Reliable Spanners for Metric Spaces

Sariel Har-Peled
Department of Computer Science, University of Illinois, Urbana, IL USA

Manor Mendel
Department of Mathematics and Computer Science, The Open University of Israel, Raanana, Israel

Dániel Oláh
Department of Mathematics and Computing Science, TU Eindhoven, The Netherlands

Abstract

A spanner is reliable if it can withstand large, catastrophic failures in the network. More precisely, any failure of some nodes can only cause a small damage in the remaining graph in terms of the dilation, that is, the spanner property is maintained for almost all nodes in the residual graph. Constructions of reliable spanners of near linear size are known in the low-dimensional Euclidean settings. Here, we present new constructions of reliable spanners for planar graphs, trees and (general) metric spaces.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Spanners, reliability

Digital Object Identifier 10.4230/LIPIcs.SoCG.2021.43

Related Version *Full Version*: `https://arxiv.org/abs/2007.08738`

Funding *Sariel Har-Peled*: Work on this paper was partially supported by a NSF AF award CCF-1907400.
Manor Mendel: Supported by BSF Grant no. 2018223.
Dániel Oláh: Supported by the Netherlands Organisation for Scientific Research (NWO) through Gravitation-grant NETWORKS-024.002.003.

Acknowledgements The authors thank Kevin Buchin for useful discussions and references.

1 Introduction

Let $\mathcal{M} = (P, \mathrm{d})$ be a finite metric space. Let $G = (P, E)$ be a sparse graph on the points of $\mathcal{M}$ whose edges are weighted with the distances of their endpoints. The graph G is a *t-spanner* (or *t-emulator*) if for any pair of vertices $u, v \in V$ we have $\mathrm{d}_G(u, v) \leqslant t \cdot \mathrm{d}(u, v)$, where $\mathrm{d}_G(u, v)$ is the length of the shortest path between u and v in G, and $\mathrm{d}(u, v)$ is the distance in the metric space between u and v. Spanners were first introduced by Peleg and Schäffer [20] as a tool in distributed computing, but have since found use in other areas of algorithms, networking, data structure and metric geometry, see [19, 18].

Fault tolerant spanners. A desired property of spanners is a resilience to failures of their vertices. The basic notion that captures this is fault tolerance [7, 13, 14, 16, 21]. A graph G is a *k-fault tolerant* t-spanner, if for any subset of vertices B, with $|B| \leqslant k$, the graph $G \setminus B$ is a t-spanner. The disadvantage of k-fault tolerant graphs, is that there is no guarantee if more than k vertices fail, and furthermore, the size of a fault tolerant graph grows (linearly) with the parameter k. In particular, for fixed $t \geqslant 1$, the optimal size of k-fault tolerant $(2t-1)$-spanners on n vertices is $\mathcal{O}(k^{1-1/t} n^{1+1/t})$ [4]. Note that vertex degrees must be at least $\Omega(k)$ to avoid the possibility of isolating a vertex. Thus, it is not suitable for massive failures in the network.

Table 1.1 Our results. Polylog factors are polynomial factors in $\log n$ and $\log \Phi$, where Φ is the spread of the metric. For trees and planar graphs, these results are for graphs with weights on the edges. Here in expectation denotes that the spanner works against an oblivious adversary (here, the expectation is over the randomization in the construction), and the guarantee is on the expected size of the damaged set. Similarly, deterministic implies an adaptive adversary.

Constructions of ϑ-reliable Δ-spanners				
metric	Δ	guarantee	size	ref
Uniform	2	expectation	$\mathcal{O}(n\vartheta^{-1}\log\vartheta^{-1})$	Lemma 4
	t	det. lower bound	$\Omega(n^{1+1/t})$	Lemma 5
	$2t-1$	deterministic	$\mathcal{O}\left(\vartheta^{-2}n^{1+1/t}\right)$	Theorem 9
	$2t$	deterministic	$\mathcal{O}\left(\vartheta^{-1}n^{1+1/t}\right)$	Theorem 9
Finite metrics	$\mathcal{O}(\log n)$	expectation	$\mathcal{O}(\vartheta^{-1}n \operatorname{polylog})$	Lemma 39
	$\mathcal{O}(t)$	expectation	$\mathcal{O}\left(\vartheta^{-1}n^{1+1/t}\operatorname{polylog}\right)$	Lemma 39
	$\mathcal{O}(t\log n)$	deterministic	$\mathcal{O}\left(\vartheta^{-1}n^{1+1/t}\operatorname{polylog}\right)$	Lemma 40
	$\mathcal{O}(t^2)$	deterministic	$\mathcal{O}\left(\vartheta^{-1}n^{1+1/t}\operatorname{polylog}\right)$	Lemma 40
Ultrametrics	$2+\varepsilon$	expectation	$\mathcal{O}(\vartheta^{-1}\varepsilon^{-2}n \operatorname{polylog})$	Lemma 41
	$(2+\varepsilon)t-1$	deterministic	$\mathcal{O}\left(\vartheta^{-2}\varepsilon^{-3}n^{1+1/t}\operatorname{polylog}\right)$	Lemma 42
Trees	$3+\varepsilon$	expectation	$\mathcal{O}(\vartheta^{-1}\varepsilon^{-2}n \operatorname{polylog})$	Lemma 43
	$(4+\varepsilon)t-3$	deterministic	$\mathcal{O}\left(\vartheta^{-2}\varepsilon^{-3}n^{1+1/t}\operatorname{polylog}\right)$	Lemma 44
Planar graphs	$3+\varepsilon$	expectation	$\mathcal{O}(\vartheta^{-1}\varepsilon^{-4}n \operatorname{polylog})$	Lemma 45
	$(4+\varepsilon)t-3$	deterministic	$\mathcal{O}(\vartheta^{-2}\varepsilon^{-6}n^{1+1/t}\operatorname{polylog})$	Lemma 46

Reliable spanners. An alternative is reliable spanners. For a parameter $\vartheta > 0$, a *ϑ-reliable t-spanner* has the property that for any failure (or attack) set B, the residual graph $G \setminus B$ has a t-spanner path between all pairs of points of $V \setminus B^+$, where $B^+ \supseteq B$ is some set such that $|B^+| \leqslant (1+\vartheta)|B|$. We consider two variants. In the standard model (i.e., adaptive adversarial model) the adversary "knows" the spanner G, and the set B is chosen as a worst case for G. In the oblivious (or randomized) case the spanner G is drawn from a probability distribution χ (over the same number of vertices). The adversary knows χ in advance, but not the sampled spanner. In this oblivious model, we require that $\mathbb{E}[|B^+|] \leqslant (1+\vartheta)|B|$. See Section 2 for precise definitions.

