

Reliable Spanners for Metric Spaces*

Sariel Har-Peled[†]

Manor Mendel[‡]

Dániel Oláh[§]

August 4, 2023

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. In other words, 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.

1. Introduction

Let $\mathcal{M} = (P, d)$ 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* if for any pair of vertices $u, v \in P$ we have $d_G(u, v) \leq t \cdot d(u, v)$, where $d_G(u, v)$ is the length of the shortest path between u and v in G , and $d(u, v)$ is the distance in the metric space between u and v . Spanners were first introduced by Peleg and Schäffer [PS89] as a tool in distributed computing, but have since found use in many other areas, such as algorithms, networking, data structures, and metric geometry, see [Pel00, NS07].

Fault tolerant spanners. A desired property of spanners is the ability to withstand failures of some of their vertices. One such notion is provided by fault tolerance [CLNS15, LNS98, LNS02, Luk99, Sol14]. A graph G is a *k-fault tolerant t-spanner*, if for any subset of vertices B , with $|B| \leq k$, the graph $G \setminus B$ is a *t-spanner*. However, for *k-fault tolerant* graphs 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 $t \geq 1$, the optimal size of *k-fault tolerant* $(2t - 1)$ -spanners on n vertices is $\mathcal{O}(k^{1-1/t} n^{1+1/t})$ [BDPW18]. 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.

*A preliminary version of this paper appeared in SoCG 2021 [HMO21].

[†]Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@illinois.edu; <http://sarielhp.org/>. Work on this paper was partially supported by a NSF AF award CCF-1907400.

[‡]Department of Mathematics and Computer Science, The Open University of Israel, manorme@openu.ac.il. Supported by BSF Grant no. 2018223.

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

Reliable spanners. An alternative notion was introduced by Bose *et al.* [BDMS13]. Here, for a parameter $\vartheta > 0$, a ϑ -reliable t -spanner has the property that for any failure (or attack) set B , the residual graph $G \setminus B$ is a t -spanner path for the points of $V \setminus B^+$, where $B^+ \supseteq B$ is some set, such that $|B^+| \leq (1 + \vartheta)|B|$. We consider two variants:

- (A) **Adaptive adversary** (i.e., standard or “deterministic” model): The adversary knows the spanner G , and the set B is chosen as a worst case for G .
- (B) **Oblivious adversary** (i.e., “randomized” model): Here, 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 $\mathbb{E}[|B^+|] \leq (1 + \vartheta)|B|$.

See [Section 2](#) for precise definitions.

Previous work. Bose *et al.* [BDMS13] provided some lower bounds and constructions in the general settings, but the bounds on the size of the damaged set (i.e., B^+) are much larger. In the Euclidean settings, 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 $\mathcal{O}(n \log n \log \log^6 n)$ edges [BHO19] (an alternative, but slightly inferior construction, was provided independently by Bose *et al.* [BCDM18]). The number of edges can be further reduced in the oblivious adversary case, where one can construct an oblivious $(1 + \varepsilon)$ -spanner that is ϑ -reliable in expectation and has $\mathcal{O}(n \log \log^2 n \log \log \log n)$ edges [BHO20].

Covers. A cover is a set of clusters (i.e., subsets of the point set) that covers the metric space with certain desirable properties. Awerbuch and Peleg [AP90] showed a cover, where (i) each cluster has diameter $O(k)$, (ii) every vertex participates in $O(kn^{1/k})$ clusters, and (iii) for every vertex v , the ball of radius k centered at v , is contained in a single cluster. Busch *et al.* [BLT14] show how to compute a cover of a planar graph, with diameter $\leq 16k$ per cluster, such that each pair in distance k from each other belongs to some cluster, and every vertex participates in at most 18 clusters. For graphs that excludes a minor of fixed size, they get a similar result, except that each vertex might participate in $O(\log n)$ vertices, where n is the number of vertices in the input graph. Abraham *et al.* [AGMW10] presented a result with better sparsity when the graphs do not have $K_{r,r}$ as a minor.

Our results.

We provide new constructions of reliable spanners for finite uniform metrics, ultrametrics, trees, planar graphs and finite metrics. 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 – that is, star graphs. In the oblivious model one can simply use “constellation of stars” with a constant number of random centers, and the resulting 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, and 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 – these are expanders that remain expanding under the type of attacks described above. See [Section 3](#) and [Section 6](#) for details.

Constructions of ϑ -reliable Δ -spanners				
metric	Δ	guarantee	size	ref
Uniform	2	expectation	$\mathcal{O}(n\vartheta^{-1} \log \vartheta^{-1})$	Lemma 3.1
	t	det. lower bound	$\Omega(n^{1+1/t})$	Lemma 3.2
	$2t - 1$	deterministic	$\mathcal{O}(\vartheta^{-2} n^{1+1/t})$	Theorem 3.6
	$2t$	deterministic	$\mathcal{O}(\vartheta^{-1} n^{1+1/t})$	Theorem 3.6
Finite metrics	$\mathcal{O}(\log n)$	expectation	$\mathcal{O}(\vartheta^{-1} n \text{ polylog})$	Lemma 5.3
	$\mathcal{O}(t)$	expectation	$\mathcal{O}(\vartheta^{-1} n^{1+1/t} \text{ polylog})$	Lemma 5.3
	$\mathcal{O}(t \log n)$	deterministic	$\mathcal{O}(\vartheta^{-1} n^{1+1/t} \text{ polylog})$	Lemma 5.4
	$\mathcal{O}(t^2)$	deterministic	$\mathcal{O}(\vartheta^{-1} n^{1+1/t} \text{ polylog})$	Lemma 5.4
Ultrametrics	$2 + \varepsilon$	expectation	$\mathcal{O}(\vartheta^{-1} \varepsilon^{-2} n \text{ polylog})$	Lemma 5.5
	$(2 + \varepsilon)t - 1$	deterministic	$\mathcal{O}(\vartheta^{-2} \varepsilon^{-3} n^{1+1/t} \text{ polylog})$	Lemma 5.6
Trees	$3 + \varepsilon$	expectation	$\mathcal{O}(\vartheta^{-1} \varepsilon^{-2} n \text{ polylog})$	Lemma 5.7
	$(4 + \varepsilon)t - 3$	deterministic	$\mathcal{O}(\vartheta^{-2} \varepsilon^{-3} n^{1+1/t} \text{ polylog})$	Lemma 5.8
Planar graphs	$3 + \varepsilon$	expectation	$\mathcal{O}(\vartheta^{-1} \varepsilon^{-4} n \text{ polylog})$	Lemma 5.9
	$(4 + \varepsilon)t - 3$	deterministic	$\mathcal{O}(\vartheta^{-2} \varepsilon^{-6} n^{1+1/t} \text{ polylog})$	Lemma 5.10

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.

Covers. A [t-cover](#) of a finite metric space $\mathcal{M} = (P, d)$ 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 $t \cdot d(p, q)$. 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 the problem of constructing good covers. For general n -point spaces with spread at most Φ , we observe that the Ramsey partitions of [\[MN07\]](#) 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 we provide. In more specific cases, like ultrametrics, trees and planar graphs, one can do better. For example, 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 finite uniform, ultrametric, tree, planar, and general 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 $\geq 1 - 1/m^{O(1)}$. As such, the constructions described can

be done efficiently, if one wants constructions of spanners for n vertices that succeeds with probability $1 - 1/n^{O(1)}$. See [Remark 6.3](#) for details.

2. Preliminaries

2.1. Metric spaces

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

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

For a finite set $X \subseteq \mathcal{X}$, the [diameter](#) of X is

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

and the [spread](#) of X is $\Phi(X) = \frac{\text{diam}(X)}{\min_{p, q \in X, p \neq q} d(p, q)}$. A metric space $\mathcal{M} = (\mathcal{X}, d)$ 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 $(V(G), d_G)$. 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 2.1. For a metric space $\mathcal{M} = (P, d)$, a graph $H = (P, E)$ is a [t-spanner](#), if for any $p, q \in P$, $d(p, q) \leq d_H(p, q) \leq t \cdot d(p, q)$. Here d_H is the shortest path distance on H whose edges are weighted according to d .

Given a weighted graph $G = (V, E)$, and a set $B \subseteq V$, we denote by $G|_B$ the subgraph induced on B . We also use the notation $G \setminus B = G|_{V \setminus B}$. A randomized graph G is a probability distribution over the edge set E for a given set of vertices V .

An [attack](#) on a graph $G = (V, E)$ is a set of vertices B that fails and no longer can be used. An attack (on a randomized graph) is [oblivious](#), if the set B is picked stochastically independent of the edge set of the graph.

Definition 2.2 (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 set of smallest possible size, such that for any pair of vertices $p, q \in P \setminus B^+$, we have

$$d_{G \setminus B}(p, q) \leq t \cdot d(p, q), \tag{2.1}$$

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 [\$\vartheta\$ -reliable](#) (in the deterministic or non-oblivious setting), if $\lambda(G, B) \leq \vartheta$ holds for any attack $B \subseteq P$. Furthermore, the graph G is [\$\vartheta\$ -reliable in expectation](#) (or in the oblivious model), if $\mathbb{E}[\lambda(G, B)] \leq \vartheta$ holds for any oblivious attack $B \subseteq P$.

Remark 2.3. 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 $G \setminus B$.

2.3. Miscellaneous

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

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

denote the **neighbors** of Y in Z . The neighbors of Y in G is denoted by $\Gamma(Y) = \Gamma_{V(G)}(Y)$.

Definition 2.4. 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 $d(p, q)$ is the same quantity (e.g., 1). 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 $* = \bigcup_{p \in C} *_p$, which is the union of the star graphs induced by centers in C .

Lemma 3.1. *The constellation $*$, defined above, is a ϑ -reliable 2-spanner in expectation. The number of its edges is $\mathcal{O}(n\vartheta^{-1} \log \vartheta^{-1})$.*

Proof: Let $B \subseteq P$ be an oblivious attack, and let $b = |B|$. If there is a point of C that is not in B , then there is a center point outside of the attack set, which provides 2-hop paths between all pairs of points in the residual graph, and therefore we choose $B^+ = B$. On the other hand, if $C \subseteq B$, then the residual graph contains only isolated vertices, and therefore we choose $B^+ = P$. If $(1 + \vartheta)b \geq n$ then there is nothing to prove. Thus, since $b/n < 1/(1 + \vartheta)$ and $1/(1 + \vartheta) \leq 1 - \vartheta/2$, we have

$$\mathbb{E}[\lambda(G, B)] = 0 \mathbb{P}[C \not\subseteq B] + \frac{n - 1 - b}{b} \mathbb{P}[C \subseteq B] \leq \frac{n}{b} \left(\frac{b}{n}\right)^k \leq \frac{1}{(1 + \vartheta)^{k-1}} \leq \exp\left(-\frac{k-1}{2}\vartheta\right) \leq \vartheta.$$

As for the number of edges, $*$ has at most $k(n - 1)$ edges, since $*$ is the union of k stars and each star has $n - 1$ edges. Thus, by the choice of k , the size of $*$ 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 3.2. *Let $G = (P, E)$ be a ϑ -reliable t -spanner on P for the uniform metric, where $\vartheta \in (0, 1)$ and $t \geq 1$. Then, in the non-oblivious settings, G must have $\Omega(n^{1+1/t})$ edges.*

Proof: We assume that the distance between any pair of points of P is one. Let the attack set B be the set of all nodes of degree at least Δ , where $\Delta = n^{1/t}/4$. Assume, toward a contradiction, that $|B| \leq n/4$. By the reliability condition, there exists a set $Q \subseteq P \setminus B$ of nodes of the residual graph of size at least $n - (1 + \vartheta)|B| \geq n/2$, that has t -hop paths in $G \setminus B$ between all pairs of vertices of Q . Let $p \in Q$ be an arbitrary vertex. Let $\mathbf{b}(p, t)$ be a ball of radius t centered at p in the shortest path metric in the residual graph $G \setminus B$. On the one hand, from the above, $\mathbf{b}(p, t) \supseteq Q$. On the other hand, the maximum degree is at most Δ , and therefore $|\mathbf{b}(p, t)| \leq \Delta^t \leq n/4$. As such, we have $n/4 \geq |\mathbf{b}(p, t)| \geq |Q| \geq n/2$. A contradiction.

