Sometimes Reliable Spanners of Almost Linear Size

Kevin Buchin[1], Sariel Har-Peled[2], and Dániel Oláh[1]

1 **Department of Mathematics and Computing Science, TU Eindhoven, The Netherlands**
`k.a.buchin@tue.nl | d.olah@tue.nl`

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

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 $O(n \log n \log\log^6 n)$ 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 $O(n \log\log^2 n \log\log\log n)$ edges.

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. 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 $t \cdot \|p - q\|$.

We focus our attention to construct spanners that can survive massive failures of vertices. The most studied notion is fault tolerance [6, 7, 8], 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. 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 [2], 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

36th European Workshop on Computational Geometry, Würzburg, Germany, March 16–18, 2020.
This is an extended abstract of a presentation given at EuroCG'20. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

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 for only a tiny additional fraction of points beyond t. The fraction is relative to the number of failed vertices and controlled by the parameter ϑ.

Recently, the authors [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 $\mathcal{O}\big(n \log n \log\log^6 n\big)$ in higher dimensions (the constant in the $\mathcal{O}$ depends both on the dimension, ε, and the reliability parameter). An alternative construction, with slightly worse bounds, was given by Bose *et al.* [1].

Limitations of previous constructions. The construction of Buchin *et al.* [4] (and also the construction of Bose *et al.* [1]) relies on using expanders to get a monotone 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.* [2].

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. Note, that B^+ is not necessarily unique. 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 $\mathbb{E}[\lambda(G, B)] \le \vartheta$, or alternatively, that the probability $\mathbb{P}[\lambda(G, B) \ge \vartheta]$ is small.

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 $\mathbb{E}[\lambda(G, B)] \le \vartheta$. This construction can also be modified so that $\lambda(G, B) \le \vartheta$ 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 $O\big(n \log\log^2 n\big)$.

In this abstract, we only present the one-dimensional construction of reliable spanners in expectation. For the missing proofs and further results, we refer to the full version [3].

2. Preliminaries

▶ **Definition 2.1** (Reliable spanner). Let $G = (P, E)$ be a t-spanner for some $t \geq 1$ constructed by a (possibly) randomized algorithm. Given an oblivious attack B, its ***damaged set*** B^+ is the 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 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 $B \subseteq P$.

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 borrow the notion of shadow from our previous work [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.

▶ **Definition 2.2.** 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 $\big|[i \ldots j] \cap B\big| \geq \alpha \big|[i \ldots j]\big|$. 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 $\mathcal{S}_{\to}(B)$ and $\mathcal{S}_{\leftarrow}(B)$, respectively. The combined shadow is denoted by $\mathcal{S}(\alpha, B) = \mathcal{S}_{\to}(B) \cup \mathcal{S}_{\leftarrow}(B)$.

▶ **Lemma 2.3** ([4]). *For any set $B \subseteq [n]$, and $\alpha \in (0, 1)$, we have that $|\mathcal{S}(\alpha, B)| \leq \big(1 + 2\lceil 1/\alpha \rceil\big) |B|$. Further, if $\alpha \in (2/3, 1)$, we have that $|\mathcal{S}(\alpha, B)| \leq |B| / (2\alpha - 1)$.*

▶ **Definition 2.4.** 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 $|j - i|$. 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 integer numbers $x, y > 0$, let $x_{\uparrow y} = \lceil x/y \rceil\, y$.

3. Construction of reliable spanners on the line

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, see Figure 3.1 as an example. 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

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

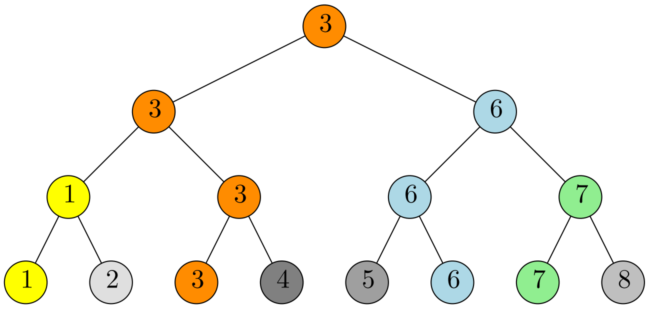

Figure 3.1 An example of a tournament tree with $n = 8$.

where $c > 1$ is a sufficiently large constant. Let M be the smallest integer for which $|P_M| \le 2^{M/2}/\varepsilon$ 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, $E(G) = \cup_{i=0}^{M} E_i$.

4. Analysis

▶ **Lemma 4.1.** *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

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

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.