Previously, results for reliable spanners were only known in the geometric setting. For any point set $P \subseteq \mathbb{R}^d$, and for any constants $\vartheta, \varepsilon \in (0,1)$, one can construct a ϑ-reliable $(1+\varepsilon)$-spanner with only $\mathcal{O}\left(n \log n \operatorname{loglog}^6 n\right)$ edges [6]. The number of edges can be further reduced by using the relaxed notion of reliable spanners. For any point set $P \subseteq \mathbb{R}^d$ and parameters $\vartheta, \varepsilon \in (0,1)$, one can construct an oblivious $(1+\varepsilon)$-spanner that is ϑ-reliable in expectation and has $\mathcal{O}(n \operatorname{loglog}^2 n \operatorname{logloglog} n)$ edges [5].

Our results

We provide new constructions of reliable spanners for uniform metrics, finite metrics, ultrametrics, trees and planar graphs. Our new results are summarized in Table 1.1.

Technique. Our approach for constructing reliable spanners is in two steps: We first construct reliable spanners for uniform metrics and then reduce the problem of constructing reliable spanners for general metrics to uniform metrics using **covers**.

Spanners for uniform metrics. Uniform metrics have trivial classical 2-spanners: star graphs. It turns out that in the oblivious model one can simply use "constellation of stars" with a constant number of random centers. I.e. the spanner is linear in size. In the adaptive settings we present a lower bound of $\Omega(n^{1+1/t})$ edges for a reliable t-spanner. We present an asymptotically matching construction of a deterministic $(2t-1)$-reliable spanner with $\mathcal{O}(n^{1+1/t})$ edges. The construction is based on reliable expanders, i.e., expanders that remain expanding under attacks as defined above for spanners. See Section 3.

Covers. A ***t*-cover** of a finite metric space $\mathcal{M} = (P, \mathsf{d})$ is a family of subsets $\mathcal{C} = \{S \mid S \subseteq P\}$, such that for each pair $p, q \in P$ of points there exists a subset in $\mathcal{C}$ that contains both points and whose diameter is at most $t \cdot \mathrm{d}(p, q)$. Covers are used here to extend reliable spanners for uniform metrics into reliable spanners for general metrics. This is done by using spanners for uniform in each set of the cover and then taking a union of the edges of those graphs. See Section 5.

Naturally, the size $\sum_{S \in \mathcal{C}} |S|$ of a t-cover $\mathcal{C}$ is an important parameter in the resulting size of the spanner, so in Section 4 we study good covers. For general n-point spaces with spread at most Φ, we observe that the Ramsey partitions of [17] provide $\mathcal{O}(t)$ covers of size $\mathcal{O}(n^{1+1/t} \log \Phi)$, which is close to optimal, because of an $\Omega(n(n^{1/t} + \log_t \Phi))$ lower bound on the size that we prove. In more specific cases like ultrametrics, trees and planar graphs one can do better. Specifically, for trees and planar graphs one gets $(2+\varepsilon)$-covers of near linear size. For planar graphs, known partitions have much larger gap, which makes these results quite interesting.

New reliable spanners. Plugging the constructions of spanners for uniform metrics with the construction of covers yields reliable spanners for uniform, ultrametric, tree, planar, and general finite metrics. The results are summarized in Table 1.1.

Efficient construction. All our constructions relies on randomized constructions of expanders (over m vertices), that succeeds with probability $\geqslant 1 - 1/m^{O(1)}$. As such, the constructions described can be done efficiently, with potentially an extra $O(\log t)$ factor in the number of edges, if one wants a constructions of spanners, for n vertices, that succeeds with probability $1 - 1/n^{O(1)}$. See Remark 13 for details.

2 Preliminaries

2.1 Metric spaces

For a set $\mathcal{X}$, a function $\mathrm{d} : \mathcal{X}^2 \to [0, \infty)$, is a **metric** if it is symmetric, complies with the triangle inequality, and $\mathrm{d}(p, q) = 0 \iff p = q$. A **metric space** is a pair $\mathcal{M} = (\mathcal{X}, \mathrm{d})$, where d is a metric. For a point $p \in \mathcal{X}$, and a radius r, the **ball** of radius r is the set

$$\mathsf{b}(p, r) = \{q \in \mathcal{X} \mid \mathrm{d}(p, q) \leqslant r\}.$$

For a finite set $X \subseteq \mathcal{X}$, the **diameter** of X is

$$\mathrm{diam}(X) = \mathrm{diam}_{\mathcal{M}}(X) = \max_{p,q \in X} \mathrm{d}(p, q),$$

and the **spread** of X is $\Phi(X) = \frac{\mathrm{diam}(X)}{\min_{p,q \in X, p \neq q} \mathrm{d}(p,q)}$. A metric space $\mathcal{M} = (\mathcal{X}, \mathrm{d})$ is **finite**, if $\mathcal{X}$ is a finite set. In this case, we use $\Phi = \Phi(\mathcal{X})$ to denote the spread of the (finite) metric.

A natural way to define a metric space is to consider an undirected connected graph $G = (P, E)$ with positive weights on the edges. The *shortest path metric* of G, denoted by d_G, assigns for any two points $p, q \in P$ the length of the shortest path between p and q in the graph. Thus, any graph G readily induces the finite metric space $(V(G), \mathrm{d}_G)$. If the graph is unweighted, then all the edges have weight 1.

A **tree metric** is a shortest path metric defined over a graph that is a tree.