Hence, $|B| > n/4$. The claim follows since $\Delta|B| \leq 2|E|$, which implies $|E| \geq \Delta|B|/2 = \Omega(n^{1+1/t})$. ■

Remark 3.3. 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 3.2](#) 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 3.4. Denote by $\lambda(G)$ the second eigenvalue of the matrix M/d , where $M = \text{Adj}(G)$ is the adjacency matrix of a d -regular graph G . A *proper expander* specifies a constant $C > 1$, and functions $\mathcal{C}_\delta, \mathbf{c}_\delta > 0$, such that for every $\delta \in (0, 1/4)$ and even integers $d \geq \mathcal{C}_\delta$, $n \geq d^2$, there exists an n -vertex, d -regular graph $G = (V, E)$, such that:

- (P1) $\forall S \subseteq V, |S| \geq 12n/(\delta d) \implies |\Gamma(S)| > (1 - \delta)n$,
- (P2) $\forall S \subseteq V, |S| \leq \mathbf{c}_\delta n/d \implies |\Gamma(S)| \geq (1 - \delta)d|S|$.
- (P3) $\lambda(G) \leq C/\sqrt{d}$.

For each one of the properties above, it is known that there exists an expander satisfying it: Property (P1) is essentially proved in [\[BHO19\]](#), Property (P2) is folklore, and Property (P3) appears in [\[DJPP13\]](#). 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, [Appendix A](#) reprove (P1) and (P2) in the same random model in which (P3) was proved. We thus get the following:

Theorem 3.5. *The random graph constructed in [Section A.1](#) is a proper expander (see [Definition 3.4](#)), asymptotically almost surely. Specifically, the probability the graph has the desired properties is $\geq 1 - n^{-O(1)}$.*

With the appropriate choice of parameters, these expanders are reliable spanners for uniform metrics.

Theorem 3.6. For every $t \in \mathbb{N}$, $\theta \in (0, 1)$, and $n \in 2\mathbb{N}$, such that $n \geq e^{\Omega(t)}$, 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.

The proof is somewhat cumbersome and is deferred to [Section 6](#).

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

Corollary 3.7. Let $\mathcal{M} = (\mathcal{X}, d)$ 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 $|E(G)| = \mathcal{O}(\vartheta^{-2} \cdot n^{1+1/t})$ 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 $|Q| \geq |P| - (1 + \vartheta)|B \cap P|$. Furthermore, for any two points $p, q \in Q$, we have

$$d_{\mathcal{M}}(p, q) \leq d_{G|_Q}(p, q) \leq (2t - 1) \cdot \text{diam}_{\mathcal{M}}(P),$$

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 $(2t - 1) \cdot \text{diam}_{\mathcal{M}}(P)$.

Corollary 3.8. Let $\mathcal{M} = (\mathcal{X}, d)$ 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 $|E(G)| = \mathcal{O}(\vartheta^{-1} \cdot n^{1+1/t})$ 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 $|Q| \geq |P| - (1 + \vartheta)|B \cap P|$. Furthermore, for any two points $p, q \in Q$, we have

$$d_{\mathcal{M}}(p, q) \leq d_{G|_Q}(p, q) \leq 2t \cdot \text{diam}_{\mathcal{M}}(P),$$

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

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

Remark 3.9. [Corollary 3.7](#) and [Corollary 3.8](#) are quite weak as far as the guarantee on the length of the resulting path (i.e., they are not good spanners). A construction that provides a spanner guarantee is provided below in [Lemma 5.4](#) below.

As aside, proper expanders are reliable spanners because their expansion property is robust, as testified by the following.

Theorem 3.10 (reliable vertex expander). For every $\vartheta \in (0, 1)$ there exists a constant $\mathfrak{c} = \mathfrak{c}_{\vartheta} > 0$, such that for any expansion constant $h \geq \mathfrak{c}^{-2}$, there exists a family of vertex expander graphs $\{G = (V, E)\}_n$ on n vertices of degree at most $4h/\vartheta$, with the following resiliency property: For any $B \subset V$, there exists $B^+ \supseteq B$, $|B^+| \leq (1 + \vartheta)|B|$ such that the graph $G \setminus B^+$ is a vertex expander in the sense that

- (i) $\text{diam}(G \setminus B^+) \leq 2\lceil \log_h n \rceil$, and
- (ii) For any $U \subset V \setminus B^+$ of size $|U| \leq \mathfrak{c}n/h$, we have $|\Gamma_{G \setminus B^+}(U)| \geq h|U|$.

The proof is deferred to [Section 6.3.4](#).

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

Definition 4.1. For a finite metric space $\mathcal{M} = (P, d)$, a *t -cover*, is a family of subsets $\mathcal{C} = \{S_i \subseteq P \mid i = 1, \dots, m\}$, such that for any $p, q \in P$, there exists an index i , such that $p, q \in S_i$, and

$$\text{diam}(S_i)/t \leq d(p, q) \leq \text{diam}(S_i).$$

The *size* of a cover $\mathcal{C}$ is $\text{size}(\mathcal{C}) = \sum_{S \in \mathcal{C}} |S|$. Recalling **Definition 2.4**, the *degree* in $\mathcal{C}$ of a point $p \in P$ is the number of sets of $\mathcal{C}$ that contain it. The *depth* of $\mathcal{C}$ is the maximum degree of any element of P , and is denoted by $D(\mathcal{C})$.

4.1. Lower bounds

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

Proposition 4.2. *For any parameter $t > 1$, any integer $h > 1$, $\Phi = t^h$, and any $n \geq h$, there exists a metric $\mathcal{M} = (P, d)$ 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 $\geq h/2$, and depth h .

Proof: For simplicity of exposition, assume that h divides n . Let P_i be a set of n/h points, such that the distance between any pair of points of P_i is $\ell_i = (t + \varepsilon)^i$, for some fixed small $\varepsilon > 0$, for $i = 1, \dots, h$. Furthermore, assume that the distance between any point of P_i and any point of P_j , for $i < j$, is ℓ_j .

Let $P = \cup P_i$, and observe that the distance function defined above is a metric (it is the union of uniform metrics of different resolutions). Now, consider any t -cover $\mathcal{C}$ of P . The *rank* of a cluster $C \in \mathcal{C}$, is the highest j , such that $P_j \cap C \neq \emptyset$. We can assume that all the clusters of $\mathcal{C}$ have at least two points, as otherwise they can be removed. For any index $j \in [h]$, any point $p \in P_j$ and any point $q \in \cup_{i=1}^j P_i - p$, by definition, there exists a cluster $C \in \mathcal{C}$, such that $p, q \in C$, and $\text{diam}(C) \leq t \cdot \ell_j$, since $d(p, q) = \ell_j$. It follows that C can not contain any point of $\cup_{i=j+1}^h P_i$. Namely, the rank of C is j .

If there are two clusters of rank j in $\mathcal{C}$, then we can merge them, since merging does not increase their diameter, and such an operation does not increase the size of the cover, and the degrees of its elements. As such, in the end of this process, the cover $\mathcal{C}$ has h clusters, and for any $j \in [h]$, there is a cluster $C_j \in \mathcal{C}$ that is of rank j , and contains (exactly) all the elements of $\cup_{i=1}^j P_i$. Namely, $\mathcal{C} = \{C_1, \dots, C_h\}$, where $C_j = \cup_{i=1}^j P_i$. It is easy to verify that this cover has the desired properties. $\blacksquare$

Proposition 4.3. *For any $t \in 2, 3, \dots$, and any sufficiently large n there exists an n -point metric space for which any t -cover must be of size at least $\Omega(n^{1+1/2t})$.*

Proof: Let $g = 2t + 2$. By a standard probabilistic argument, for any sufficiently large $n \in \mathbb{N}$ there exists a simple graph $G = (V, E)$ on n vertices and $m = \Omega(n^{1+\frac{1}{g-2}})$ edges whose girth is at least g . (This is not the best known bound, but it is sufficient for our purposes.) Consider G as a metric space with the shortest (unweighted) path metric. Let $\mathcal{C}$ be a t -cover for G , and let $\mathcal{C}' = \{S \in \mathcal{C} \mid \text{diam}(S) \leq t\}$. For $S \subset V$ let $E(S) = \{uv \in E \mid u, v \in S\}$.

We claim that for every $S \in \mathcal{C}'$, $E(S)$ is a forest, and hence $|E(S)| < |S|$. Indeed, Suppose $E(S)$ contains a cycle $(v_0, v_1, \dots, v_h, v_0)$ such that $v_i v_{i+1}, v_0 v_h \in E(S) \subseteq E$. Denote $d_i = d(v_0, v_i)$. Since $\{v_i, v_{i+1}\} \in E$, we have $d_{i+1} - d_i \in \{-1, 0, 1\}$. Let j be the smallest i such that $d_{i+1} \leq d_i$. Hence, $d_j = j$. Let $P = (v_{j+1}, u_1, \dots, u_k, v_0)$ be a shortest path in G between v_{j+1} and v_0 . P 's length is at most $d_j = j$.

Thus, the sequence $v_0, v_1, \dots, v_j, v_{j+1}, u_1, \dots, u_k, v_0$ is closed path of length at most $2j + 1$, and hence contains a cycle of length at most $2j + 1$. By the girth condition, $2j + 1 \geq 2t + 2$, and hence $d_j = j > t$. But since $v_0, v_j \in S$, this means that $\text{diam}(S) \geq d_j > t$ which contradicts the definition of $\mathcal{C}'$.

For every edge $pq \in E$, $d(p, q) = 1$, and by the t -covering condition there exists $S \in \mathcal{C}'$ such that $p, q \in S$. Therefore $pq \in E(S)$. Hence $E \subset \bigcup_{S \in \mathcal{C}'} E(S)$, and therefore

$$cn^{1+\frac{1}{2t}} \leq |E| \leq \sum_{S \in \mathcal{C}'} |E(S)| < \sum_{S \in \mathcal{C}'} |S| = \text{size}(\mathcal{C}') \leq \text{size}(\mathcal{C}). \quad \blacksquare$$

4.2. Cover for ultrametrics

Definition 4.4. 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 an associated label $\Delta_u \geq 0$. This label is zero for all the leaves of T , and it is a positive number for all the interior nodes. The labels satisfy for every non-root vertex $v \in T$, $\Delta_v \leq \Delta_{\bar{p}(v)}/k$, where $\bar{p}(v)$ is the parent of v in T . The distance between two leaves $x, y \in T$ is defined as $\Delta_{\text{lca}(x, y)}$, where $\text{lca}(x, y)$ is the least common ancestor of x and y in T . An HST T is a *k-HST* if for every non-root vertex $v \in T$, $\Delta_v \leq \Delta_{\bar{p}(v)}/k$.

HSTs are one of the simplest non-trivial metrics, and surprisingly, general metrics can be embedded randomly into HSTs with expected distortion of $\mathcal{O}(\log n)$ [Bar98, FRT04].

Definition 4.5. A metric $\mathcal{M} = (P, d)$ is an *ultrametric*, if for any $x, y, z \in P$, we have that $d(x, z) \leq \max(d(x, y), d(y, z))$. Notice, that this is a stronger version of the triangle inequality, which states that $d(x, z) \leq d(x, y) + d(y, z)$.

The following is folklore, and it is also easy to verify (see, e.g., [BLMN05, Lemma 3.5]).

Lemma 4.6. For $k \geq 1$, every finite ultrametric can be k -approximated by a k -HST.

Lemma 4.7. For $k > 1$, every k -HST with spread Φ has 1-cover of depth at most $\log_k \Phi$.

Proof: Let T be the tree of the HST. With every vertex $u \in T$ we associate a cluster C_u the leaves of the subtree rooted at u . The covers is defined as $\mathcal{C} = \{C_u \mid u \in T\}$. The properties of the cover are immediate. $\blacksquare$

Corollary 4.8. Let $\mathcal{M} = (P, d)$ 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 4.9 ([MN07]). Let (P, d) be an n -point metric space and $k \geq 1$. Then there exists a distribution over decreasing sequences of subsets $P = P_0 \supsetneq P_1 \supsetneq \dots \supsetneq P_s = \emptyset$ (s itself is a random variable), such that for all $\mu > -1/k$, we have $\mathbb{E} \left[\sum_{j=1}^{s-1} |P_j|^\mu \right] \leq \max \left(\frac{k}{1+\mu k}, 1 \right) \cdot n^{\mu+1/k}$, and such that for each $j \in [s]$ there exists an ultrametric ρ_j on P_{j-1} satisfying for every $p, q \in P$, that $\rho_j(p, q) \geq d(p, q)$, and if $p \in P_{j-1}$ and $q \in P_{j-1} \setminus P_j$ then $\rho_j(p, q) \leq \mathcal{O}(k) \cdot d(p, q)$.

