

Locality-Sensitive Orderings and Applications to Reliable Spanners*

Arnold Filtser^{†1} and Hung Le^{‡2}

¹Bar-Ilan University

²University of Massachusetts at Amherst

Abstract

Chan, Har-Peled, and Jones [2020] recently developed locality-sensitive ordering (LSO), a new tool that allows one to reduce problems in the Euclidean space $\mathbb{R}^d$ to the 1-dimensional line. They used LSO's to solve a host of problems. Later, Buchin, Har-Peled, and Oláh [2019,2020] used the LSO of Chan *et al.* to construct very sparse *reliable spanners* for the Euclidean space. A highly desirable feature of a reliable spanner is its ability to withstand a massive failure: the network remains functioning even if 90% of the nodes fail. In a follow-up work, Har-Peled, Mendel, and Oláh [2021] constructed reliable spanners for general and topologically structured metrics. Their construction used a different approach, and is based on sparse covers.

In this paper, we develop the theory of LSO's in non-Euclidean metrics by introducing new types of LSO's suitable for general and topologically structured metrics. We then construct such LSO's, as well as constructing considerably improved LSO's for doubling metrics. Afterwards, we use our new LSO's to construct reliable spanners with improved stretch and sparsity parameters. Most prominently, we construct $\tilde{O}(n)$ -size reliable spanners for trees and planar graphs with the optimal stretch of 2. Along the way to the construction of LSO's and reliable spanners, we introduce and construct ultrametric covers, and construct 2-hop reliable spanners for the line.

*This paper appeared in arXiv under the title “Reliable Spanners: Locality-Sensitive Orderings Strike Back”. We changed the title in order to emphasize that the main theme of the paper is Locality-Sensitive Orderings.

[†]Email: arnold273@gmail.com.

[‡]Email: hungle@cs.umass.edu. The research was supported by the start-up grant of Umass Amherst and by the National Science Foundation under Grant No. CCF-2121952.

Contents

1	Introduction	1
1.1	Our contribution	3
1.2	Related work	8
2	Preliminaries	9
3	Ultrametric Covers	10
3.1	From Pairwise Partition Cover to Ultrametric Cover: Proof of Lemma 3	12
3.2	Pairwise Partition Cover for General Metrics: Proof of Lemma 4	14
3.3	Ultrametric cover for doubling spaces	15
4	Locality Sensitive Ordering	17
5	Triangle Locality Sensitive Ordering	19
5.1	Proof of Theorem 4	20
5.2	2-hop spanner: Proof of Lemma 12	21
6	Left-sided Locality Sensitive Ordering	23
6.1	k -separable graphs	24
6.2	SPD Graphs	25
6.3	Reliable Spanners from Left-sided LSO	26
7	Subgraph Reliable Connectivity Preserves	28
8	Conclusions	30

1 Introduction

The Algorithmist’s toolkit consists of diverse “tools” frequently utilized for many different problems. In the geometric context, some tools apply to general metric spaces such as metric embeddings [Bou85, FRT04] and padded decompositions [KPR93, Bar96, Fil19a], while many tools apply mainly to Euclidean spaces, such as dimension reduction [JL84], locality-sensitive hashing [IM98], well-separated pair decomposition (WSPD) [CK95], and many others. Recently, Chan, Har-Peled, and Jones [CHJ20] developed a new and exciting tool for Euclidean spaces called *Locality-Sensitive Ordering* (LSO).

Definition 1 $((\tau, \rho)$ -LSO). *Given a metric space (X, d_X) , we say that a collection Σ of orderings is a (τ, ρ) -LSO (locality-sensitive ordering) if $|\Sigma| \leq \tau$, and for every $x, y \in X$, there is a linear ordering $\sigma \in \Sigma$ such that (w.l.o.g.) $x <_\sigma y$ and the points between x and y w.r.t. σ could be partitioned into two consecutive intervals I_x, I_y where $I_x \subseteq B_X(x, \rho \cdot d_X(x, y))$ and $I_y \subseteq B_X(y, \rho \cdot d_X(x, y))$. Parameter ρ is called the stretch parameter.*

The main reason that LSO has become an extremely useful tool is that it reduces the problem at hand in the d -dimensional Euclidean space to the same problem in a much simpler space: the 1-dimensional line. [CHJ20] constructed an $(O(\epsilon)^{-d} \log \frac{1}{\epsilon}, \epsilon)$ -LSO for any given set of points in the d -dimensional Euclidean space $\mathbb{R}^d$ (more generally, Chan *et al.* [CHJ20] constructed $(O(\epsilon^{-1} \cdot \log n)^{O(d)}, \epsilon)$ -LSO for metric spaces with doubling dimension¹ d). They used their LSO to design simple dynamic algorithms for approximate nearest neighbor search, approximate bichromatic closest pair, approximate MST, spanners, and fault-tolerant spanners. Afterwards, Buchin, Har-Peled, and Oláh [BHO19, BHO20] used the LSO of Chan *et al.* [CHJ20] to construct *reliable spanners* (see Definition 2) for Euclidean spaces following the same methodology: reducing the problem to the construction on the line. In this work, we introduce new notions of LSO and apply them to construct reliable spanners for non-Euclidean metrics.

Given a metric space (X, d_X) , a t -spanner is a weighted graph $H = (X, E, w)$ over² X where for every pair of points $x, y \in X$, $d_X(x, y) \leq d_H(x, y) \leq t \cdot d_X(x, y)$, with d_H being the shortest path metric of H . The parameter t is called the stretch of the spanner. A highly desirable property of a t -spanner is the ability to withstand extensive vertex failures. Levcopoulos, Narasimhan, and Smid [LNS02] introduced the notion of a fault-tolerant spanner. A subgraph $H = (V, E_H, w)$ is an f -vertex-fault-tolerant t -spanner of a weighted graph $G = (V, E, w)$, if for every set $F \subset V$ of at most f vertices, it holds that $\forall u, v \notin F$, $d_{H \setminus F}(u, v) \leq t \cdot d_{G \setminus F}(u, v)$. A major limitation of fault-tolerant spanners is that the number of failures must be determined in advance; in particular, such spanners cannot withstand a massive failure. One can imagine a scenario where a significant portion (even 90%) of a network fails and ceases to function (due to, e.g., close-down during a pandemic), it is important that the remaining parts of the network (or at least most of it) will remain highly connected and functioning. To this end, Bose *et al.* [BDMS13] introduced the notion of a *reliable spanner*. Here, given a failure set $B \subseteq X$, the residual spanner $H \setminus B$ is a t -spanner for $X \setminus B^+$, where $B^+ \supseteq B$ is a set slightly larger than B . Buchin *et al.* [BHO20] relaxed the notion of reliable spanners by allowing the size of B^+ to be bounded only in expectation.

¹A metric space (X, d) has doubling dimension d if every ball of radius $2r$ can be covered by 2^d balls of radius r .

²Often in the literature, the metric space (X, d_X) is the shortest path metric of a graph G , and there is a requirement that H will be a subgraph of G . We will not have such a requirement in this paper.

Definition 2 (Reliable spanner). *A weighted graph H over point set X is a deterministic ν -reliable t -spanner of a metric space (X, d_X) if d_H dominates³ d_X , and for every set $B \subseteq X$ of points, called an attack set, there is a set $B^+ \supseteq B$, called a faulty extension of B , such that:*

1. $|B^+| \leq (1 + \nu)|B|$.
2. For every $x, y \notin B^+$, $d_{H[X \setminus B]}(x, y) \leq t \cdot d_X(x, y)$.

An oblivious ν -reliable t -spanner is a distribution $\mathcal{D}$ over dominating graphs H , such that for every attack set $B \subseteq X$ and $H \in \text{supp}(\mathcal{D})$, there exist a superset B^+ of B such that, for every $x, y \notin B^+$, $d_{H[X \setminus B]}(x, y) \leq t \cdot d_X(x, y)$, and $\mathbb{E}_{H \sim \mathcal{D}}[|B^+|] \leq (1 + \nu)|B|$. We say that the oblivious spanner $\mathcal{D}$ has m edges if every graph $H \in \text{supp}(\mathcal{D})$ has at most m edges.

We call the distribution $\mathcal{D}$ in Definition 2 an oblivious ν -reliable t -spanner because the adversary is oblivious to the specific spanner produced by the distribution (it may be aware to the distribution itself).

For constant dimensional Euclidean spaces, Bose *et al.* [BDMS13] constructed a deterministic reliable $O(1)$ -spanner, such that for every attack B , the faulty extension B^+ contains at most $O(|B|^2)$ vertices. The construction of reliable spanners where the size of B^+ is a linear function of B was left as an open question. For every $\nu, \epsilon \in (0, 1)$, and n points in d -dimensional Euclidean space $(\mathbb{R}^d, \|\cdot\|_2)$, Buchin *et al.* [BHO19] used the LSO of Chan *et al.* [CHJ20] to construct a *deterministic* ν -reliable $(1 + \epsilon)$ -spanner with $n \cdot \nu^{-6} \cdot \tilde{O}(\epsilon)^{-7d} \cdot \tilde{O}(\log n)$ edges (see also [BCDM18]). Later, for the *oblivious* case, Buchin *et al.* [BHO20] applied the same LSO to construct an oblivious ν -reliable $(1 + \epsilon)$ -spanner with $n \cdot \tilde{O}(\epsilon)^{-2d} \cdot \tilde{O}(\nu^{-1}(\log \log n)^2)$ edges.

Very recently, Har-Peled, Mendel, and Oláh [HMO21] constructed reliable spanners for general metric spaces, as well as for topologically structured spaces (e.g. trees and planar graphs). They showed that for every integer k , every general n -point metric space admits an *oblivious* ν -reliable $(512 \cdot k)$ -spanner⁴ with $n^{1+1/k} \cdot O(\nu^{-1}k \log^2 \Phi \log \frac{n}{\nu})$ edges, where $\Phi = \frac{\max_{x,y} d_X(x,y)}{\min_{x,y} d_X(x,y)}$ is the *aspect ratio* of the metric space (also known as the *spread*, which a priori is unbounded). Additionally, they showed that ultrametrics (see Definition 3) admit oblivious ν -reliable $(2 + \epsilon)$ -spanners with $n \cdot \tilde{O}(\nu^{-1}\epsilon^{-2} \log^2 \Phi)$ edges, tree metrics admit oblivious ν -reliable $(3 + \epsilon)$ -spanners with $n \cdot \tilde{O}(\nu^{-1}\epsilon^{-2} \log^2 n \log^2 \Phi)$ edges, and planar metrics admit oblivious ν -reliable $(3 + \epsilon)$ -spanners with $n \cdot \tilde{O}(\nu^{-1}\epsilon^{-4} \log^2 \Phi)$ edges (see Table 2).

The reliable spanner constructions of Har-Peled *et al.* [HMO21] are based on *sparse covers*. A (τ, ρ) -sparse cover is a collection $\mathcal{C}$ of clusters such that every point belongs to at most τ clusters, and for every pair $x, y \in X$, there is a cluster $C \in \mathcal{C}$ containing both $x, y \in C$ where $\text{diam}(C) \leq \rho \cdot d_X(x, y)$; ρ is called the *stretch* of the cover $\mathcal{C}$. They then treat each cluster in $\mathcal{C}$ as a uniform metric, construct a reliable spanner for each cluster, and return the union of all the constructed spanners. Thus the main task becomes constructing a reliable spanner for the uniform metric. Specifically, instead of the oblivious ν -reliable 1-spanner for the line constructed in [BHO20], Har-Peled *et al.* [HMO21] constructed an oblivious ν -reliable 2-spanner for the uniform metric, which is the best stretch possible for subquadratic size spanners (see Observation 1). Indeed, this additional factor 2 appears in the stretch parameter in all the spanners in [HMO21]. Most prominently, for trees they

³Metric space (X, d_H) dominates metric space (X, d_X) if $\forall u, v \in X, d_X(u, v) \leq d_H(u, v)$.

⁴[HMO21] did not compute the constant explicitly. Their construction is based on the Ramsey-trees of Mendel-Naor [MN07], which have stretch $128k$. Using state of the art Ramsey trees [NT12] of stretch $2ek$ instead (see also [ACE⁺20]), the approach of [HMO21] provides stretch $8ek$.

constructed an $(O(\epsilon^{-1} \log \Phi \log n), 2 + \epsilon)$ -sparse cover, resulting in a stretch $4 + \epsilon$ spanner,⁵ while the natural lower bound is 2 (Observation 1). A similar phenomenon occurs for planar graphs. An additional drawback in the sparse cover based approach of [HMO21] is its dependency on the aspect ratio Φ (which a priori can be unbounded). This dependency on the aspect ratio is inherent in their technique and cannot be avoided (see Lemma 20 in [HMO21]).