2.2 Reliable spanners

▶ **Definition 1.** *For a metric space $\mathcal{M} = (P, \mathrm{d})$, a graph $H = (P, E)$ is a **t-spanner**, if for any $p, q \in P$, we have that $\mathrm{d}(p, q) \leqslant \mathrm{d}_H(p, q) \leqslant t \cdot \mathrm{d}(p, q)$.*

Given a weighted graph $G = (V, E)$, and $B \subseteq V$ we denoted by $G|_B$ the subgraph induced on B. We also use the notation $G \setminus B = G|_{V \setminus B}$.

An **attack** on a graph 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 G.

▶ **Definition 2** (Reliable spanner). *Let $G = (P, E)$ be a t-spanner for some $t \geqslant 1$ constructed by a (possibly) randomized algorithm. Given an attack B, its **damaged set** B^+ is a set of smallest possible size, such that for any pair of vertices $p, q \in P \setminus B^+$, we have*

$$\mathrm{d}_{G \setminus B}(p, q) \leqslant t \cdot \mathrm{d}(p, q), \tag{2.1}$$

*that is, distances are preserved (up to a factor of t) 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** (in the deterministic or non-oblivious setting), if $\lambda(G, B) \leqslant \vartheta$ holds for any attack $B \subseteq P$. Furthermore, the graph G is **ϑ-reliable in expectation** (or in the oblivious model), if $\mathbb{E}[\lambda(G, B)] \leqslant \vartheta$ holds for any oblivious attack $B \subseteq P$.*

▶ Remark. The damaged set B^+ is not necessarily unique, since there might be freedom in choosing the point to include in B^+ for a pair that does not have a t-path in $G \setminus B$.

Miscellaneous. For a graph G, and a set of vertices $Y \subseteq V(G)$, let

$$\Gamma(Y) = \{x \in V(G) \mid xy \in E(G) \text{ and } y \in Y\}$$

denote the **neighbors** of Y in G.

▶ **Definition 3.** *For a collection of sets $\mathcal{F}$, and an element x, let*

$$\deg(x, \mathcal{F}) = |\{X \in \mathcal{F} \mid x \in X\}|$$

*denote the **degree** of x in $\mathcal{F}$. The maximum degree of any element of $\mathcal{F}$ is the **depth** of $\mathcal{F}$.*

Notations. We use $P + p = P \cup \{p\}$ and $P - p = P \setminus \{p\}$. Similarly, for a graph G, and a vertex p, let $G - p$ denote the graph resulting from removing p.

3 Reliable spanners for uniform metric

Let P be a set of n points and let (P, d) be a metric space equipped with the uniform metric, that is, for all distinct pairs $p, q \in P$, we have that $\mathrm{d}(p, q)$ is the same quantity. Note that $n - 1$ edges are enough to achieve a 2-spanner for the uniform metric by using the star graph.

3.1 A randomized construction for the oblivious case

Construction. Let $\vartheta \in (0,1)$ be a fixed parameter. Set $k = 2\lceil \vartheta^{-1}\log\vartheta^{-1}\rceil + 1$ and sample k points from P uniformly at random (with replacement). Let $C \subseteq P$ be the resulting set of *center points*. For each point $p \in C$, build the star graph $*_p = (P, \{pq \mid q \in P - p\})$, where p is the center of the star. The **constellation** of C is the graph $\circledast = \bigcup_{p\in C} *_p$, which is the union of the star graphs induced by centers in C.

▶ **Lemma 4** (For the proof, see [11]). *The constellation $\circledast$, defined above, is a ϑ-reliable 2-spanner in expectation. The number of its edges is $\mathcal{O}(n\vartheta^{-1}\log\vartheta^{-1})$.*

3.2 Lower bound for a deterministic construction

In the *non-oblivious settings*, the attacker knows the constructed graph G when choosing the attack set B.

▶ **Lemma 5** (For the proof, see [11]). *Let $G = (P, E)$ be a ϑ-reliable t-spanner on P for the uniform metric, where $\vartheta \in (0,1)$ and $t \geqslant 1$. Then, in the non-oblivious settings, G must have $\Omega(n^{1+1/t})$ edges.*

▶ Remark 6. Erdős' girth conjecture states that there exists a graph G with n vertices and $\Omega(n^{1+1/k})$ edges, and girth at least $2k+1$, where the *girth* of G is the length of the shortest cycle in G. The argument in the proof of Lemma 5 is closely related to the standard argument for proving a tight counterpart – any graph with $\Omega(n^{1+1/k})$ edges has girth at most $2k+1$.

3.3 Reliable spanners of the uniform metric for adaptive adversary

Here, we present a construction of reliable spanner that is close to being tight. The spanner is simply a high-degree expander whose properties are described in the following definition.

▶ **Definition 7.** *Denote by $\lambda(G)$ the second eigenvalue of the matrix M/d, where $M = \mathrm{Adj}(G)$ is the adjacency matrix of a d-regular graph G. A* ***proper expander*** *specifies a constant $C > 1$, and functions $\mathscr{C}_\delta, c_\delta > 0$, such that for every $\delta \in (0, 1/4)$ an even integers $d \geqslant \mathscr{C}_\delta$, $n \geqslant d^2$, there exists an n-vertex, d-regular graph $G = (V, E)$, such that:*

(P1) $\forall S \subseteq V,\ |S| \geqslant 12n/(\delta d) \implies |\Gamma_V(S)| > (1-\delta)n$,

(P2) $\forall S \subseteq V,\ |S| \leqslant c_\delta n/d \implies |\Gamma_V(S)| \geqslant (1-\delta)d|S|$.

(P3) $\lambda(G) \leqslant C/\sqrt{d}$.

For each one of the properties above, it is known that there exists an expander satisfying it: Property (P1) is essentially proved in [6], Property (P2) is folklore, and Property (P3) appears in [8]. Since they hold almost surely for "random regular graphs", they also hold simultaneously. However, we were unable to find in the literature proofs of almost sure existence in the same model of random regular graphs, and contiguity of the different random models does not necessarily hold in the high-degree regime (which is what we need here). Therefore, for completeness, in the full version we reprove (P1) and (P2) in the same random model in which (P3) was proved. We thus get the following in the full version of the paper [11].

▶ **Theorem 8.** *A random graph construction leads to is a proper expander (see Definition 7), asymptotically almost surely.*

With the appropriate choice of parameters, these expanders are reliable spanners for uniform metrics. The proof is somewhat cumbersome and is included in the full version of the paper.

