Sometimes Reliable Spanners of Almost Linear Size

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

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

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

Abstract

Reliable spanners can withstand huge failures, even when a linear number of vertices are deleted from the network. In case of failures, some of the remaining vertices of a reliable spanner may no longer admit the spanner property, but this collateral damage is bounded by a fraction of the size of the attack. It is known that $\Omega(n \log n)$ edges are needed to achieve this strong property, where n is the number of vertices in the network, even in one dimension. Constructions of reliable geometric $(1+\varepsilon)$-spanners, for n points in $\mathbb{R}^d$, are known, where the resulting graph has $\mathcal{O}(n \log n \operatorname{loglog}^6 n)$ edges.

Here, we show randomized constructions of smaller size spanners that have the desired reliability property in expectation or with good probability. The new construction is simple, and potentially practical – replacing a hierarchical usage of expanders (which renders the previous constructions impractical) by a simple skip list like construction. This results in a 1-spanner, on the line, that has linear number of edges. Using this, we present a construction of a reliable spanner in $\mathbb{R}^d$ with $\mathcal{O}(n \operatorname{loglog}^2 n \operatorname{logloglog} n)$ edges.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Geometric spanners, vertex failures, reliability

Digital Object Identifier 10.4230/LIPIcs.ESA.2020.27

Funding *Sariel Har-Peled*: Work on this paper was partially supported by a NSF AF awards CCF-1421231 and CCF-1907400.
Dániel Oláh: Supported by the Netherlands Organisation for Scientific Research (NWO) through Gravitation-grant NETWORKS-024.002.003.

1 Introduction

Geometric graphs are such that their vertices are points in the d-dimensional Euclidean space $\mathbb{R}^d$ and edges are straight line segments. The quality or efficiency of a geometric graph is often measured in terms of the ratio of shortest path distances and geometric distances between its vertices. Let $G = (P, E)$ be a geometric graph, where $P \subset \mathbb{R}^d$ is a set of n points and E is the set of edges. The shortest path distance between two points $p, q \in P$ in the graph G is denoted by $\mathsf{d}_G(p, q)$ (or just $\mathsf{d}(p, q)$). The graph G is a t-spanner for some constant $t \geq 1$, if $\mathsf{d}(p, q) \leq t \cdot \|p - q\|$ holds for all pairs of points $p, q \in P$, where $\|p - q\|$ stands for the Euclidean distance of p and q. The spanning ratio, stretch factor, or dilation of a graph G is the minimum number $t \geq 1$ for which G is a t-spanner. A path between p and q is a *t-path* if its length is at most $t \cdot \|p - q\|$.

Table 1.1 Comparison of the size of constructions of reliable spanners and reliable spanners in expectation. The reliability parameter is $\vartheta > 0$, and, for dimensions $d \geq 2$, the graphs are $(1+\varepsilon)$-spanners for $\varepsilon > 0$.

	dim	# edges	constants
Reliable spanners			
Buchin et al. [4]	$d = 1$	$\mathcal{O}(n \log n)$	$\mathcal{O}\big(\vartheta^{-6}\big)$
	$d \geq 2$	$\mathcal{O}\big(n \log n \log \log^6 n\big)$	$\mathcal{O}\big(\varepsilon^{-7d}\vartheta^{-6}\log^7 \varepsilon^{-1}\big)$
Bose et al. [2]	$d \geq 1$	$\mathcal{O}\big(n \log^2 n \log\log n\big)$	?
Reliable spanners *in expectation*			
New results	$d = 1$	$\mathcal{O}(n)$	$\mathcal{O}\big(\vartheta^{-1}\log \vartheta^{-1}\big)$
	$d \geq 2$	$\mathcal{O}(n \log\log^2 n \log\log\log n)$	$\mathcal{O}\big(\varepsilon^{-2d}\vartheta^{-1}\log^3 \varepsilon^{-1}\log \vartheta^{-1}\big)$

We focus our attention to construct spanners that can survive massive failures of vertices. The most studied notion is fault tolerance [5, 7, 8, 9, 10], which provides a properly functioning residual graph if there are no more failures than a predefined parameter k. It is clear, that a k-fault tolerant spanner must have $\Omega(kn)$ edges to avoid small degree nodes, which can be isolated by deleting their neighbors. Therefore, fault tolerant spanners must have quadratic size to be able to survive a failure of a constant fraction of vertices. Another notion is robustness [3], which gives more flexibility by allowing the loss of some additional nodes by not guaranteeing t-paths for them. For a function $f : \mathbb{N} \to \mathbb{R}^+$ a t-spanner G is f-robust, if for any set of failed points B there is an extended set B^+ with size at most $f(|B|)$ such that the residual graph $G \setminus B$ has a t-path for any pair of points $p, q \in P \setminus B^+$. The function f controls the robustness of the graph - the slower the function grows the more robust the graph is. The benefit of robustness is that a near linear number of edges are enough to achieve it, even for the case when f is linear, there are constructions with nearly $\mathcal{O}(n \log n)$ edges. For $\vartheta \in (0, 1)$, a spanner that is f-robust with $f(k) = (1 + \vartheta)k$ is a ϑ-*reliable* spanner [4]. This is the strongest form of robustness, since the dilation can increase beyond t only for a tiny additional fraction of points. The fraction is relative to the number of failed vertices and controlled by the parameter ϑ.

Recently, Buchin et al. [4] showed a construction of reliable 1-spanners of size $\mathcal{O}(n \log n)$ in one dimension, and of reliable $(1+\varepsilon)$-spanners of size $\mathcal{O}\big(n \log n \operatorname{loglog}^6 n\big)$ in higher dimensions (the constant in the $\mathcal{O}$ depends on the dimension, ε, and the reliability parameter). An alternative construction, with slightly worse bounds, was given by Bose et al. [2].

Limitations of previous constructions. The construction of Buchin et al. [4] (and also the construction of Bose et al. [2]) relies on using expanders to get a 1-spanner for points on the line, and then extending it to higher dimensions. The spanner (in one dimension) has $\mathcal{O}(n \log n)$ edges. Unfortunately, even in one dimension, such a reliable spanner requires $\Omega(n \log n)$ edges, as shown by Bose et al. [3]. Furthermore, the constants involved in these constructions [2, 4] are quite bad, because of the usage of expanders. See Table 1.1 for a summary of the sizes of different constructions (together with the new results).

The problem. As such, the question is whether one can come up with simple and practical constructions of spanners that have linear or near linear size, while still possessing some reliability guarantee – either in expectation or with good probability.

Some definitions. Given a graph G, an *attack* $B \subseteq V(G)$ is a set of vertices that are being removed. The *damaged set* B^+, is the set of all the vertices which are no longer connected to the rest the graph, or are badly connected to the rest of the graph – that is, these vertices no longer have the desired spanning property. The *loss* caused by B, is the quantity $|B^+ \setminus B|$ (where we take the minimal damaged set). The *loss rate* of B is $\lambda(G, B) = |B^+ \setminus B| \,/\, |B|$. A graph G is ϑ*-reliable* if for any attack B, the loss rate $\lambda(G, B)$ is at most ϑ.

Randomness and obliviousness. As mentioned above, reliable spanners must have size $\Omega(n \log n)$. A natural way to get a smaller spanner, is to consider randomized constructions, and require that the reliability holds in expectation (or with good probability). Randomized constructions are (usually) still sensitive to adversarial attacks, if the adversary is allowed to pick the attack set after the construction is completed (and it is allowed to inspect it). A natural way to deal with this issue is to restrict the attacks to be *oblivious* – that is, the attack set is chosen before the graph is constructed (or without any knowledge of the edges).

In such an oblivious model, the loss rate is a random variable (for a fixed attack B). It is thus natural to construct the graph G randomly, in such a way that $\mathbb{E}[\lambda(G, B)] \leq \vartheta$, or alternatively, that the probability $\mathbb{P}[\lambda(G, B) \geq \vartheta]$ is small.