By computing a cover (using [Corollary 4.8](#)) for each ultrametric generated by the above lemma, we get the following.

Lemma 4.10. *For an n -point metric space $\mathcal{M} = (P, d)$ with spread Φ , and a parameter $k > 1$, one can compute, in polynomial time, an $\mathcal{O}(k)$ -cover of $\mathcal{M}$ of size $\mathcal{O}(n^{1+1/k} \log \Phi)$ and depth $\mathcal{O}(kn^{1/k} \log \Phi)$.*

Proof: Using [Lemma 4.9](#), for the parameter k , compute the sequence $P = P_0 \supsetneq P_1 \supsetneq \dots \supsetneq P_s = \emptyset$, and the associated ultrametrics $\rho_1, \dots, \rho_{s-1}$. We build a 1-HST T_i for ρ_i , when restricted to the set P_{i-1} , for $i = 1, \dots, s-1$. For every HST in this collection, we compute a 2-cover using [Corollary 4.8](#). Let $\mathcal{C}$ be the union of all these covers.

Since for every HST T_i the resulting cover has size $\mathcal{O}(|P_i| \log \Phi)$, then by [Lemma 4.9](#) (applied with $\mu = 1$), we have

$$\mathbb{E} \left[\sum_i |P_i| \log \Phi \right] = \mathcal{O}(n^{1+1/k} \log \Phi).$$

As for the quality of the cover, let p, q be any two points in P , and assume (without loss of generality) that $p \in P_{i-1}$ and $q \in P_{i-1} \setminus P_i$ for some $i \in [s]$. We have that $\rho_i(p, q) = \mathcal{O}(k) \cdot d(p, q)$, and since we computed a 2-cover for this tree, there is a cluster in the computed cover that contains both points and its diameter is at most twice the distance between those points.

The maximum depth, follows by using [Lemma 4.9](#) with $\mu = 0$. This implies that a point of P participates in $s = \mathcal{O}(kn^{1/k})$ HSTs, and each cover generated by such an HST might a point an element at most $\mathcal{O}(\log \Phi)$ times.

The bounds on the size and depth hold in expectation, and one can repeat the construction if they exceed the desired size by (say) a factor of eight. In expectation, after a constant number of iterations, the algorithm would compute with high probability a cover with the desired bounds. $\blacksquare$

4.4. Covers for trees

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

Lemma 4.11. *For a weighted tree metric $T = (P, d)$, 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 $n = |P|$.*

Proof: The proof is by induction on n . When $n = 1$ the trivial cover is sufficient. Assume next that $n \geq 1$ and the minimum distance in T is one. Find a separator node $\mathbf{s}$ in T such that there is no connected component in T larger than $n/2$ after removing $\mathbf{s}$. For $i \in \{0, 1, 2, \dots, m = \lceil \log_{1+\varepsilon/2} \Phi \rceil\}$, define the sets $P(\mathbf{s}, i) = \{p \in P \mid d(\mathbf{s}, p) \leq (1 + \varepsilon/2)^i\}$. By the inductive hypothesis, for each connected component Q of $T - \mathbf{s}$ there is a $(2 + \varepsilon)$ -cover $\mathcal{C}_Q$ of Q of depth $\mathcal{O}(\varepsilon^{-1} \log \Phi \log(n/2))$ and size $\mathcal{O}(n\varepsilon^{-1} \log \Phi \log(n/2))$. The cover for P is

$$\mathcal{C}_P = \{P(\mathbf{s}, i) \mid i = 0, 1, 2, \dots, m\} \cup \bigcup_Q \mathcal{C}_Q.$$

Since every element of P participates in at most m sets in each level of the recursion, the bound on the depth and size is immediate.

As for the quality of the cover, by the inductive hypothesis, we need only to check pairs of points $p, q \in P$, that are separated by $\mathbf{s}$. Assume $d(\mathbf{s}, p) \geq d(\mathbf{s}, q)$, and let j be the minimum index, such that $d(\mathbf{s}, p) \leq (1 + \varepsilon/2)^j$. We have that $p, q \in P(\mathbf{s}, j)$, and

$$d(p, q) \geq d(\mathbf{s}, p) \geq (1 + \varepsilon/2)^{j-1} \quad \text{and} \quad \text{diam}(P(\mathbf{s}, j)) \leq 2(1 + \varepsilon/2)^j.$$

As such, we have $\frac{\text{diam}(P(\mathbf{s}, j))}{d(p, q)} \leq \frac{2(1 + \varepsilon/2)^j}{(1 + \varepsilon/2)^{j-1}} = 2 + \varepsilon$. $\blacksquare$

4.5. Covers for planar graphs

The next lemma can be traced back to the work of Lipton and Tarjan [LT79].

Lemma 4.12. *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) $|X| \leq (2/3)n$ and $|Y| \leq (2/3)n$,
- (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 4.13. For a metric space $(\mathcal{X}, d)$ and a parameter r , an r -net N is a maximal set of points in $\mathcal{X}$ satisfying:

- (i) For any two net points $p, q \in N$, $p \neq q$, we have $d(p, q) > r$.
- (ii) For any $p \in \mathcal{X}$, $d(p, N) = \min_{q \in N} d(p, q) \leq r$.

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 $\text{net}(\mathcal{X}, r)$.

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 4.14. *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 $N = \text{net}(\pi, r) \subseteq \pi$ be a net computed for some distance $r > 0$. For some $R > 0$, consider the set of balls $\mathcal{B} = \{b(p, R) \mid p \in N\}$. For any point $q \in V(G)$, we have that the degree of q in $\mathcal{B}$ is at most $2R/r + 1$.*

Proof: Let $p_1, \dots, p_k$ be the points of $N \cap b(q, R)$ sorted by their order along π – these are the only points that their balls in $\mathcal{B}$ would contain q . By the definition of the net, $d(p_i, p_{i+1}) > r$ for all i . Since π is a shortest path, we also have that

$$2R \geq \text{diam}(b(q, R)) \geq d(p_1, p_k) = \sum_{i=1}^{k-1} d(p_i, p_{i+1}) > (k-1)r. \quad \blacksquare$$

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

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

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 4.15. *For any two vertices $p, q \in V(G)$, there exists a cluster $C \in \mathcal{C}$, such that $p, q \in C$, and $\text{diam}(C) \leq (2 + \varepsilon)d_G(p, q)$. That is, $\mathcal{C}$ is a $(2 + \varepsilon)$ -cover of G .*

Proof: Assume, for the simplicity of exposition, that p and q get separated in the top level of the recursion (otherwise, apply the argument to the inductive step in which they get separated). The shortest path between p and q must intersect the separator Z , say at a vertex v . Assume that $d_G(p, v) \geq d_G(v, q)$ and that $r_{j-1} \leq d_G(p, v) \leq r_j = (1 + \varepsilon/8)^j$. There is a point $u \in N_j$ within distance $\varepsilon r_j/8$ from v . As such,

$$d_G(p, u) \leq d_G(p, v) + d_G(v, u) \leq (1 + \varepsilon/8)^j + (\varepsilon/8)(1 + \varepsilon/8)^j \leq (1 + \varepsilon/8)^{j+1},$$

which implies that $p \in \mathbf{b} = \mathbf{b}(u, (1 + \varepsilon/8)r_j)$. A similar argument shows that q is also in $\mathbf{b}$. Furthermore, we have that $d(p, q) \geq d(p, v) \geq r_{j-1}$. Note that $\mathbf{b} \in \mathcal{B}_j \subseteq \mathcal{C}$. We have that

$$\frac{\text{diam}(\mathbf{b})}{d(p, q)} \leq \frac{2(1 + \varepsilon/8)r_j}{r_{j-1}} = 2(1 + \varepsilon/8)^2 \leq 2 + \varepsilon. \quad \blacksquare$$

Lemma 4.16. *The depth of $\mathcal{C}$ is $\mathcal{O}(\varepsilon^{-2} \log n \log \Phi)$.*

Proof: Fix a vertex p . By [Lemma 4.14](#), for each i , at most

$$3 + 2 \left(\frac{2(1 + \varepsilon/8)r_i}{\varepsilon r_i/8} + 1 \right) = \mathcal{O}(1/\varepsilon).$$

balls of $\mathcal{B}_i$ contains p . The number of such sets is $\mathcal{O}(\log_{1+\varepsilon/8} \Phi) = \mathcal{O}(\varepsilon^{-1} \log \Phi)$. The vertex p get sent down to at most one recursive subproblem, and the recursion depth is $\mathcal{O}(\log n)$. It follows that the depth of any point (and the degree of p specifically) is at most $\mathcal{O}(\varepsilon^{-2} \log \Phi \log n)$. $\blacksquare$

Theorem 4.17. *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 4.18. It is possible to generalize [Theorem 4.17](#) to the shortest path metric on families of graphs excluding a fixed minor. Specifically, by [\[KLMN05, Lemma 3.3\]](#), there exists $\mathcal{O}(s^2)$ -cover of depth $\mathcal{O}(3^s \log \Phi)$ for every metric space supported on a graph excluding K_s minor and spread Φ . It may be possible to improve the approximation parameter to $\mathcal{O}(s)$ using [\[AGG+19\]](#). This approach does not have a $\log n$ factor in the depth parameter, but it can not provide a $(2 + \varepsilon)$ -approximation as in [Theorem 4.17](#). As communicated to us by an anonymous referee, the approach used here to prove [Theorem 4.17](#) can also be extended to any family of graphs excluding fixed minor and obtain $(2 + \varepsilon)$ -cover using the shortest paths separators of [\[AG06\]](#). We have not pursued those avenues.

5. From covers to reliable spanners

5.1. The oblivious construction

Lemma 5.1. *Let $\mathcal{M} = (P, d)$ be a finite metric space, and suppose there exists a ξ -cover $\mathcal{C}$ 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 $\mathcal{O}(s \frac{D}{\vartheta} \log \frac{D}{\vartheta})$.*

Proof: For each cluster $C \in \mathcal{C}$, let $*_C$ be a random constellation graph on C as defined in [Section 3.1](#), with reliability parameter $\psi = \vartheta/D$. The resulting spanner is the union $\bigcup_{C \in \mathcal{C}} *_C$. The number of edges in the resulting graph is at most

$$\sum_{C \in \mathcal{C}} \mathcal{O}(|C| \psi^{-1} \log \psi^{-1}) = \mathcal{O}(s \psi^{-1} \log \psi^{-1}) = \mathcal{O}\left(s \frac{D}{\vartheta} \log \frac{D}{\vartheta}\right).$$

Fix an attack set $B \subset P$. A cluster $C \in \mathcal{C}$ is *failed* if $|C| \leq (1 + \psi) |C \cap B|$. Denote the set failed clusters by $\mathcal{F}$. For $C \in \mathcal{C} \setminus \mathcal{F}$, let $\text{dmg}(C, B)$ be the (random) set of damaged points in C as defined in the proof of [Lemma 3.1](#), i.e., $\text{dmg}(C, B) = C$ if B contains all the constellation's centers, and $\text{dmg}(C, B) = B$ otherwise. by [Lemma 3.1](#), $\mathbb{E}[|\text{dmg}(C, B)|] \leq (1 + \psi) |B \cap C|$. The damaged set is defined as

$$B^+ = \left(\bigcup_{C \in \mathcal{F}} C \right) \cup \left(\bigcup_{C \in \mathcal{C} \setminus \mathcal{F}} \text{dmg}(C, B) \right).$$

We next bound expected size of the *loss* $B^+ \setminus B$:

$$\begin{aligned} \mathbb{E}[|B^+ \setminus B|] &\leq \mathbb{E}\left[\sum_{C \in \mathcal{F}} |C \setminus B|\right] + \sum_{C \in \mathcal{C} \setminus \mathcal{F}} \mathbb{E}[|\text{dmg}(C, B) \setminus B|] \\ &\leq \psi \sum_{C \in \mathcal{F}} |B \cap C| + \psi \sum_{C \in \mathcal{C} \setminus \mathcal{F}} |B \cap C| = \psi \sum_{p \in B} \sum_{C \in \mathcal{C}} 1_{p \in C} \leq \psi D |B| = \vartheta |B|. \end{aligned}$$

Finally, for any two points $p, q \in P \setminus B^+$, there exists a non-failed cluster $C \in \mathcal{C}$ that contains both points, such that $\text{diam}(C) \leq \xi d(p, q)$. As such, the two hops in the resulting graph are going to be of length at most $2\text{diam}(C) \leq 2\xi d(p, q)$. $\blacksquare$

5.2. The deterministic construction

Lemma 5.2. *Let $\mathcal{M} = (P, d)$ 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 $t \geq 1$, there exists:*

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

Proof: (A) For a cluster $C \in \mathcal{C}$, let $G(C)$ be a ψ -reliable spanner constructed using [Corollary 3.7](#) on C , with the reliability parameter $\psi = \vartheta/D$. Let $G = (P, E)$ be a graph whose edge set E is the union of the edge sets of $G(C)$ for $C \in \mathcal{C}$.