▶ **Theorem 9.** *For every $t \in \mathbb{N}$, $\theta \in (0,1)$, and $n \in 2\mathbb{N}$ such that $n \geqslant e^{\Omega(t)}$, there exist:*

(i) *ϑ-reliable $2t$-spanner with $\mathcal{O}(\vartheta^{-1} n^{1+1/t})$ edges for n-point uniform space, and*

(ii) *ϑ-reliable $(2t-1)$-spanner with $\mathcal{O}(\vartheta^{-2} n^{1+1/t})$ edges for n-point uniform space.*

Applying the above theorem directly on non-uniform metric we obtain the following corollaries.

▶ **Corollary 10.** *Let $\mathcal{M} = (\mathcal{X}, \mathrm{d})$ be a metric space, and let $P \subseteq \mathcal{X}$ be a finite subset of size n. Given parameters $t \in \mathbb{N}$, and $\vartheta \in (0,1)$, there exists a weighted graph G on P, such that:*

(A) *The graph G has $|E(G)| = \mathcal{O}\big(\vartheta^{-2} \cdot n^{1+1/t}\big)$ edges.*

(B) *The graph G is ϑ-reliable. Namely, given any attack set $B \subset \mathcal{X}$, there exists a subset $Q \subseteq P$, such that $|Q| \geqslant |P| - (1+\vartheta)|B \cap P|$. Furthermore, for any two points $p, q \in Q$, we have*

$$\mathrm{d}_{\mathcal{M}}(p,q) \leqslant \mathrm{d}_{G|_Q}(p,q) \leqslant (2t-1) \cdot \mathrm{diam}_{\mathcal{M}}(P),$$

and the path realizing it has at most $(2t-1)$ hops.

In particular, G has hop diameter at most $2t-1$, and diameter at most $(2t-1) \cdot \mathrm{diam}_{\mathcal{M}}(P)$.

▶ **Corollary 11.** *Let $\mathcal{M} = (\mathcal{X}, \mathrm{d})$ be a metric space, and let $P \subseteq \mathcal{X}$ be a finite subset of size n. Given parameters $t \in \mathbb{N}$, and $\vartheta \in (0,1)$, there exists a weighted graph G on P, such that:*

(A) *The graph G has $|E(G)| = \mathcal{O}\big(\vartheta^{-1} \cdot n^{1+1/t}\big)$ edges.*

(B) *The graph G is ϑ-reliable. Namely, given any attack set $B \subset \mathcal{X}$, there exists a subset $Q \subseteq P$, such that $|Q| \geqslant |P| - (1+\vartheta)|B \cap P|$. Furthermore, for any two points $p, q \in Q$, we have*

$$\mathrm{d}_{\mathcal{M}}(p,q) \leqslant \mathrm{d}_{G|_Q}(p,q) \leqslant 2t \cdot \mathrm{diam}_{\mathcal{M}}(P),$$

and the path realizing it has at most $2t$ hops.

In particular, G has hop diameter at most $2t$, and diameter at most $2t \cdot \mathrm{diam}_{\mathcal{M}}(P)$.

3.3.1 A simpler construction of reliable spanner

Here, we prove a somewhat inferior construction of reliable spanner. In particular, the ensuing construction cannot produce reliable expanders with a constant degree but simpler to prove.

The proof will only use Property (P1) of proper expanders from Definition 7. Specifically, it will use the following lemma which is a minor variant of known constructions.

▶ **Lemma 12** ([6]). *Let L, R be two disjoint sets, with a total of $n \in 2\mathbb{N}$ elements, and let $\xi \in (0,1)$ be a parameter. There exists 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 $|X| \geqslant \xi|L|$, we have that $|\Gamma(X)| > (1-\xi)|R|$, and*

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

▶ Remark 13. The randomized construction of Lemma 12 succeeds with probability $1 - 1/n^{O(1)}$. Since we use the construction below on sets that are polynomially large (i.e., $n^{1/t}$), one can use Lemma 12 constructively in this case (potentially losing an additional $\log t$ factor). This applies to the other expander constructions used in this paper. Note, that while the construction works with high probability, verifying that it indeed works seems computationally intractable.

Construction of the reliable spanner. We define a tree T that contains the n points in its leaves and has exactly $t+1$ levels. Level 0 corresponds to the leaves, and in level k each node has exactly $m = n^{1/t}$ children from level $k-1$, for $k = 1, \ldots, t$. Note that each level corresponds to a partition of P. At level k, each cluster (i.e., tree node) is of size $n^{k/t}$, and there are $n^{1-k/t}$ clusters (i.e., nodes) at this level. For a node $u \in T$, let $v_1, \ldots, v_m$ be the children of u, and let P_u be the set of leaves of the subtree rooted at u. The graph G_u is defined to be the union of bipartite expanders between P_{v_i} and P_{v_j}, for $1 \leqslant i < j \leqslant m$, using Lemma 12 with parameter $\xi = \vartheta/8$. The graph G is the union of graphs G_u for all internal nodes $u \in T$.

Analysis.

▶ **Lemma 14** (For the proof, see [11]). *The graph G, defined above, has $\mathcal{O}(\vartheta^{-2}t \cdot n^{1+1/t})$ edges.*

Fix an arbitrary attack set $B \subseteq P$. A node $u \in T$ is **damaged**, if $|P_u \cap B| \geqslant (1-\vartheta/2)\,|P_u|$ (i.e., most of the nodes in the subtree of u are in the attack set). Let B^+ be all the points of P, that get removed when one removes all the subtrees rooted at damaged nodes.

▶ **Lemma 15** (For the proof, see [11]). *Let B^+ be the set defined above. We have $|B^+| \leqslant (1+\vartheta)\,|B|$.*

Let $\ominus(p, \ell)$ denote the ball of radius ℓ centered at p in the graph $G \backslash B$. For a point $p \in P$, let $u(p, \ell)$ be the node in the tree that is the ancestor of p that is ℓ hops upward toward the root. For a point $p \in P$ and $\ell \geqslant 1$, let

$$\beta(p,\ell) = \left|P_{u(p,\ell)} \cap B\right| \Big/ \left|P_{u(p,\ell)}\right|, \qquad \text{and} \qquad \omega(p,\ell) = \left|P_{u(p,\ell)} \cap \ominus(p,\ell)\right| \Big/ \left|P_{u(p,\ell)}\right|,$$

be the ratio of attack points in $P_{u(p,\ell)}$, and the ratio of ℓ-reachable survivors from p in $P_{u(p,\ell)}$, respectively.