1-spanner. Surprisingly, the one-dimensional problem is the key for building reliable spanners. Here, the graph G is constructed over the set of vertices $[n] = \{1, \ldots, n\}$. An attack is a subset $B \subseteq [n]$. Given an attack B, the requirement is that for all $i, j \in [n] \setminus B^+$, such that $i < j$, there is a monotonically increasing path from i to j in $G \setminus B$ – here, the length of the path between i and j is exactly $j - i$. Since there is no distortion in the length of the path, such graphs are 1-spanners.

Our results. We give a randomized construction of a 1-spanner in one dimension, that is ϑ-reliable in expectation, and has size $\mathcal{O}(n)$. Formally, the construction has the property that $\mathbb{E}[\lambda(G, B)] \leq \vartheta$. This construction can also be modified so that $\lambda(G, B) \leq \vartheta$ holds with some desired probability. This is the main technical contribution of this work.

Next, following in the footsteps of the construction of reliable spanners, we use the one-dimensional construction to get $(1 + \varepsilon)$-spanners that are ϑ-reliable either in expectation or with good probability. The new constructions have size roughly $\mathcal{O}\big(n \operatorname{loglog}^2 n\big)$.

Main idea. We borrow the notion of shadow from the work of Buchin et al. [4]. A point p is in the α-shadow if there is a neighborhood of p, such that an α-fraction of it belongs to the attack set. One can think about the maximum α such that p is in the α-shadow of B as the *depth* of p (here, the depth is in the range $[0, 1]$). A point with depth close to one, are intuitively surrounded by failed points, and have little hope of remaining well connected. Fortunately, only a few points have depth truly close to one 1. The flip side is that the attack has little impact on shallow points (i.e., points with depth close to 0). Similar to people, shallow points are surrounded by shallow points. As such, only a small fraction of the shallow points needs to be strongly connected to other points in the graph, as paths from (shallow) points around them can then travel via these hub points.

To this end, similar in spirit to skip lists, we define a random gradation of the points $P = P_0 \supseteq P_1 \supseteq \ldots \supseteq P_{\log n}$, where $|P_i| = n/2^i$ – this is done via a random tournament tree. In each level, each point of P_i is connected to all its neighbors within a certain distance (which increases as i increases). Intuitively, because of the improved connectivity, the probability that a point is well-connected (after the attack) increases if they belong to higher level of the gradation. Thus, the probability of a shallow point to remain well connected is, intuitively, good. Specifically, we can quantify the probability of a vertex to lose its connectivity as a function of its depth. Combining this with bounds on the number of points of certain depths, results in bounds on the expected size of the damaged set.

Reliable skip lists. Our construction can be interpreted as a reliable construction of skip lists. Here, an attack removes certain cells in the skip list, which are no longer available. This can happen, for example, if the skip list is stored in a distributed fashion in a network, and certain nodes of the network are down. Our construction implies that one can withstand an attack with small expected loss. The previous work on skip graphs [1], or [4], presented constructions of variants of skip lists with somewhat similar properties, but using $\mathcal{O}(n \log n)$ pointers. The current construction requires only $\mathcal{O}(n)$ pointers.

Comparison to previous work. While we borrow some components of Buchin et al. [4], the basic scheme in the one-dimensional case, is new, and significantly different – the previous construction used expanders in a hierarchical way. The new construction requires different analysis and ideas. The extension to higher dimension is relatively straightforward and follows the ideas of Buchin et al. [4], although some modifications and care are necessary.

Paper organization. We review some necessary machinery in Section 2. The one-dimensional construction is described in Section 3. We describe the extension to higher dimensions in Section 4.

2 Preliminaries

Let $G = (P, E)$ be a t-spanner for some $t \geq 1$. An ***attack*** on G is a set of vertices B that fail, and no longer can be used. An attack is ***oblivious***, if the set B is picked without any knowledge of E.

▶ **Definition 1** (Reliable spanner). *Let $G = (P, E)$ be a t-spanner for some $t \geq 1$ constructed by a (possibly) randomized algorithm. Given an attack B, its* **damaged set** *B^+ is a smallest set, such that for any pair of vertices $u, v \in P \setminus B^+$, we have*

$$\mathsf{d}_{G \setminus B}(u, v) \leq t \cdot \|u - v\|,$$

that is, t-paths are preserved for all pairs of points not contained in B^+. The quantity $|B^+ \setminus B|$ is the **loss** *of G under the attack B. The* **loss rate** *of G is $\lambda(G, B) = |B^+ \setminus B| \,/\, |B|$. For $\vartheta \in (0, 1)$, the graph G is ϑ-***reliable** *if $\lambda(G, B) \leq \vartheta$ holds for any attack $B \subseteq P$.*

*Further, we say that the random graph G is ϑ-***reliable in expectation** *if $\mathbb{E}[\lambda(G, B)] \leq \vartheta$ holds for any oblivious attack $B \subseteq P$. For $\vartheta, \rho \in (0, 1)$, we say that the graph G is ϑ-***reliable with probability** *$1 - \rho$ if $\mathbb{P}[\lambda(G, B) \leq \vartheta] \geq 1 - \rho$ holds for any oblivious attack $B \subseteq P$.*

▶ Remark 2. We emphasize that in the latter case the graph is random and the expectation and the probability is taken with respect to the distribution of graphs.

Another remark is that the set B^+ is not unique, since one can (possibly) choose the point to include in B^+ for a pair that does not have a t-path in $G \setminus B$. However, this does not cause a problem in defining the loss rate.

▶ **Definition 3.** *Let $[n]$ denote the* interval *$\{1, \ldots, n\}$. Similarly, for x and y, let $[x \ldots y]$ denote the interval $\{x, x+1, \ldots, y\}$.*

We use the shadow notion as it was introduced by Buchin et al. [4].

▶ **Definition 4.** *Consider an arbitrary set $B \subseteq [n]$ and a parameter $\alpha \in (0,1)$. A number i is in the* **left α-shadow** *of B, if and only if there exists an integer $j \geq i$, such that $\left|[i \ldots j] \cap B\right| \geq \alpha \left|[i \ldots j]\right|$. 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 $|[h \ldots i] \cap B| \geq \alpha \,|[h \ldots i]|$. The left and right α-shadow of B is denoted by $\mathcal{S}_{\rightarrow}(\alpha, B)$ and $\mathcal{S}_{\leftarrow}(\alpha, B)$, respectively. The combined shadow is denoted by $\mathcal{S}(\alpha, B) = \mathcal{S}_{\rightarrow}(\alpha, B) \cup \mathcal{S}_{\leftarrow}(\alpha, B)$.*

▶ **Lemma 5** **([4]).** *For any set $B \subseteq [n]$, and $\alpha \in (0,1)$, we have that $|\mathcal{S}(\alpha, B)| \leq \left(1 + 2\lceil 1/\alpha \rceil\right)|B|$. Further, if $\alpha \in (2/3, 1)$, we have that $|\mathcal{S}(\alpha, B)| \leq |B|/(2\alpha - 1)$.*

▶ **Definition 6.** *Given a graph G over $[n]$, a* **monotone path** *between $i, j \in [n]$, such that $i < j$, is a sequence of vertices $i = i_1 < i_2 < \cdots < i_k = j$, such that $i_{\ell-1} i_\ell \in E(G)$, for $\ell = 2, \ldots, k$.*

A monotone path between i and j has length $|j - i|$. Throughout the paper we use $\log x$ and $\ln x$ to denote the base 2 and natural base logarithm of x, respectively. For any set $A \subseteq P$, let $A^c = P \setminus A$ denote the complement of A. For two integers $x, y > 0$, let $x_{\uparrow y} = \lceil x/y \rceil\, y$.

3 Reliable spanners in one dimension

We show how to build a random graph on $[n]$ that still has monotone paths almost for all vertices that are not directly attacked. First, in Section 3.2, we show that our construction is ϑ-reliable in expectation. Then, in Section 3.3, we show how to modify the construction to obtain a 1-spanner that is ϑ-reliable with probability $1 - \rho$.