1.1 Our contribution

Our major contribution is to the theory of locality-sensitive orderings. Specifically, we significantly improve the parameters of LSO in doubling metrics¹, and extend the idea of LSO to general metrics, as well as to topologically structured metrics. This is done by introducing left-sided LSO and triangle-LSO (see Table 1). LSO's are a powerful tool enabling one to reduce many problems to the line. LSO's already have many applications in computational geometry [CHJ20]; we expect that our LSO for doubling metrics, as well as those for general and topologically structured graphs, will find many additional applications in the future. Next, we use these newly introduced LSO's (or improved in the case of doubling) to construct oblivious reliable spanners. Our constructions have smaller stretch (optimal in the case of topologically structured metrics) and smaller sparsity (see Table 2). Below we describe each type of LSO in detail, and which spanners it was used to construct. Our constructions of LSO for general and doubling metrics are going through the construction of *ultrametric covers*. An ultrametric cover is a collection of dominating ultrametrics such that the distance between every pair of points is well approximated by some ultrametric in the collection. We construct the first ultrametric cover for doubling metrics with stretch $1 + \epsilon$ (previously only tree covers were known), and improve the stretch parameter in the ultrametric covers of general metrics (see Table 3). Finally, a crucial ingredient when one constructs reliable spanners using LSO is reliable spanners for the line. Buchin *et al.* [BHO19, BHO20] constructed such spanners; however their spanners have $\Omega(\log n)$ hops, which will incur additional $\log n$ factor in the stretch (in all cases other than Euclidean/doubling). To avoid this overhead, we construct a 2-hop reliable spanner for the line, and a 2-hop left spanner, which is a newly defined type of spanner suitable for our left-LSO (see Table 4). See Figure 1 for a graphic illustration of how the different parts in the paper are related. Finally, we answer an open question by Har-Peled [Har20] regarding sub-graph reliable spanners, by providing matching upper and lower bounds for reliable connectivity preservers.

Classic LSO. Chan *et al.* [CHJ20] constructed a $((\epsilon^{-1} \cdot \log n)^{O(d)}, \epsilon)$ -LSO for metric spaces of doubling dimension d . Applying the (implicit) framework of [BHO19, BHO20] yields a reliable spanner with $n \cdot (\epsilon^{-1} \cdot \log n)^{O(d)}$ edges. In this work, we completely remove the dependency on n of the number of orderings. Specifically, we construct an $(\epsilon^{-O(d)}, \epsilon)$ -LSO for doubling metrics (Corollary 1); this immediately implies reliable spanners for metric spaces of doubling dimension d with the same performance, up to the dependency on ϵ , as for Euclidean spaces (Corollary 2).

Triangle LSO. A (τ, ρ) -triangle LSO for a metric space (X, d_X) is a collection Σ of at most τ linear orderings over X , such that for every $x, y \in X$, there is an ordering $\sigma \in \Sigma$ such that for every two points $z, q \in X$ satisfying $x \leq_\sigma z \leq_\sigma q \leq_\sigma y$ (or $y \leq_\sigma z \leq_\sigma q \leq_\sigma x$) it holds that $d_X(z, q) \leq \rho \cdot d_X(x, y)$

⁵With an additional effort, [HMO21] reduced the stretch of the spanner to $3 + \epsilon$. This analysis is tight, and their technique cannot give a reliable spanner with a stretch factor smaller than 3.




Figure 1: Relationships between different concepts; new concepts introduced in this papers are green-shaded.

LSO type	Metric Space	# of orderings (τ)	Stretch	Ref
(Classic) LSO	Euclidean space $\mathbb{R}^d$	$O(\epsilon)^{-d} \cdot \log \frac{1}{\epsilon}$	ϵ	[CHJ20]
	Doubling dimension d	$O(\epsilon^{-1} \cdot \log n)^{O(d)}$	ϵ	[CHJ20]
		$\epsilon^{-O(d)}$	ϵ	Corollary 1
Triangle-LSO	General metric	$\tilde{O}(n^{1/k} \cdot \epsilon^{-1})$	$2k + \epsilon$	Corollary 3
	Ultrametric	1	1	Lemma 11
Left-sided LSO	Tree	$\log n$	1	Theorem 6
	Treewidth k	$k \cdot \log n$	1	Theorem 6
	Planar graph	$\frac{1}{\epsilon} \cdot \log^2 n$	$1 + \epsilon$	Theorem 7
	Minor Free	$\frac{1}{\epsilon} \cdot \log^2 n$	$1 + \epsilon$	Theorem 7

Table 1: Summery of all known results, on all the different types of locality sensitive orderings (LSO). $k \in \mathbb{N}$ is an integer, $\epsilon \in (0, 1)$ is an arbitrarily small parameter.

(Definition 6). Note that every (τ, ρ) -triangle LSO is also a (τ, ρ) -LSO; however, a (τ, ρ) -LSO is only a $(\tau, 2\rho + 1)$ -triangle LSO (by the triangle inequality). Hence a triangle-LSO is preferable to the classic LSO. We observe that ultrametrics admit a $(1, 1)$ -triangle-LSO (Lemma 11), and show that general n -point metric spaces admit an $(\tilde{O}(n^{\frac{1}{k}} \cdot \epsilon^{-1}), 2k + \epsilon)$ -triangle-LSO (Corollary 3).

We then prove a meta-theorem stating that every metric space admitting a (τ, ρ) -triangle LSO has an oblivious ν -reliable 2ρ -spanner with $n\tau \cdot O(\log^2 n + \nu^{-1}\tau \log n \cdot \log \log n)$ edges (Theorem 4). This gives oblivious reliable spanners for ultrametrics and general metric spaces. Our spanners for general metrics have significantly smaller stretch compared to [HMO21] ($8k$ compared to $512k$ ⁴), this constant is highly important as it governs the parameter in the power of n . An additional advantage is that we remove the dependency on the aspect ratio (which a priori can be unbounded).

Left-sided LSO. A (τ, ρ) -left-sided LSO for a metric space (X, d_X) is a collection Σ of linear orderings over *subsets* of X , called *partial orderings*, such that every point x belongs to at most τ partial orderings, and for every $x, y \in X$, there is a partial ordering $\sigma \in \Sigma$ such that for every two

Family	stretch	guarantee	size	ref
Euclidean ($\mathbb{R}^d, \ \cdot\ _2$)	$O(1)$	Deterministic	$\Omega(n \log n)$	[BDMS13]
	$1 + \epsilon$	Deterministic	$n \cdot \tilde{O}(\epsilon)^{-7d} \nu^{-6} \cdot \tilde{O}(\log n)$	[BHO19]
	$1 + \epsilon$	Oblivious	$n \cdot \tilde{O}(\epsilon)^{-2d} \cdot \tilde{O}(\nu^{-1}(\log \log n)^2)$	[BHO20]
Doubling dimension d	$1 + \epsilon$	Deterministic	$n \cdot \epsilon^{-O(d)} \nu^{-6} \cdot \tilde{O}(\log n)$	Corollary 2
	$1 + \epsilon$	Oblivious	$n \cdot \epsilon^{-O(d)} \nu^{-1} \log \nu^{-1} \cdot \tilde{O}(\log \log n)^2$	Corollary 2
General metric	$512 \cdot k$ ⁴	Oblivious	$\tilde{O}(n^{1+1/k} \cdot \nu^{-1}) \cdot \log^2 \Phi$	[HMO21]
	$8k + \epsilon$	Oblivious	$\tilde{O}(n^{1+1/k} \cdot \epsilon^{-2}) \cdot \nu^{-1}$	Theorem 5
Ultrametric, Tree, planar	k	Deterministic	$\Omega(n^{1+1/k})$	[HMO21]
	$k < 2$	Oblivious	$\Omega(n^2)$	Observation 1
Ultrametric	$2 + \epsilon$	Oblivious	$n \cdot \tilde{O}(\nu^{-1} \epsilon^{-2} \log^2 \Phi)$	[HMO21]
	2 (tight)	Oblivious	$n \cdot \tilde{O}(\log^2 n + \nu^{-1} \log n)$	Corollary 4
Tree	$3 + \epsilon$	Oblivious	$n \cdot \tilde{O}(\nu^{-1} \epsilon^{-2} \log^2 n \log^2 \Phi)$	[HMO21]
	2 (tight)	Oblivious	$n \cdot O(\nu^{-1} \log^3 n)$	Theorem 9
Treewidth k	2 (tight)	Oblivious	$n \cdot O(\nu^{-1} k^2 \log^3 n)$	Theorem 9
Planar	$3 + \epsilon$	Oblivious	$n \cdot \tilde{O}(\nu^{-1} \epsilon^{-4} \log^2 n \log^2 \Phi)$	[HMO21]
	$2 + \epsilon$ (tight)	Oblivious	$n \cdot O(\nu^{-1} \epsilon^{-2} \log^5 n)$	Theorem 10
Minor-free	$2 + \epsilon$ (tight)	Oblivious	$n \cdot O(\nu^{-1} \epsilon^{-2} \log^5 n)$	Theorem 10

Table 2: Comparison between previous and new constructions of reliable spanners. All spanners (except [BDMS13]) constructed on n -point metric spaces with reliability parameter ν . For doubling metrics, we recover the same strong results previously known only for Euclidean space (up to a polynomial dependence on ϵ). Both lower bounds hold for the uniform metric (which is sub-metric of a star metric). For all other metric spaces, we improve both stretch and sparsity, and remove the undesirable dependence on the aspect ratio (spread) Φ . Most notably, for trees and planar graphs, the stretch was improved from $3 + \epsilon$, to the best possible stretch 2. Finally, for general graphs, our spanner has stretch $8k$, considerably improving the constant hiding in [HMO21]. This constant is highly important as it governs the parameter in the power of n .

points $x', y' \in X$ satisfying $x \leq_\sigma x'$ and $y' \leq_\sigma y$, it holds that $d_X(x', y') \leq \rho \cdot d_X(x, y)$ (Definition 8). Note that the *stretch guarantee* of a (τ, ρ) -left-sided LSO implies that of a (τ, ρ) -LSO (but not the vice versa). However, there could be $\Omega(n)$ (partial) orderings in a (τ, ρ) -left-sided LSO. By lifting the restriction on the total number of partial orderings, we can construct a left-sided LSO with an optimal stretch of 1 or a nearly optimal stretch of $1 + \epsilon$; see Table 1. This small stretch ultimately leads to the (nearly) optimal stretch for the reliable spanners of tree and planar metrics constructed in this work, which is not attainable in previous work [HMO21].

We then prove a meta-theorem stating that every metric space admitting a (τ, ρ) -left-sided LSO has an oblivious ν -reliable 2ρ -spanner with $n \cdot O(\nu^{-1} \tau^2 \log n)$ edges (Theorem 8). We show that n -vertex trees admit a $(\log n, 1)$ -left-sided LSO and conclude that trees have oblivious ν -reliable 2-spanners with $n \cdot O(\nu^{-1} \log^3 n)$ edges (Theorem 9). Note that the stretch parameter 2 is optimal. Later, we show that planar graphs admit a $(\frac{1}{\epsilon} \log^2 n, 1 + \epsilon)$ -left-sided LSO for every $\epsilon \in (0, 1)$ (Theorem 7). An oblivious ν -reliable $(2 + \epsilon)$ -spanner with $n \cdot O(\nu^{-1} \epsilon^{-2} \log^5 n)$ edges follows

Space	type	stretch	# of trees	ref
Euclidean $\mathbb{R}^d$	tree	$1 + \epsilon$	$O((\frac{d}{\epsilon})^d \log \frac{d}{\epsilon})$	[ADM ⁺ 95]
Doubling dimension d	ultrametric	$O(d^2)$	$O(d \log d)$	[CGMZ16]
	tree	$1 + \epsilon$	$\epsilon^{-O(d)}$	[BFN19b]
	ultrametric	$1 + \epsilon$	$\epsilon^{-O(d)}$	Theorem 2
General metric	ultrametric	$2e \cdot k$	$O(k \cdot n^{1/k})$	[MN07, NT12]
	ultrametric	$2k + \epsilon$	$\tilde{O}(n^{\frac{1}{k}} \cdot \epsilon^{-1})$	Theorem 1

Table 3: New and previous constructions of tree and ultrametric covers.

(Theorem 10). The same results also hold for bounded treewidth graphs and graphs excluding a fixed minor.

Ultrametric cover. A (τ, ρ) -tree cover for a metric space (X, d_X) is a set $\mathcal{T}$ of τ dominating trees³ such that the distance between every pair of points is preserved up to a factor ρ in at least one tree ($\forall u, v, \min_{T \in \mathcal{T}} d_T(u, v) \leq \rho \cdot d_X(u, v)$). When all trees in the cover are ultrametrics, we call it an ultrametric cover (Definition 4). The first study on tree covers was for Euclidean spaces by Arya *et al.* [ADM⁺95] who constructed the so-called Dumbbell trees. For general metrics, Mendel and Naor [MN07] (implicitly) constructed an ultrametric cover from Ramsey type embeddings. These covers actually have a stronger guarantee, where every vertex v is guaranteed to have an ultrametric in the cover approximating its shortest path tree ($\forall v \exists T \forall u, d_T(v, u) \leq \rho \cdot d_X(v, u)$). There is a long line of work on Ramsey-type embeddings [BFM86, BBM06, BLMN05, MN07, NT12, BGS16, ACE⁺20, FL21, Bar21, Fil21]. The state of the art covers follow from Naor and Tao [NT12], and implies a $(2e \cdot k, O(k \cdot n^{1/k}))$ ultrametric cover. For doubling metrics, Bartal, Fandina, and Neiman [BFN19a] constructed a $(1 + \epsilon, \epsilon^{-O(d)})$ tree cover. We refer to [BFN19a] for further results and background on tree covers.