▶ **Lemma 16** (For the proof, see [11]). *For any $p \in P \backslash B^+$, and any $\ell = 1, \ldots, t$, we have $\omega(p, \ell) \geqslant 3\vartheta/8$, where t is the height of T.*

▶ **Proposition 17** (For the proof, see [11]). *For every $t \in \mathbb{N}$, $\theta \in (0,1)$, and $n \in \mathbb{N}$ such that $n^{1/t} \in \{2, 3, 4, \ldots\}$, there exists ϑ-reliable $(2t-1)$-spanner with $\mathcal{O}\big(\vartheta^{-2}tn^{1+1/t}\big)$ edges for n-point uniform space.*

▶ Remark 18. A constructive version of Proposition 17 requires repeating the construction of the underlying expanders $O(\log t)$ times, so that it succeeds (per such expander) with probability $> 1 - 1/n$. This would yield an extra $O(\log t)$ factor in the number of edges in the graph, see Remark 13.

4 Covers for trees, bounded spread metrics, and planar graphs

▶ **Definition 19.** *For a finite metric space $\mathcal{M} = (P, \mathsf{d})$, a **t-cover**, is a family of subsets $\mathcal{C} = \{S_i \subseteq P \mid i = 1, \ldots m\}$, such that for any $p, q \in P$, there exists an index i, such that $p, q \in S_i$, and*

$$\mathrm{diam}(S_i)/t \leqslant \mathrm{d}(p,q) \leqslant \mathrm{diam}(S_i).$$

*The **size** of a cover $\mathcal{C}$ is $\mathrm{size}(\mathcal{C}) = \sum_{S \in \mathcal{C}} |S|$. For a point $p \in P$, its **degree** in $\mathcal{C}$, is the number of sets of P that contain it. The **depth** of $\mathcal{C}$ is the maximum degree of any element of P, and is denoted by $\mathsf{D}(\mathcal{C})$.*

4.1 Lower bounds

Unfortunately, in the worst case, the depth of any cover and its size must depend on the spread of the metric.

▶ **Lemma 20** (For the proof, see [11]). *For any parameter $t > 1$, any integer $h > 1$, and $\Phi = t^h$, there exists a metric $\mathcal{M} = (P, \mathrm{d})$ of n points, such that*

(I) $\Phi(P) = \Phi$, *and*

(II) *any t-cover $\mathcal{C}$ of P has size $\Omega(n \log_t \Phi) = \Omega(nh)$, average degree $\geqslant h/2$, and depth h.*

▶ **Proposition 21** (For the proof, see [11]). *For any $t \in 2, 3, \ldots$, and any sufficiently large n there exists an n-point metric space for which any t-cover must be of size at least $\Omega(n^{1+1/2t})$.*

4.2 Cover for ultrametrics

▶ **Definition 22.** *A* **hierarchically well-separated tree** *(***HST** *) is a metric space defined on the leaves of a rooted tree T. To each vertex $u \in T$ there is associated a label $\Delta_u \geqslant 0$. This label is zero for all the leaves of T, and it is a positive number for all the interior nodes. The labels are such that if a vertex u is a child of a vertex v, then $\Delta_u \leqslant \Delta_v$. The distance between two leaves $x, y \in T$ is defined as $\Delta_{\mathrm{lca}(x,y)}$, where $\mathrm{lca}(x, y)$ is the least common ancestor of x and y in T. An HST T is a k-***HST** *if for all vertices $v \in T$, we have that $\Delta_v \leqslant \Delta_{\overline{\mathrm{p}}(v)}/k$, where $\overline{\mathrm{p}}(v)$ is the parent of v in T.*

HSTs are one of the simplest non-trivial metrics, and surprisingly, general metrics can be embedded (randomly) to trees with expected distortion $\mathcal{O}(\log n)$ [2, 9].

▶ **Definition 23.** *A metric $\mathcal{M} = (P, \mathrm{d})$ is an* ***ultrametric****, if for any $x, y, z \in P$, we have that $\mathrm{d}(x, z) \leqslant \max(\mathrm{d}(x, y), \mathrm{d}(y, z))$. Notice, that this is a stronger version of the triangle inequality, which states that $\mathrm{d}(x, z) \leqslant \mathrm{d}(x, y) + \mathrm{d}(y, z)$.*

The following is folklore, easy to verify (see, e.g., [3, Lemma 3.5].

▶ **Lemma 24.** *For $k \geqslant 1$, every finite ultrametric can be k-approximated by a k-HST.*

▶ **Lemma 25** (For the proof, see [11]). *For $k > 1$, every k-HST with spread Φ has 1-cover of depth at most $\log_k \Phi$.*

The following corollary is immediate from the above two lemmas.

▶ **Corollary 26.** *Let $\mathcal{M} = (P, \mathrm{d})$ be an ultrametric over n points with spread Φ. For any $\varepsilon \in (0, 1)$, one can compute a $(1 + \varepsilon)$-cover of $\mathcal{M}$ of depth $\mathcal{O}(\varepsilon^{-1} \log \Phi)$.*

4.3 Cover for general finite metrics

We need the following result.

▶ **Lemma 27** ([17]). *Let (P, d) be an n-point metric space and $k \geqslant 1$. Then there exists a distribution over decreasing sequences of subsets $P = P_0 \supsetneq P_1 \supsetneq \cdots \supsetneq P_s = \varnothing$ (s itself is a random variable), such that for all $\mu > -1/k$, we have $\mathbb{E}\Big[\sum_{j=1}^{s-1} |P_j|^{\mu}\Big] \leqslant \max\Big(\frac{k}{1+\mu k}, 1\Big) \cdot n^{\mu+1/k}$, and such that for each $j \in [s]$ there exists an ultrametric ρ_j on P satisfying for every $p, q \in P$, that $\rho_j(p, q) \geqslant \mathrm{d}(p, q)$, and if $p \in P$ and $q \in P_{j-1} \backslash P_j$ then $\rho_j(p, q) \leqslant \mathcal{O}(k) \cdot \mathrm{d}(p, q)$.*

By computing a cover for each ultrametric generated by the above lemma, we get the following.