3.1 Construction

The input consists of a parameter $\vartheta > 0$ and the point set $P = [n] = \{1, \ldots, n\}$. The backbone of the construction is a random elimination tournament. We assume that n is a power of 2 as otherwise one can construct the graph for the next power of two, and then throw away the unneeded vertices.

The tournament is a full binary tree, with the leafs storing the values from 1 to n, say from left to right. The value of a node is computed randomly and recursively. For a node, once the values of the nodes were computed for both children, it randomly copies the value of one of its children, with equal probability to choose either child. Let P_i be the values stored in the ith bottom level of the tree. As such, $P_0 = P$, and $P_{\log n}$ is a singleton. Each set P_i can be interpreted as an ordered set (from left to right, or equivalently, by value).

Let

$$\alpha = 1 - \frac{\vartheta}{8} \qquad \text{and} \qquad \varepsilon = \frac{8(1-\alpha)}{c \ln \vartheta^{-1}} = \frac{\vartheta}{c \ln \vartheta^{-1}}, \tag{3.1}$$

where $c > 1$ is a sufficiently large constant. Let M be the smallest integer for which $|P_M| \leq 2^{M/2}/\varepsilon$ holds (i.e., $M = \lceil (2/3)\log(\varepsilon n) \rceil$). For $i = 0, 1, \ldots, M$, and for all $p \in P_i$ connect p with the first

$$\ell(i) = \left\lceil \frac{2^{i/2}}{\varepsilon} \right\rceil \tag{3.2}$$

successors (and hence predecessors) of p in P_i. Let E_i be the set of all edges in level i. The graph G on P is defined as the union of all edges over all levels – that is, $E(G) = \cup_{i=0}^{M} E_i$. Note, that top level of the graph G is a clique.

▶ Remark 7. Before dwelling on the correctness of the construction, note that the obliviousness of the attack is critical. Indeed, it is quite easy to design an attack if the structure of G is known. To this end, let B_i be the set of $\ell(M) = O(n^{1/3}/\varepsilon)$ values of P_i closest to $n/2$ – namely, we are taking out the middle-part of the graph, that belongs to the ith level. Consider the attack $B = \cup B_i$. It is easy to verify that this attack breaks G into at least two disconnected graphs, each of size at least $n/2 - O(n^{1/3}\varepsilon^{-1}\log n)$.

3.2 Analysis

▶ **Lemma 8.** *The graph G has $\mathcal{O}\left(n\vartheta^{-1}\log\vartheta^{-1}\right)$ edges.*

Proof. The number of edges contributed by a point in P_i is at most $\ell(i)$ at level i, and $|P_i| = n/2^i$. Thus, we have

$$|E(G)| \leq \sum_{i=0}^{M} |P_i| \cdot \ell(i) \leq \sum_{i=0}^{M} \frac{n}{2^i} \cdot \left\lceil \frac{2^{i/2}}{\varepsilon} \right\rceil \leq \sum_{i=0}^{M} \frac{n}{2^i} \cdot \frac{2 \cdot 2^{i/2}}{\varepsilon} \leq \frac{n}{\varepsilon} \cdot \sum_{i=0}^{\infty} \frac{2}{2^{i/2}} = \mathcal{O}\left(\frac{n}{\varepsilon}\right). \quad ◀$$

Fix an attack $B \subseteq P$. The high-level idea is to show that if a point $p \in P \setminus B$ is far enough from the faulty set, then, with high probability, there exist monotone paths reaching far from p in both directions. For two points $p < q$, we show that if both p and q have far reaching monotone paths, then the path going to the right from p, and the path going to the left from q must cross each other, which in turn implies, that there is a monotone path between p and q. Therefore, it is enough to bound the number of points that does not have far reaching monotone paths.

▶ **Definition 9** (Stairway). *Let $p \in P$ be an arbitrary point. The path $p = p_0, p_1, \ldots, p_j$ is a right (resp., left)* **stairway** *of p to level j, if*

- **(i)** $p = p_0 \leq p_1 \leq \cdots \leq p_j$ *(resp., $p \geq p_1 \geq \cdots \geq p_j$),*
- **(ii)** *if $p_i \neq p_{i+1}$, then $p_i p_{i+1} \in E_i$, for $i = 0, 1, \ldots, j-1$,*
- **(iii)** $p_i \in P_i$*, for $i = 1, \ldots, j$.*

Furthermore, a stairway is **safe** *if none of its points are in the attack set B. A right (resp., left) stairway is* **usable***, if $[p_j \ldots n] \cap P_j$ (resp., $[1 \ldots p_j] \cap P_j$) forms a clique in G. Let $T \subseteq P$ denote the set of points that have a safe and usable stairway to both directions.*

Let $\alpha_k = \alpha/2^k$, for $k = 0, 1, \ldots, \log n$. Let $\mathcal{S}_k = \mathcal{S}(\alpha_k, B)$ be the α_k-shadow of B, for $k = 0, 1, \ldots, \log n$. Observe that $\mathcal{S}_0 \subseteq \mathcal{S}_1 \subseteq \cdots \subseteq \mathcal{S}_{\log n}$, and there is an index j such that $\mathcal{S}_j = P$, if $B \neq \emptyset$. A point is classified according to when it gets "buried" in the shadow. A point p, for $k \geq 1$, is a ***kth round*** point, if $p \in \mathcal{S}_k \setminus \mathcal{S}_{k-1}$. Intuitively, a kth round point is more likely to have a safe stairway the larger the value of k is.

▶ **Definition 10.** *A point is* **bad** *if it belongs to B, or it does not have a right or left stairway that is safe and usable. Formally, a point $p \in P$ is bad, if and only if $p \in P \setminus T$.*

▶ **Lemma 11.** *For any two points $p, q \in T$ that are not bad, there is a monotone path connecting p and q in the residual graph $G \setminus B$.*

Proof. Suppose we have $p < q$. Let $(p, p_1, \ldots, p_{j(p)})$ be a safe usable right stairway starting from p and $(q, q_1, \ldots, q_{j(q)})$ be a safe usable left stairway from q. These stairways exist, since $p, q \in T$. Let $j = \min(j(p), j(q))$ and consider the stairways $(p, p_1, \ldots, p_j)$ and $(q, q_1, \ldots, q_j)$. Notice that both are safe and at least one of them is usable.

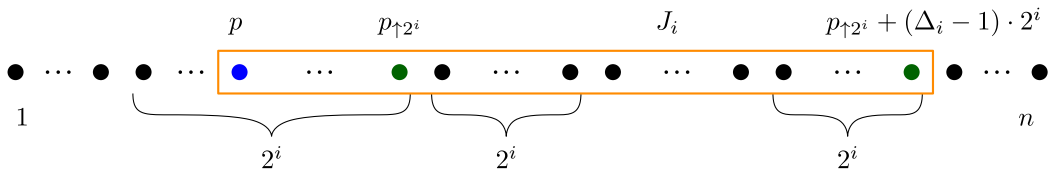

Figure 3.1 The interval $J_i = \left[p \ldots p_{\uparrow 2^i} + (\Delta_i - 1) \cdot 2^i\right]$.

Let i be the first index such that $p_i \geq q_i$, if there is any. We distinguish two cases based on whether $p_i < q_{i-1}$ holds or not. In case $p_i < q_{i-1}$, the path $(p, p_1, \ldots, p_{i-1}, p_i, q_{i-1}, \ldots, q_1, q)$ is a monotone path from p to q, since $q_i q_{i-1} \in E_{i-1}$ implies $p_i q_{i-1} \in E_{i-1}$. On the other hand, if we have $p_i \geq q_{i-1}$, the path $(p, p_1, \ldots, p_{i-1}, q_{i-1}, \ldots, q_1, q)$ is a monotone path between p and q, since $p_{i-1} p_i \in E_{i-1}$ implies $p_{i-1} q_{i-1} \in E_{i-1}$.

