Gamma(X)| > (1 - \xi)|R|$, and

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

Proof This is a variant of an expander graph. See [21, Sect. 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, fix a subset $X \subseteq L$ of size at least $\xi|L|$, and fix a subset on the right, $Z \subseteq R$ of size $\leq (1 - \xi)|R|$. Notice that there are at most 2^n choices for both X and Z . The probability that all the edges we picked for the vertices of X stay inside Z , is at most

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

since $c \geq 3/\xi^2$ 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. \square

Remark The above construction is randomized, but simple algebraic deterministic constructions are known [3,15,23]. One can improve their expansion ratio by repeatedly squaring them. These constructions are implicit, and are easily constructed quickly on the fly. As such, in the following, we assume that the expanders of Lemma 2.4 are readily available for use in our constructions.

Remark It is not hard to show that regular expanders (i.e., not bipartite, as constructed above) are reliable graphs. That is, deleting a set of vertices leaves almost all the remaining vertices in a single connected component.

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 , such that $h \leq i$ and $|[h : i] \cap B| \geq \alpha|[h : i]|$. The α -shadow of




Fig. 1 Consider a set $B \subseteq [n]$, the points $p_i = (i, 1 + \lceil [1, i - 1] \cap B \rceil)$, for $i = 1, \dots, n$, and the associated polyline σ resulting from connecting these points from left to right (the squares are the points that belong to B). A number $i \in [n] \setminus B$ is in the left α -shadow of B , if the ray emanating from p_i to the right, with slope α , hits σ

B , denoted by $\mathcal{S}(\alpha, B)$, is the union of the left α -shadow and the right α -shadow. See Fig. 1 for a visual interpretation of the shadow.

Lemma 3.2 For any $B \subseteq [n]$ and $\alpha \in (0, 1)$, $|\mathcal{S}(\alpha, B)| \leq (1 + 2\lceil 1/\alpha \rceil)|B|$.

Proof Fix a set $B = \{q_1 < \dots < q_b\}$. Let $P_0 = \{p_1 < \dots < p_{n-b}\}$ be all the points of $[n] \setminus B$, where $b = |B|$. Let the set H be initially empty. Consider the process where in the i th iteration, for $i = 1, \dots, b$, L_i (resp. R_i) is the set of $\lceil 1/\alpha \rceil$ largest (resp. smallest) numbers of P_{i-1} that are smaller (resp. larger) than q_i . We set $H = H \cup L_i \cup R_i \cup \{q_i\}$, and $P_i = P_{i-1} \setminus (L_i \cup R_i)$.

The claim is that if $k \in [n]$ is in (say) the left α -shadow of B , then $k \in H$. If $k \in B$ then the claim is immediate, so assume that $k \notin B$. Next, fix a witness interval $J = [k : k']$, such that $|J \cap B| \geq \alpha|J|$. Observe that the above process handles all the elements of $J^* = B \cap J$ in turn. When handling an element $q \in J^*$, as long as $k \notin H$, the process either adds k to H , or adds at least $1 + \lceil 1/\alpha \rceil$ new points of J (that were not previously in H) to H —indeed, the last $1 + \lceil 1/\alpha \rceil$ “free” points in the interval $[k : q]$ are added to H . Suppose that $k \notin H$ at the end of the process. Then, we have

$$|J| \geq |H \cap J| \geq (1 + \lceil 1/\alpha \rceil)\alpha|J| > |J|,$$

which is a contradiction.

Applying the above argument symmetrically implies that, at the end of the process, H contains all the points in the right and left shadows of B . □

One can get a better bound if α is close to one, as testified by the following. This is crucial to achieve ϑ -reliability, since Lemma 3.2 is not sharp when α is close to one.

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 starting at j and adding 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

$$\frac{\ell - 1}{\ell} \geq \alpha \iff \ell - 1 \geq \alpha\ell \iff \ell \geq \frac{1}{1 - \alpha} \iff \ell \geq \left\lceil \frac{1}{1 - \alpha} \right\rceil \geq 3,$$

as $\alpha \geq 2/3$. In particular, by the minimality of I , we have 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 avoid 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 correspond 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

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

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 $[i : j]$, which we refer to as the *block* of v , see Fig. 2. Let \mathcal{I} be the resulting set of all blocks. On 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 $\text{next}(I)$ and $\text{prev}(I)$ be the blocks (on 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

Here, we show that the resulting graph H is an $\mathcal{O}(k)$ -robust 1-spanner with $\mathcal{O}(n \log n)$ edges.