▶ **Lemma 28** (For the proof, see [11]). *For an n-point metric space $\mathcal{M} = (P, \mathsf{d})$ with spread Φ, and a parameter $k > 1$, one can compute, in polynomial time, a $\mathcal{O}(k)$-cover of $\mathcal{M}$ of size $\mathcal{O}(n^{1+1/k} \log \Phi)$ and depth $\mathcal{O}(kn^{1/k} \log \Phi)$.*

4.4 Covers for trees

Using a tree separator, and applying it recursively, implies the following construction of covers for trees.

▶ **Lemma 29** (For the proof, see [11]). *For a weighted tree metric $T = (P, \mathsf{d})$, with spread Φ, and a parameter $\varepsilon \in (0,1)$, one can compute in polynomial time a $(2+\varepsilon)$-cover of T of depth $\mathcal{O}(\varepsilon^{-1}\log\Phi\log n)$, and size $\mathcal{O}(n\varepsilon^{-1}\log\Phi\log n)$, where $n = |P|$.*

4.5 Covers for planar graphs

Preliminaries. The next lemma can be traced back to the work of Lipton and Tarjan [15], and we include a sketch of the proof for completeness.

▶ **Lemma 30** (For the proof, see [11]). *Let $H = (P, E)$ be a planar triangulated graph with non-negative edge weights. There is a partition of P to three sets X, Y, Z, such that*
(i) *$|X| \leqslant (2/3)n$ and $|Y| \leqslant (2/3)n$,*
(ii) *there is no edge between X and Y, and*
(iii) *Z is composed of two interior disjoint shortest paths that share one of their endpoints, and an edge connecting their other two endpoints.*

▶ **Definition 31.** *For a metric space $(\mathcal{X}, \mathrm{d})$ and a parameter r, an r-net N is a maximal set of points in $\mathcal{X}$, such that*
(i) *for any two net points $p, q \in N$, we have $\mathrm{d}(p,q) > r$, and*
(ii) *for any $p \in \mathcal{X}$, we have $\mathsf{d}(p, N) = \min_{q\in N} \mathrm{d}(p,q) \leqslant r$.*

A net can be computed by repeatedly picking the point furthest away from the current net N, and adding it to the net if this distance is larger than r, and stopping otherwise. We denote a net computed by this algorithm by $\mathrm{net}(\mathcal{X}, r)$.

The following lemma testifies that if we restrict the net to lay along a shortest path in the graph, locally the cover it induces has depth as if the graph was one dimensional.

▶ **Lemma 32** (For the proof, see [11]). *Let G be a weighted graph, and let d be the shortest path metric it induces. Let π be a shortest path in G and let $N = \mathrm{net}(\pi, r) \subseteq \pi$ be a net computed for some distance $r > 0$. For some $R > 0$, consider the set of balls $\mathcal{B} = \{\mathsf{b}(p, R) \mid p \in N\}$. For any point $q \in V(G)$, we have that the degree of q in $\mathcal{B}$ is at most $2R/r + 1$.*

Construction. Let $\varepsilon \in (0,1)$ be an input parameter, and let G be a weighted planar graph. We assume that G is triangulated, as otherwise it can be triangulated (we also assume that we have its planar embedding). Any new edge pq is assigned as weight the distance between its endpoints in the original graph. This can be done in linear time. As usual, we assume that the minimum distance in G is one, and its spread is Φ.

Let Z be the cycle separator given by Lemma 30 made out of two shortest paths π_1 and π_2. Let p_1, p_2, p_3 be the endpoints of these two paths.

For $i = 0, \ldots, m = \lceil \log_{1+\varepsilon/8} \Phi \rceil$, let $N_i = \mathrm{net}(\pi_1, \varepsilon r_i/8) \cup \mathrm{net}(\pi_2, \varepsilon r_i/8) \cup \{p_1, p_2, p_3\}$, where $r_i = (1+\varepsilon/8)^i$. The associated set of balls is

$$\mathcal{B}_i = \{\mathsf{b}(p, (1+\varepsilon/8)r_i) \mid p \in N_i\}.$$

The resulting set of balls is $\mathcal{B}(Z) = \bigcup_i \mathcal{B}_i$. We add the sets of $\mathcal{B}(Z)$ to the cover, and continue recursively on the connected components of $G - Z$. Let $\mathcal{C}$ denote the resulting cover.

Analysis.

▶ **Lemma 33** (For the proof, see [11]). *For any two vertices $p, q \in V(G)$, there exists a cluster $C \in \mathcal{C}$, such that $p, q \in C$, and $\operatorname{diam}(C) \leqslant (2+\varepsilon)\mathrm{d}_G(p,q)$. That is, $\mathcal{C}$ is a $(2+\varepsilon)$-cover of G.*

▶ **Lemma 34** (For the proof, see [11]). *The depth of $\mathcal{C}$ is $\mathcal{O}(\varepsilon^{-2}\log n \log \Phi)$.*

4.5.1 The result

▶ **Theorem 35.** *Let G be a weighted planar graph over n vertices with spread Φ. Then, given a parameter $\varepsilon \in (0,1)$, one can construct a $(2+\varepsilon)$-cover of G with depth $\mathcal{O}(\varepsilon^{-2}\log n \log \Phi)$ in polynomial time.*

▶ Remark 36. It is possible generalize Theorem 35 to the shortest path metric on families of graphs excluding a fixed minor. Specifically, by [12, Lemma 3.3], there exists $O(s^2)$-cover of depth $O(3^s \log \Phi)$ for every metric spaces supported on graph excluding K_s minor and spread Φ. It may be possible to improve the approximation parameter to $O(s)$ using [1], but we have not pursued this avenue. However, this approach can not recover approximation parameter $2+\varepsilon$ for planar graphs as in Theorem 35.