Finally, if $p_i < q_i$ holds for all $i = 1, \ldots, j$, then the path $(p, p_1, \ldots, p_j, q_j, \ldots, q_1, q)$ is a monotone path between p and q. We have $p_j q_j \in E_j$, since at least one of the stairways is usable. This concludes the proof that there is a monotone path from p to q. ◀

▶ **Lemma 12.** *For a fixed set $Q \subseteq [n]$, we have that $\mathbb{P}[Q \cap P_i = \emptyset] \leq \exp(-|Q|/2^i)$.*

Proof. Let $Q = \{q_1, \ldots, q_r\}$, and observe that knowing that certain points of Q are not in P_i, increases the probability of another point to be in P_i. That is, $\mathbb{P}[q_j \in P_i \mid q_1, \ldots, q_{j-1} \notin P_i] \geq \mathbb{P}[q_j \in P_i] = 1/2^i$. As such, we have

$$\begin{aligned} \mathbb{P}\big[Q \cap P_i = \emptyset\big] &= \mathbb{P}\Big[\bigcap_j (q_j \notin P_i)\Big] = \prod_{j=1}^{r} \mathbb{P}[q_j \notin P_i \mid q_1, \ldots, q_{j-1} \notin P_j] \\ &\leq \left(1 - 1/2^i\right)^r \leq \exp(-r/2^i). \end{aligned}$$ ◀

▶ **Lemma 13.** *Assume that $\vartheta \in (0, 1/2)$ and let $p \in \mathcal{S}_k \setminus \mathcal{S}_{k-1}$ be a kth round point for some $k \geq 1$. The probability that p is bad is at most $(\vartheta/2)^k/32$.*

Proof. For any integer $i \geq 1$, let $\Delta_i = \left\lfloor 2^{(i-1)/2}/(2\varepsilon)\right\rfloor$ and let $J_i = \left[p \ldots p_{\uparrow 2^i} + (\Delta_i - 1) \cdot 2^i\right]$, see Figure 3.1. Recall that $p \in [n]$, so $p_{\uparrow 2^i} = \left\lceil p/2^i \right\rceil 2^i$ is the next multiple of 2^i. Let ξ be the largest integer such that $J_\xi \subseteq P$. For $i = 0, \ldots, \xi$, the points of $J_{i+1} \cap P_i$ form a clique in G, since

$$|J_{i+1} \cap P_i| \leq \left\lceil |J_{i+1}| / 2^i \right\rceil \leq \left\lceil 2^{i+1} \Delta_{i+1} / 2^i \right\rceil = 2\Delta_{i+1} \leq \left\lceil 2^{i/2}/\varepsilon \right\rceil = \ell(i).$$

Indeed, any two vertices of P_i with distance at most $\ell(i)$ are connected by an edge of E_i. As such, it is enough to prove that there is a right safe stairway from p, that climbs on the levels to level ξ. Since $J_{\xi+1} \cap P_\xi$ forms a clique, it follows that such a stairway would be usable.

Let $\mathcal{E}_i$ be the event that $(J_i \setminus B) \cap P_i$ is empty, for $i = 1, \ldots, \xi$. Since $p \notin \mathcal{S}_{k-1}$, we have that $|J_i \cap B| < \alpha_{k-1} |J_i| \leq 2^i \alpha_{k-1} \Delta_i$. On the other hand, we have $|J_i \cap P_i| \geq 2^i (\Delta_i - 1)/2^i = \Delta_i - 1$. As such, if $|J_i \cap B| < |J_i \cap P_i|$ then $\mathbb{P}[\mathcal{E}_i] = 0$. This happens if $2^i \alpha_{k-1} \Delta_i \leq \Delta_i - 1 \iff 2^{i-k+1} \alpha \leq (\Delta_i - 1)/\Delta_i$, which happens if $i \leq k-2$, given that $\Delta_i \geq 2$. Notice that $\Delta_i \geq 2$ holds for all $i \geq 1$, if $\varepsilon \leq \frac{1}{4}$.

So assume that $i \geq k - 1$. Let $q_1, \ldots, q_r$ be all points of $J_i \setminus B$, which are the possible candidates to be contained in $(J_i \setminus B) \cap P_i$. By Eq. (3.1), there are at least

$$r = |J_i| - |J_i \cap B| \geq (1 - \alpha_{k-1}) |J_i| \geq (1 - \alpha_{k-1}) 2^i (\Delta_i - 1)$$

$$\geq (1-\alpha_{k-1})2^i\Big(\frac{2^{(i-1)/2}}{2\varepsilon}-2\Big) = \frac{c(1-\alpha_{k-1})\ln\vartheta^{-1}}{16(1-\alpha)}2^{3i/2-1/2}-(1-\alpha_{k-1})2^{i+1}$$
$$\geq c2^{3i/2-9/2}\ln\vartheta^{-1}-2^{i+1}$$

such points. Observe, that by the structure of the construction, a point is more likely to be contained in P_i conditioned on the event there are some other points which are not contained in P_i. Therefore, by Lemma 12, we have $\mathbb{P}\big[\mathcal{E}_i\big] \leq \exp\big(-r/2^i\big) \leq \tau_i$, for $\tau_i = \exp\big(2 - c2^{i/2-9/2}\ln\vartheta^{-1}\big)$. The sequence τ_i has a fast decay in i, since

$$\frac{\tau_{i+1}}{\tau_i} = \exp\Big(-(\sqrt{2}-1)c2^{i/2-9/2}\ln\vartheta^{-1}\Big) \leq \exp\big(-c2^{-6}\ln 2\big) = 2^{-c2^{-6}} \leq \frac{1}{2},$$

if $c \geq 2^6$ holds. Thus, we have

$$\begin{aligned}\mathbb{P}\Big[\cup_{i=1}^{\xi}\mathcal{E}_i\Big] &\leq \sum_{i=1}^{\xi}\mathbb{P}\big[\mathcal{E}_i\big] \leq \sum_{i=k-1}^{\xi}\tau_i \leq 2\tau_{k-1} = 2\exp\Big(2-c2^{(k-1)/2-9/2}\ln\vartheta^{-1}\Big)\\ &\leq 16\exp\Big(-\frac{c}{32}2^{k/2}\ln\vartheta^{-1}\Big) = 16\cdot\vartheta^{\frac{c}{32}\cdot 2^{k/2}} \leq 2^4\cdot\vartheta^{\frac{c}{2^6}\cdot k}\\ &\leq 2^4\cdot\left(\frac{1}{2}\right)^{\frac{c}{2^7}\cdot k}\cdot\left(\vartheta^{\frac{c}{2^7}}\right)^k \leq \frac{(\vartheta/2)^k}{64}\end{aligned}$$

for $c \geq 2^{11}$, using the conditions $0 < \vartheta \leq \frac{1}{2}$, $k \geq 1$ and the fact that $x \leq 2^x$.

Let p_i be the leftmost point in $(J_i \setminus B) \cap P_i$, for $i \geq 0$. Since $P_i \subseteq P_{i-1}$, for all i, it follows that $p = p_0 \leq p_1 \leq \cdots \leq p_\xi$. Furthermore, since $J_{i+1} \cap P_i$ is a clique in level i of G, and $p_i, p_{i+1} \in J_{i+1} \cap P_i$, it follows that $p_i p_{i+1} \in E_i$, if $p_i \neq p_{i+1}$, for all i. We conclude that $p, p_1, \ldots, p_\xi$ is a safe and usable right stairway in G.