▶ **Definition 4.2** (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) $p = p_0 \le p_1 \le \cdots \le p_j$ (resp., $p \ge p_1 \ge \cdots \ge p_j$),
(ii) if $p_i \ne p_{i+1}$, then $p_i p_{i+1} \in E$, 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. Finally, a point p is ***bad*** if it belongs to B, or it does not have safe and usable stairways to both directions, that is, $p \in P \setminus T$.

Let $p, q \in T$ be two points such that $p < q$. Intuitively, it is clear that the right stairway of p and the left stairway of q must cross each other at some level. Combining these stairways, with some care at the point where they cross, we obtain a monotone path between p and q.

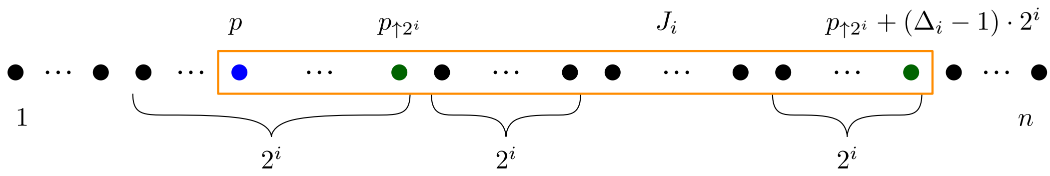

Figure 4.1 The interval $J_i = \left[p \ldots p_{\uparrow 2^i} + (\Delta_i - 1) \cdot 2^i\right]$.

▶ **Lemma 4.3.** *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 $G \setminus B$.*

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 ***k*th 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.

▶ **Lemma 4.4.** *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 (idea). By symmetry, it is enough to consider right stairways. We define a sequence of intervals $J_1, J_2, \ldots$, see Figure 4.1, such that each interval starts at p, their length increases exponentially, and $J_i \cap P_i$ contains exactly Δ_i or $\Delta_i - 1$ points. We set Δ_i to ensure that any pair of points $p_i \in J_i \cap P_i$ and $p_{i+1} \in J_{i+1} \cap P_{i+1}$ are connected. Thus, if the sets $J_i \cap P_i$, for all $i \geq 1$, contain at least one point outside of B, then we have a possible candidate for a safe and usable stairway. It is not hard to see that, for example, by choosing the leftmost available point in each set, we obtain a monotone path. ◀

We obtain bounds on the expected number of bad kth round points by using Lemma 2.3 and Lemma 4.4 to bound the number of such points and the probability of a kth round point being bad, respectively. Then, we sum up for all rounds to obtain the desired bound.

▶ **Lemma 4.5.** *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 $\mathbb{E}[|T^c|] \leq (1 + \vartheta)\,|B|$.*

▶ **Theorem 4.6.** *Let $\vartheta \in (0, 1/2)$ and $P = [n]$ be fixed. The graph G, constructed in Section 3, 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 4.1, the size of G is $\mathcal{O}\big(n\vartheta^{-1} \log \vartheta^{-1}\big)$. Let $B \subseteq P$ be an oblivious attack and consider the bad set $P \setminus T$. By Lemma 4.3, for any two points outside the bad set, there is a monotone path connecting them. Further, by Lemma 4.5, we have $\mathbb{E}[|P \setminus T|] \leq (1 + \vartheta)\,|B|$ for any oblivious attack. Thus, we obtain $\mathbb{E}[\lambda(G, B)] \leq \mathbb{E}[|T^c \setminus B| \,/\, |B|] \leq \vartheta$. ◀

Using Theorem 4.6 and a result of Chan *et al.* [5] on orderings of a set of points in $\mathbb{R}^d$, we can construct spanners for higher dimensional point sets that are reliable in expectation.

▶ **Theorem 4.7.** *Let $\vartheta, \varepsilon \in (0, 1)$ be fixed and $P \subseteq \mathbb{R}^d$ be a set of n points. We can construct a $(1 + \varepsilon)$-spanner of P that is ϑ-reliable in expectation and has size $\mathcal{O}\big(n \log \log^2 n \log \log \log n\big)$.*

References

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

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

3 K. Buchin, S. Har-Peled, and D. Oláh. Sometimes reliable spanners of almost linear size, 2019. arXiv:1912.01547.

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. URL: https://doi.org/10.4230/LIPIcs.SoCG.2019.19, doi:10.4230/LIPIcs.SoCG.2019.19.

5 T. M. Chan, S. Har-Peled, and M. Jones. On Locality-Sensitive Orderings and Their Applications. In *Proc. 34th Int. Annu. Sympos. Comput. Geom.* (SoCG), pages 21:1–21:17, 2018. URL: http://drops.dagstuhl.de/opus/volltexte/2018/10114, doi:10.4230/LIPIcs.ITCS.2019.21.

6 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. URL: http://doi.acm.org/10.1145/276698.276734, doi:10.1145/276698.276734.

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

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