5 From covers to reliable spanners

5.1 The oblivious construction

▶ **Lemma 37** (For the proof, see [11]). *Let $\mathcal{M} = (P, \mathrm{d})$ be a finite metric space, and suppose there exists $\mathcal{C}$, a ξ-cover of $\mathcal{M}$ of size s and depth D. Then, there exists an oblivious ϑ-reliable 2-hop 2ξ-spanner for $\mathcal{M}$, of size $\mathcal{O}\big(s\frac{\mathsf{D}}{\vartheta}\log\frac{\mathsf{D}}{\vartheta}\big)$.*

5.2 The deterministic construction

▶ **Lemma 38** (For the proof, see [11]). *Let $\mathcal{M} = (P, \mathrm{d})$ be a finite metric space over n points, and let $\mathcal{C}$ be a ξ-cover of it of depth D and size s. Then, for any integer $t \geqslant 1$, there exists:*

(A) *A ϑ-reliable $(2t-1)$-hop $(2t-1)\xi$-spanner for $\mathcal{M}$, of size $\mathcal{O}(\vartheta^{-2}\mathsf{D}^2 s n^{1/t})$.*

(B) *A ϑ-reliable $2t$-hop $2t\xi$-spanner for $\mathcal{M}$, of size $\mathcal{O}(\vartheta^{-1}\mathsf{D} s n^{1/t})$.*

5.3 Applications

5.3.1 General metrics

▶ **Lemma 39** (For the proof, see [11]). *Let $\mathcal{M} = (P, \mathrm{d})$ be an n-point metric space of spread at most Φ, and let $\vartheta \in (0,1)$ and $k \in \mathbb{N}$ be parameters. Then, one can build an oblivious ϑ-reliable $\mathcal{O}(k)$-spanner for $\mathcal{M}$ with*

$$\mathcal{O}\Big(\vartheta^{-1} k n^{1+1/k} \log^2 \Phi \log \frac{k n^{1/k} \log \Phi}{\vartheta}\Big)$$

edges. In particular, for $k = \log n$, we obtain ϑ-reliable $\mathcal{O}(\log n)$-spanner for $\mathcal{M}$ with

$$\mathcal{O}\big(\vartheta^{-1} n \log n \log^2 \Phi (\log\log n + \log\log \Phi + \log(\vartheta^{-1}))$$

edges.

▶ **Lemma 40** (For the proof, see [11]). *Let $\mathcal{M} = (P, \mathrm{d})$ be a finite metric over n points of spread Φ, and let $\vartheta \in (0,1)$ and $k, h \in \mathbb{N}$ int be parameters. Then, one can build a ϑ-reliable $\mathcal{O}(kt)$-spanner for $\mathcal{M}$ with $\mathcal{O}(\vartheta^{-1}kn^{1+1/k+1/t}\log^2\Phi)$ edges. In particular, when $t = \log n$, we obtain ϑ-reliable $\mathcal{O}(k\log n)$-spanner for $\mathcal{M}$ with $\mathcal{O}(\vartheta^{-1}kn^{1+1/(2k)}\log^2\Phi)$ edges, and when $t = k$ we obtain ϑ-reliable $\mathcal{O}(t^2)$-spanner for $\mathcal{M}$ with $\mathcal{O}(\vartheta^{-1}tn^{1+1/t}\log^2\Phi)$ edges.*

5.3.2 Ultrametrics

▶ **Lemma 41** (For the proof, see [11]). *Let $\mathcal{M} = (P, \mathrm{d})$ be an ultrametric over n points with spread Φ, and let $\vartheta, \varepsilon \in (0,1)$ be parameters. Then, one can build an* oblivious *ϑ-reliable $(2+\varepsilon)$-spanner for $\mathcal{M}$ with $\mathcal{O}\big(\vartheta^{-1}\varepsilon^{-2}n\log^2\Phi\log\frac{\log\Phi}{\vartheta\varepsilon}\big)$ edges.*

▶ **Lemma 42** (For the proof, see [11]). *Let $\mathcal{M} = (P, \mathrm{d})$ be an ultrametric over n points with spread Φ, and let $\vartheta, \varepsilon \in (0,1)$, and $t \in \mathbf{N}$ be parameters. Then, one can build a ϑ-reliable $((2+\varepsilon)t-1)$-spanner for $\mathcal{M}$ of size $\mathcal{O}(\vartheta^{-2}\varepsilon^{-3}t \cdot n^{1+1/t}\log^3\Phi)$.*

5.3.3 Tree metrics

▶ **Lemma 43** (For the proof, see [11]). *Let $\mathcal{M} = (P, \mathrm{d})$ be a tree metric over n points with spread Φ, and let $\vartheta, \varepsilon \in (0,1)$ be parameters. Then, one can build an* oblivious *ϑ-reliable $(3+\varepsilon)$-spanner for $\mathcal{M}$ with $\mathcal{O}\big(\vartheta^{-1}\varepsilon^{-2}n\,\mathrm{polylog}(n,\Phi)\big)$ edges, where* $\mathrm{polylog}(n,\Phi) = \log^2 n\log^2\Phi\log\frac{\log\Phi\log n}{\vartheta\varepsilon}$.

▶ **Lemma 44** (For the proof, see [11]). *Let $\mathcal{M} = (P, \mathrm{d})$ be a tree metric over n points with spread Φ, and let $\vartheta, \varepsilon \in (0,1)$, and $t \in \mathbf{N}$ be parameters. Then, one can build a ϑ-reliable $((4+\varepsilon)t-3)$-spanner for $\mathcal{M}$ of size $\mathcal{O}(\vartheta^{-2}\varepsilon^{-3} \cdot n^{1+1/t}\log^3 n\log^3\Phi)$.*

5.3.4 Planar graphs

▶ **Lemma 45** (For the proof, see [11]). *Let G be a weighted planar graph with n vertices and spread Φ. Furthermore, let $\vartheta, \varepsilon \in (0,1)$ be parameters. Then, one can build an oblivious ϑ-reliable $(3+\varepsilon)$-spanner for G with $\mathcal{O}\big(\vartheta^{-1}\varepsilon^{-4}n\,\mathrm{polylog}(n,\Phi)\big)$ edges, where* $\mathrm{polylog}(n,\Phi) = \log^2 n\log^2\Phi\log\frac{\log\Phi\log n}{\vartheta\varepsilon}$.