The bound now follows by applying the same argument symmetrically for the left stairway. Indeed, using the union bound, we obtain $\mathbb{P}[p \text{ is bad}] \leq 2(\vartheta/2)^k/64 = (\vartheta/2)^k/32$. ◀

▶ **Lemma 14.** *Let $\vartheta \in (0, 1/2)$ and $B \subseteq P$ be an oblivious attack. Recall, that T^c is the set of bad points. Then, we have $\mathbb{E}[|T^c|] \leq (1+\vartheta)\,|B|$.*

Proof. We may assume that all the points of $\mathcal{S}_0$ are bad. Fortunately, by Lemma 5, we have $|\mathcal{S}_0| \leq |B|/(2\alpha-1) = |B|/(1-\vartheta/4) \leq (1+\vartheta/2)\,|B|$, since $\alpha = 1-\vartheta/8$ and $1/(1-x/4) \leq 1 + x/2$ for $0 \leq x \leq 2$. Again, using Lemma 5, we have

$$|\mathcal{S}_k \setminus \mathcal{S}_{k-1}| \leq |\mathcal{S}_k| \leq \big(1 + 2\left\lceil 2^k/\alpha\right\rceil\big)\,|B| \leq \left(3 + \frac{2^{k+1}}{\alpha}\right)|B| \leq 2^{k+3}\,|B|\,.$$

For $k \geq 1$, we have, by Lemma 13, that

$$b_k = \mathbb{E}[|(\mathcal{S}_k \setminus \mathcal{S}_{k-1}) \cap T^c|] \leq \sum_{p \in \mathcal{S}_k \setminus \mathcal{S}_{k-1}} \mathbb{P}[p \text{ is bad}] \leq 2^{k+3}\,|B| \cdot \frac{(\vartheta/2)^k}{32} \leq \frac{\vartheta^k}{4}\,|B|\,.$$

Since, $T^c = (\mathcal{S}_0 \cap T^c) \cup \bigcup_{k\geq 1}\big[(\mathcal{S}_k \setminus \mathcal{S}_{k-1}) \cap T^c\big]$, we have, by linearity of expectation, that

$$\frac{\mathbb{E}\Big[|T^c|\Big]}{|B|} \leq \frac{1}{|B|}\left(|\mathcal{S}_0| + \sum_{k=1}^{\infty} b_k\right) \leq 1 + \frac{\vartheta}{2} + \sum_{k=1}^{\infty}\frac{\vartheta^k}{4} \leq 1 + \frac{\vartheta}{2} + \frac{\vartheta}{4(1-\vartheta)} \leq (1+\vartheta),$$

since $\vartheta < 1/2$. ◀

▶ **Theorem 15.** *Let $\vartheta \in (0, 1/2)$ and $P = [n]$ be fixed. The graph G, constructed in Section 3.1, has $\mathcal{O}\big(n\vartheta^{-1} \log \vartheta^{-1}\big)$ edges, and it is a ϑ-reliable 1-spanner of P in expectation. Formally, for any oblivious attack B, we have $\mathbb{E}[\lambda(G, B)] \leq \vartheta$.*

Proof. By Lemma 8 the size of the construction is $|E(G)| = \mathcal{O}\big(n\vartheta^{-1} \log \vartheta^{-1}\big)$. Let $B \subseteq P$ be an oblivious attack and consider the bad set $P \setminus T$. By Lemma 11, for any two points outside the bad set, there is a monotone path connecting them. Further, by Lemma 14, we have $\mathbb{E}[|P \setminus T|] \leq (1 + \vartheta) |B|$ for any oblivious attack. Therefore, we obtain $\mathbb{E}[\lambda(G, B)] \leq \mathbb{E}[|T^c \setminus B| / |B|] \leq \vartheta$. ◀

3.3 Probabilistic bound

One can replace the guarantee, in Theorem 15, on the bound of the loss rate (which holds in expectation), by an upper bound that holds with probability at least $1 - \rho$, for some prespecified $\rho > 0$. A straightforward application of Markov's inequality implies that taking the union of $\log \rho^{-1}$ independent copies (G') of the construction of Theorem 15 with parameter $\vartheta/2$, results in a graph with the desired property. Indeed, we have

$$\mathbb{P}[\lambda(G, B) > \vartheta] \leq \mathbb{P}[\lambda(G', B) > \vartheta]^{\log \rho^{-1}} \leq \left(\frac{\mathbb{E}[\lambda(G', B)]}{\vartheta}\right)^{\log \rho^{-1}} \leq \left(\frac{1}{2}\right)^{\log \rho^{-1}} = \rho.$$

Here we show how one can do better to avoid the multiplicative factor $\log \rho^{-1}$.

Construction. The input consists of two parameters $\vartheta, \rho > 0$ and the set $P = [n]$. Let G be the graph constructed in Section 3.1 with parameters

$$\alpha = 1 - \frac{\vartheta}{8} \qquad \text{and} \qquad \varepsilon = \frac{8(1-\alpha)}{c(\ln \vartheta^{-1} + \ln \rho^{-1})} = \frac{\vartheta}{c(\ln \vartheta^{-1} + \ln \rho^{-1})},$$

where $c > 1$ is a sufficiently large constant. First, we need a variant of Lemma 13 to bound the probability of a kth round point being bad, using the new value of ε.

▶ **Lemma 16.** *Assume that $\vartheta \in (0, 1/2)$, $\rho \in (0, 1)$ and let $p \in \mathcal{S}_k \setminus \mathcal{S}_{k-1}$ be a kth round point for some $k \geq 1$. The probability that p is bad is at most $\vartheta \cdot \rho / 2^{3k+4}$.*

Proof. The proof is the same as the proof of Lemma 13. The only difference is due to the new value of ε, which results in $\tau_i = \exp\big(2 - c2^{i/2-9/2}(\ln \vartheta^{-1} + \ln \rho^{-1})\big)$, using the same notation. Therefore, we have

$$\begin{aligned}\mathbb{P}[p \in \mathcal{S}_k \setminus \mathcal{S}_{k-1} \text{ is bad}] &\leq 4\tau_{k-1} = 4\exp\Big(2 - c2^{k/2-5}(\ln \vartheta^{-1} + \ln \rho^{-1})\Big)\\ &\leq 2^5 \exp\Big(-\frac{c}{2^6} k(\ln \vartheta^{-1} + \ln \rho^{-1})\Big) = 2^5 \cdot \vartheta^{\frac{c}{2^6}k} \cdot \rho^{\frac{c}{2^6}k}\\ &\leq 2^5 \cdot \left(\frac{1}{2}\right)^{\frac{c}{2^6}k-1} \cdot \vartheta \cdot \rho = 2^{-\frac{c}{2^6}k+6} \cdot \vartheta \cdot \rho \leq 2^{-3k-4} \cdot \vartheta \cdot \rho,\end{aligned}$$

for $c \geq 2^{10}$. See Lemma 13 for a complete proof. ◀

▶ **Lemma 17.** *Let $\vartheta \in (0, 1/2)$, $\rho \in (0, 1)$ be fixed and $B \subseteq P$ be an oblivious attack. Then, with probability $\geq 1 - \rho$, the number of bad points is at most $(1 + \vartheta) |B|$. That is, we have $\mathbb{P}[|T^c| \leq (1 + \vartheta) |B|] \geq 1 - \rho$.*

Proof. The idea is to give bounds on the number of bad kth round points for all $k \geq 1$. Let $\mathcal{E}_k$ be the event that $|(\mathcal{S}_k \setminus \mathcal{S}_{k-1}) \cap T^c| > \frac{\vartheta}{2^{k+1}} |B|$ happens, for $k \geq 1$. Recall, by the choice of α, we have $|\mathcal{S}_0 \cap T^c| \leq |\mathcal{S}_0| \leq \left(1 + \frac{\vartheta}{2}\right) |B|$. Notice, that at least one of the events $\mathcal{E}_k$ must happen, for $k \geq 1$, in order to have $|T^c| > (1 + \vartheta) |B|$, since

$$|T^c| = |\mathcal{S}_0 \cap T^c| + \sum_{k=1}^{\infty} |(\mathcal{S}_k \setminus \mathcal{S}_{k-1}) \cap T^c| \leq \left(1 + \frac{\vartheta}{2}\right) |B| + \sum_{k=1}^{\infty} \frac{\vartheta}{2^{k+1}} |B| = (1 + \vartheta) |B| .$$