We observe that every ultrametric admits a $(1, 1)$ -triangle LSO, which implies that given a (τ, ρ) -ultrametric cover, one can construct a (τ, ρ) -triangle LSO (Lemma 11). Indeed, the main step in our construction of a triangle LSO for general metrics is a construction of a $(\tilde{O}(n^{\frac{1}{k}} \cdot \epsilon^{-1}), 2k + \epsilon)$ -ultrametric cover (Theorem 1). Our construction provides a constant improvement in the stretch parameter (equivalently, a polynomial improvement in the number of ultrametrics in the cover) compared to previous results.

A more structured case is that of a (τ, ρ, k, δ) -ultrametric cover, where in addition to being a (τ, ρ) -ultrametric cover, we require that each ultrametric will be a k -HST of degree at most δ (see Definition 3). We show that every $\Omega(\frac{1}{\epsilon})$ -HST of degree bounded by δ admits a (classic) $(\frac{\delta}{2}, \epsilon)$ -LSO (Lemma 9). It follows that a $(\tau, \rho, \Omega(\frac{1}{\epsilon}), \delta)$ -ultrametric cover implies a $(\tau \cdot \frac{\delta}{2}, (1 + \epsilon)\rho)$ -LSO (Lemma 10). The trees in the tree cover for doubling metrics of [BFN19a] are far from being ultrametrics and cannot be used in our framework. We then construct an $(\epsilon^{-O(d)}, 1 + \epsilon, \frac{1}{\epsilon}, \epsilon^{-O(d)})$ -ultrametric cover for spaces with doubling dimension d (Theorem 2), which implies the respective LSO. Interestingly, having such an ultrametric cover is a characterizing property for metric spaces of bounded doubling dimension (Theorem 2). See Table 3 for a summary.

Type	guarantee	size	hops	ref
1 spanner	Deterministic	$n \cdot O(\log n \cdot \nu^{-6})$	$O(\log n)$	[BHO19]
	Oblivious	$n \cdot O(\nu^{-1} \cdot \log \nu^{-1})$	$O(\log n)$	[BHO20]
	Oblivious	$n \cdot O(\log^2 n + \nu^{-1} \log n \cdot \log \log n)$	2	Lemma 12
Left spanner	Oblivious	$n \cdot O(\nu^{-1} \log n)$	2	Lemma 16

Table 4: Construction of reliable spanners for the line. [BHO19] and [BHO20] constructed sparse (both deterministic and oblivious) reliable 1-spanners for points on the line. However, their spanners have $O(\log n)$ hops, which will incur distortion $O(\rho \cdot \log n)$ when applied on a (τ, ρ) -triangle LSO (with $\rho > 1$). We construct a 1-spanner with only 2-hops, which we later use to construct a reliable spanner from a triangle-LSO. In addition, we construct a 2-hop left spanner for the line, which is later used to construct a reliable spanner from a left-sided LSO.

2-hop reliable spanners for the path graph. Using (different types of) LSO, we can reduce the problem of constructing reliable spanners for different complicated metric spaces to that of constructing reliable spanners for the 1-dimensional path graph. Buchin *et al.* [BHO20] constructed an oblivious ν -reliable 1-spanner with $n \cdot \tilde{O}(\nu^{-1}(\log \log n)^2)$ edges for the path graph. However, the shortest path between two given vertices in their spanner could contain $\Omega(\log n)$ edges (called *hops*). While $(\log n)$ -hop spanners are acceptable when applying them upon a (τ, ϵ) -LSO, using a h -hop spanner of the path graph for (τ, ρ) -triangle-LSO will result in distortion $h \cdot \rho$. It is therefore desirable to minimize the number of hops used by the spanner. Having a 1-hop spanner will require $\Omega(n^2)$ edges; we thus settle for the next best thing: a 2-hop reliable spanner. Specifically, we construct an oblivious ν -reliable, 2-hop 1-spanner with $n \cdot O(\log^2 n + \nu^{-1} \log n \cdot \log \log n)$ edges for the path graph (Lemma 12). This spanner is later used in our meta Theorem 4 to construct reliable spanners for metric spaces admitting a (τ, ρ) -triangle-LSO.

For the left-sided LSO case, we also need a 2-hop spanner for the path graph. However, the shortest 2-hop path between the i 'th and the j 'th vertices for $i < j$ must go through a vertex to the left of i , i.e., a vertex in $[1, i]$ (as opposed to a vertex in $[i, j]$ in the triangle-LSO case). This requirement inspires us to define a *left-spanner* (Definition 11); we then construct an oblivious ν -reliable 2-hop left-spanner with $n \cdot O(\nu^{-1} \log n)$ edges (Lemma 16). The left-spanner is later used in meta Theorem 8 to construct a reliable spanner from a left-sided LSO.

Connectivity preservers. While research on reliable spanners for metric spaces has been fruitful, nothing is known for reliable spanners of *graphs*, where we require the spanner to be a subgraph of the input graph. In a recent talk, Har-Peled [Har20] asked a “probably much harder question”: whether it is possible to construct a non-trivial subgraph reliable spanner. We show that, even for a much simpler problem where one seeks a subgraph to only preserve *connectivity* for vertices outside B^+ , the faulty extension of B , the subgraph must have $\Omega(n^2)$ edges in the worst case. Indeed, our lower bound is much more general: it applies to g -reliable connectivity preservers for some function g . A g -reliable connectivity preserver of a graph $G = (V, E)$ is a subgraph H of G such that for every attack $B \subseteq V$, there is a superset $B^+ \supseteq B$ of size at most $g(|B|)$, such that for every $u, v \in V \setminus B^+$, if u and v are connected in $G \setminus B$, then they are also connected in $H \setminus B$. Observe that a ν -reliable spanner defined in Definition 2 is a g -reliable (non-subgraph) spanners for the linear function $g(x) = (1 + \nu)x$. We showed that there is an n -vertex graph G such that every oblivious g_k -reliable connectivity preserver has $\Omega(n^{1+1/k})$ edges for any function $g = O(x^k)$.

(see Theorem 12). Taking $k = 1$ gives a lower bound $\Omega(n^2)$ on the number of edges of *subgraph* ν -reliable spanners. On the positive side, we provide a construction of a *deterministic* connectivity preserver matching the lower bound (Theorem 11).

1.2 Related work

The tradeoff between stretch and sparsity (number of edges) of (regular) t -spanners has been extensively studied [Yao82, Cla87, Kei88, PS89, ADD⁺93, CDNS95, HIS13, CGMZ16, FN18, LS19]; see the recent survey of Ahmed *et al.* [ABS⁺20], and the book [NS07] and references therein for more details. The bottom line is that n -point metric spaces admit $(2k - 1)$ -spanners (for every integer k) with $O(n^{1+1/k})$ edges [CDNS95], while the metric induced by n points in d dimensional Euclidean space admits a $(1 + \epsilon)$ -spanner with $n \cdot O(\epsilon)^{1-d}$ edges [ADD⁺93]. Similarly, n -point metric spaces with doubling dimension d admit $(1 + \epsilon)$ -spanners with $n \cdot \epsilon^{-O(d)}$ edges [CGMZ16].

For vertex-fault-tolerant spanner, it was shown that every n -point set in $\mathbb{R}^d$, or more generally in a space of doubling dimension d , admits an f -vertex-fault-tolerant $(1 + \epsilon)$ -spanner with $\epsilon^{-O(d)} \cdot f \cdot n$ edges [LNS02, Luk99, Sol14].

AF: To do before journal version: Chan *et al.* [CHJ20] observed that an extremely simple construction using LSO's is possible for fault-tolerant spanners: namely take each ordering and add edges to the $f + 1$ predecessors and successors. This given optimal bounds on sparsity and max degree (other constructions much more involved). The same follows now for doubling metrics using our LSO's. Specifically an f -fault tolerant spanner with $n \cdot f \cdot \epsilon^{-O(d)}$ edges and max degree $f \cdot \epsilon^{-O(d)}$. See also [Sol14] for the state of the art. We should state it explicitly and take the dividend....

For general graphs, after a long line of works [CLPR10, DK11, BDPW18, BP19, DR20, BDR21], it was shown that every n -vertex graph admits an efficiently constructible f -vertex-fault-tolerant $(2k - 1)$ -spanner with $O(f^{1-1/k} \cdot n^{1+1/k})$ edges, which is optimal assuming the Erdős' Girth Conjecture [Erd64]. A related notion is that of a vertex-fault-tolerant (VFT) emulator. Unlike spanners, emulators are not required to be subgraphs, and the weight of an emulator edge is determined w.r.t. the faulty set. It was recently shown that vertex-fault-tolerant (VFT) emulators are asymptotically sparser from their spanner counterparts [BDN22].

In addition to vertex-fault-tolerant (VFT) spanners, also edge-fault-tolerant (EFT) spanners were studied, where the guarantee is to withstand up to f -edge faults (as opposed to f vertex faults in VFT). Bodwin, Dinitz, and Robelle [BDR22] constructed an f -EFT $2k - 1$ -spanners with $O(k^2 f^{\frac{1}{2} - \frac{1}{2k}} \cdot n^{1+\frac{1}{k}} + kfn)$ / $O(k^2 f^{\frac{1}{2}} \cdot n^{1+\frac{1}{k}} + kfn)$ edges for odd/even values of k respectively. There is also a lower bound of $\Omega(f^{\frac{1}{2} - \frac{1}{2k}} n^{1+1/k})$ [BDPW18].

Abam *et al.* [AdBFG09] introduced the notion of *region* fault-tolerant spanners for the Euclidean plane. They showed that one can construct a t -spanner with $O(n \log n)$ edges in such a way that if points belonging to a convex region are deleted, the residual graph is still a spanner for the remaining points.

Spanners with low hop diameter for Euclidean spaces of fixed dimension were studied in the pioneering work of Arya *et al.* [AMS94]. State of the art is a $(1 + \epsilon)$ -spanner constructible in $O(n \log n)$ time by Solomon [Sol13] that has $O(n\alpha_k(n))$ ⁶ edges and hop diameter k .

In addition to having a small number of edges, it is desirable to have a spanner with a small total edge weight, called a *light spanner*. Light spanners have been thoroughly studied in the spanner literature [CDNS95, ES13, ENS15, Got15, CW18, FS20, BLW17, BLW19, ADF⁺19, LS19, CFKL20]. Sparse (and light) spanners were constructed efficiently in different computational models such as

⁶ $\alpha_k(n)$ is the inverse function of a very fast growing function at level k of the primitive recursive hierarchy; see [Sol13] for a more formal description.

LOCAL [KPX08], CONGEST [EFN20], streaming [Elk11, Bas08, BS07, KW14, AGM12, FKN20], massive parallel computation (MPC) [BDG⁺20] and dynamic graph algorithms [BKS12, BFH19].

2 Preliminaries

Let (X, d_X) be a metric space. The aspect ratio, or spread, denoted by Φ , is defined as follows: $\Phi = \frac{\max_{x,y \in X} d_X(x,y)}{\min_{x \neq y \in X} d_X(x,y)}$. We denote by $[n]$ the set of integers $\{1, 2, \dots, n\}$. For two integers $a \leq b$, we define $[a : b] = \{a, a+1, \dots, b\}$.

We use $\tilde{O}$ notation hides poly-logarithmic factors. That is $\tilde{O}(f) = O(f) \cdot \log^{O(1)}(f)$.

Let G be a graph. We denote the vertex set and edge set of G by $V(G)$ and $E(G)$, respectively. When we want to explicitly specify the vertex set V and edge set E of G , we write $G = (V, E)$. If G is a weighted graph, we write $G = (V, E, w)$ with $w : E \rightarrow \mathbb{R}_+$ being the weight function on the edges of G . For every pair of vertices $x, y \in V$, we denote by $d_G(x, y)$ the shortest path distance between x and y in $G = (V, E, w)$. Given a path $P \subseteq G$, we define the *hop length* of P to be the number of edges on the path.

A t -*spanner* for a metric space (X, d_X) is a weighted graph $H(V, E, w)$ that has $V = X$, $w(u, v) = d_X(u, v)$ for every edge $(u, v) \in E$ and $d_X(x, y) \leq d_H(x, y) \leq t \cdot d_X(x, y)$ for every pair of points $x, y \in X$. We say that a t -spanner H has hop number h if for every pair of vertices x, y , there is an $x - y$ path P in H of at most h hops such that $w_H(P) \leq t \cdot d_X(x, y)$.