▶ **Lemma 46** (For the proof, see [11]). *Let G be a weighted planar graph with n vertices and spread Φ. Furthermore, let $\vartheta, \varepsilon \in (0,1)$ and $t \in \mathbf{N}$ be parameters. Then, one can build a deterministic ϑ-reliable $((4+\varepsilon)t-3)$-spanner for G of size $\mathcal{O}(\vartheta^{-2}\varepsilon^{-6} \cdot n^{1+1/t}\log^3 n\log^3\Phi)$.*

6 Concluding remarks and open problems

Subsequent work. Recently Filtser and Le [10] improved the results here, getting bounds that do not depend on the spread of the metrics in some cases, for the oblivious adversary case.

Tradeoffs in deterministic constructions for general spaces. Classical spanners are known to have an approximation–size trade-off for general n-point metrics: To achieve $\Theta(t)$ approximation it is sufficient and necessary to have $n^{1+1/t}$ edges in the worst case. In contrast, for reliable spanners we were able only to give an upper bound on the trade-off, with no asymptotically matching lower bound: To achieve $\mathcal{O}(t^2)$ approximation it is sufficient to have $\widetilde{\mathcal{O}}(n^{1+1/t})$ edges. Classically, the uniform metric is 2-approximated by a star graph with only $n-1$ edges. In contrast, we have shown here reliable spanners for uniform metric have approximation–size trade similar to the classical spanner for general metrics. The connection between the two problems is quite intriguing, and is worthy of further research.

The dependence of the size on the spread. The size of spanners constructed in this paper depends on the spread of the metric space. This is because of the reduction to uniform spaces via covers, in which the dependence on the spread is unavoidable in general. However, in some setting this dependence is avoidable. For example [6, 5] achieves this for fixed dimensional Euclidean spaces, and [10] achieves it for doubling spaces, and general finite spaces in the oblivious adversary model. Getting spread-free bounds for the non-oblivious adversary is an interesting problem for further research.

Explicit constructions. To the best of our knowledge, there is no known polynomial time deterministic algorithm for constructing expanders with Property (P1) or Property (P2). Getting such a construction is an interesting open problem.

References

1 Ittai Abraham, Cyril Gavoille, Anupam Gupta, Ofer Neiman, and Kunal Talwar. Cops, robbers, and threatening skeletons: padded decomposition for minor-free graphs. *SIAM J. Comput.*, 48(3):1120–1145, 2019. `doi:10.1137/17M1112406`.

2 Y. Bartal. On approximating arbitrary metrics by tree metrics. In *Proc. 30th Annu. ACM Sympos. Theory Comput.* (STOC), pages 161–168, 1998. `doi:10.1145/276698.276725`.

3 Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. On metric Ramsey-type phenomena. *Ann. of Math. (2)*, 162(2):643–709, 2005. `doi:10.4007/annals.2005.162.643`.

4 Greg Bodwin, Michael Dinitz, Merav Parter, and Virginia Vassilevska Williams. *Optimal Vertex Fault Tolerant Spanners (for fixed stretch)*, pages 1884–1900. Society for Industrial and Applied Mathematics, 2018. `doi:10.1137/1.9781611975031.123`.

5 K. Buchin, S. Har-Peled, and D. Oláh. Sometimes reliable spanners of almost linear size. *ArXiv*, 2019. (accepted to ESA 2020). `arXiv:1912.01547`.

6 K. Buchin, S. Har-Peled, and D. Oláh. A spanner for the day after. In *Proceedings of the 35th Symposium on Computational Geometry*, pages 19:1–19:15, 2019. `doi:10.4230/LIPIcs.SoCG.2019.19`.

7 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`.

8 Ioana Dumitriu, Tobias Johnson, Soumik Pal, and Elliot Paquette. Functional limit theorems for random regular graphs. *Probab. Theory Related Fields*, 156(3-4):921–975, 2013. `doi:10.1007/s00440-012-0447-y`.

9 J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *Journal of Computer and System Sciences*, 69(3):485–497, 2004. `doi:10.1016/j.jcss.2004.04.011`.

10 Arnold Filtser and Hung Le. Reliable spanners: Locality-sensitive orderings strike back. *CoRR*, abs/2101.07428, 2021. `arXiv:2101.07428`.

11 Sariel Har-Peled, Manor Mendel, and Dániel Oláh. Reliable spanners for metric spaces. *CoRR*, abs/2007.08738v2, 2021. `arXiv:2007.08738v2`.

12 R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: A new embedding method for finite metric spaces. *Geom. funct. anal. (GAFA)*, 15(4):839–858, 2005.

13 C. Levcopoulos, G. Narasimhan, and M. Smid. Efficient algorithms for constructing fault-tolerant geometric spanners. In *Proceedings of the 30th ACM Symposium on Theory of Computing*, pages 186–195, 1998. `doi:10.1145/276698.276734`.

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

15 R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979. `doi:10.1137/0136016`.

16 T. Lukovszki. New results of fault tolerant geometric spanners. In *Proceedings of the 6th Workshop on Algorithms and Data Structures*, pages 193–204, 1999. doi:10.1007/3-540-48447-7_20.

17 M. Mendel and A. Naor. Ramsey partitions and proximity data structures. *J. Euro. Math. Soc.*, 009(2):253–275, 2007. URL: http://eudml.org/doc/277787.

18 G. Narasimhan and M. Smid. *Geometric spanner networks*. Cambridge University Press, New York, NY, USA, 2007.

19 D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. Society for Industrial and Applied Mathematics, 2000. doi:10.1137/1.9780898719772.

20 D. Peleg and A. Schäffer. Graph spanners. *J. Graph Theory*, 13:99–116, 1989.

21 S. Solomon. From hierarchical partitions to hierarchical covers: Optimal fault-tolerant spanners for doubling metrics. In *Proceedings of the 46th ACM Symposium on Theory of Computing*, STOC '14, page 363–372, 2014. doi:10.1145/2591796.2591864.