Using Markov's inequality and Lemma 16 we get

$$\mathbb{P}[\mathcal{E}_k] \leq \frac{\mathbb{E}[|(\mathcal{S}_k \setminus \mathcal{S}_{k-1}) \cap T^c|]}{\frac{\vartheta}{2^{k+1}} |B|} \leq \frac{|\mathcal{S}_k| \cdot \mathbb{P}[p \in \mathcal{S}_k \setminus \mathcal{S}_{k-1} \text{ is bad}]}{\frac{\vartheta}{2^{k+1}} |B|} \leq \frac{2^{k+3} |B| \cdot \frac{\vartheta \cdot \rho}{2^{3k+4}}}{\frac{\vartheta}{2^{k+1}} |B|} = \frac{\rho}{2^k}.$$

Therefore, we obtain

$$\mathbb{P}[|T^c| > (1 + \vartheta) |B|] \leq \mathbb{P}[\cup_{k \geq 1} \mathcal{E}_k] \leq \sum_{k=1}^{\infty} \mathbb{P}[\mathcal{E}_k] \leq \sum_{k=1}^{\infty} \frac{\rho}{2^k} \leq \rho,$$

which is equivalent to $\mathbb{P}[|T^c| \leq (1 + \vartheta) |B|] \geq 1 - \rho$. ◀

▶ **Theorem 18.** *Let $\vartheta \in (0, 1/2)$, $\rho \in (0, 1)$ and $P = [n]$ be fixed. The graph G, constructed above, is a ϑ-reliable 1-spanner of P, with probability at least $1 - \rho$. Formally, we have $\mathbb{P}[\lambda(G, B) \leq \vartheta] \geq 1 - \rho$ for any oblivious attack B. Furthermore, the graph G has $\mathcal{O}\left(n\vartheta^{-1}(\log \vartheta^{-1} + \log \rho^{-1})\right)$ edges.*

Proof. The bound on the size follows directly from Lemma 8. Let $B \subseteq P$ be an oblivious attack and consider the bad set $P \setminus T$. By Lemma 11, for any two points outside the bad set, there is a monotone path connecting them. Further, by Lemma 17, we have $\mathbb{P}[\lambda(G, B) \leq \vartheta] \geq \mathbb{P}[|T^c| \leq (1 + \vartheta) |B|] \geq 1 - \rho$ for any oblivious attack. ◀

4 Reliable spanners in higher dimensions

Now we turn to the higher-dimensional setting, and show that one can construct spanners with near linear size that are reliable in expectation or with some fixed probability (which can be provided as part of the input). We use the same technique as Buchin et al. [4], that is, we use our one-dimensional construction as a black box in combination with a result of Chan et al. [6]. Let the dimension $d > 1$ be fixed. In the following we assume $P \subset [0, 1)^d$, which can be achieved by an appropriate scaling and translation of the d-dimensional Euclidean space $\mathbb{R}^d$. For an ordering σ of $[0, 1)^d$, and two points $p, q \in [0, 1)^d$, such that $p \prec q$, let $(p, q)_\sigma = \left\{z \in [0, 1)^d \;\middle|\; p \prec z \prec q\right\}$ be the set of points between p and q in the order σ.

▶ **Theorem 19** ([6]). *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) $(p, z)_\sigma \subseteq \text{ball}\left(p, \varsigma\ell\right)$,

(iii) $(z, q)_\sigma \subseteq \text{ball}\left(q, \varsigma\ell\right)$, *and*

(iv) $z \in \text{ball}\left(p, \varsigma\ell\right)$ *or* $z \in \text{ball}\left(q, \varsigma\ell\right)$.

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.

The above theorem ensures that it is enough to maintain only a "few" linear orderings, and for any pair of points $p, q \in P$ there exists an ordering where all points that lie between p and q are either very close to p or q. It is natural to build the one-dimensional construction for each of these orderings with some carefully chosen parameter. Then, since there is a reliable path in the one-dimensional construction, there is an edge $p'q'$ along the path between p and q that connects the locality of p and the locality of q. We fix the edge $p'q'$ and apply recursion on the subpaths from p to p' and q to q'.

4.1 Construction

Let $\vartheta, \varepsilon \in (0,1)$ be fixed parameters and $P \subseteq [0,1)^d$ be a set of n points. Set $\varsigma = \varepsilon/16$ in Theorem 19 and let $\Pi^+ = \Pi^+(\varsigma)$ be the set of $M = M(\varsigma)$ orderings that fulfills the conditions of the theorem. We define $\vartheta' = \frac{\vartheta}{3MN}$, where $N = \lceil \operatorname{loglog} n \rceil$. Now, for each ordering $\sigma \in \Pi^+$, we build N independent spanners $G^1_\sigma, \ldots, G^N_\sigma$, using the construction in Section 3.1 with parameter ϑ'. The (random) graph G is defined as the union of graphs G^i_σ for all $\sigma \in \Pi^+$ and $i \in [N]$, that is, $E(G) = \cup_{\sigma\in\Pi^+, i\in[N]} E\big(G^i_\sigma\big)$.

4.2 Analysis

▶ **Lemma 20.** *The graph G, constructed above, has $\mathcal{O}\big(c\, n \operatorname{loglog}^2 n \operatorname{logloglog} n\big)$ edges, where the $\mathcal{O}$ hides constant that depends on the dimension d, and $c = \mathcal{O}(\varepsilon^{-2d}\vartheta^{-1}\log^3\varepsilon^{-1}\log\vartheta^{-1})$.*

Proof. There are $M = \mathcal{O}\big(\varepsilon^{-d}\log\varepsilon^{-1}\big)$ orderings, and for each ordering there are N copies, for which we build the one-dimensional construction with parameter ϑ'. The size of the one-dimensional construction is $\mathcal{O}\big(n \cdot \vartheta'^{-1} \cdot \log \vartheta'^{-1}\big)$, by Lemma 8. Therefore, G has size

$$
\begin{aligned}
|E(G)| &= \big|\cup_{\sigma\in\Pi^+, i\in[N]} E\big(G^i_\sigma\big)\big| \le \sum_{\sigma\in\Pi^+, i\in[N]} |E(G'_\sigma)| \le NM \cdot \mathcal{O}\big(n \cdot \vartheta'^{-1} \cdot \log\vartheta'^{-1}\big) \\
&= \mathcal{O}\big(n \cdot N^2M^2\vartheta^{-1} \cdot \big(\log\vartheta^{-1} + \log N + \log M\big)\big) \\
&= \mathcal{O}\big(n \cdot \operatorname{loglog}^2 n \cdot \varepsilon^{-2d}\log^2\varepsilon^{-1} \cdot \vartheta^{-1} \cdot (\log\vartheta^{-1} + \\
&\qquad\qquad + \operatorname{logloglog} n + d\log\varepsilon^{-1} + \operatorname{loglog}\varepsilon^{-1})\big) \\
&= \mathcal{O}\big(c\, n \operatorname{loglog}^2 n \operatorname{logloglog} n\big), \qquad \text{where } c = \mathcal{O}(\varepsilon^{-2d}\vartheta^{-1}\log^3\varepsilon^{-1}\log\vartheta^{-1}).
\end{aligned}
$$ ◀

Fix an attack set $B \subseteq P$. In order to bound $\lambda(G, B)$ in expectation, we define a sequence of sets $B_0 \subseteq B_1 \subseteq \cdots \subseteq B_N$ as follows. First, we set $B_0 = B$. Then, for $i = 1, \ldots, N$, we define B^σ_i for each $\sigma \in \Pi^+$ to contain all points that do not have a right or left stairway in G^i_σ that is safe and usable with respect to B_{i-1}, that is, B^σ_i contains the bad points with respect to B_{i-1}. We set $B_i = \cup_{\sigma\in\Pi^+} B^\sigma_i$. Our goal is to show that the expected size of B_N is small, and there is a $(1+\varepsilon)$-path for all pairs of points outside of B_N.