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

Proof Let $h = \log n$ be the depth of the tree T . For $i = 1, 2, \dots, h$, on the i th level 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 on 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. Summing this up for all levels implies the bound. □

There is a natural path between two leaves in the tree T going through their lowest common ancestor. However, we need something somewhat different here, as the path has to move forward (from left to right) 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 on 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.




Fig. 2 The binary tree built over $[n]$. The block of node v is the interval $[i : j]$

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, 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


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

be the resulting walk on the blocks where we removed repeated blocks. Figure 3 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 , see Fig. 4. 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,



(a) The canonical path in the tree.



(b) The canonical path on the blocks.

Fig. 3 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



Fig. 4 Contamination

$I_j \subseteq I_{j+1}$. Thus we have that either $i \in I_t$, or $i \in \text{prev}(I_t)$, or $i \in \text{prev}(\text{prev}(I_t))$, since there are at most $|I_t| + |I_t|/2 + \dots + 2 + 1 = 2|I_t| - 1$ 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 .

So, let r be the maximum number in I_t , and observe that $||i : r|| \leq 3|I_t|$. Furthermore, since I_t is α -contaminated, we have

$$||i : r| \cap B| \geq |I_t \cap B| \geq \alpha|I_t| \geq \frac{\alpha}{3}||i : r||.$$

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. \square

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 proven 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 $|B^+| \leq (1 + 2\lceil 3/\alpha \rceil)|B| \leq 200|B|$.

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 (after removing the vertices of B). 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 $3/4$ of the vertices have 1-paths from i (i.e., $|I_t \cap \mathcal{S}| \geq 3|I_t|/4$).

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 $3|I_t|/4$ vertices of I_t are in \mathcal{S} . As such, at least $15|I_{t+1}|/16$ 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 $|I_{t+1} \cap \mathcal{S}| \geq (15/16 - \alpha)|I_{t+1}| \geq 3|I_{t+1}|/4$, 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 $|I_{t+1} \cap \mathcal{S}| \geq 3|I_{t+1}|/4$.

The symmetric claim for the descent part of the path is handled in a similar fashion. Therefore, at least $3/4$ 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. \square

Note that it is easy to generalize the construction for arbitrary n . Let h be the 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 $[n]$ do 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

Here, we show how to extend Theorem 3.6, to build a 1-spanner on $[n]$, such that for any fixed $\vartheta \in (0, 1)$ 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. There are two main new ingredients: (i) a shifting scheme, and (ii) using much larger intervals when ascending from a level upward. Unfortunately, the details are somewhat tedious.

3.3.1 The Construction

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

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

where c is a sufficiently large constant (e.g., $c \geq 512$). We first connect any $i \in [n]$ to all the vertices that are within 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

$$\Delta(i, j) = 1 + \frac{(j - 1)2^i}{N} - 2^i.$$

For a fixed i , the $\Delta(i, j)$'s are N equally spaced numbers in the interval $[1 - 2^i : 1 - 2^i/N]$, starting at its left endpoint. Here, i is the *resolution* of $\Delta(i, j)$, the *shift* corresponding to resolution i is $2^i/N$, and the number of different shifts is N . For $k = 0, \dots, n/2^i$, and i, j as above, the corresponding *block* is

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




Fig. 5 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 $n/2^i + 1 = 9$ blocks per each shift

Such a block is an interval of length 2^i that starts at $\Delta(i, j) + k2^i$, see Fig. 5. The set of all intervals/blocks of interest is

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

Constructing the graph H_ϑ . Let $G_E(i, j, k)$ denote the expander graph of Lemma 2.4, constructed over $\mathbb{I}(i, j, k)$ and $\mathbb{I}(i, j, k + 1)$, with the value of the parameter ξ as specified in (3.1). We define H_ϑ to be the union of all the graphs G_E over all choices of i, j, k , also including the graph G_0 (described above). The last step is to delete vertices from H_ϑ that are outside the range of interest $[n]$.

Remark As before, if n is not a power of two, repeat the construction on $\lceil \text{pow}_2(n) \rceil$, and remove redundant vertices.