Let n_i be the size of the i th cluster, for $i \in \{1, \dots, m = |\mathcal{C}|\}$. The number of edges in G is bounded by $\mathcal{O}(\psi^{-2} N) = \mathcal{O}(D^2 \vartheta^{-2} N)$, where

$$N = \sum_{i=1}^m n_i^{1+1/t} = \sum_{i=1}^m n_i^{1/t} n_i \leq (\max_i n_i^{1/t}) \sum_{i=1}^m n_i \leq n^{1/t} s,$$

since $\sum_i n_i = s$, and $\max_i n_i \leq n$. We conclude that the total number of edges in G is $\mathcal{O}(\vartheta^{-2} D^2 s n^{1/t})$.

Let $B \subset P$ be an attack set. The damage set $B^+ \supseteq B$ is constructed as in the proof of [Lemma 5.1](#). Thus, as argued there, $|B^+| \leq (1 + \vartheta) |B|$. For every two points $p, q \in P \setminus B^+$, there exists a cluster $C \in \mathcal{C}$, such that $p, q \in C$, and

$$\text{diam}(C)/\xi \leq d(p, q) \leq \text{diam}(C).$$

Furthermore, by [Corollary 3.7](#), we have

$$d_{G \setminus B}(p, q) \leq d_{G(C) \setminus B}(p, q) \leq (2t - 1) \text{diam}(C) \leq (2t - 1) \xi \cdot d(p, q).$$

and shortest path in $G(C) \setminus B$ realizing $d_{G(C) \setminus B}(p, q)$ has at most $2t - 1$ hops.

(B) Similar to the above, except for using [Corollary 3.8](#) instead of [Corollary 3.7](#). ■

5.3. Applications

5.3.1. General metrics

Lemma 5.3. *Let $\mathcal{M} = (P, d)$ 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*

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

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

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

edges.

Proof: By [Lemma 4.10](#), $\mathcal{M}$ has a $\xi = \mathcal{O}(k)$ -cover of size $\mathcal{O}(n^{1+1/2k} \log \Phi)$ and depth $D = \mathcal{O}(k n^{1/2k} \log \Phi)$. Plugging this into [Lemma 5.1](#) yields an oblivious ϑ -reliable $\mathcal{O}(k)$ -spanner with $\mathcal{O}(\vartheta^{-1} k n^{1+1/k} \log^2 \Phi \log \frac{k n^{1/k} \log \Phi}{\vartheta})$ edges. ■

Lemma 5.4. *Let $\mathcal{M} = (P, d)$ be a finite metric over n points of spread Φ , and let $\vartheta \in (0, 1)$ and $k, t \in \mathbb{N}$ be parameters. Then, one can build a ϑ -reliable $\mathcal{O}(kt)$ -spanner for $\mathcal{M}$ with $\mathcal{O}(\vartheta^{-1} k n^{1+1/k+1/t} \log^2 \Phi)$ edges. In particular, when $t = \log n$, we obtain a ϑ -reliable $\mathcal{O}(k \log n)$ -spanner for $\mathcal{M}$ with*

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

edges, and when $t = k$ we obtain ϑ -reliable $\mathcal{O}(t^2)$ -spanner for $\mathcal{M}$ with $\mathcal{O}(\vartheta^{-1} t n^{1+1/t} \log^2 \Phi)$ edges.

Proof: By [Lemma 4.10](#), $\mathcal{M}$ has a $\mathcal{O}(k)$ -cover of size $s = \mathcal{O}(n^{1+1/2k} \log \Phi)$ and depth $D = \mathcal{O}(k n^{1/2k} \log \Phi)$. Plugging this into [Lemma 5.2 \(B\)](#), yields a ϑ -reliable $\mathcal{O}(kt)$ -spanner with the number of edges bounded by

$$\mathcal{O}(\vartheta^{-1} D s n^{1/t}) = \mathcal{O}(\vartheta^{-1} k n^{1+1/k+1/t} \log^2 \Phi). \quad \blacksquare$$

5.3.2. Ultrametrics

Lemma 5.5. *Let $\mathcal{M} = (P, d)$ 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*

$$\mathcal{O}\left(\vartheta^{-1} \varepsilon^{-2} n \log^2 \Phi \log \frac{\log \Phi}{\vartheta \varepsilon}\right)$$

edges.

Proof: By [Corollary 4.8](#), one can build a $(1 + \varepsilon/2)$ -cover of $\mathcal{M}$ of depth $D = \mathcal{O}(\varepsilon^{-1} \log \Phi)$ and size $\mathcal{O}(nD)$. Plugging this into [Lemma 5.1](#) yields an oblivious ϑ -reliable 2-hop $(2 + \varepsilon)$ -spanner for $\mathcal{M}$, of size

$$\mathcal{O}\left(\frac{nD^2}{\vartheta} \log \frac{D}{\vartheta}\right) = \mathcal{O}\left(\vartheta^{-1} \varepsilon^{-2} n \log^2 \Phi \log \frac{\log \Phi}{\vartheta \varepsilon}\right). \quad \blacksquare$$

Lemma 5.6. *Let $\mathcal{M} = (P, d)$ be an ultrametric over n points with spread Φ , and let $\vartheta, \varepsilon \in (0, 1)$, and $t \in \mathbb{N}$ be parameters. Then, one can build a ϑ -reliable $((2 + \varepsilon)t - 1)$ -spanner for $\mathcal{M}$ of size $\mathcal{O}(\vartheta^{-2} \varepsilon^{-3} t \cdot n^{1+1/t} \log^3 \Phi)$.*

Proof: By [Corollary 4.8](#), one can build a $\xi = (1 + \varepsilon/2)$ -cover of $\mathcal{M}$ of depth $D = \mathcal{O}(\varepsilon^{-1} \log \Phi)$ and size $\mathcal{O}(nD)$. Plugging this into [Lemma 5.2 \(A\)](#) yields a deterministic ϑ -reliable $(2t - 1)\xi$ -spanner for $\mathcal{M}$, of size

$$\mathcal{O}(\vartheta^{-2} t D^3 n^{1+1/t}) = \mathcal{O}(\vartheta^{-2} \varepsilon^{-3} t \cdot n^{1+1/t} \log^3 \Phi). \quad \blacksquare$$

5.3.3. Tree metrics

Lemma 5.7. *Let $\mathcal{M} = (P, d)$ 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*

$$\mathcal{O}(\vartheta^{-1} \varepsilon^{-2} n \text{polylog}(n, \Phi))$$

edges, where $\text{polylog}(n, \Phi) = \log^2 n \log^2 \Phi \log \frac{\log \Phi \log n}{\vartheta \varepsilon}$.

Proof: By [Lemma 4.11](#), one can build a $(2 + \varepsilon/2)$ -cover of T of depth $D = \mathcal{O}(\varepsilon^{-1} \log \Phi \log n)$ and size $\mathcal{O}(nD)$. Plugging this into [Lemma 5.1](#) yields an oblivious ϑ -reliable 2-hop $(4 + \varepsilon)$ -spanner for $\mathcal{M}$, of size

$$\mathcal{O}\left(\frac{nD^2}{\vartheta} \log \frac{D}{\vartheta}\right) = \mathcal{O}\left(\vartheta^{-1} \varepsilon^{-2} n \log^2 n \log^2 \Phi \log \frac{\log \Phi \log n}{\vartheta \varepsilon}\right).$$

To get the improved bound on the dilation, let $p, q \in P$ be two points and let $C \in \mathcal{C}$ be the cluster such that $p, q \in C$ and $\text{diam}(C) \leq (2 + \varepsilon)d(p, q)$. Assume that $|C| > 2$ (otherwise $C = \{p, q\}$ and there is nothing to prove). By construction, for the separator node $\mathbf{s} \in C$, we have $d(\mathbf{s}, z) \leq \text{diam}(C)/2$ for all $z \in C$. Observe, that the (shortest) path between p and q in T passes through $\mathbf{s}$. Thus, using the triangle inequality, the length of a 2-hop path between p and q via $z \in C$ can be bounded by

$$d(p, z) + d(z, q) \leq d(p, \mathbf{s}) + 2d(\mathbf{s}, z) + d(\mathbf{s}, q) \leq d(p, q) + 2\text{diam}(C)/2 \leq (3 + \varepsilon)d(p, q). \quad \blacksquare$$

Lemma 5.8. *Let $\mathcal{M} = (P, d)$ be a tree metric over n points with spread Φ , and let $\vartheta, \varepsilon \in (0, 1)$, and $t \in \mathbb{N}$ be parameters. Then, one can build a ϑ -reliable $((4 + \varepsilon)t - 3)$ -spanner for $\mathcal{M}$ of size $\mathcal{O}(\vartheta^{-2} \varepsilon^{-3} \cdot n^{1+1/t} \log^3 n \log^3 \Phi)$.*

Proof: By [Lemma 4.11](#), one can build a $\xi = (2 + \varepsilon/2)$ -cover of T of depth $D = \mathcal{O}(\varepsilon^{-1} \log \Phi \log n)$ and size $\mathcal{O}(nD)$. Plugging this into [Lemma 5.2 \(A\)](#) yields a deterministic ϑ -reliable $(2t - 1)\xi$ -spanner for $\mathcal{M}$, of size

$$\mathcal{O}(\vartheta^{-2} D^3 n^{1+1/t}) = \mathcal{O}(\vartheta^{-2} \varepsilon^{-3} \cdot n^{1+1/t} \log^3 n \log^3 \Phi).$$

To get the improved bound on the dilation, consider a $(2t - 1)$ -hop path $p = p_0, p_1, \dots, p_{2t-1} = q$, such that $p_i \in C$, for all i , for some cluster C . We have

$$\begin{aligned} \sum_{i=0}^{2t-2} d(p_i, p_{i+1}) &= d(p, p_1) + d(p_{2t-2}, q) + \sum_{i=1}^{2t-3} d(p_i, p_{i+1}) \\ &\leq d(p, \mathbf{s}) + d(\mathbf{s}, p_1) + d(p_{2t-2}, \mathbf{s}) + d(\mathbf{s}, q) + (2t - 3)\text{diam}(C) \\ &\leq d(p, q) + (2t - 2)\text{diam}(C) \leq (1 + \xi(2t - 2)) \cdot d(p, q) \end{aligned}$$

for the length of the path. Thus, by using $2 + \varepsilon/2$ for the cover quality, we get

$$1 + (2t - 2)(2 + \varepsilon/2) = (4 + \varepsilon)t - 3 - \varepsilon \leq (4 + \varepsilon)t - 3. \quad \blacksquare$$

5.3.4. Planar graphs

Lemma 5.9. *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 $\mathcal{O}(\vartheta^{-1}\varepsilon^{-4}n \text{polylog}(n, \Phi))$ edges, where $\text{polylog}(n, \Phi) = \log^2 n \log^2 \Phi \log \frac{\log \Phi \log n}{\vartheta \varepsilon}$.*

Proof: By [Theorem 4.17](#), one can build a $(2 + \varepsilon/2)$ -cover of G of depth $D = \mathcal{O}(\varepsilon^{-2} \log \Phi \log n)$ and size $\mathcal{O}(nD)$. Plugging this into [Lemma 5.1](#) yields an oblivious ϑ -reliable 2-hop $(4 + \varepsilon)$ -spanner for G , of size $\mathcal{O}\left(\frac{nD^2}{\vartheta} \log \frac{D}{\vartheta}\right) = \mathcal{O}\left(\vartheta^{-1}\varepsilon^{-4}n \log^2 n \log^2 \Phi \log \frac{\log \Phi \log n}{\vartheta \varepsilon}\right)$.