The path graph P_n contains n vertices $v_1, v_2, \dots, v_n$ and there is (unweighted) edge between v_i and v_j iff $|i - j| = 1$. A path $v_{i_1}, v_{i_2}, \dots, v_{i_s}$ is monotone iff for every j , $i_j < i_{j+1}$. Note that if a spanner H contains a monotone path between v_i, v_j then $d_H(v_i, v_j) = d_{P_n}(v_i, v_j) = |i - j|$. We sometimes identify vertices of P_n with numbers in $\{1, 2, \dots, n\}$, and refer to $\{1, 2, \dots, n\}$ as the vertex set of P_n .

A metric (X, d_X) has doubling dimension d if every ball of radius r can be covered by at most 2^d balls of radius $r/2$. The following lemma gives the standard packing property of doubling metrics (see, e.g., [GKL03]).

Lemma 1 (Packing Property). *Let (X, d) be a metric space with doubling dimension d . If $S \subseteq X$ is a subset of points with minimum interpoint distance r that is contained in a ball of radius R , then $|S| = \left(\frac{2R}{r}\right)^{O(d)}$.*

In the following lemma, we show that when constructing oblivious spanners, it is enough to bound the number of edges in expectation to obtain a worst-case guarantee.

Lemma 2. *Consider an n -vertex graph $G = (V, E, w)$ that admits an oblivious ν -reliable t -spanner with m edges in expectation. Then G admits an oblivious 2ν -reliable t -spanner with $2m$ edges in the worst case.*

Proof. Formally, there is a distribution $\mathcal{D}$ over spanners H such that for every attack $B \subseteq V$, $\mathbb{E}[|B^+ \setminus B|] \leq \nu|B|$, and $\mathbb{E}[|H|] \leq m$. Let $\mathcal{D}'$ be the distribution over spanners H obtained by conditioning $\mathcal{D}$ on the event $|H| \leq 2m$. Clearly, all the spanners in $\text{supp}\{\mathcal{D}'\}$ have at most $2m$

edges. Furthermore, for every attack $B \subseteq V$, it holds that

$$\begin{aligned}\mathbb{E}_{H \sim \mathcal{D}'}[|B^+ \setminus B|] &= \mathbb{E}_{H \sim \mathcal{D}}[|B^+ \setminus B| \mid |H| \leq 2m] \\ &= \frac{1}{\Pr[|H| \leq 2m]} \cdot (\mathbb{E}_{H \sim \mathcal{D}'}[|B^+ \setminus B| \mid |H| \leq 2m] \cdot \Pr[|H| \leq 2m]) \\ &\leq \frac{1}{\Pr[|H| \leq 2m]} \cdot \mathbb{E}_{H \sim \mathcal{D}}[|B^+ \setminus B|] \leq 2\nu \cdot |B| ,\end{aligned}$$

where in the last inequality, we use Markov's inequality. $\square$

3 Ultrametric Covers

Ultrametric. An ultrametric (X, d) is a metric space satisfying a strong form of the triangle inequality, that is, for all $x, y, z \in X$, $d(x, z) \leq \max\{d(x, y), d(y, z)\}$. A related notion is a k -hierarchical well-separated tree (k -HST).

Definition 3 (k -HST). *A metric (X, d_X) is a k -hierarchical well-separated tree (k -HST) if there exists a bijection φ from X to leaves of a rooted tree T in which:*

1. *Each node $v \in T$ is associated with a label Γ_v such that $\Gamma_v = 0$ if v is a leaf and $\Gamma_v \geq k\Gamma_u$ if v is an internal node and u is any child of v .*
2. *$d_X(x, y) = \Gamma_{\text{lca}(\varphi(x), \varphi(y))}$ where $\text{lca}(u, v)$ is the least common ancestor of any two given nodes u, v in T .*

It is well known that any ultrametric is a 1-HST, and any k -HST is an ultrametric (see [BLMN05]).

Ultrametric cover. Consider a metric space (X, d_X) , a distance measure d_Y is said to be dominating if $\forall x, y \in X$, $d_X(x, y) \leq d_Y(x, y)$. A tree/ultrametric over X is said to be dominating if their metric is dominating. Bartal, Fandina, and Neiman [BFN19a] studied *tree covers*: a metric space (X, d_X) admits a (τ, ρ) -tree cover if there are at most τ dominating trees $\{T_1, T_2, \dots, T_\tau\}$ such that $X \subseteq V(T_i)$ for every $i \in [\tau]$ and for every pair of points $x, y \in X$, there is some tree T_i where $d_{T_i}(x, y) \leq \rho \cdot d_X(x, y)$. Bartal *et al.* [BFN19a] observed that the previous constructions of Ramsey trees⁷ [MN07, NT12, ACE⁺20] give an $(\tilde{O}(n^{1/k}), 2ek)$ -tree cover for general metrics, and explicitly constructed an $(\epsilon^{-O(d)}, 1+\epsilon)$ -tree cover for metric spaces with doubling dimension d . Here we initiate the study of ultrametric covers.

Definition 4 (Ultrametric Cover). *A (τ, ρ) -ultrametric cover for a space (X, d) is a collection of at most τ dominating ultrametrics $\mathcal{U} = \{(U_i, d_{U_i})\}_{i=1}^\tau$ over X , such that for every $x, y \in X$ there is an ultrametric U_i for which $d_{U_i}(x, y) \leq \rho \cdot d_X(x, y)$.*

If every metric $(U, d_U) \in \mathcal{U}$ is a k -HST and the corresponding tree T_U of U has maximum degree at most δ , we say that $\mathcal{U}$ is a (τ, ρ, k, δ) -ultrametric cover of (X, d_X) .

Note that ultrametrics are much more structured than general trees. For example, every ultrametric embeds isometrically into ℓ_2 , while trees require distortion $\sqrt{\log \log n}$ [Bou86, Mat99]. Later,

⁷Ramsey trees have additional desired property compared to general tree covers: for every vertex v , there is a single tree in the cover satisfying all its pairwise distances, as oppose to union of all the trees in a general tree cover.

we will show how to use ultrametric covers to construct locality-sensitive orderings (see Lemma 10 and Lemma 11).

The first main result of this section is Theorem 1 where we construct an ultrametric cover for general metrics.

Theorem 1 (Ultrametric Cover For General Metrics). *For every $k \in \mathbb{N}$ and $\epsilon \in (0, \frac{1}{2})$, every n -point metric space admits an $(O(n^{\frac{1}{k}} \cdot \log n \cdot \frac{k^2}{\epsilon} \cdot \log \frac{k}{\epsilon}), 2k + \epsilon)$ -ultrametric cover.*

Interestingly, the tree cover in [BFN19a] for general metrics actually consists of ultrametrics; in other words, Bartal *et al.* [BFN19a] observed that Ramsey trees constitute an $(\tilde{O}(n^{1/k}), 2ek)$ -ultrametric cover. Thus, we obtain a polynomial improvement in the number of ultrametrics in the cover. Specifically, to guarantee stretch $\approx 2(k+1)$, our cover uses $\tilde{O}(n^{1/k})$ ultrametrics, while previous covers have $\Omega(n^{\frac{e}{k+1}})$ ultrametrics.

Next, in Theorem 2 below, we show that every metric space with doubling dimension d admits an $(\epsilon^{-O(d)}, 1 + \epsilon, \frac{1}{\epsilon}, \epsilon^{-O(d)})$ -ultrametric cover for any parameter $\epsilon \in (0, \frac{1}{6})$. It turns out that this property is actually a characterization of doubling spaces. The proof of Theorem 2 is provided in Section 3.3.

Theorem 2 (Ultrametric Cover For Doubling Metrics). *Every metric space (X, d_X) with doubling dimension d admits an $(\epsilon^{-O(d)}, 1 + \epsilon, \frac{1}{\epsilon}, \epsilon^{-O(d)})$ -ultrametric cover for any parameter $\epsilon \in (0, \frac{1}{6})$.*

Conversely, if a metric space (X, d_X) admits a (τ, ρ, k, δ) -ultrametric cover for $k \geq 2\rho$, then it has doubling dimension $d \leq \log(\tau\delta)$.

The main tool in proving Theorems 1 and 2 is *pairwise partition cover*, a newly introduced notion, which is closely related to the previously introduced stochastic/padded decompositions and sparse covers [AP90, KPR93, Bar96, GKL03, AGG⁺19, Fil19a]. A partition $\mathcal{P}$ of a metric space (X, d_X) is Δ -bounded if every cluster $C \in \mathcal{P}$ has diameter at most Δ .

Definition 5 (Pairwise Partition Cover Scheme). *A collection of partitions $\mathbb{P} = \{\mathcal{P}_1, \dots, \mathcal{P}_s\}$ is $(\tau, \rho, \epsilon, \Delta)$ -pairwise partition cover if (a) $s \leq \tau$, (b) every partition $\mathcal{P}_i$ is Δ -bounded, and (c) for every pair x, y such that $\frac{\Delta}{2\rho} \leq d_X(x, y) \leq \frac{\Delta}{\rho}$, there is a cluster C in one of the partitions $\mathcal{P}_i$ such that C contains both closed balls $B(x, \epsilon\Delta), B(y, \epsilon\Delta)$.*

A space (X, d_X) admits a (τ, ρ, ϵ) -pairwise partition cover scheme if for every Δ , it admits a $(\tau, \rho, \epsilon, \Delta)$ -pairwise partition cover.

We will show that given a pairwise partition cover scheme, one can construct an ultrametric cover. The proof appears in Section 3.1.

Lemma 3. *Suppose that a metric space (X, d_X) admits a (τ, ρ, ϵ) -pairwise partition cover scheme for $\tau \in \mathbb{N}$, $\rho \geq 1$, and $\epsilon \in (0, \frac{1}{2})$. Then X admits an $(O(\frac{\tau}{\epsilon} \log \frac{\rho}{\epsilon}), \rho(1 + 7\epsilon))$ -ultrametric cover. Furthermore, every ultrametric in the cover is a $\Theta(\frac{\rho}{\epsilon})$ -HST.*

In Section 3.2 we construct a pairwise partition cover for general metrics:

Lemma 4. *Every n -point metric space (X, d_X) admits an $(O(n^{\frac{1}{k}} \log n), 2k + \delta, \frac{\delta}{4k(2k+\delta)})$ -pairwise partition cover scheme for any $\delta \in [0, 1]$ and integer $k \geq 1$.*

We are now ready to prove Theorem 1:

Proof of Theorem 1. Let (X, d_X) be an n -point metric space, and fix $\delta = \frac{\epsilon}{8}$. By Lemma 4, X admits an $(O(n^{\frac{1}{k}} \log n), 2k + \delta, \frac{\delta}{4k(2k+\delta)})$ -pairwise partition cover. By Lemma 3, X admits an ultrametric cover with $O(\frac{n^{\frac{1}{k}} \log n}{\delta} \cdot 4k(2k + \delta) \cdot \log \frac{(2k+\delta)4k(2k+\delta)}{\delta}) = O(n^{\frac{1}{k}} \log n \cdot \frac{k^2}{\delta} \cdot \log \frac{k}{\delta}) = O(n^{\frac{1}{k}} \log n \cdot \frac{k^2}{\epsilon} \cdot \log \frac{k}{\epsilon})$ ultrametrics, and stretch $(2k + \delta)(1 + \frac{7\delta}{2k(2k+\delta)}) < 2k + 8\delta = 2k + \epsilon$. $\square$

3.1 From Pairwise Partition Cover to Ultrametric Cover: Proof of Lemma 3

Lemma 3 is a reduction from pairwise partition cover scheme to ultrametric cover. In essence, an ultrametric is simply a hierarchical partition. Thus, this reduction takes unrelated partitions in all possible scales, and combines them into hierarchical/laminar partitions. Reductions similar in spirit were constructed in the context of the Steiner point removal problem [Fil20], stochastic Steiner point removal [EGK⁺14], universal Steiner tree [BDR⁺12], and others. We follow here a bottom-up approach, where the ratio between consecutive scales in a single hierarchical partition (a.k.a. ultrametric) is $O(\frac{\rho}{\epsilon})$. When constructing the next level in the hierarchical partition, we take partitions from a pairwise partition cover of the current scale, and slightly “round” them around the “borders” so that no previously created cluster will be divided (see Figure 2). The argument is that due to the large ratio between consecutive scales, the effects of this rounding are marginal.

Proof of Lemma 3. Assume w.l.o.g. that the minimal pairwise distance in X is 1, while the maximal pairwise distance is Φ . Fix $c \geq 1$ to be determined later. For $i \geq 0$, set $\Delta_i = c \cdot (\frac{4\rho}{\epsilon})^i$, and let $\mathbb{P}_i = \{\mathcal{P}_1^i, \dots, \mathcal{P}_\tau^i\}$ be a (τ, ρ, Δ_i) -padded partition cover (we assume that $\mathbb{P}_i$ has exactly τ partitions; we can enforce this assumption by duplicating partitions if necessary). Fix some j , let $\mathcal{P}_j^{-1}$ be the partition where each vertex is a singleton, and consider $\{\mathcal{P}_j^i\}_{i \geq -1}$. We will inductively define a new set of partitions, enforcing it to be a laminar system. The basic idea of doing this is to produce a tree of partitions where the lower level is a refinement of the higher level, and we do so by grouping a cluster at a lower level to one of the clusters at a higher level separating it.