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 $\mathcal{O}(\log n)$ resolutions. For every such resolution the number of different shifts is $N = \mathcal{O}(1/\vartheta^2)$. For every shift, the number of edges created is bounded by $\mathcal{O}(n\xi^{-2}) = \mathcal{O}(n/\vartheta^4)$, by Lemma 2.4. Thus, H_ϑ has $\mathcal{O}(\vartheta^{-6}n \log n)$ edges. \square

In the following, let $\llbracket s, \ell \rrbracket = [s : s + \ell - 1]$ be the set of consecutive integers starting at s and 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 H_ϑ , if $x < y$ and there is a 1-path between x and y in H_ϑ . For a set $B \subseteq [n]$, and a vertex s , let $\mathcal{D} = \mathcal{D}(H_\vartheta, s, B)$ be the set of all descendants of s in the graph $H_\vartheta \setminus B$. Similarly, for a vertex t , let $\mathcal{A} = \mathcal{A}(H_\vartheta, t, B)$ be the set of ancestors of t in $H_\vartheta \setminus B$. 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 $H_\vartheta \setminus B$. In a symmetric fashion, the set of ancestors in I that can reach a node t is denoted by $I \cap \mathcal{A}$.

Now, we show that if a point outside of the shadow has a reasonably large fraction of descendants in an interval, then, one can find an extended interval, which is $\Theta(1/\vartheta)$ times longer and has the same property. The crucial part is to carefully choose two

consecutive blocks of \mathcal{I} , so that the expander between them can be used to extend the set of descendants.

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 $\mathcal{D} = \mathcal{D}(H_\vartheta, s, B)$, and assume that $|\llbracket s, h \rrbracket \cap \mathcal{D}| \geq \vartheta h/32$. Then, for some number ℓ , $8h/\vartheta \leq \ell \leq ch/(8\vartheta)$, we have $|\llbracket s, \ell \rrbracket \cap \mathcal{D}| \geq \vartheta \ell/32$.*

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

$$\Delta = \text{pow}_2 \left(\frac{\vartheta h}{64} \right) \implies \frac{\vartheta h}{64} \leq \Delta \leq \frac{\vartheta h}{32}$$

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 $L = \mathcal{I}(i, j, k)$ lies in the interval $\llbracket s + h, \Delta \rrbracket$. Notice that $\llbracket s, h \rrbracket \subseteq L$, since

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

holds. Let $R = \mathcal{I}(i, j, k + 1)$ and $\ell = \text{right}(R) - s + 1$, where $\text{right}(R)$ is the right endpoint of the interval R , see Fig. 6. Observe that $\vartheta h/64 \leq \Delta \leq \vartheta h/32$ and

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

since $c \geq 512$. Similarly,

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

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

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




Fig. 6 The intervals L and R and their relation to s, h, Δ , and ℓ

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

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

Let $U' \subseteq R$ 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 $|U'| \geq (1 - \xi)|R|$, where $G_E(i, j, k)$ is the expander graph built over L and R . We have that

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

Since $\xi \leq \vartheta/16$, we have

$$\frac{|\llbracket s, \ell \rrbracket \cap \mathcal{D}|}{\ell} \geq \frac{(\tau - \xi)2^i}{2 \cdot 2^i} = \frac{\tau - \xi}{2} \geq \frac{\vartheta/8 - \vartheta/16}{2} = \frac{\vartheta}{32}. \quad \square$$

Remark One can state a symmetric version of Lemma 3.9 about the number of ancestors that can reach a target node t .

Lemma 3.10 *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 L and R be two adjacent consecutive blocks of the same size in \mathcal{I} (see (3.2)) such that $s \in L$ and $t \in R$, 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 $\min(\text{right}(L) - s, t - \text{left}(R))$ is maximized (i.e., the common boundary between the two intervals is as close to the middle $(s + t)/2$ as possible). Let $2^i = |R| = |L|$. It is easy to verify that $2^i/2 \leq \llbracket s : t \rrbracket \leq 2 \cdot 2^i$. Indeed, the lower bound holds by the minimality of L and R . Otherwise, the right half of L and the left half of R would also be a valid choice and would have smaller size. The upper bound follows from the fact that $|L| + |R| = 2 \cdot 2^i$.