We next show the improved bound on the dilation. Using the notation from [Lemma 4.15](#), for a pair of points $p, q \in P$, there is a cluster $C \in \mathcal{C}$, such that $p, q \in C$ and $\text{diam}(C) \leq (2 + \varepsilon)d(p, q)$. Let v be the point where the cycle separator and the shortest path between p and q intersect. Notice, that for the center point $u \in C$ we have $d(u, z) \leq \text{diam}(C)/2$, for all $z \in C$. Furthermore, using the notation from the proof of [Lemma 4.15](#), for some j , we have

$$\frac{d(u, v)}{d(p, q)} \leq \frac{\varepsilon r_j / 8}{r_{j-1}} = \frac{\varepsilon}{8} \left(1 + \frac{\varepsilon}{8}\right) \leq \frac{\varepsilon}{4}.$$

Thus, the length of a 2-hop path between p and q via $z \in C$ can be bounded by

$$\begin{aligned} d(p, z) + d(z, q) &\leq d(p, v) + 2d(v, u) + 2d(u, z) + d(v, q) \leq d(p, q) + 2\frac{\varepsilon}{4}d(p, q) + 2\text{diam}(C)/2 \\ &\leq \left(1 + \frac{\varepsilon}{2} + 2 + \varepsilon\right) d(p, q) \leq (3 + 2\varepsilon)d(p, q). \quad \blacksquare \end{aligned}$$

Lemma 5.10. *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 $\mathcal{O}(\vartheta^{-2}\varepsilon^{-6} \cdot n^{1+1/t} \log^3 n \log^3 \Phi)$.*

Proof: Let $\xi = 2 + \varepsilon/2$. By [Theorem 4.17](#), one can build a ξ -cover of G with depth $D = \mathcal{O}(\varepsilon^{-2} \log \Phi \log n)$ and size $\mathcal{O}(nD)$. Plugging this into [Lemma 5.2 \(A\)](#) yields a deterministic ϑ -reliable $(2t - 1)\xi$ -spanner for $\mathcal{M}$, of size $\mathcal{O}(\vartheta^{-2}D^3 n^{1+1/t}) = \mathcal{O}(\vartheta^{-2}\varepsilon^{-6} \cdot n^{1+1/t} \log^3 n \log^3 \Phi)$.

The improved dilation follows by using the same argument as in the proof of [Lemma 5.9](#) and [Lemma 5.8](#). \blacksquare

6. Proper expanders as reliable spanners for uniform metric

The purpose of this section is to prove that with the appropriate parameters, proper expanders (as defined in [Definition 3.4](#)) and edge-union of proper expanders constitute good reliable spanners for uniform metrics. That is, they satisfy the requirements of [Theorem 3.6](#).

6.1. Preliminaries

In the following, (n, d) -graph denotes a d -regular, n -vertex graph. Recall that $\lambda(G)$ denotes the normalized eigenvalue of G – that is, the second largest in absolute value (see [Definition 3.4](#)).

For a given graph $G = (V, E)$ and $S, H \subseteq V$, denote by $\Gamma_H(S) = \{v \in H \mid \exists u \in S, uv \in E\}$ the *neighbors* of S in H . For $S, T \subset V$, denote $E(S, T) = \{uv \in E \mid u \in S, v \in T\}$.

The following is well known result on expanders, attributed to Alon and Chung [[AC88](#)] in [[HLW06](#), Sec. 2.4].

Lemma 6.1 (Expander mixing lemma). *Let $G = (V, E)$ be an (n, d) graph. Then for every $S, T \subset V$,*

$$\left| |E(S, T)| - \frac{d|S||T|}{n} \right| \leq \lambda(G)d\sqrt{|S||T|}. \quad (6.1)$$

We also need the following lemma which is a minor variant of known constructions.

Lemma 6.2 ([[BHO19](#)]). *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 $|X| \geq \xi|L|$, we have that $|\Gamma(X)| > (1 - \xi)|R|$, and*
- (II) *for any subset $Y \subseteq R$, with $|Y| \geq \xi|R|$, we have that $|\Gamma(Y)| > (1 - \xi)|L|$.*

Remark 6.3. The randomized construction of [Lemma 6.2](#) 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 6.2](#) constructively in this case (potentially losing an additional $\log t$ factor). This also applies to the other expander constructions used in this paper. But while the randomized construction works with high probability, verifying it seems computationally intractable.

6.2. Construction of reliable spanners from proper expanders

[Theorem 3.6](#) states the existence two different spanners and accordingly, we present two different graphs, $\mathcal{G}_{n,\vartheta,2t-1}$ and $\mathcal{G}_{n,\vartheta,2t}$.

We begin with $\mathcal{G}_{n,\vartheta,2t}$. Recall the definition of proper expander ([Definition 3.4](#)) with parameter δ . Fix $n, t \in \mathbb{N}$, $n > t$, and $\delta \in (0, 1/4)$ such that

$$n^{1/t} \geq (c_\delta \delta)^{-1}. \quad (6.2)$$

The graph $\mathcal{G}_{n,\vartheta,2t}$ is defined to be an (n, d) -graph that is a proper expander with $\delta = \vartheta$, and

$$d = \lceil \max\{2\vartheta^{-1}n^{1/t}, 36C^2\vartheta^{-3}c_\vartheta^{-2}\} \rceil \quad (6.3)$$

To define $\mathcal{G}_{n,\vartheta,2t-1}$ we follow an idea we used slightly inferior construction in a preliminary version of this paper, see [[HMO21](#), Section 3.3.1]: Let $n' = n^{1-1/t}$. Partition the space to $n^{1/t}$ subsets, $A_1, \dots, A_{n^{1/t}}$,

each of size n' , and let $t' = t - 1$. For every A_i construct a copy of the graph $\mathcal{G}_{n', \vartheta, 2t'}$ with A_i being the vertices. The degree in those graphs is $d' = \mathcal{O}(\vartheta^{-1} n'^{1/t'}) = \mathcal{O}(\vartheta^{-1} n^{1/t})$.

In addition, for every $i \neq j$ connect A_i with A_j with a bipartite expander $(A_i \cup A_j, E_{ij})$ according to [Lemma 6.2](#), with $\xi = \vartheta$. This increases the degree by $\mathcal{O}(\vartheta^{-2} n^{1/t})$. Thus, the total degree of $\mathcal{G}_{n, \vartheta, 2t-1}$ is $\mathcal{O}(\vartheta^{-2} n^{1/t})$.

6.3. Analysis of $\mathcal{G}_{n, \vartheta, 2t}$

For the rest of [Section 6.3](#), let $G = \mathcal{G}_{n, \vartheta, 2t}$, and $\lambda = \lambda(G)$.

6.3.1. The shadow of a bad set

Let $B \subset V$ an arbitrary subset, such that $(1 + 5\vartheta)|B| < n$ (otherwise, we can choose $B^+ = V$ and there is nothing to prove). Choose

$$\varepsilon = (1 + \vartheta)(|B|/n + \lambda/\sqrt{\vartheta}) \stackrel{(P3)}{\leq} (1 + \vartheta)(|B|/n + C/\sqrt{d\vartheta}) \stackrel{\text{Eq. (6.3)}}{\leq} (1 + \vartheta)|B|/n + \mathfrak{c}_\vartheta \vartheta, \quad (6.4)$$

so (recalling that $\delta = \vartheta$, and $\mathfrak{c}_\delta = \mathfrak{c}_\vartheta$)

$$1 - \delta - \varepsilon \geq 1 - \vartheta - |B|/n - \vartheta|B|/n - \mathfrak{c}_\vartheta \vartheta \geq 1 - \vartheta - (1 - 4\vartheta) - \vartheta - \mathfrak{c}_\vartheta \vartheta \geq \vartheta. \quad (6.5)$$

Define the “shadows of B ” as follows. Let $S_0 = \emptyset$, and for $i > 0$ let

$$S_i = S_{i-1} \cup \{u \in V \setminus (B \cup S_{i-1}) \mid |E(u, B \cup S_{i-1})| \geq \varepsilon d\}. \quad (6.6)$$

These are all the “bad” vertices that have a lot of neighbors inside the (growing) bad set $B \cup S_{i-1}$. Lastly, set $S = \cup_i S_i$ the limit of S_i , and $B^+ = B \cup S$.

6.3.2. Bounding the size of the shadow

By the construction above of the damaged set B^+ , for every $u \in V \setminus B^+$,

$$|E(u, B^+)| < \varepsilon d. \quad (6.7)$$

Claim 6.4. $|S| \leq \vartheta|B|$.

Proof: The argument we use is similar to the one used in the proof of [\[BLMN05, Lemma 5.3\]](#). Let $\Delta_i = S_i \setminus S_{i-1}$. We have that

$$\begin{aligned} \varepsilon d|S_i| &\leq \varepsilon d \sum_{j=1}^i |\Delta_j| \stackrel{\text{Eq. (6.6)}}{\leq} \sum_{j=1}^i |E(\Delta_j, B \cup S_{j-1})| = |E(S_i, B \cup S_{i-1})| \\ &\stackrel{\text{Eq. (6.1)}}{\leq} \frac{d(|B| + |S_{i-1}|)|S_i|}{n} + \lambda d \sqrt{(|B| + |S_{i-1}|)|S_i|}. \end{aligned}$$

This implies that $|S_i| \left(\varepsilon - \frac{|B| + |S_{i-1}|}{n} \right) \leq \lambda \sqrt{(|B| + |S_{i-1}|)|S_i|}$. Squaring, and dividing by $|S_i|$, we have

$$|S_i| \left(\varepsilon - \frac{|B| + |S_{i-1}|}{n} \right)^2 \leq \lambda^2 (|B| + |S_{i-1}|) \leq \lambda^2 (|B| + |S_i|),$$

which implies, for $\rho_i = \left(\varepsilon - \frac{|B| + |S_{i-1}|}{n}\right)^2$, that

$$|S_i| \leq \frac{\lambda^2 |B|}{\rho_i - \lambda^2}.$$

We claim that $|S_i| \leq \vartheta |B|$ for every i . Indeed, otherwise let i the smallest index such that $|S_i| > \vartheta |B|$. So $|S_{i-1}| \leq \vartheta |B| < |S_i|$. By definition, see [Eq. \(6.4\)](#), we have that $\varepsilon = (1 + \vartheta)(|B|/n + \lambda/\sqrt{\vartheta})$ and as such

$$\rho_i = \left(\varepsilon - \frac{|B| + |S_{i-1}|}{n}\right)^2 \geq \left((1 + \vartheta)\left(\frac{|B|}{n} + \frac{\lambda}{\sqrt{\vartheta}}\right) - \frac{(1 + \vartheta)|B|}{n}\right)^2 \geq \frac{1 + \vartheta}{\vartheta} \lambda^2$$

We thus have that

$$\vartheta |B| < |S_i| \leq \frac{\lambda^2 |B|}{\rho_i - \lambda^2} \leq \frac{\lambda^2 |B|}{\frac{1 + \vartheta}{\vartheta} \lambda^2 - \lambda^2} = \vartheta |B|,$$

which is a contradiction. ■

6.3.3. The expansion happens outside the bad set

Claim 6.5. *Let $U \subseteq V \setminus B^+$. If $|U| \leq \mathfrak{c}_\vartheta \vartheta n^{1-1/t}$ then $|\Gamma_{V \setminus B^+}(U)| \geq n^{1/t} |U|$.*

Proof: We have

$$|U| \leq \mathfrak{c}_\vartheta \vartheta n^{1-1/t} = \frac{\mathfrak{c}_\vartheta n}{n^{1/t}/\vartheta} \leq \frac{\mathfrak{c}_\vartheta n}{d/2} \leq \frac{\mathfrak{c}_\vartheta n}{d},$$

since $d \geq 2n^{1/t}/\vartheta$ by [Eq. \(6.3\)](#). As such, by the expansion property [\(P2\)](#), $|\Gamma_V(U)| \geq (1 - \delta)d|U|$ (as $\delta = \vartheta$). Furthermore

$$|\Gamma_{B^+}(U)| \leq |E(U, B^+)| \stackrel{\text{Eq. (6.7)}}{\leq} \varepsilon d |U|,$$

so

$$|\Gamma_{V \setminus B^+}(U)| \geq |\Gamma_V(U)| - |\Gamma_{B^+}(U)| \geq (1 - \delta - \varepsilon)d|U| \stackrel{\text{Eq. (6.5)}}{\geq} \vartheta d |U| \stackrel{\text{Eq. (6.3)}}{\geq} n^{1/t} |U|. \quad \blacksquare$$

Let

$$\mathbb{B}_{G \setminus B^+}(u, i) = \{v \in V \setminus B^+ \mid \mathbf{d}_{G \setminus B^+}(u, v) \leq i\},$$

the ball of radius i around u in the shortest path metric of the graph $G \setminus B^+$. Define $U_0 = \{u\}$, and $U_i = \Gamma_{V \setminus B^+}(U_{i-1})$. Observe that $\mathbb{B}_{G \setminus B^+}(u, i) = \bigcup_{j=0}^i U_j$.