The lowest level $\mathcal{P}_j^{-1}$ where each set in the partition is a singleton, stays as-is. Inductively, for any $i \geq 0$, after constructing $\tilde{\mathcal{P}}_j^{i-1}$ from $\mathcal{P}_j^{i-1}$, we will construct $\tilde{\mathcal{P}}_j^i$ from $\mathcal{P}_j^i$ using $\tilde{\mathcal{P}}_j^{i-1}$. Let $\mathcal{P}_j^i = \{C_1, \dots, C_\phi\}$ be the clusters in the partition $\mathcal{P}_j^i$. For each $q \in [1, \phi]$, let $Y_q = X \setminus \cup_{a < q} \tilde{C}_a$ be the set of unclustered points (w.r.t. level i , before iteration q). Let $C'_q = C_q \cap Y_q$ be the cluster C_q restricted to vertices in Y_q , and let $S_{C'_q} = \{C \in \tilde{\mathcal{P}}_j^{i-1} \mid C \cap C'_q \neq \emptyset\}$ be the set of new level- $(i-1)$ clusters with non empty intersection with C'_q . We set the new cluster $\tilde{C}_q = \cup_{C \in S_{C'_q}} C$ to be the union of these clusters, and continue iteratively. See Figure 2 for illustration. Clearly, $\tilde{\mathcal{P}}_j^{i-1}$ is a refinement of $\mathcal{P}_j^i$. We conclude that $\{\tilde{\mathcal{P}}_j^i\}_{i \geq -1}$ is a laminar hierarchical set of partitions that refine each other.

We next argue by induction that $\tilde{\mathcal{P}}_j^i$ has diameter $\Delta_i(1 + \epsilon)$. Consider $\tilde{C}_q \in \tilde{\mathcal{P}}_j^i$; it consists of $C'_q \subseteq C_q \in \mathcal{P}_j^i$ and of clusters in $\tilde{\mathcal{P}}_j^{i-1}$ intersecting C'_q . As the diameter of C'_q is bounded by $\text{diam}(C_q) \leq \Delta_i$, and by the induction hypothesis, the diameter of each cluster $C \in \tilde{\mathcal{P}}_j^{i-1}$ is bounded by $(1 + \epsilon)\Delta_{i-1}$, we conclude that the diameter of $\tilde{C}_q$ is bounded by

$$\Delta_i + 2 \cdot (1 + \epsilon)\Delta_{i-1} = \Delta_i \left(1 + \frac{2(1 + \epsilon)}{4\rho} \cdot \epsilon\right) \leq \Delta_i(1 + \epsilon),$$

since $\rho \geq 1$ and $\epsilon < 1$.




Figure 2: Illustration of the construction of the partition $\tilde{\mathcal{P}}_j^i$ given $\mathcal{P}_j^i$ and $\tilde{\mathcal{P}}_j^{i-1}$. The black lines in both the left and right parts of the figure border clusters in $\tilde{\mathcal{P}}_j^{i-1}$. On the left illustrated the partition $\mathcal{P}_j^i = \{C_1, C_2, C_3, C_4\}$, where different clusters colored by different colors. On the right illustrated the modified partition $\mathcal{P}_j^i = \{\tilde{C}_1, \tilde{C}_2, \tilde{C}_3, \tilde{C}_4\}$. $\tilde{C}_1$ contains all the clusters in $\tilde{\mathcal{P}}_j^{i-1}$ intersecting C_1 . C_2 contains all the clusters in $\tilde{\mathcal{P}}_j^{i-1} \setminus \tilde{C}_1$ intersecting C_2 , and so on.

Next we argue that $\tilde{\mathbb{P}}_i = \{\tilde{\mathcal{P}}_1^i, \dots, \tilde{\mathcal{P}}_\tau^i\}$ is a $(\tau, (1+\epsilon)\rho, 0, (1+\epsilon)\Delta)$ -pairwise partition cover. Observe that it contains τ partitions, and we have shown that all the clusters have diameter at most $(1+\epsilon)\Delta$. Thus, it remains to prove that for every pair x, y at distance $d_X(u, v) \in \left[\frac{(1+\epsilon)\Delta_i}{2(1+\epsilon)\rho}, \frac{(1+\epsilon)\Delta_i}{(1+\epsilon)\rho}\right] = \left[\frac{\Delta_i}{2\rho}, \frac{\Delta_i}{\rho}\right]$ contained in some cluster. As $d_X(u, v) \in \left[\frac{\Delta_i}{2\rho}, \frac{\Delta_i}{\rho}\right]$, there is some index j such that $B_X(u, \epsilon\Delta_i), B_X(v, \epsilon\Delta_i) \subseteq C_i \in \mathcal{P}_j^i$. That is, the balls of radius $\epsilon\Delta_i$ around u, v are contained in a cluster of $\mathcal{P}_j^i$. We argue that $u, v \in \tilde{C}_i \in \tilde{\mathcal{P}}_j^i$. Let $\tilde{C}_v, \tilde{C}_u \in \tilde{\mathcal{P}}_j^{i-1}$ be the clusters containing v, u respectively at $(i-1)$ -th level. Note that they both have diameter at most $(1+\epsilon)\Delta_{i-1} = \frac{(1+\epsilon)\epsilon}{4\rho}\Delta_i < \epsilon\Delta_i$. Hence $\tilde{C}_v \subseteq B_X(v, \epsilon\Delta_i) \subseteq C_i$, and similarly $\tilde{C}_u \subseteq C_i$. By the partitioning algorithm, it follows that $\tilde{C}_u, \tilde{C}_v \subseteq \tilde{C}_i$ (as $\tilde{C}_u, \tilde{C}_v$ do not intersect any other clusters), and in particular $u, v \in \tilde{C}_i$ as required.

Finally, we construct an ultrametric cover. Fix an index $j \in [1, \tau]$; we construct a $(\frac{4\rho}{\epsilon})$ -HST U_j as follows. Leaves of U_j bijectively correspond to points in X and have label 0. For each $i \in [0, I]$ where $I = \lceil \log_{4\rho/\epsilon} \Phi/c \rceil$, internal nodes at level i bijectively correspond to the clusters $\tilde{\mathcal{P}}_j^i$ (leaves of U_j is at level -1), and have label $(1+\epsilon)\Delta_i$. There is an edge from each node corresponding to a cluster $\tilde{C}_{i-1} \in \tilde{\mathcal{P}}_j^{i-1}$ to the node corresponding the unique cluster $\tilde{C}_i \in \tilde{\mathcal{P}}_j^i$ containing $\tilde{C}_{i-1}$. The root of U_j is the unique single cluster in $\tilde{\mathcal{P}}_j^I$. Clearly, the ultrametric cover $\{U_j\}_{j=1}^\tau$ is dominating, and every ultrametric is a $\frac{4\rho}{\epsilon}$ -HST.

To bound the stretch, we will construct such an ultrametric cover with $c = (1+\epsilon)^l$ for every $l \in [0, \lceil \log_{1+\epsilon} \frac{4\rho}{\epsilon} \rceil]$. The final ultrametric cover will be a union of these $O(\log_{1+\epsilon} \frac{4\rho}{\epsilon})$ ultrametric covers. Clearly, their cardinality is bounded by $\tau \cdot O(\log_{1+\epsilon} \frac{4\rho}{\epsilon}) = O(\frac{\tau}{\epsilon} \log \frac{\rho}{\epsilon})$.

Consider a pair $x, y \in X$. Let $l \in [0, \lceil \log_{1+\epsilon} \frac{4\rho}{\epsilon} \rceil]$, and $i \geq 0$ be the unique indices such that $(1+\epsilon)^{l-1}(\frac{4\rho}{\epsilon})^i \leq (1+\epsilon)\rho \cdot d_X(x, y) \leq (1+\epsilon)^l(\frac{4\rho}{\epsilon})^i$. For $c = (1+\epsilon)^l$, there is some index j , and a cluster $\tilde{C}_i \in \tilde{\mathcal{P}}_j^i$ such that $x, y \in \tilde{C}_i \in \tilde{\mathcal{P}}_j^i$. Thus in the corresponding ultrametric, x, y both decedents of an internal node with label $(1+\epsilon)^{l+1}(\frac{4\rho}{\epsilon})^i \leq (1+\epsilon)^3\rho \cdot d_X(x, y)$, the stretch guarantee follows.

In summary, we have constructed an $(O(\frac{\tau}{\epsilon} \log \frac{\rho}{\epsilon}), \rho(1+7\epsilon))$ -ultrametric cover, consisting of $(\frac{4\rho}{\epsilon})$ -HST's. $\square$

3.2 Pairwise Partition Cover for General Metrics: Proof of Lemma 4

Fix parameter $\delta \in (0, 1]$. We begin by creating a distribution over partitions, such that for every pair of points u, v at distance $\frac{\Delta}{2k+\delta}$, there is a non trivial probability that the some balls around u, v contained in a single cluster. Later, Lemma 4 will follow by taking the union of many independently sampled such partitions.

Lemma 5. *For every n -point metric space (X, d_X) , integer $k \geq 1$, $\delta \in [0, 1]$, and $\Delta > 0$ there is a distribution over Δ -bounded partitions such that for every pair of points u, v where $d_X(u, v) \leq \frac{\Delta}{2k+\delta}$, with probability at least $n^{-\frac{1}{k}}$, the balls $B_X(u, \frac{\delta}{4k(2k+\delta)}\Delta), B_X(v, \frac{\delta}{4k(2k+\delta)}\Delta)$ contained in a single cluster.*

For the case where $\delta = 0$, is a distribution formerly constructed by the first author [Fil19b].⁸ Our proof here follows the steps of [Fil19b] (which is based on the partition in [CKR04]).

Proof of Lemma 5. Pick u.a.r. a radius $r \in \{\frac{1}{k}, \frac{2}{k}, \dots, \frac{k}{k}\}$, and a random permutation $\pi = \{v_1, v_2, \dots, v_n\}$ over the points. Then set $C_i = B_X(v_i, r \cdot \frac{\Delta}{2}) \setminus \cup_{j < i} B_X(v_j, r \cdot \frac{\Delta}{2})$. As a result we obtain a Δ bounded partition $\{C_i\}_{i=1}^n$.

For a pair u, v where $d_X(u, v) \leq \frac{\Delta}{2k+\delta}$, let $T = B_X(u, \frac{\delta}{4k(2k+\delta)}\Delta) \cup B_X(v, \frac{\delta}{4k(2k+\delta)}\Delta)$. Note that for every pair of points $x, y \in T$, by triangle inequality it holds that

$$d_X(x, y) \leq d_X(u, v) + \frac{\delta}{2k(2k+\delta)} \cdot \Delta \leq \left(\frac{1}{2k+\delta} + \frac{\delta}{2k(2k+\delta)} \right) \cdot \Delta = \frac{\Delta}{2k}.$$

Let $A_s = \{v_j \mid d_X(v_j, T) \leq \frac{s}{k} \cdot \frac{\Delta}{2}\}$. Then $A_0 = T$. Suppose that $r = \frac{s}{k}$, and let v_i be the vertex with minimal index such that $d_X(T, v_i) \leq \frac{s}{k} \cdot \frac{\Delta}{2}$. Then no vertex in T will join the clusters $C_1, \dots, C_{i-1}$, and some vertex in T will join C_i . Let $z \in T \cap C_i$, and suppose further that $v_i \in A_{s-1}$. By the triangle inequality, for every $y \in T$, $d_X(y, v_i) \leq d_X(y, z) + d_X(z, v_i) \leq \frac{\Delta}{2k} + \frac{s-1}{k} \cdot \frac{\Delta}{2} = \frac{s}{k} \cdot \frac{\Delta}{2}$. Hence all the points in T will join the cluster of v_i . Denote by Ψ the event that all the vertices in T are contained in a single cluster. Using the law of total probability, we conclude

$$\Pr[\Psi] = \frac{1}{k} \cdot \sum_{s=1}^k \Pr[\Psi \mid r = \frac{s}{k}] \geq \frac{1}{k} \cdot \sum_{s=1}^k \frac{|A_{s-1}|}{|A_s|} \geq \left(\prod_{s=1}^k \frac{|A_{s-1}|}{|A_s|} \right)^{\frac{1}{k}} = \left(\frac{|A_0|}{|A_k|} \right)^{\frac{1}{k}} \geq n^{-\frac{1}{k}},$$

where the second inequality follows by the inequality of arithmetic and geometric means. $\square$

We now ready to prove Lemma 4 (restated for convenience).

Lemma 4. *Every n -point metric space (X, d_X) admits an $(O(n^{\frac{1}{k}} \log n), 2k + \delta, \frac{\delta}{4k(2k+\delta)})$ -pairwise partition cover scheme for any $\delta \in [0, 1]$ and integer $k \geq 1$.*

⁸This is the full version, see also the conference version [Fil19a].