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 $|L_0 \setminus B| \geq \vartheta |L_0|/4$ and $|R_0 \setminus B| \geq \vartheta |R_0|/4$. 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 a fraction of at least $\vartheta/32$ 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 $L_{j+1} \not\subseteq L$. By the choice of L and 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

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

Applying the same argumentation, using Lemma 3.9 for the reachable ancestors, we have that

$$\frac{|R \cap \mathcal{A}|}{|R|} \geq \xi$$

(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 $8|L_i| \leq 8|L_i|/\vartheta \leq |L_{i+1}|$ 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 $2 \log n$. \square

Theorem 3.11 *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 3.10. \square

4 Building a Reliable Spanner in \mathbb{R}^d

4.1 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. We use a recent result of Chan et al. [14], which

introduced *locality-sensitive orderings*. These orderings (which are the same as total orders and are extensions of the \mathcal{Z} -order) 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 $p \prec_\sigma q$, let

$$(p, q)_\sigma = \{z \in [0, 1)^d \mid p \prec_\sigma z \prec_\sigma q\}$$

be the open interval between p and q in the order σ . Further, let $\text{ball}(p, r) = \{z \in [0, 1)^d \mid \|p - z\| \leq r\}$ denote the ball centered at p with radius r .

Theorem 4.1 [14] *For $\delta \in (0, 1)$, there is a set $\Pi^+(\delta)$ of at most $M(\delta) = \mathcal{O}(\delta^{-d} \log \delta^{-1})$ 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*

- $p \prec_\sigma q$,
- $(p, z)_\sigma \subseteq \text{ball}(p, \delta\ell)$,
- $(z, q)_\sigma \subseteq \text{ball}(q, \delta\ell)$, and
- $z \in \text{ball}(p, \delta\ell)$ or $z \in \text{ball}(q, \delta\ell)$.

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

First, we give a very simple construction and analysis, for building reliable spanners, using the theorem above and the 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, with similar ideas but used 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 $\delta = \varepsilon / (c \log n)$,

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

and c be some sufficiently large constant. Next, let $\vartheta' = \vartheta / M$ and $\Pi^+ = \Pi^+(\delta)$ be the set of orderings from Theorem 4.1. For each ordering $\sigma \in \Pi^+$, compute the ϑ' -reliable exact spanner G_σ of P , see Theorem 3.11, according to σ . Let G be the graph resulting from 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 $|B^+| \leq (1 + M \cdot \vartheta')|B| = (1 + \vartheta)|B|$.

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 $(p, z)_\sigma \subseteq \text{ball}(p, \delta\ell)$ and $(z, q)_\sigma \subseteq \text{ball}(q, \delta\ell)$ (and z is in one of these balls).

By Lemma 3.10, 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:

- For all $j \neq i$, $\|p_j - p_{j+1}\| \leq 2\delta\ell$, since p_j and p_{j+1} are contained in a ball of radius $\delta\ell$.
- $\|p_i - p_{i+1}\| \leq \|p_i - p\| + \|p - q\| + \|q - p_{i+1}\| \leq \ell + 2\delta\ell$.

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

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

4.2 An Improved Construction

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

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

and c be some sufficiently large constant. Next, let $\vartheta' = \vartheta/(3NM)$ where $N = \lceil \log \log n \rceil + 1$ and $\Pi^+ = \Pi^+(\delta)$ be the set of orderings from Theorem 4.1. For each ordering $\sigma \in \Pi^+$, compute the ϑ' -reliable exact spanner G_σ of P , see Theorem 3.11, according to σ . Let G be the graph resulting from 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 3.11. Each of these graphs has $\mathcal{O}\left((\vartheta')^{-6} n \log n\right)$ edges. Therefore, the size of G is

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

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

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

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

The claim is 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 3.11, for the right choice of σ , 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:

- $\|p' - q'\| \leq \|p - q\| + 2\delta\|p - q\| = (1 + 2\varepsilon/c)\|p - q\|$.
- $\|p - p'\| \leq 2\delta\|p - q\| = 2\varepsilon\|p - q\|/c$, and similarly $\|q - q'\| \leq 2\varepsilon\|p - q\|/c$.
- $p', q' \in P \setminus B_{N-1}$.