Claim 6.6. $|\mathbb{B}_{G \setminus B^+}(u, t-1)| \geq n^{1-1/t}$.

Proof: Assume the contrary. Then $|U_{t-1}| \leq |\mathbb{B}_{G \setminus B^+}(u, t-1)| < n^{1-1/t}$, which implies that there exists $j < t-1$ such that $|U_{j+1}| < |U_j| n^{1/t}$. From [Claim 6.5](#), this means that $|U_j| > \mathfrak{c}_\vartheta \vartheta n^{1-1/t}$. So take $K \subset U_j$ such that $|K| = \mathfrak{c}_\vartheta \vartheta n^{1-1/t}$. Then again by [Claim 6.5](#),

$$|\mathbb{B}_{G \setminus B^+}(u, t-1)| \geq |\Gamma_{V \setminus B^+}(K)| \geq \mathfrak{c}_\vartheta \vartheta n^{1-1/t} n^{1/t} \stackrel{\text{Eq. (6.2)}}{\geq} n^{1-1/t}. \quad \blacksquare$$

We summarize the relevant properties of $\mathcal{G}_{n, \vartheta, 2t}$.

Lemma 6.7. *The graph $\mathcal{G}_{n,\vartheta,2t}$ has the following properties. For any $B \subset V$, there exists a set B^+ , $B \subseteq B^+ \subset V$, $|B^+| \leq (1 + 5\vartheta)|B|$, such that for any $u \in V \setminus B^+$, we have*

$$|\mathbb{B}_{G \setminus B^+}(u, t-1)| \geq n^{1-1/t}, \quad (6.8)$$

$$\Gamma_V(\mathbb{B}_{G \setminus B^+}(u, t-1)) \geq (1 - \vartheta)n, \quad (6.9)$$

$$|\mathbb{B}_{G \setminus B^+}(u, t)| \geq \vartheta n. \quad (6.10)$$

Proof: Recall that when $(1 + 5\vartheta)|B| \geq n$, we can take $B^+ = V$ and there is nothing to prove. So assume from now on that $(1 + 5\vartheta)|B| < n$.

When $t = 1$, $\mathcal{G}_{n,\vartheta,2t}$ is the complete graph, and we take $B^+ = B$. Eq. (6.8) is equivalent to $1 \geq 1$; Eq. (6.9) follows since $\Gamma_V(\{u\}) = V$ in complete graph; and Eq. (6.10) follows simply because $|B| \leq n/(1 + 5\vartheta) \leq (1 - \vartheta)n$.

Assume next that $t \geq 2$. Eq. (6.8) is just a repetition of Claim 6.6. To prove Eq. (6.9). Observe that

$$|\mathbb{B}_{G \setminus B^+}(u, t-1)| \stackrel{\text{Eq. (6.8)}}{\geq} n^{1-1/t} \stackrel{\text{Eq. (6.3)}}{\geq} 2n/(\vartheta d).$$

We conclude using Property (P1) of proper expanders that

$$\Gamma_V(\mathbb{B}_{G \setminus B^+}(u, t-1)) \geq (1 - \vartheta)n.$$

By Claim 6.4,

$$|B^+| \leq (1 + \vartheta)|B| \leq (1 + \vartheta) \frac{n}{1 + 5\vartheta} \leq (1 - 3\vartheta)n.$$

So,

$$|\mathbb{B}_{G \setminus B^+}(u, t)| \geq |\Gamma_V(\mathbb{B}_{G \setminus B^+}(u, t-1))| \geq |\Gamma_V(U_{t-1})| - |B^+| \geq \vartheta n.$$

This completes the proof of Lemma 6.7. ■

Proposition 6.8. *The graph $\mathcal{G}_{n,\vartheta,2t}$ is a ϑ -reliable $2t$ -spanner for n -point uniform space with $\mathcal{O}(\vartheta^{-1}n^{1+1/t})$ edges.*

Proof: Fix $u, v \in V \setminus B^+$. By Eq. (6.9),

$$\min\left\{|\Gamma_V(\mathbb{B}_{G \setminus B^+}(u, t-1))|, |\Gamma_V(\mathbb{B}_{G \setminus B^+}(v, t-1))|\right\} \geq (1 - \vartheta)n.$$

So,

$$\begin{aligned} |\mathbb{B}_{G \setminus B^+}(u, t) \cap \mathbb{B}_{G \setminus B^+}(v, t)| &\geq |\Gamma_V(\mathbb{B}_{G \setminus B^+}(u, t-1)) \cap \Gamma_V(\mathbb{B}_{G \setminus B^+}(v, t-1))| - |B^+| \\ &\geq (1 - 2\vartheta)n - (1 - 5\vartheta)n > 0. \end{aligned}$$

We conclude that $\mathbb{B}_{G \setminus B^+}(u, t) \cap \mathbb{B}_{G \setminus B^+}(v, t) \neq \emptyset$, which means that there is a path of length at most $2t$ in $G \setminus B^+$ between u and v . ■

6.3.4. Proof of Theorem 3.10

Restatement of Theorem 3.10. *For every $\vartheta \in (0, 1)$ there exists a constant $\mathfrak{c} = \mathfrak{c}_\vartheta > 0$, such that for any expansion constant $h \geq \mathfrak{c}^{-2}$, there exists a family of vertex expander graphs $\{G = (V, E)\}_n$ on n vertices of degree at most $4h/\vartheta$, with the following resiliency property: For any $B \subset V$, there exists $B^+ \supseteq B$, $|B^+| \leq (1 + \vartheta)|B|$ such that the graph $G \setminus B^+$ is a vertex expander in the sense that*

(i) $\text{diam}(G \setminus B^+) \leq 2\lceil \log_h n \rceil$, and

(ii) For any $U \subset V \setminus B^+$ of size $|U| \leq \mathfrak{c}n/h$, we have $|\Gamma_{G \setminus B^+}(U)| \geq h|U|$.

Proof: The graphs are simply $\mathcal{G}_{n,\vartheta,2t}$ for $t = \log_h n$. The claims follow immediately from Claim 6.5 and Proposition 6.8. ■

6.4. Analysis of $\mathcal{G}_{n,\vartheta,2t-1}$

Proposition 6.9. *The graph $\mathcal{G}_{n,\vartheta,2t-1}$ is ϑ -reliable $(2t - 1)$ -spanner for n -point uniform space with $\mathcal{O}(\vartheta^{-2}n^{1+1/t})$ edges.*

Proof: Recall the definition of $\mathcal{G}_{n,\vartheta,2t-1}$ from [Section 6.2](#). Let $t' = t - 1$, and $n' = n^{1-1/t}$. Given an attack set B , construct B_i^+ from $A_i \cap B$ according to [Lemma 6.7](#), and define $B^+ = \cup_i B_i^+$. Fix $u, v \in V \setminus B^+$. If $u, v \in A_i$ then by [Proposition 6.8](#), there is a $2t' = 2t - 2$ path between them.

If $u \in A_i, v \in A_j, i \neq j$. Then by [Eq. \(6.10\)](#), $|\mathbb{B}_{A_i \setminus B^+}(u, t')| \geq \vartheta n'$, and $|\mathbb{B}_{A_j \setminus B^+}(v, t')| \geq \vartheta n'$. By [Lemma 6.2](#), there is an edge in E_{ij} between $\mathbb{B}_{A_i \setminus B^+}(u, t')$ and $\mathbb{B}_{A_j \setminus B^+}(v, t')$. and hence a path of length $2t' + 1 = 2t - 1$ in $G \setminus B^+$. $\blacksquare$

7. Concluding remarks and open problems

Subsequent work. Recently Filtser and Le [\[FL21\]](#) improved some of the results here. They obtained bounds that do not depend on the spread of the metrics in some cases, for the oblivious adversary case. They also obtained reliable spanners (in the oblivious adversary model) for trees with tight stretch of 2 and for planar graphs with tight stretch of $2 + \varepsilon$.

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 only able to show 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 $\tilde{\mathcal{O}}(n^{1+1/t})$ 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 [\[BHO19, BHO20\]](#) achieves this for fixed dimensional Euclidean spaces, and [\[FL21\]](#) 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.

Acknowledgments. The authors thank Kevin Buchin for useful discussions and references, and the anonymous referees for their suggestions.

References

- [AC88] N. Alon and F. R. K. Chung. Explicit construction of linear sized tolerant networks. *Proceedings of the First Japan Conference on Graph Theory and Applications (Hakone, 1986)*, vol. 72. 15–19, 1988.
- [AG06] I. Abraham and C. Gavoille. Object location using path separators. *Proc. 25th ACM Sympos. Prin. Dist. Comput.*, 188–197, 2006.
- [AGG+19] I. Abraham, C. Gavoille, A. Gupta, O. Neiman, and K. Talwar. Cops, robbers, and threatening skeletons: padded decomposition for minor-free graphs. *SIAM J. Comput.*, 48(3): 1120–1145, 2019.
- [AGMW10] I. Abraham, C. Gavoille, D. Malkhi, and U. Wieder. Strong-diameter decompositions of minor free graphs. *Theory Comput. Syst.*, 47(4): 837–855, 2010.
- [AP90] B. Awerbuch and D. Peleg. Sparse partitions. *Proc. 31st Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, 503–513, 1990.
- [Bar98] Y. Bartal. On approximating arbitrary metrics by tree metrics. *Proc. 30th Annu. ACM Sympos. Theory Comput. (STOC)*, 161–168, 1998.
- [BCDM18] P. Bose, P. Carmi, V. Dujmovic, and P. Morin. Near-optimal $O(k)$ -robust geometric spanners. *CoRR*, abs/1812.09913, 2018. arXiv: 1812.09913.
- [BDMS13] P. Bose, V. Dujmović, P. Morin, and M. Smid. Robust geometric spanners. *SIAM Journal on Computing*, 42(4): 1720–1736, 2013.
- [BDPW18] G. Bodwin, M. Dinitz, M. Parter, and V. V. Williams. *Optimal Vertex Fault Tolerant Spanners (for fixed stretch)*. *Proc. 29th ACM-SIAM Sympos. Discrete Algs. (SODA)*. 3600 University City Science Center, Philadelphia, PA, USA: SIAM, 2018, pp. 1884–1900. eprint: <https://epubs.siam.org/doi/pdf/10.1137/1.9781611975031.123>.
- [BHO19] K. Buchin, S. Har-Peled, and D. Oláh. A spanner for the day after. *Proc. 35th Int. Annu. Sympos. Comput. Geom. (SoCG)*, vol. 129. 19:1–19:15, 2019.
- [BHO20] K. Buchin, S. Har-Peled, and D. Oláh. Sometimes reliable spanners of almost linear size. *Proc. 29th Annu. Euro. Sympos. Alg. (ESA)*, vol. 173. 27:1–27:15, 2020.
- [BLMN05] Y. Bartal, N. Linial, M. Mendel, and A. Naor. On metric Ramsey-type phenomena. *Ann. of Math. (2)*, 162(2): 643–709, 2005.
- [BLT14] C. Busch, R. LaFortune, and S. Tirthapura. Sparse covers for planar graphs and graphs that exclude a fixed minor. *Algorithmica*, 69(3): 658–684, 2014.
- [CLNS15] 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.
- [DJPP13] I. Dumitriu, T. Johnson, S. Pal, and E. Paquette. Functional limit theorems for random regular graphs. *Probab. Theory Related Fields*, 156(3-4): 921–975, 2013.
- [FKS89] J. Friedman, J. Kahn, and E. Szemerédi. On the second eigenvalue in random regular graphs. *Proceedings of the 21st Annual ACM Symposium on Theory of Computing, May 14-17, 1989, Seattle, Washington, USA*, 587–598, 1989.
- [FL21] A. Filtser and H. Le. Reliable spanners: locality-sensitive orderings strike back. *Arxiv CoRR*, abs/2101.07428, 2021. arXiv: 2101.07428.