Proof. Fix Δ . Sample $s = n^{\frac{1}{k}} \cdot 2 \ln n$ i.i.d. partitions using Lemma 5. Consider a pair of points u, v such that $d_X(u, v) \leq \frac{\Delta}{2k+\delta}$. Then in each sampled partition, the probability that the balls $B_X(u, \frac{\delta}{4k(2k+\delta)}\Delta), B_X(v, \frac{\delta}{4k(2k+\delta)}\Delta)$ contained in a single cluster is at least $p = n^{-\frac{1}{k}}$. The probability that u, v are not satisfied by any partition, is at most $(1-p)^s \leq e^{-ps} = e^{-2 \ln n} = n^{-2}$. As there are at most $\binom{n}{2} \leq \frac{n^2}{2}$ pairs at distance at most $\frac{\Delta}{2k+\delta}$, by union bound, with probability at least $\frac{1}{2}$, every pair is satisfied by some partition. It follows that the union of s random partitions is, with a probability at least $\frac{1}{2}$, an $(O(n^{\frac{1}{k}} \log n), 2k + \delta, \frac{\delta}{4k(2k+\delta)})$ -pairwise partition cover as required. $\square$

3.3 Ultrametric cover for doubling spaces

In this section, we will construct a pairwise partition cover for doubling spaces, and then use them to construct ultrametric covers, and thus proving Theorem 2. We begin with the following combinatorial lemma.

Lemma 6. *Consider a graph $G = (V, E_b \cup E_r)$ with disjoint sets of blue edges E_b and red edges E_r , such the maximal blue degree is $\delta_b \geq 1$, and the maximal red degree is $\delta_r \geq 1$. Then there is a set of at most $\gamma = O(\delta_r \delta_b)$ matching $\mathcal{M} = \{M_1, M_2, \dots, M_\gamma\}$ of G such that (a) $E_b \subseteq \cup_{i=1}^\gamma M_i$, and (b) for every matching $M \in \mathcal{M}$, there is no red edge whose both endpoints are matched by M .*

Proof. We construct $\mathcal{M}$ greedily. Initially, $\mathcal{M} = \emptyset$. Let E'_b be the set of blue edges of G that are not added to any matching in $\mathcal{M}$. Let $M \subseteq E'_b$ be a maximal matching such that there is no red edge whose endpoints are both matched by M (such maximal matching could be found greedily in linear time); we add M to $\mathcal{M}$ and repeat.

We argue by contradiction that the greedy algorithm adds at most $4\delta_r \delta_b$ matching to $\mathcal{M}$. Consider a vertex v such that after $\delta_b(2\delta_r + 2)$ maximal matchings added to $\mathcal{M}$, there remains at least one blue edge incident to v that is not covered by any matching in $\mathcal{M}$. Since there is at most δ_b blue edges incident to v , there must be a set $\mathcal{M}_v \subseteq \mathcal{M}$ of at least $\delta_b(2\delta_r + 1)$ matchings where v is not matched by any of the matchings in $\mathcal{M}_v$. By the maximality, in each matching $M \in \mathcal{M}_v$, either:

- (a) A red neighbor of v is matched by M .
- (b) For every blue neighbor u of v , either u is matched, or a red neighbor of u is matched by M , which prevents u from being matched.

Since v has at most δ_r red neighbors, and each of them can be matched at most δ_b times, case (a) happens at most $\delta_b \delta_r$ times. The blue neighbors of v could be matched at most $\delta_b - 1$ times, while their red neighbors could be matched at most $\delta_r \delta_b$ times. Thus, case (b) happens at most $\delta_b - 1 + \delta_r \delta_b < \delta_b(\delta_r + 1) - 1$ times. We conclude that $|\mathcal{M}_v| \leq \delta_r \delta_b + \delta_b(\delta_r + 1) - 1 = \delta_b(2\delta_r + 1) - 1$, a contradiction. $\square$

Lemma 7. *Every metric space (X, d_X) with doubling dimension d admits an $(\epsilon^{-O(d)}, (1 + \epsilon), \epsilon)$ -pairwise partition cover scheme for any $\epsilon \in (0, 1/16)$.*

Proof. Let $\Delta > 0$ be any given real number. We show that (X, d_X) admits an $(\epsilon^{-O(d)}, (1 + 8\epsilon), \frac{\epsilon}{2}, (1 + 8\epsilon)\Delta)$ -pairwise partition cover $\mathbb{P}$, the lemma then follows by rescaling ϵ and Δ .

Let N be an $(\epsilon\Delta)$ -net of (X, d_X) . We construct a graph G with N as the vertex set; there is a blue edge $(u, v) \in E_b$ in G iff $d_X(u, v) \in [(1 - 4\epsilon)\frac{\Delta}{2}, (1 + 2\epsilon)\Delta]$, and there is a red edge $(u, v) \in E_r$




Figure 3: Illustration of the partition $\mathcal{P}_M$. The black points represent metric points, while the red points represent the net points N . The blue edges are the matching M . For each edge $\{u, v\} \in M$, the cluster $B_X(u, 2\epsilon\Delta) \cup B_X(v, 2\epsilon\Delta)$ is added to $\mathcal{P}_M$. These clusters are illustrated with colored filled boxes. The remaining points are clustered around unclustered net points. These clusters are encircled by green lines. Distances shown in the figure are not properly scaled for a better visualization.

iff $d_X(u, v) \leq 4\epsilon\Delta$. As $\epsilon < \frac{1}{12}$, the set of blue and red edges are disjoint. By the packing property of doubling metrics (Lemma 1), every vertex in G has blue degree $\epsilon^{-O(d)}$ and red degree $2^{O(d)}$. Let $\mathcal{M}$ be the set of matching of $G(N, E_b \cup E_r)$ guaranteed by Lemma 6; $|\mathcal{M}| = O(\epsilon^{-O(d)} 2^{O(d)}) = \epsilon^{-O(d)}$.

For each matching $M \in \mathcal{M}$, we construct a partition $\mathcal{P}_M$ as follows: for every edge $\{u, v\} \in M$, we add $B_X(u, 2\epsilon\Delta) \cup B_X(v, 2\epsilon\Delta)$ as a cluster to $\mathcal{P}_M$. Denote by N_M the set of net points that remain unclustered. For every net point $x \in N_M$, we initiate a new cluster C_x containing x only. Then, every remaining unclustered point $z \in X$ joins the cluster of its closest net point x_z (from either N_M or $N \setminus N_M$). See Figure 3 for an illustration.

We observe that for any two edges (u, v) and (u', v') in matching M , $B_X(u, 2\epsilon\Delta) \cap B_X(u', 2\epsilon\Delta) = \emptyset$ since otherwise, there is a red edge between u and u' , contradicting item (b) in Lemma 6. Thus, $\mathcal{P}_M$ is indeed a partition of X . We next bound the diameter of each cluster in $\mathcal{P}_M$. Clearly every cluster C_x for $x \in N_M$ has diameter at most $2\epsilon\Delta$. On the other hand, by the construction and the triangle inequality, the diameter of every cluster resulting from the matching is bounded by $2 \cdot (\epsilon\Delta + 2\epsilon\Delta) + (1 + 2\epsilon)\Delta = (1 + 8\epsilon)\Delta$. Thus $\mathcal{P}_M$ is $((1 + 8\epsilon)\Delta)$ -bounded.

Let $\mathbb{P} = \{\mathcal{P}_M\}_{M \in \mathcal{M}}$. It remains to show that for every $x, y \in X$ such that $d_X(x, y) \in \left[\frac{(1+8\epsilon)\Delta}{(1+8\epsilon)^2}, \frac{(1+8\epsilon)\Delta}{(1+8\epsilon)}\right] = \left[\frac{\Delta}{2}, \Delta\right]$, there is a cluster C in a partition $\mathcal{P} \in \mathbb{P}$ containing both $B_X(x, \frac{\epsilon}{2} \cdot (1 + 8\epsilon)\Delta)$ and $B_X(y, \frac{\epsilon}{2} \cdot (1 + 8\epsilon)\Delta)$. Note that as $\epsilon \leq \frac{1}{16}$, $\frac{\epsilon}{2} \cdot (1 + 8\epsilon)\Delta \leq \epsilon\Delta$. Let $x', y' \in N$ be net points such that $d_X(x, x'), d_X(y, y') \leq \epsilon\Delta$. Then by the triangle inequality $|d_X(x', y') - d_X(x, y)| \leq 2\epsilon\Delta$, implying that $d_X(x', y') \in \left[(1 - 4\epsilon)\frac{\Delta}{2}, (1 + 2\epsilon)\Delta\right]$. Hence, G contains a blue edge between x', y' . It follows that there is a matching M containing the edge $\{x', y'\}$, and a partitions $\mathcal{P}_M$ containing the cluster $C = B_X(x', 2\epsilon\Delta) \cup B_X(y', 2\epsilon\Delta)$. In particular, $B_X(x, \epsilon\Delta) \cup B_X(y, \epsilon\Delta) \subseteq C$ as required. $\square$

We are finally ready to prove Theorem 2 that we restate below for convenience.

Theorem 2 (Ultrametric Cover For Doubling Metrics). *Every metric space (X, d_X) with doubling dimension d admits an $(\epsilon^{-O(d)}, 1 + \epsilon, \frac{1}{\epsilon}, \epsilon^{-O(d)})$ -ultrametric cover for any parameter $\epsilon \in (0, \frac{1}{6})$.*

Conversely, if a metric space (X, d_X) admits a (τ, ρ, k, δ) -ultrametric cover for $k \geq 2\rho$, then it has doubling dimension $d \leq \log(\tau\delta)$.

Proof. We begin with the first assertion (doubling metrics admit ultrametric covers). After appropriate rescaling, by Lemma 3 and Lemma 7, we obtain an $(\epsilon^{-O(d)}, 1 + \epsilon)$ -ultrametric cover where every ultrametric in the cover is a $\frac{1}{\epsilon}$ -HST. It remains to show that every ultrametric in the cover returned by Lemma 3 w.r.t. the pairwise partition cover scheme constructed in Lemma 7 has bounded degree.

We will use the terminology of Lemmas 3 and 7. Consider some ultrametric $\mathcal{U}_j$ and some cluster $\tilde{C}$ at level i with label $(1 + \epsilon)\Delta_i$. The clusters at level $i - 1$ correspond to points in $(\epsilon\Delta_{i-1})$ -net. The cluster $\tilde{C}$ has diameter $(1 + \epsilon)\Delta_i$, and hence it has at most $\epsilon^{-O(d)}$ $(\epsilon\Delta_{i-1})$ -net points. In particular $\tilde{C}$ can contain at most $\epsilon^{-O(d)}$ level- $(i - 1)$ clusters. The bound on the degree follows.

Next, we prove the second assertion. Consider a metric space (X, d_X) admitting a (τ, ρ, k, δ) -ultrametric cover with $k \geq 2\rho$. Let $B_X(x, r)$ be some ball of radius r . In each ultrametric U_i in the cover, let L_i be the node closest to root that is an ancestor of x and has label at most $\rho \cdot r$. Let $\{L_{i,1}, L_{i,2}, \dots\}$ be the set of at most δ children of L_i in U_i . For each $L_{i,j}$, we pick an arbitrary leaf $u_{i,j} \in X$ descendent of $L_{i,j}$. We argue that

$$B_X(x, r) \subseteq \cup_{i,j} B_X(u_{i,j}, \frac{r}{2}),$$

as the number of balls in the union is at most $\tau\delta$, the theorem will follow. Consider a vertex $y \in B_X(x, r)$. There is necessarily an ultrametric U_i such that $d_{U_i}(x, y) \leq \rho \cdot r$. In particular, in U_i , x, y are both decedents of a node with label at most $\rho \cdot r$. Recall that L_i is such a node with maximal label. Let $L_{i,j}$ be the child of L_i such that y is decedent of $L_{i,j}$. As U_i is a k -HST, the label of $L_{i,j}$ is bounded by $\frac{\rho r}{k} \leq \frac{r}{2}$ since $k \geq 2\rho$. In particular, $y \in B_X(u_{i,j}, \frac{r}{2})$ as required. $\square$

4 Locality Sensitive Ordering

Locality-Sensitive Ordering. Chan *et al.* [CHJ20] introduced and studied the notion of locality-sensitive ordering (Definition 1). In the same paper, Chan *et al.* [CHJ20] showed that the Euclidean metric of dimension d has an $(O(\epsilon)^{-d} \log \frac{1}{\epsilon}, \epsilon)$ -LSO. They also presented various applications of the LSO to solve fundamental geometry problems in Euclidean spaces. The proof relies on the following lemma by Walecki [Als08].

Lemma 8. *Given a set of n elements $[n] = \{1, \dots, n\}$, there exists a set Σ of $\lceil \frac{n}{2} \rceil$ orderings such that for any two elements $i \neq j \in [n]$, there exists an ordering σ in which i and j are adjacent.*

We show that metrics admitting an ultrametric cover of bounded degree have an LSO with a small number of orders. Our proof relies on the following lemma.