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 $p, q \in P \setminus B_N$, to 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 both $\|p - p'\|$ and $\|q - q'\|$ are much smaller than $\|p - q\|$. 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 on the i th level of the recursion. Thus, we have that $Q_0 = \{(p, q)\}$, $Q_1 = \{(p, p'), (q, q')\}$, and so on. We repeat this $N - 1$ times. On 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

- $\|x' - y'\| \leq \|x - y\|(1 + 2\varepsilon/c) \leq (2\varepsilon/c)^i(1 + 2\varepsilon/c)\|p - q\|$;
- $\|x - x'\| \leq 2\varepsilon\|x - y\|/c \leq (2\varepsilon/c)^{i+1}\|p - q\|$ and $\|y - y'\| \leq (2\varepsilon/c)^{i+1}\|p - q\|$;
- $x', y' \in P \setminus B_{N-(i+1)}$.

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 Theorems 4.1 and 3.11 to obtain a path of length at most $2\|x - y\| \log n$ 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

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

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. □

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 $d(B, C) \geq s \cdot \max(\text{diam}(B), \text{diam}(C))$, where $d(B, C) = \min_{p \in B, q \in C} \|p - q\|$.

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 $\mathcal{W} = \{(B_i, C_i)\}_{i=1}^m$ of pairs of subsets of P , such that

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

The well-separated pair decomposition was introduced by Callahan and Kosaraju [12]. 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. Here, 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 [2, 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} of P , such that each pair consists of two cells of T , 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

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

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 so that the expander property holds with $\xi = \vartheta/8$. Furthermore, for every two nodes 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 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 $\mathcal{O}(\xi^{-2}\varepsilon^{-d}n \log \Phi(P))$. \square

We note that there are point sets, with unbounded spread, such that any WSPD on them has weight $\Omega(n^2)$. A simple example is a sequence of points along a line with exponentially increasing distances. Thus, requiring the bounded spread on the point set is unavoidable to achieve optimal size by using the above construction.

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,

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

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

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 $B^+ = \mathcal{S}(\gamma, B)$, and the residual graph $H = G_\Phi \setminus B$. For a point $p \in P_u \setminus B^+$, we define the set*

$$R_u(p) = \{q \in P_u \setminus B \mid d_H(p, q) \leq 2 \cdot \text{diam}(\square_u)\}.$$

Then, $|R_u(p)| \geq 3\xi|P_u|$.

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 $pq_1, q_1q_2, \dots, q_{m-1}q_m$ 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

$$|Q_i| \geq (1 - \xi)n(u_i) - b(u_i) \geq (1 - \xi - \gamma)n(u_i) = \left(\frac{\vartheta}{2} - \xi\right)n(u_i) = 3\xi n(u_i),$$

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

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

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

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

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

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, 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 Fig. 7. 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 $|Y \cap R_v(q)| = |R_v(q) \setminus (P_v \setminus Y)| \geq |R_v(q)| - |P_v \setminus Y| \geq 3\xi n(v) - \xi n(v) > 0$.




Fig. 7 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 $G_\Phi \setminus B$

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

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

6 Conclusions

A natural open question is whether ϑ -reliable spanners can be constructed with $\mathcal{O}(n \log n)$ edges for general point sets. There are different approaches that lead to near optimal bounds. While we use the one-dimensional construction with a careful application of locality-sensitive orderings, Bose et al. [7] uses WSPD, centroid decomposition and an idea of Willard [25] for order maintenance. Another natural open question is how to construct reliable spanners that are required to be subgraphs of a given graph.

Acknowledgements The authors thank Pat Morin for his useful comments on this manuscript, and keeping us updated about his own work [7]. The authors also thank the anonymous referees for the detailed and useful comments.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Abam, M.A., de Berg, M., Farshi, M., Gudmundsson, J.: Region-fault tolerant geometric spanners. *Discrete Comput. Geom.* **41**(4), 556–582 (2009)