- [Fri03] J. Friedman. A proof of alon’s second eigenvalue conjecture. *Proc. 35th Annu. ACM Sympos. Theory Comput.* (STOC), 720–724, 2003.
- [FRT04] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Sys. Sci.*, 69(3): 485–497, 2004.
- [HLW06] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S.)*, 43(4): 439–561, 2006.
- [HMO21] S. Har-Peled, M. Mendel, and D. Oláh. Reliable spanners for metric spaces. *Proc. 37th Int. Annu. Sympos. Comput. Geom.* (SoCG), vol. 189. 43:1–43:13, 2021.
- [KLMN05] 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.
- [LNS02] C. Levcopoulos, G. Narasimhan, and M. Smid. Improved algorithms for constructing fault-tolerant spanners. *Algorithmica*, 32(1): 144–156, 2002.
- [LNS98] C. Levcopoulos, G. Narasimhan, and M. H. M. Smid. Efficient algorithms for constructing fault-tolerant geometric spanners. *Proc. 30th Annu. ACM Sympos. Theory Comput.* (STOC), 186–195, 1998.
- [LT79] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2): 177–189, 1979. eprint: <https://doi.org/10.1137/0136016>.
- [Luk99] T. Lukovszki. New results of fault tolerant geometric spanners. *Proc. 6th Int. Workshop Alg. and Data Struct.*, 193–204, 1999.
- [MN07] M. Mendel and A. Naor. Ramsey partitions and proximity data structures. *eng. J. Euro. Math. Soc.*, 009(2): 253–275, 2007.
- [NS07] G. Narasimhan and M. Smid. *Geometric spanner networks*. New York, NY, USA: Cambridge University Press, 2007.
- [Pel00] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. 3600 University City Science Center, Philadelphia, PA, USA: SIAM, 2000. eprint: <https://epubs.siam.org/doi/pdf/10.1137/1.9780898719772>.
- [PS89] D. Peleg and A. Schäffer. Graph spanners. *J. Graph Theory*, 13: 99–116, 1989.
- [Pud15] D. Puder. Expansion of random graphs: new proofs, new results. *Invent. Math.*, 201(3): 845–908, 2015.
- [Sol14] S. Solomon. From hierarchical partitions to hierarchical covers: optimal fault-tolerant spanners for doubling metrics. *Proceedings of the 46th ACM Symposium on Theory of Computing*, 363–372, 2014.
- [Wor99] N. C. Wormald. Models of random regular graphs. *Surveys in combinatorics, 1999 (Canterbury)*. Vol. 267. London Math. Soc. Lecture Note Ser. UK: Cambridge Univ. Press, Cambridge, 1999, pp. 239–298.

A. Random regular graphs

In [Theorem 3.5](#) above we prove that random construction using a union of random permutations yields a proper expander, see [Definition 3.4](#). The properties are well known and are held by random regular graphs asymptotically almost surely. However, the literature on random regular graphs is more concerned

with the setting of constant d and n tends to infinity, especially regarding Property (P3)). Here we need a slightly different range, where d is sufficiently large, and $n \geq d^2$. For completeness, we gather here proofs and references that prove Theorem 3.5.

A.1. Construction

We use the *permutation model* $\mathcal{G}_{n,d}$ for constructing random regular graph (see [Wor99] for other models). Assuming d is even, sample $d/2$ independent and identically distributed random permutations $\pi_1, \dots, \pi_{d/2} \in S_n$, where S_n is the set of all permutations of $[n] = \{1, \dots, n\}$. The resulting graph $G = (V, E)$, has $V = [n]$ and

$$E = \left\{ \{i, \pi_j(i)\} \mid i \in [n] \text{ and } j \in [d/2] \right\}.$$

A.2. Analysis

A.2.1. Property (P3)

All the proofs that random d -regular graphs have second eigenvalue at most $C/\sqrt{d}$ (that we are aware of) are non-trivial. It was first proved in [FKS89], and by now there are many proofs of this. Notably, Friedman [Fri03] showed that random regular graphs are “almost Ramanujan”, i.e., $\lambda \leq \frac{2\sqrt{d-1}+\varepsilon}{d}$ with probability $1 - o_n(1)$, see also [Pud15] for a recent and simpler proof. Unfortunately, most of those papers are interested in the settings where the degree d is constant and only the number of vertices $n \rightarrow \infty$, which is not suitable for our needs. However, the argument of Kahn and Szemerédi [FKS89] does work when d is allowed to tend to infinity (together with n). Specifically, Dumitriu *et al.* [DJPP13] used it to prove the following.

Theorem A.1 ([DJPP13, Theorem 24]). *Fix $C = 41000$. There exists $K > 0$ such that for any even d and $n > \max\{K, d\}$, we have $\mathbb{P}[\lambda(G) \leq C/\sqrt{d}] \geq 1 - 2/n^2$, where $\mathbb{P}[\cdot]$ is the probability in the permutation model $\mathcal{G}_{n,d}$.*

This implies (P3) holds a.a.s. in the permutation model. Properties (P1) and (P2) are much easier to prove and are considered folklore. They are usually proved in different, more convenient, models of random regular graphs. However, contrary to the case when the degree is fixed, in the high-degree regime the different random models are not necessarily contiguous, i.e., asymptotically almost surely properties are not necessarily equivalent among the different models (for more on this, see [Wor99]). Therefore, for completeness, we next provide proofs that Properties (P1) and (P2) hold asymptotically almost surely in the permutation model.

A.2.2. Property (P1)

Lemma A.2. *For integers $s, t, n \geq 0$, such that $s \leq t \leq n$, we have $\binom{t}{s} / \binom{n}{s} \leq \left(\frac{t}{n}\right)^s$.*

Similarly, if $s \leq t \leq n$ and $n \geq 3s$, we have $\binom{t}{s} / \binom{n}{2s} \leq \left(\frac{t}{n}\right)^s \left(\frac{2s}{n-s}\right)^s$.

Proof: Observe that for $k \leq m' \leq m$, we have

$$\frac{\binom{m'}{k}}{\binom{m}{k}} = \frac{m'}{m} \cdot \frac{m'-1}{m-1} \cdots \frac{m'-k+1}{m-k+1} \leq \left(\frac{m'}{m}\right)^k,$$

as $(m' - i)/(m - i) \leq m'/m$. As such, using the (easy to verify) identity $\binom{n}{2s} = \binom{n}{s} \binom{n-s}{s} / \binom{2s}{s}$, we have

$$\frac{\binom{t}{s}}{\binom{n}{2s}} = \frac{\binom{t}{s} \binom{2s}{s}}{\binom{n}{s} \binom{n-s}{s}} \leq \left(\frac{t}{n}\right)^s \left(\frac{2s}{n-s}\right)^s. \quad \blacksquare$$

Lemma A.3. Fix sets $S, T \subset V$, with $s = |S|$ and $t = |T|$, such that $|S| \leq |T|$. Then $\mathbb{P}[\Gamma(S) \subseteq T] \leq \left(\frac{t}{n}\right)^{sd/2}$.

Proof: Fix $T' \subseteq T$ of cardinality $|T'| = s$. Then $\mathbb{P}[\pi(S) = T'] = 1/\binom{n}{s}$. By the union bound, and [Lemma A.2](#), we have

$$\mathbb{P}[\pi(S) \subseteq T] = \sum_{T' \in \binom{T}{s}} \mathbb{P}[\pi(S) = T'] = \frac{\binom{t}{s}}{\binom{n}{s}} \leq \left(\frac{t}{n}\right)^s.$$

Let $\pi_1, \dots, \pi_{d/2}$ be the permutations used in the construction. We have that

$$\mathbb{P}[\Gamma(S) \subseteq T] \leq \mathbb{P}\left[\bigwedge_{k=1}^{d/2} \pi_k(S) \subseteq T\right] = \left(\mathbb{P}[\pi(S) \subseteq T]\right)^{d/2} \leq \left(\frac{t}{n}\right)^{sd/2}. \quad \blacksquare$$

Lemma A.4. Fix $\varepsilon, \delta \in (0, 1/2)$. If $\sqrt{n} \geq d \geq 12/(\varepsilon\delta)$ is an integer. Then, asymptotically almost surely over the random (n, d) -graph $G = (V, E)$, for any $S \subset V$ with $|S| \geq \varepsilon n$, we have $|\Gamma(S)| \geq (1 - \delta)n$. That is, property (P1) of [Definition 3.4](#) holds.

Proof: Using the union bound on the bound established in [Lemma A.3](#) we conclude that the probability that there exists a subset S of size $s = \varepsilon n$, such that $|\Gamma(S)| \leq (1 - \delta)n = t$, is at most

$$\begin{aligned} \binom{n}{s} \binom{n}{t} \left(\frac{t}{n}\right)^{sd/2} &\leq \left(\frac{\varepsilon n}{s}\right)^s 2^n \left(\frac{t}{n}\right)^{sd/2} \leq \left(\frac{e}{\varepsilon}\right)^{\varepsilon n} 2^n (1 - \delta)^{\varepsilon nd/2} \leq \left(\frac{e}{\varepsilon}\right)^{\varepsilon n} 2^n \exp(-\delta \varepsilon nd/2) \\ &\leq 2^n \exp\left(\varepsilon n + \varepsilon n \ln \frac{1}{\varepsilon} - 6n\right) \leq \exp(2n + n/e - 6n) \leq \exp(-3n), \end{aligned}$$

since $\max_{x>0} x \ln(1/x) = 1/e$, as easy calculation shows¹. \blacksquare

A.2.3. Property (P2)

The argument above used only the forward edges associated with the $d/2$ random permutations. Since there are only $d/2$ such edges associated with every vertex, that argument can not prove vertex-expansion close to d as is stated in (P2). For this we have to also use the backward edges associated with the permutation as well. Since the backward edges and the forward edges associated with the same random permutation are not stochastically independent, more care is needed to prove this property, as we shall now see.

Claim A.5. Fix $\varepsilon, \eta \in (0, 1)$ such that $0 < \varepsilon < \frac{1}{3d} < \frac{3}{d} < \eta < 1$. Then

$$\mathbb{P}\left[\exists S \subset V, |S| \leq \varepsilon n, |E(S, S)| > \eta d |S|\right] = O(n^{-0.4}).$$

¹Indeed, for $f(x) = x \ln(1/x)$, we have $f'(x) = -1 + \ln x^{-1}$. Setting $f'(x) = 0$, implies $1 = \ln x^{-1}$. Namely, the maximum of f is achieved at $x = 1/e$, where $f(1/e) = 1/e$.

Proof: Fix $S \subset V$, of cardinality $s \leq \varepsilon n$. Observe that

$$E(S, S) = \{ \{u, \pi_i(u)\} \mid u \in S, \pi_i(u) \in S, \text{ and } i \in [d/2] \}.$$

Fix $R \subseteq S \times [d/2]$. For any $(u, i) \in R$, let $\mathcal{E}_{u,i}$ be the event that $\pi_i(u) \in S$. The events induced by R , that is $\{\mathcal{E}_{u,i} \mid (u, i) \in R\}$ are not independent, but they are non-positively correlated. Indeed, for a set $R \subset S \times [d/2]$ consider the event $\mathcal{E}_R = \bigwedge_{(z,j) \in R} \mathcal{E}_{z,j}$. For $(u, i) \notin R$ we have,

$$\mathbb{P}[\pi_i(u) \in S \mid \mathcal{E}_R] \leq \mathbb{P}[\pi_i(u) \in S] = \frac{s}{n}.$$

Indeed, observe that all the sub-events in $\mathcal{E}_R$ involving $j \neq i$, are irrelevant (that is, stochastically independent of the events $\mathcal{E}_{x,i}$). As such, let $X = \{x \mid (x, i) \in R\}$. Imagine now picking the permutation π_i by first picking the locations of the elements of X , and then choosing the location of u . Clearly, if the event $\mathcal{E}_R$ happened, then there are only $s - |X|$ empty slots in S , and $n - |X|$ slots available overall. As such, we have that

$$\mathbb{P}[\pi_i(u) \in S \mid \mathcal{E}_R] \leq \frac{s - |X|}{n - |X|} \leq \frac{s}{n} = \mathbb{P}[\pi_i(u) \in S].$$

Now, order the elements of R in arbitrary order, and let $R(i)$ be the subset with the first i elements in that order. We have

$$\mathbb{P}[\mathcal{E}_R] = \prod_{i=1}^{|R|} \mathbb{P}[\mathcal{E}_{R(i)} \mid \mathcal{E}_{R(i-1)}] \leq (s/n)^{|R|}.$$