Lemma 9. *Every α -HST (U, d_U) of degree δ admits a $(\lceil \frac{\delta}{2} \rceil, \frac{1}{\alpha})$ -LSO.*

Proof. For simplicity, we will assume that the number of children in each node is exactly δ . This could be achieved by adding dummy nodes. By Lemma 8, every set of δ vertices can be ordered into $\lceil \frac{\delta}{2} \rceil$ orderings such that every two vertices are adjacent in at least one of them. Denote these orderings by $\sigma_1, \dots, \sigma_{\lceil \frac{\delta}{2} \rceil}$. We construct the set of orderings for (U, d_U) inductively.

Let A be the root of the HST with children $A_1, \dots, A_\delta$. By the induction hypothesis, each A_i admits a set $\sigma_1^i, \dots, \sigma_{\lceil \frac{\delta}{2} \rceil}^i$ orderings. We construct $\lceil \delta/2 \rceil$ orderings as follows: for each $j \in [\lceil \delta/2 \rceil]$, order the vertices inside each A_i w.r.t. σ_j^i , and order the sets in between them w.r.t. σ_j . The resulting ordering is denoted by $\tilde{\sigma}_j$. This finishes the construction.

Next, we argue that this is a $(\lceil \frac{\delta}{2} \rceil, \frac{1}{\alpha})$ -LSO. Clearly, we used exactly $\lceil \frac{\delta}{2} \rceil$ orderings. Let Δ be the label of the root. Consider a pair of leaves x, y . If $d_U(x, y) < \Delta$, then there is some i such that $x, y \in A_i$. By the induction hypothesis, there is some order σ_j^i of A_i such that (w.l.o.g.) $x <_{\sigma_j^i} y$, and the points between x and y w.r.t. σ_j^i could be partitioned into two consecutive intervals I_x, I_y such that $I_x \subseteq B_U(x, d_U(x, y)/\alpha)$ and $I_y \subseteq B_U(y, d_U(x, y)/\alpha)$. The base case is trivial since every leaf has label 0. Note that σ_j^i is a sub-ordering of $\tilde{\sigma}_j$. In particular, the lemma holds.

The next case is when $d_U(x, y) = \Delta$. Then there are $i \neq i'$ such that $x \in A_i$ and $y \in A_{i'}$. There is some ordering σ_j such that A_i and $A_{i'}$ are consecutive. In particular all the vertices between x to y in $\tilde{\sigma}_j$ can be partitioned to two sets, the first belonging to A_i , and the second to $A_{i'}$. The lemma follows as all the vertices in A_i ($A_{i'}$) are at distance at most $\frac{\Delta}{\alpha} = \frac{d_U(x, y)}{\alpha}$ from x (y). $\square$

Lemma 10. *If a metric (X, d_X) admits a (τ, ρ, k, δ) -ultrametric cover, then it has a $(\tau \cdot \lceil \frac{\delta}{2} \rceil, \frac{\rho}{k})$ -LSO.*

Proof. Let $\mathcal{U}$ be an ultrametric cover for (X, d_X) . For each ultrametric (U, d_U) , let Σ_U be the set of orderings obtained by applying Lemma 9. Let $\Sigma = \cup_{U \in \mathcal{U}} \Sigma_U$. We show that Σ is the LSO claimed by the lemma. Clearly, it contains at most $\tau \cdot \lceil \frac{\delta}{2} \rceil$ orderings.

Consider two points $x \neq y \in X$. Let U be an ultrametric in $\mathcal{U}$ such that $d_U(x, y) \leq \rho \cdot d_X(x, y)$. By Lemma 9, there is an ordering $\rho \in \Sigma_U$ such that (w.l.o.g.) $x <_\rho y$ and points between x and y w.r.t. ρ can be partitioned into two consecutive intervals I_x and I_y where $I_x \subseteq B_U(x, \frac{d_U(x, y)}{k})$ and $I_y \subseteq B_U(y, \frac{d_U(x, y)}{k})$. Since $d_U(x, y) \leq \rho \cdot d_X(x, y)$, we conclude that $I_x \subseteq B_X(x, \frac{\rho}{k} \cdot d_X(x, y))$ and $I_y \subseteq B_X(y, \frac{\rho}{k} \cdot d_X(x, y))$ as desired. $\square$

Using Theorem 2 and Lemma 10 with $\rho = (1 + \epsilon)$ and $k = O(\frac{1}{\epsilon})$, we conclude

Corollary 1 (LSO For Doubling Metrics). *For every ϵ sufficiently smaller than 1, every metric space (X, d_X) of doubling dimension d admits an $(\epsilon^{-O(d)}, \epsilon)$ -LSO.*

Reliable $(1 + \epsilon)$ -Spanners from LSO. Buchin *et al.* [BHO19, BHO20] constructed reliable $(1 + \epsilon)$ spanners for Euclidean metrics using $(\tau(\epsilon), \epsilon)$ -LSO by Chan *et al.* [CHJ20]. Specifically, their spanner for the deterministic case has $n \cdot O(\epsilon)^{-7d} \log^7(\frac{1}{\epsilon}) \cdot \nu^{-6} \log n (\log \log n)^6$ edges, while for the oblivious case, they constructed a spanner with an almost linear number of edges: $n \cdot O(\epsilon)^{-2d} \log^3 \frac{1}{\epsilon} \cdot \nu^{-1} \log \nu^{-1} (\log \log n)^2 \log \log \log n$. Their key idea is to reduce the problem to the construction of reliable $(1 + \epsilon)$ -spanners for the (unweighted) path graph P_n with n vertices. We observe that their construction of the reliable $(1 + \epsilon)$ -spanners did not use any property of the metric space other than the existence of an LSO.

Theorem 3 ([BHO19, BHO20], implicit). *Suppose that for any $\epsilon \in (0, 1)$, an n -point metric space (X, d_X) admits a $(\tau(\epsilon), \epsilon)$ -LSO for some function $\tau : (0, 1) \rightarrow \mathbb{N}$. Then for every $\nu \in (0, 1)$ and $\epsilon \in (0, 1)$:*

1. (X, d_X) admits a deterministic ν -reliable $(1 + \epsilon)$ -spanner with $n \cdot O\left(\left(\tau\left(\frac{\epsilon}{c_d}\right)\right)^7 \nu^{-6} \log n (\log \log n)^6\right)$ edges for some universal constant c_d .
2. (X, d_X) admits an oblivious ν -reliable $(1 + \epsilon)$ -spanner with $n \cdot O\left(\left(\tau\left(\frac{\epsilon}{c_0}\right)\right)^2 \nu^{-1} (\log \log n)^2 \cdot \log\left(\frac{\tau\left(\frac{\epsilon}{c_0}\right) \log \log n}{\nu}\right)\right)$ edges for some universal constant c_0 .

By Theorem 3 and Corollary 1, we have:

Corollary 2. Consider a metric space (X, d_X) with doubling dimension d . Then for every $\nu \in (0, 1)$ and $\epsilon \in (0, 1)$, (X, d_X) admits a deterministic ν -reliable $(1 + \epsilon)$ -spanner with $n \cdot \epsilon^{-O(d)} \nu^{-6} \log n (\log \log n)^6$ edges, and an oblivious ν -reliable $(1 + \epsilon)$ -spanner with $n \cdot \epsilon^{-O(d)} \cdot \nu^{-1} (\log \log n)^2 \cdot \log\left(\frac{\log \log n}{\epsilon \nu}\right) = n \cdot \epsilon^{-O(d)} \cdot \tilde{O}\left(\nu^{-1} (\log \log n)^2\right)$ edges.

5 Triangle Locality Sensitive Ordering

A triangle locality-sensitive ordering (triangle-LSO) is defined as follows.

Definition 6 $((\tau, \rho)$ -Triangle-LSO). Given a metric space (X, d_X) , we say that a collection Σ of orderings is a (τ, ρ) -triangle-LSO if $|\Sigma| \leq \tau$, and for every $x, y \in X$, there is an ordering $\sigma \in \Sigma$ such that (w.l.o.g.) $x <_\sigma y$, and for every $a, b \in X$ such that $x \leq_\sigma a \leq_\sigma b \leq_\sigma y$ it holds that $d_X(a, b) \leq \rho \cdot d_X(x, y)$.

Note that every (τ, ρ) -triangle LSO is also a (τ, ρ) -LSO; however, a (τ, ρ) -LSO is a $(\tau, 2\rho + 1)$ -triangle LSO (by triangle inequality). Hence for stretch parameter $\rho > 1$, triangle-LSO is preferable to the classic LSO. Similar to Lemma 10, we show that a metric space admitting an ultrametric cover has a triangle-LSO with a small number of orderings.

Lemma 11. If a metric (X, d_X) admits a (τ, ρ) -ultrametric cover $\mathcal{U}$, then it has a (τ, ρ) -triangle-LSO.

Proof. Let (U, d_U) be an ultrametric in the cover $\mathcal{U}$. Note that U is also a 1-HST. We create a single ordering for σ_U by simply putting the leaves (only) in the pre-order fashion. That is, given a node v with children $v_1, \dots, v_\delta$, we recursively put all the descendants of v_1 , then v_2 , and so on until we put all the descendants of v to σ_U . It holds that for every $x <_{\sigma_U} y$ with lca z , all the vertices a, b such that $x <_{\sigma_U} a <_{\sigma_U} b <_{\sigma_U} y$ are also descendants of z . Thus, $d_U(a, b) \leq d_U(x, y)$.

We now show that $\{\sigma_U\}_{U \in \mathcal{U}}$ is a (τ, ρ) -triangle-LSO. For any given x, y , there exists $U \in \mathcal{U}$ such that $d_X(x, y) \leq d_U(x, y) \leq \rho \cdot d_X(x, y)$. Then for every pair of vertices a, b such that $x <_{\sigma_U} a <_{\sigma_U} b <_{\sigma_U} y$, $d_X(a, b) \leq d_U(a, b) \leq d_U(x, y) \leq \rho \cdot d_X(x, y)$ as desired. $\square$

Using Theorem 1 and Lemma 11, we conclude.

Corollary 3. For every $k \in \mathbb{N}$, and $\epsilon \in (0, 1)$, every n -point metric space admits an $\left(O\left(n^{\frac{1}{k}} \cdot \log n \cdot \frac{k^2}{\epsilon} \cdot \log \frac{k}{\epsilon}\right), 2k + \epsilon\right)$ -triangle-LSO.

Reliable Spanners from triangle-LSO. We show that if a metric space admits a (τ, ρ) -triangle-LSO, it has an oblivious 2ρ -spanners with about $n \cdot \tau^2$ edges. We use this result to construct reliable $(8k - 2)(1 + \epsilon)$ -spanners for general metrics (and reliable (2ρ) -spanners for ultrametrics).

Theorem 4. *Suppose that a metric space (X, d_X) admits a (τ, ρ) -Triangle-LSO. Then for every $\nu \in (0, 1)$, X admits an oblivious ν -reliable, (2ρ) -spanner with $n\tau \cdot O(\log^2 n + \nu^{-1}\tau \log n \cdot \log \log n)$ edges.*

The proof of Theorem 4 is deferred to Section 5.1. Using Corollary 3 with parameters $2k$ and $\frac{\epsilon}{2}$, and Theorem 4 we conclude:

Theorem 5 (Oblivious Reliable Spanner for General Metric). *For every n -point metric space (X, d) and parameters $\nu \in (0, 1)$, $\epsilon \in (0, \frac{1}{2})$, $k \in \mathbb{N}$, (X, d) admits an oblivious ν -reliable, $8k + \epsilon$ -spanner for X with $n^{1+\frac{1}{k}} \cdot \nu^{-1} \cdot \log^3 n \cdot \log \log n \cdot \frac{k^4}{\epsilon^2} O(\log \frac{k}{\epsilon})^2 = n^{1+\frac{1}{k}} \cdot \nu^{-1} \cdot \tilde{O}(\log^3 n \cdot \frac{k^4}{\epsilon^2})$ edges.*

By Lemma 11 and Theorem 4, we obtain:

Corollary 4. *For every parameter $\nu \in (0, 1)$, every n -point ultrametric space (X, d) admits an oblivious ν -reliable, 2-spanner with $n \cdot O(\log^2 n + \nu^{-1} \log n \cdot \log \log n)$ edges.*

The stretch parameter in Corollary 4 is tight; see Remark 1.

For the rest of this section, we show how to construct reliable spanners for metric spaces admitting a triangle LSO. Following the approach of Buchin *et al.* [BHO19, BHO20], we reduce the problem to the construction of reliable spanners for the (unweighted) path graph P_n . However, in our setting, we face a very different challenge: when $\rho > 1$, the *stretch* of the reliable spanner for (X, d_X) grows linearly w.r.t. the *number of hops* of the reliable spanner for the path graph. In the Euclidean setting studied by Buchin *et al.* [BHO20], the stretch parameter is $\rho = \epsilon < 1$, and as a result, the stretch is not significantly affected by the number of hops of the spanner for the path graph.