▶ **Lemma 21.** *Let B be an oblivious attack and let $B_0 \subseteq B_1 \subseteq \cdots \subseteq B_N$ be the sequence defined above. Then, for $i = 1, \ldots, N$, we have $\mathbb{E}[|B^\sigma_i| \mid B_{i-1}] \le (1+\vartheta')\,|B_{i-1}|$, for all $\sigma \in \Pi^+$.*

Proof. The set B_{i-1} has information only about graphs G^j_σ for $j \le i-1$. Thus, the attack B_{i-1} on the graph G^i_σ is oblivious and we have $\mathbb{E}[|B^\sigma_i| \mid B_{i-1}] \le (1+\vartheta')\,|B_{i-1}|$ by Lemma 14. ◀

▶ **Lemma 22.** *Let B_N be the set defined above. For any oblivious attack B, the expected size of B_N is at most $(1+\vartheta) \cdot |B|$.*

Proof. By Lemma 21 we have $\mathbb{E}[|B_i^\sigma| \mid B_{i-1}] \le (1+\vartheta')\,|B_{i-1}|$ for all $\sigma \in \Pi^+$. Therefore,

$$\mathbb{E}[|B_i| \mid B_{i-1}] \le \big((1+\vartheta')\,|B_{i-1}| - |B_{i-1}|\big) \cdot M + |B_{i-1}| = \left(1 + \frac{\vartheta}{3N}\right)|B_{i-1}|$$

holds, for $i = 1, \ldots, N$, which gives

$$\begin{aligned}\mathbb{E}[|B_N|] \le \mathbb{E}[\mathbb{E}[|B_N| \mid B_{N-1}]] &\le \left(1 + \frac{\vartheta}{3N}\right) \cdot \mathbb{E}[|B_{N-1}|] \\ &\le \left(1 + \frac{\vartheta}{3N}\right)^N \cdot \mathbb{E}[|B_0|] = \left(1 + \frac{\vartheta}{3N}\right)^N \cdot |B|\,.\end{aligned}$$

Using $1 + x \le e^x \le 1 + 3x$, for $x \in [0,1]$, we obtain

$$\mathbb{E}[|B_N|] \le \left(1 + \frac{\vartheta}{3N}\right)^N \cdot |B| \le \exp\left(N\frac{\vartheta}{3N}\right) \cdot |B| = e^{\frac{\vartheta}{3}} \cdot |B| \le (1+\vartheta) \cdot |B|\,. \qquad ◀$$

▶ **Lemma 23.** *Let B_N be the set defined above. Then, for any two points $p, q \in P \setminus B_N$, there is a $(1+\varepsilon)$-path in the graph $G \setminus B$.*

Proof. The proof is essentially the same as the proof of Theorem 15 in [4].

Let $p, q \in P \setminus B_N$ be fixed. According to Theorem 19, there is an ordering $\sigma \in \Pi^+$, such that all the points $z \in (p,q)_\sigma$ lie in one of the balls of radius $\varsigma\,\|p-q\|$ around p and q. Recall that the graph G contains G_σ^N as a subgraph. Since $p, q \notin B_N$ and G_σ^N is reliable, there is a path connecting p and q that is monotone with respect to σ and avoids any point in B_{N-1} by Theorem 15. Therefore, there is a unique edge $p'q'$ along this path such that p' is in the close neighborhood of p and q' is in the close neighborhood of q. Furthermore, we also have that $p', q' \in P \setminus B_{N-1}$. We fix the edge $p'q'$ in path π and find subpaths between the pairs pp' and qq' in a recursive manner. The bounds on the distances are

(i) $\|p'-q'\| \le (1+2\varsigma)\,\|p-q\|$,
(ii) $\|p-p'\| \le \varsigma\,\|p-q\|$ and similarly $\|q-q'\| \le \varsigma\,\|p-q\|$.

We repeat this process $N-1$ times. Let Q_i be the set of pairs that needs to be connected in the ith round, that is, $Q_0 = \{pq\}, Q_1 = \{pp', qq'\}$ and so on. There are at most 2^i pairs in Q_i and for any pair $xy \in Q_i$ we have $x, y \in P \setminus B_{N-i}$. For each pair $xy \in Q_i$, there is an ordering σ such that the argument above can be repeated. That is, there is a monotone path in the graph $G_\sigma^{N-i} \setminus B_{N-i-1}$ according to σ and there is an edge $x'y'$ along this path such that

(i) $\|x'-y'\| \le (1+2\varsigma)\,\|x-y\| \le (1+2\varsigma)\varsigma^i\,\|p-q\|$,
(ii) $\|x-x'\| \le \varsigma\,\|x-y\| \le \varsigma^{i+1}\,\|p-q\|$ and similarly $\|y-y'\| \le \varsigma^{i+1}\,\|p-q\|$.

The edge $x'y'$ is added to path π and the pairs xx' and yy' are added to Q_{i+1}, unless they are trivial (i.e., $x = x'$ or $y = y'$). After $N-1$ rounds, Q_{N-1} is the set of active pairs that still needs to be connected. Notice that $x, y \in P \setminus B_1$ holds for any pair $xy \in Q_{N-1}$. Again, for each pair in Q_{N-1}, we apply Theorem 19 and Theorem 15 to obtain a monotone path according to some ordering σ in the graph G_σ^1. None of these paths use any points in B. In order to complete the path π we add the whole paths obtained in the last step. It is not hard to see that the number of edges of each of the paths added in the last step is at most $2\log n$. Indeed, it is clear from the analysis of our one-dimensional construction that a path using the stairways can have at most two points per level. Since the number of levels in the construction is fewer than $\log n$, we get the bound $2\log n$.

Now, that we have a path π that connects the points p and q without using any points in the failed set B, we give an upper bound on the length of π. First, we calculate the total length added in the last step. There are $|Q_{N-1}| \le 2^{N-1}$ pairs in the last step and for each pair $xy \in Q_{N-1}$ we have $\|x - y\| \le \|p - q\| \varsigma^{N-1}$. Thus, we obtain

$$\begin{aligned}\sum_{\{x,y\}\in Q_{N-1}} \operatorname{length}(\pi[x,y]) &\le 2^{N-1}\big((1+2\varsigma)\|p-q\|\varsigma^{N-1} + 2\log n\,\|p-q\|\varsigma^{N}\big)\\ &\le 2\cdot 2\varsigma\,\|p-q\| + (2\varsigma)^N \log n\,\|p-q\| = \Big(4\varsigma + (2\varsigma)^N\log n\Big)\|p-q\|\\ &\le \left(\frac{\varepsilon}{4} + \left(\frac{\varepsilon}{8}\right)^{\operatorname{loglog} n}\log n\right)\|p-q\|\\ &\le \left(\frac{\varepsilon}{4} + \frac{\varepsilon}{4}\cdot\left(\frac{1}{2}\right)^{\operatorname{loglog} n}\log n\right)\|p-q\| = \frac{\varepsilon}{2}\|p-q\|,\end{aligned}$$

where we simply use $2\varsigma \le 1$ in the second line and $\varsigma = \varepsilon/16$ and $N = \lceil \operatorname{loglog} n\rceil$ in the third line. Second, we bound the total length of the edges that were added to path π in any round except the last. This contributes at most

$$\begin{aligned}\sum_{i=0}^{N-2} 2^i\cdot(1+2\varsigma)\varsigma^i\,\|p-q\| &\le (1+2\varsigma)\,\|p-q\|\cdot\sum_{i=0}^{\infty}(2\varsigma)^i = (1+2\varsigma)\,\|p-q\|\cdot\frac{1}{1-2\varsigma}\\ &= \left(1+\frac{4\varsigma}{1-2\varsigma}\right)\|p-q\| = \left(1+\frac{\varepsilon/4}{1-\varepsilon/8}\right)\|p-q\| \le \left(1+\frac{\varepsilon}{2}\right)\|p-q\|\end{aligned}$$

to the length of π. Therefore the total length of the path π connecting p and q, without using any points of B, is at most $(1+\varepsilon)\,\|p-q\|$. ◀