Therefore,

$$\begin{aligned} \mathbb{P}[|E(S, S)| > \eta ds] &\leq \sum_{r=\eta ds+1}^{ds/2} \sum_{R \in \binom{S \times [d/2]}{r}} \mathbb{P}[\mathcal{E}_R] \leq \sum_{r=\eta ds}^{ds/2} \binom{ds/2}{r} \cdot \left(\frac{s}{n}\right)^r \leq \sum_{r=\eta ds}^{ds/2} \left(\frac{eds/2}{r} \cdot \frac{s}{n}\right)^r \\ &\leq \sum_{r=\eta ds}^{ds/2} \left(\frac{eds^2}{2rn}\right)^r \leq \left(\frac{es}{\eta n}\right)^{\eta ds}, \end{aligned}$$

since the summations behaves like a geometric series, using $\frac{eds^2}{2rn} \leq \frac{eds^2}{2\eta ds n} = \frac{es}{2\eta n} \leq \frac{1}{2}$. Therefore

$$\Pr[\exists S \subset V, |S| \leq \varepsilon n, |E(S, S)| > \eta ds] \leq \sum_{s=1}^{\varepsilon n} \binom{n}{s} \left(\frac{es}{\eta n}\right)^{\eta ds} \leq \sum_{s=1}^{\varepsilon n} \left(e^2 \left(\frac{s}{n}\right)^{\eta d-1} \eta^{-\eta d}\right)^s = O(n^{-0.4}).$$

The last inequality is derived by partitioning the sum on s up-to $2 \log n$ – this part is bounded by $O\left(\frac{\log^{3/2} n}{\sqrt{n}}\right)$. The second part of the summation from $2 \log n$ to εn be bounded by $O((\varepsilon/\eta)^{2 \log n})$. ■

Lemma A.6. *For every $\delta \in (0, 1/4)$, there exists $c_\delta > 0$, such that for any sufficiently large integer $d \geq e^{32/\delta}$, and any $n \geq d^2$, asymptotically (in n) almost surely, a random (n, d) graph $G = (V, E) \sim \mathcal{G}_{n,d}$ have the following vertex expansion property: For any $S \subset V$ with $|S| \leq e^{-2.2/\delta} n/d$, we have $|\Gamma(S)| \geq (1 - \delta)d|S|$.*

That is, Property (P2) of Definition 3.4 holds.

Proof: Fix $\eta = \delta/2$. Define $\mathcal{B}$ to be the event, that for all subsets $S \subset V$ of cardinality at most εn , we have that

$$|E(S, S)| \leq \eta d|S|.$$

By Claim A.5, $\Pr[\mathcal{B}] \geq 1 - O(n^{-0.4})$. So, fix $S \subset V$ of cardinality $s \leq \varepsilon n$. The following stochastic process samples a random bijection π on V uniformly.

- (i) Order the vertices of S arbitrarily $v_1, \dots, v_s$.
- (ii) For $i = 1, \dots, s$, pick $\pi(v_i)$ uniformly at random from $V \setminus \{\pi(v_1), \dots, \pi(v_{i-1})\}$.
- (iii) Let $U = S \cap \pi(S)$. Reorder the elements of S as $u_1, \dots, u_s$, such that $U = \{u_{\ell+1}, \dots, u_s\}$.
- (iv) For $i = 1, \dots, \ell$, pick $\pi^{-1}(u_i)$ uniformly at random from $V \setminus (S \cup \{\pi^{-1}(u_1), \dots, \pi^{-1}(u_{i-1})\})$.
- (v) Denote by π' the partial injection constructed so far.
- (vi) Note that, for $i = \ell + 1, \dots, s$, the element $\pi^{-1}(u_i)$ is already fixed, and it is in S , as such we do nothing.
- (vii) Complete π' to a full permutation by sampling a random bijection $(V \setminus (S \cup \pi'(S))) \mapsto (V \setminus (S \cup \pi'(S)))$.

One can sample random regular graph G from $\mathcal{G}_{n,d}$ by applying the above independently $d/2$ times, and constructing $d/2$ random permutations $\pi_1, \dots, \pi_{d/2}$.

We next do for a fixed $S \subset V$ of cardinality $s \leq \varepsilon n$ the following thought experiment. In the construction of each of the $d/2$ permutation, we replace step (vi) above with the following step generating a new partial injection τ into V , as follows:

- (vi') For $i = \ell + 1, \dots, s$, pick $\tau^{-1}(u_i)$ randomly and uniformly from

$$V \setminus (S \cup \{\pi^{-1}(u_1), \dots, \pi^{-1}(u_\ell), \tau^{-1}(u_{\ell+1}), \dots, \tau^{-1}(u_{i-1})\}).$$

Denote by $\tau_i = \tau_i^S$ that random partial injection. Denote

$$F_S = \{\{u, \tau_i^{-1}(u)\} \mid i \in [d/2], u \in S \cap \pi_i(S)\}$$

the set of auxiliary edges associated with the above process for S . Denote $G_S = (V, E \cup F_S)$ – it is the result of redirecting all the edges of $E(S, S)$ to be outside the $S \times S$ biclique. The graph G_S is randomly generated, but the content of F_S is not measurable in $\mathcal{G}_{n,d}$ – meaning that the underlying probability spaces is more refined than $\mathcal{G}_{n,d}$. We denote by $\widetilde{\mathcal{G}}_{n,d}$ a probability space from which one can sample all of G_S . I.e., it contains $d/2$ random permutations $\pi_1, \dots, \pi_{d/2}$ that constitute $\mathcal{G}_{n,d}$ and all the relevant partial injections τ_i^S .

Note that $|F_S| = |E(S, S)|$, and therefore depends only on $\mathcal{G}_{n,d}$. Furthermore, If $G \in \mathcal{B}$, then for every $S \subset V$, with cardinality at most εn , we have that $|F_S| \leq \eta d |S|$. Hence, conditioned on $\mathcal{B}$, for any such S ,

$$|\Gamma_G(S)| \geq |\Gamma_{G_S}(S)| - \eta d |S|. \quad (\text{A.1})$$

We next bound for a fixed $S \subset V$ of cardinality $s \leq \varepsilon n$, the probability that there exists a subset T of cardinality t for which $\Gamma_G(S) \subseteq T$. Begin by fixing $S \subseteq T \subseteq V$ of cardinality $|T| = t$. We claim that

$$\mathbb{P}[\pi(S) \cup \pi^{-1}(S) \cup \tau^{-1}(S \cap \pi(S)) \subseteq T] \leq \left(\frac{t}{n}\right)^{2s}. \quad (\text{A.2})$$

Indeed, fix $T' \subset T$ of cardinality s . Then

$$\mathbb{P}[\pi(S) = T'] \leq \binom{n}{s}^{-1}.$$

Fix $T'' \subset V \setminus S$ of cardinality s . Recall that we denoted by $\pi' \subset \pi$ the partial injection sampled at the end of step (v) in the sampling process described above. Observe that conditioned on $\pi(S) = T'$ the

partial injection $\pi'^{-1} \cup \tau^{-1} : S \rightarrow V \setminus S$ is also uniformly random, and therefore

$$\mathbb{P}[(\pi'^{-1} \cup \tau^{-1})(S) = T'' \mid \pi(S) = T'] \leq \binom{n-s}{s}^{-1}.$$

Hence,

$$\begin{aligned} & \mathbb{P}[\pi(S) \cup \pi^{-1}(S) \cup \tau^{-1}(S \cap \pi(S)) \subseteq T] \\ &= \mathbb{P}\left[\bigvee_{T' \subset T, |T'|=s} \bigvee_{T'' \subset T \setminus T', |T''|=s} \pi(S) = T' \wedge (\pi'^{-1} \cup \tau^{-1})(S) = T''\right] \\ &\leq \sum_{T' \subset T, |T'|=s} \sum_{T'' \subset T \setminus T', |T''|=s} \Pr[\pi(S) = T'] \cdot \Pr[(\pi'^{-1} \cup \tau^{-1})(S) = T'' \mid \pi(S) = T'] \\ &= \frac{\binom{t}{s} \binom{t-s}{s}}{\binom{n}{s} \binom{n-s}{s}} \leq \left(\frac{t}{n}\right)^s \left(\frac{t-s}{n-s}\right)^s \leq \left(\frac{t}{n}\right)^{2s}, \end{aligned}$$

which complete the proof of [Eq. \(A.2\)](#).

The probability on $\{G_S\}_{S \subset V} \sim \widetilde{\mathcal{G}}_{n,d}$ that there exists $S \subset V$ of cardinality at most εn for which $|\Gamma_{G_S}(S)| \leq t$, where $t = (1 - \eta)ds$ is at most the probability that there exists $\widetilde{T} \subset V \setminus S$ of cardinality t for which $\Gamma_{G_S}(S) \subseteq \widetilde{T} \cup S$. By the union bound and [Eq. \(A.2\)](#) this is at most

$$\sum_{s=1}^{\varepsilon n} \binom{n}{s} \binom{n-s}{t} \left(\left(\frac{t+s}{n}\right)^{2d}\right)^{d/2}.$$

We bound the summand above by

$$\begin{aligned} \binom{n}{s} \binom{n-s}{t} \left(\frac{t+s}{n}\right)^{sd} &\leq \binom{n}{s} \binom{n}{t} \left(\frac{t+s}{n}\right)^{sd} \\ &\leq \left(\frac{en}{s}\right)^s \left(\frac{en}{(1-\eta)ds}\right)^{(1-\eta)ds} \left(\frac{(1-\eta)ds + s}{n}\right)^{ds} \\ &\leq \left(e^{(1-\eta)d+1}(1-\eta)d \left(\frac{(1-\eta)d+1}{(1-\eta)d}\right)^d \left(\frac{(1-\eta)ds}{n}\right)^{\eta d-1}\right)^s \\ &\leq \left(e^d d e^{1/(1-\eta)} \left(\frac{ds}{n}\right)^{\eta d-1}\right)^s \\ &\leq \left(4e^d d \left(\frac{ds}{n}\right)^{\eta d-1}\right)^s. \end{aligned}$$

To bound the above we let $\varepsilon = e^{-1.1/\eta}/d$, and we do a case analysis according to the value of s . For $3 \log_2 n \leq s \leq e^{-1.1/\eta}n/d$, we have

$$\left(4e^d d \left(\frac{ds}{n}\right)^{\eta d-1}\right)^s \leq (4e^d d e^{-1.1d} e^{1.1/\eta})^s \leq 2^{-s} \leq n^{-2}.$$

For $1 \leq s \leq 3 \log_2 n$, we have

$$\left(4e^d d \left(\frac{ds}{n}\right)^{\eta d-1}\right)^s \leq \left(4e^d d \left(\frac{3 \log_2 n}{\sqrt{n}}\right)^{\eta d-1}\right)^s$$

$$\leq \left(4d \left(\frac{e^{1/\eta} 3 \log_2 n}{\sqrt{n}} \right)^{\eta d - 1} \right)^s \leq \left(d \left(\frac{1}{n^{6/16}} \right)^{\eta d - 1} \right)^s \leq n^{-2}.$$

Summing the above over $s = 1, \dots, e^{-1.1/\eta} n/d$, we have

$$\mathbb{P}_{\{G_S\}_{S \subset V} \sim \widetilde{\mathcal{G}}_{n,d}} [\forall S \subset V, |S| \leq e^{-1.1/\eta} n/d \implies |\Gamma_{G_S}(S)| \geq (1 - \eta)d|S|] \geq 1 - O(n^{-1}).$$

So, conditioned on $\mathcal{B}$, by [Eq. \(A.1\)](#), $\forall S \subset V, |S| \leq e^{-1.1/\eta} n/d$ we have

$$|\Gamma_G(S)| \geq |\Gamma_{G_S}(S)| - \eta d|S|.$$

Since $\Pr_{\mathcal{G}_{n,d}}[\mathcal{B}] \geq 1 - O(n^{-0.4})$, we conclude

$$\mathbb{P}_{G \sim \mathcal{G}_{n,d}} [\forall S \subset V, |S| \leq e^{-1.1/\eta} n/d \implies |\Gamma_G(S)| \geq (1 - 2\eta)d|S|] \geq 1 - O(n^{-0.4}). \quad \blacksquare$$

A.3. Summary

Restatement of [Theorem 3.5](#). *The random graph constructed in [Section A.1](#) is a proper expander (see [Definition 3.4](#)), asymptotically almost surely. Specifically, the probability the graph has the desired properties is $\geq 1 - n^{-O(1)}$.*

Proof: Property (P1) is proved in [Lemma A.4](#). Property (P2) is proved in [Lemma A.6](#). Property (P3) follows from [Theorem A.1](#). ■