2. Abam, M.A., Har-Peled, S.: New constructions of SSPDs and their applications. *Comput. Geom.* **45**(5–6), 200–214 (2012)
3. Alon, N., Schwartz, O., Shapira, A.: An elementary construction of constant-degree expanders. *Comb. Probab. Comput.* **17**(3), 319–327 (2008)
4. Aronov, B., de Berg, M., Cheong, O., Gudmundsson, J., Haverkort, H., Smid, M., Vigneron, A.: Sparse geometric graphs with small dilation. *Comput. Geom.* **40**(3), 207–219 (2008)
5. Arya, S., Mount, D.M., Smid, M.: Randomized and deterministic algorithms for geometric spanners of small diameter. In: 35th Annual Symposium on Foundations of Computer Science (Santa Fe 1994), pp. 703–712. IEEE, Los Alamitos (1994)
6. Arya, S., Mount, D.M., Smid, M.: Dynamic algorithms for geometric spanners of small diameter: randomized solutions. *Comput. Geom.* **13**(2), 91–107 (1999)
7. Bose, P., Carmi, P., Dujmović, V., Morin, P.: Near-optimal $O(k)$ -robust geometric spanners (2018). [arXiv:1812.09913](https://arxiv.org/abs/1812.09913)
8. Bose, P., Carmi, P., Farshi, M., Maheshwari, A., Smid, M.: Computing the greedy spanner in near-quadratic time. *Algorithmica* **58**(3), 711–729 (2010)
9. Bose, P., Dujmović, V., Morin, P., Smid, M.: Robust geometric spanners. *SIAM J. Comput.* **42**(4), 1720–1736 (2013)
10. Buchin, K., Har-Peled, S., Oláh, D.: A spanner for the day after. In: 35th International Symposium on Computational Geometry. *Leibniz Int. Proc. Inform.*, vol. 129, # 19. Leibniz-Zent. Inform., Wadern (2019)
11. Buchin, K., Hulshof, T., Oláh, D.: $O(k)$ -robust spanners in one dimension (2018). [arXiv:1803.08719](https://arxiv.org/abs/1803.08719)
12. Callahan, P.B., Kosaraju, S.R.: A decomposition of multidimensional point sets with applications to k -nearest-neighbors and n -body potential fields. *J. ACM* **42**(1), 67–90 (1995)
13. Carmi, P., Chaitman, L.: Stable roommates and geometric spanners. In: Canadian Conference on Computational Geometry (Winnipeg 2010), pp. 31–34 (2010)
14. Chan, T.M., Har-Peled, S., Jones, M.: On locality-sensitive orderings and their applications. In: 10th Innovations in Theoretical Computer Science. *Leibniz Int. Proc. Inform.*, vol. 124, # 21. Leibniz-Zent. Inform., Wadern (2019)
15. Gabber, O., Galil, Z.: Explicit constructions of linear-sized superconcentrators. *J. Comput. Syst. Sci.* **22**(3), 407–420 (1981)
16. Gudmundsson, J., Levcopoulos, Ch., Narasimhan, G.: Fast greedy algorithms for constructing sparse geometric spanners. *SIAM J. Comput.* **31**(5), 1479–1500 (2002)
17. Hoory, S., Linial, N., Wigderson, A.: Expander graphs and their applications. *Bull. Am. Math. Soc.* **43**(4), 439–561 (2006)
18. Levcopoulos, Ch., Narasimhan, G., Smid, M.: Efficient algorithms for constructing fault-tolerant geometric spanners. In: 30th Annual ACM Symposium on Theory of Computing (Dallas 1998), pp. 186–195. ACM, New York (1999)
19. Levcopoulos, Ch., Narasimhan, G., Smid, M.: Improved algorithms for constructing fault-tolerant spanners. *Algorithmica* **32**(1), 144–156 (2002)
20. Lukovszki, T.: New results of fault tolerant geometric spanners. In: Workshop on Algorithms and Data Structures (Vancouver 1999). *Lecture Notes in Computer Science*, vol. 1663, pp. 193–204. Springer, Berlin (1999)
21. Motwani, R., Raghavan, P.: *Randomized Algorithms*. Cambridge University Press, Cambridge (1995)
22. Narasimhan, G., Smid, M.: *Geometric Spanner Networks*. Cambridge University Press, Cambridge (2007)
23. Reingold, O., Vadhan, S., Wigderson, A.: Entropy waves, the zig-zag graph product, and new constant-degree expanders. *Ann. Math.* **155**(1), 157–187 (2002)
24. Smid, M.H.M.: Geometric spanners with few edges and degree five. In: *Theory of Computing 2006* (Hobart), pp. 7–9. Australian Computer Society, Darlinghurst (2006)
25. Willard, D.E.: Maintaining dense sequential files in a dynamic environment (extended abstract). In: 14th Annual ACM Symposium on Theory of Computing (San Francisco 1982), pp. 114–121. ACM, New York (1982)