▶ **Theorem 24.** *Let $\vartheta, \varepsilon \in (0,1)$ be fixed parameters and $P \subseteq [0,1)^d$ be a set of n points. The graph G, constructed in Section 4.1, is a ϑ-reliable $(1+\varepsilon)$-spanner of P in expectation and has size $\mathcal{O}\big(c\,n \operatorname{loglog}^2 n \operatorname{logloglog} n\big)$, where $\mathcal{O}$ hides constant that depends on the dimension d, and $c = \mathcal{O}(\varepsilon^{-2d}\vartheta^{-1}\log^3\varepsilon^{-1}\log\vartheta^{-1})$.*

Proof. The size of the construction is proved in Lemma 20. Let B_N be the set defined above. By Lemma 22, the expected size of B_N is at most $(1+\vartheta)\,|B|$. By Lemma 23, for any two points $p, q \in P \setminus B_N$, there is a $(1+\varepsilon)$-path between p and q in the graph $G \setminus B$. Thus, we have $\mathbb{E}[\lambda(G,B)] \le \vartheta$. ◀

4.3 Probabilistic bound

The same construction, as we used in Section 4.1, can be applied to construct spanners with near linear edges that are reliable with probability $1-\rho$. The idea is to use the probabilistic version of the one-dimensional construction with parameters $\rho' = \frac{\rho}{MN}$ and $\vartheta' = \frac{\vartheta}{3MN}$. Then, similarly to Lemma 22, it is not hard to show that $|B_N| \le (1+\vartheta)\,|B|$ holds with probability $1-\rho$.

▶ **Lemma 25.** *Let B_N be the set defined in Section 4.2. The probability that the size of B_N is larger than $(1+\vartheta)\cdot|B|$ is at most ρ.*

Proof. By Lemma 17, and since all attacks are oblivious, we have $\mathbb{P}[|B_i^\sigma| > (1+\vartheta')\,|B_{i-1}|] \leq \rho'$ for all $\sigma \in \Pi^+$ and $i \geq 1$. Therefore,

$$\begin{aligned}\mathbb{P}[|B_i| > (1+M\vartheta')\,|B_{i-1}|] &= \mathbb{P}[|B_i \setminus B_{i-1}| > M\vartheta'\,|B_{i-1}|]\\ &\leq \mathbb{P}[\cup_{\sigma\in\Pi^+} |B_i^\sigma \setminus B_{i-1}| > \vartheta'\,|B_{i-1}|] \leq \sum_{\sigma\in\Pi^+} \mathbb{P}[|B_i^\sigma \setminus B_{i-1}| > \vartheta'\,|B_{i-1}|]\\ &= \sum_{\sigma\in\Pi^+} \mathbb{P}[|B_i^\sigma| > (1+\vartheta')\,|B_{i-1}|] \leq M\rho'\end{aligned}$$

holds for $i = 1, \ldots, N$. Since $\left(1+\frac{\vartheta}{3N}\right)^N \leq \left(e^{\frac{\vartheta}{3N}}\right)^N \leq 1+\vartheta$, we get

$$\begin{aligned}\mathbb{P}[|B_N| > (1+\vartheta)\,|B|] &\leq \mathbb{P}\left[|B_N| > \left(1+\frac{\vartheta}{3N}\right)^N |B|\right] \leq \mathbb{P}\left[\bigcup_{i=1}^{N} |B_i| > \left(1+\frac{\vartheta}{3N}\right)|B_{i-1}|\right]\\ &\leq \sum_{i=1}^{N} \mathbb{P}\left[|B_i| > \left(1+\frac{\vartheta}{3N}\right)|B_{i-1}|\right] = \sum_{i=1}^{N} \mathbb{P}\big[|B_i| > \big(1+M\vartheta'\big)\,|B_{i-1}|\big] \leq NM\rho' = \rho.\end{aligned}$$ ◀

Therefore, using the same argument as for Theorem 24, we obtain the following result, which gives a slight improvement in the constants, compared to the trivial multiplicative factor $\mathcal{O}\big(\log \rho^{-1}\big)$ by simply repeating the construction of Section 4.1.

▶ **Theorem 26.** *Let $\vartheta, \varepsilon, \rho \in (0,1)$ be fixed parameters and $P \subseteq [0,1)^d$ be a set of n points. The graph described above is a ϑ-reliable $(1+\varepsilon)$-spanner of P with probability $1-\rho$. Furthermore, the size of the construction is $\mathcal{O}\big(c\, n \operatorname{loglog}^2 n \operatorname{logloglog} n\big)$, where $\mathcal{O}$ hides constant that depends on the dimension d, and $c = \mathcal{O}\big(\varepsilon^{-2d}\vartheta^{-1}\log^3\varepsilon^{-1}(\log\vartheta^{-1} + \log\rho^{-1})\big)$.*

5 Conclusions

Reliable spanners require $\Omega(n \log n)$ edges. In this paper, we showed that fewer edges are sufficient, if the spanner only has to be reliable against oblivious attacks (in expectation or with a certain probability). Our new construction avoids the use of expanders, and as a result has much smaller constants than previous constructions, making it potentially practical. The number of edges in the new spanner is significantly smaller – it is linear in one dimension, and roughly $\mathcal{O}(n \operatorname{loglog}^2 n)$ in higher dimensions. An open problem is whether these loglog-factors in higher dimensions can be avoided. Furthermore, similar results for reliable spanners for general metrics would be of interest.

References

1 J. Aspnes and G. Shah. Skip graphs. *ACM Transactions on Algorithms*, 3(4):37, November 2007.

2 P. Bose, P. Carmi, V. Dujmović, and P. Morin. Near-optimal O(k)-robust geometric spanners. *CoRR*, abs/1812.09913, 2018. arXiv:1812.09913.

3 P. Bose, V. Dujmović, P. Morin, and M. Smid. Robust geometric spanners. *SIAM Journal on Computing*, 42(4):1720–1736, 2013. doi:10.1137/120874473.

4 K. Buchin, S. Har-Peled, and D. Oláh. A spanner for the day after. In *Proc. 35th Int. Annu. Sympos. Comput. Geom.* (SoCG), pages 19:1–19:15, 2019. doi:10.4230/LIPIcs.SoCG.2019.19.

5 T.-H. H. Chan, M. Li, L. Ning, and S. Solomon. New doubling spanners: Better and simpler. *SIAM Journal on Computing*, 44(1):37–53, 2015. doi:10.1137/130930984.

6 T. M. Chan, S. Har-Peled, and M. Jones. On Locality-Sensitive Orderings and Their Applications. In *Proc. 10th Innovations in Theoretical Computer Science Conference (ITCS 2019)*, pages 21:1–21:17, 2018. `doi:10.4230/LIPIcs.ITCS.2019.21`.

7 C. Levcopoulos, G. Narasimhan, and M. Smid. Efficient algorithms for constructing fault-tolerant geometric spanners. In *Proc. 30th Annu. ACM Sympos. Theory Comput.* (STOC), pages 186–195, 1998. `doi:10.1145/276698.276734`.

8 C. Levcopoulos, G. Narasimhan, and M. Smid. Improved algorithms for constructing fault-tolerant spanners. *Algorithmica*, 32(1):144–156, 2002. `doi:10.1007/s00453-001-0075-x`.

9 T. Lukovszki. New results of fault tolerant geometric spanners. In *Proc. 6th Workshop Algorithms Data Struct.* (WADS), pages 193–204, 1999. `doi:10.1007/3-540-48447-7_20`.

10 S. Solomon. From hierarchical partitions to hierarchical covers: Optimal fault-tolerant spanners for doubling metrics. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, page 363–372, New York, NY, USA, 2014. Association for Computing Machinery. `doi:10.1145/2591796.2591864`.