Buchin *et al.* [BHO20] constructed an oblivious ν -reliable 1-spanner for the path graph P_n with $O(n \cdot \nu^{-1} \log \nu^{-1})$ edges. However, their spanner has hop diameter $2 \log n$; if we use their reliable spanner for the path graph, we will end up in a spanner for (X, d_X) with stretch $2\rho \log n$. To have a stretch 2ρ , we construct a reliable spanner for the path graph P_n with hop diameter 2. As a consequence, the sparsity of our spanner has some additional logarithmic factors. Note that a 2-hop spanner for the path graph P_n even without any reliability guarantee must contain $\Omega(n \log n)$ edges (see Exercise 12.10 [NS07]). Our result is summarized in the following lemma whose proof is deferred to Section 5.2.

Lemma 12 (2-hop-spanner). *For every $\nu \in (0, 1)$, the path graph P_n admits an oblivious ν -reliable, 2-hop 1-spanner H with $n \cdot O(\log^2 n + \nu^{-1} \log n \cdot \log \log n)$ edges.*

Using Lemma 12, we can construct a reliable spanner for metric spaces admitting a (τ, ρ) -triangle LSO as claimed by Theorem 4.

5.1 Proof of Theorem 4

Let Σ be a (τ, ρ) -triangle LSO as assumed by the theorem. Let $\nu' = \frac{\nu}{\tau}$. For every ordering $\sigma \in \Sigma$, we form an unweighted path graph P_σ with vertex set X and the order of vertices along the path is σ . We construct a ν' -reliable 2-hop spanner $H_\sigma(X, E_\sigma, w_\sigma)$ for P_n^σ with

$n \cdot O(\log^2 n + \nu'^{-1} \log n \cdot \log \log n)$ edges by Lemma 12. Note that for every edge $\{u, v\} \in E_\sigma$ of H_σ , $w_\sigma(u, v)$ is the distance between u and v in the (unweighted) path graph P_n^σ .

We form a new weight function w_X that assign each edge $\{u, v\} \in E_\sigma$ a weight $w_X(u, v) = d_X(u, v)$. The reliable spanner for (X, d_X) is $H = \bigcup_{\sigma \in \Sigma} H_\sigma(X, E_\sigma, w_X)$. We observe that the total number of edges in H is bounded by

$$\sum_{\sigma \in \Sigma} n \cdot O(\log^2 n + \nu'^{-1} \log n \cdot \log \log n) = n\tau \cdot O(\log^2 n + \nu^{-1} \tau \log n \cdot \log \log n) .$$

Let $B \subseteq X$ be an oblivious attack. Let B_σ^+ be the faulty extension of B in H_σ , and $B^+ = \bigcup_{\sigma \in \Sigma} B_\sigma^+$ be the faulty extension of B in H . We observe that:

$$\mathbb{E}[|B^+|] \leq |B| + \sum_{\sigma} \mathbb{E}[|B_\sigma^+ \setminus B|] \leq |B| + \tau\nu' \cdot |B| \leq (1 + \nu) \cdot |B| .$$

It remains to show the stretch guarantee of H . For every pair of points $x, y \notin B^+$, let $\sigma \in \Sigma$ be the ordering that satisfies the ordering property for x and y : for every $a, b \in X$ such that $x \leq_\sigma a \leq_\sigma b \leq_\sigma y$, $d_X(a, b) \leq \rho \cdot d_X(x, y)$. (Here we assume w.l.o.g. that $x \leq_\sigma y$.) Since $x, y \notin B^+$, $x, y \notin B_\sigma^+$. Since $H_\sigma(X, E_\sigma, w_\sigma)$ is a 2-hop 1-spanner for P_σ , there must be $z \notin B$ such that $x \leq_\sigma z \leq_\sigma y$ and $\{x, z\}, \{z, y\} \in E_\sigma$. We conclude that

$$d_H(x, y) \leq w_X(x, z) + w_X(z, y) = d_X(x, z) + d_X(z, y) \leq 2\rho \cdot d_X(x, y) ;$$

the theorem follows.

5.2 2-hop spanner: Proof of Lemma 12

In this section, we construct a 2-hop reliable spanner for the path graph P_n . Our construction uses the concept of a *shadow* introduced by Buchin, *et al.* [BHO19].

Definition 7. Let B be a subset of $[n]$. The left α -shadow of B is all the vertices b such for some $a < b$, $|[a : b] \cap B| \geq \alpha \cdot |[a : b]|$, denoted by $\mathcal{S}_L(\alpha, B)$. The right α -shadow $\mathcal{S}_R(\alpha, B)$ is defined symmetrically. The set $\mathcal{S} = \mathcal{S}_L(\alpha, B) \cup \mathcal{S}_R(\alpha, B)$ is called the α -shadow of B .

Buchin *et al.* [BHO19] showed that for every $\alpha \in (\frac{2}{3}, 1)$ and any $B \subseteq [n]$, $|\mathcal{S}_L(\alpha, B)| \leq \frac{|B|}{2\alpha-1}$. In particular, for $\beta < \frac{1}{3}$,

$$|\mathcal{S}(1 - \beta, B) \setminus B| \leq \frac{|B|}{2(1 - \beta) - 1} - |B| \leq (1 + 6\beta)|B| - |B| = 6\beta \cdot |B| . \quad (1)$$

We also use a construction of a (non-reliable) 2-hop 1-spanner for path graph P_n with $O(n \log n)$ edges [AS87, BTS94].

Lemma 13 ([AS87, BTS94]). *There is a 2-hop 1-spanner for P_n with $O(n \log n)$ edges.*

We are now ready to prove Lemma 12, which we restate below for convenience.

Lemma 12 (2-hop-spanner). *For every $\nu \in (0, 1)$, the path graph P_n admits an oblivious ν -reliable, 2-hop 1-spanner H with $n \cdot O(\log^2 n + \nu^{-1} \log n \cdot \log \log n)$ edges.*

Proof. We will assume that $\nu > \frac{1}{n}$, as otherwise, we can take all the possible edges to the spanner. We will construct a spanner with reliability parameter $O(\nu)$; the lemma will follow by rescaling.

Let $M = \log n$, and for every $j \in [0, M]$, let $p_j = c(\log n + \nu^{-1} \ln(\log n)) \cdot 2^{-j}$ for a constant c to be determined later. We construct a vertex set N_j by sampling each vertex with probability $\min\{1, p_j\}$. Vertices in N_j are called *centers*.

The spanner H is constructed as follows: For every vertex v , and parameter $j \in [0, M]$, we add edges from v to $[v - 2^{j+1} : v + 2^{j+1}] \cap N_j$. That is to every vertex $u \in N_j$ such that $|v - u| \leq 2^{j+1}$. We also add to H a 2-hop spanner of size $O(n \log n)$ guaranteed by Lemma 13.

To bound the number of edges, we charge only to the centers; each edge between two different centers is charged twice. The expected number of edges charged to a vertex a is

$$\sum_{j=0}^M \Pr[a \in N_j] \cdot 2^{j+1} \cdot 2 = c(\log n + \nu^{-1} \ln \log n) \cdot \sum_{j=1}^M 2^{-j} \cdot 2^{j+2} = O(\log n) \cdot (\log n + \nu^{-1} \ln \log n),$$

implying that the expected number of edges is $n \cdot O(\log^2 n + \nu^{-1} \log n \cdot \log \log n)$. The worst-case guarantee promised in the lemma follows by using Lemma 2.

Let B be an oblivious attack. The faulty extension B^+ will consist of B and all the vertices v such that for some j , (a) $v > 2^j$ and $[v - 2^j, v] \cap N_j \subseteq B$ or (b) $v \leq n - 2^j$ and $[v, v + 2^j] \cap N_j \subseteq B$. This way, for every pair of vertices $u, v \notin B^+$ whose distance is in $[2^i, 2^{i+1})$, there is a vertex $z \in N_i \setminus B$ such that both will have edges to z . This implies the monotone 2-hop property for u and v .

Finally we analyze the expected size of B^+ . If $B = \emptyset$, then $B^+ = \emptyset$ as we added a 2-hop spanner. Otherwise, a vertex $v \notin B$ joins B^+ iff $N_j \cap [v : v + 2^j] \subseteq B$ or $N_j \cap [v - 2^j : v] \subseteq B$ for some $j \in [0, M]$.

Fix an $i \leq \lfloor \log \frac{1}{3\nu} \rfloor$. For every vertex $a \notin \mathcal{S}_L(1 - 2^i \cdot \nu, B)$ and every $j \in [0, M]$, it holds that $[a, a + 2^j] \cap B \leq (1 - 2^i \cdot \nu) \cdot 2^j$ by the definition of a left shadow. In particular, there are at least $2^{i+j} \cdot \nu$ vertices in $[a, a + 2^j] \setminus B$. Vertex a joins B^+ iff for some j , $([a : a + 2^j] \setminus B) \cap N_j = \emptyset$ or $([a - 2^j : a] \setminus B) \cap N_j = \emptyset$. (Note that if the distance from a to either 1 or n is smaller than 2^j , then it cannot join B^+ by definition.) We observe that:

$$\begin{aligned} \Pr\left([a, a + 2^j] \cap N_j \subseteq B \mid a \notin \mathcal{S}(1 - 2^i \cdot \nu, B)\right) &\leq (1 - 2^{-j} \cdot c \cdot \nu^{-1} \cdot \ln(\log n))^{2^{i+j} \cdot \nu} \\ &\leq e^{-2^{-j} \cdot c \cdot \nu^{-1} \cdot \ln \log n \cdot 2^{i+j} \cdot \nu} = (\log n)^{-c 2^i}. \end{aligned}$$

Thus, the probability that a is added to B^+ is at most:

$$\begin{aligned} \Pr[a \in B^+ \mid a \notin \mathcal{S}(1 - 2^i \cdot \nu, B)] &\leq \sum_{j=1}^M 2 \cdot \Pr\left([a, a + 2^j] \cap N_j \subseteq B \mid a \notin \mathcal{S}(1 - 2^i \cdot \nu, B)\right) \\ &\leq \sum_{j=1}^M 2 \cdot (\log n)^{-c 2^i} \leq 2 \cdot (\log n)^{-c 2^i - 1}. \end{aligned}$$

On the other hand, for a vertex $a \notin \mathcal{S}(\frac{2}{3}, B)$, we have that

$$\Pr\left([a, a + 2^j] \cap N_j \subseteq B \mid a \notin \mathcal{S}(\frac{2}{3}, B)\right) \leq (1 - 2^{-j} \cdot c \cdot \ln n)^{\frac{1}{3} \cdot 2^j} \leq e^{-2^{-j} \cdot c \cdot \ln n \cdot \frac{1}{3} \cdot 2^j} = n^{-\frac{c}{3}},$$

implying that

$$\begin{aligned} \Pr \left[a \in B^+ \mid a \notin \mathcal{S}\left(\frac{2}{3}, B\right) \right] &\leq \sum_{j=1}^M 2 \cdot \Pr \left[[a, a+2^j] \cap N_j \subseteq B \mid a \notin \mathcal{S}\left(\frac{2}{3}, B\right) \right] \\ &\leq \sum_{j=1}^M 2 \cdot n^{-\frac{c}{3}} \leq 2 \log n \cdot n^{-\frac{c}{3}}. \end{aligned}$$

We conclude that

$$\begin{aligned} \mathbb{E}[B^+] &\leq |\mathcal{S}(1-\nu, B)| + \sum_{i=1}^{\log \frac{1}{3\nu}} |\mathcal{S}(1-2^i \cdot \nu, B) \setminus B| \cdot \Pr[a \in B^+ \mid a \notin \mathcal{S}(1-2^{i-1} \cdot \nu, B)] \\ &\quad + \left| [n] \setminus \mathcal{S}\left(\frac{2}{3}, B\right) \right| \cdot \Pr \left[a \in B^+ \mid a \notin \mathcal{S}\left(\frac{2}{3}, B\right) \right] \\ &\stackrel{(1)}{\leq} (1+4\nu)|B| + \sum_{i=1}^{\log \frac{1}{3\nu}} 2^{i+2} \nu \cdot |B| \cdot 2 \cdot (\log n)^{-c2^{i-1}} + n \cdot 2 \log n \cdot n^{-\frac{c}{3}} = (1+O(\nu))|B|, \end{aligned}$$

for a sufficiently large constant c . $\square$

6 Left-sided Locality Sensitive Ordering

We use a left-sided LSO to construct reliable spanners with *optimal stretch* for trees, planar graphs, bounded treewidth graphs, and minor free graphs. We say an ordering of (X, d_X) is *partial* if it is a linear ordering on a *subset* of points in X .

Definition 8 $((\tau, \rho)$ -left-sided LSO). *Given a metric space (X, d_X) , we say that a collection Σ of partial orderings is a (τ, ρ) -left-sided LSO if every point $x \in X$ belongs to at most τ orderings in the collection, and for every $x, y \in X$, there is an order $\sigma \in \Sigma$ such that for every $x' \preceq_\sigma x$ and $y' \preceq_\sigma y$ it holds that $d_X(x', y') \leq \rho \cdot d_X(x, y)$.*

Unlike the classic LSO and the triangle-LSO, a (τ, ρ) -left-sided LSO may contain $\Omega(n)$ (partial) orderings; it only guarantees that each point belongs to at most τ (partial) orderings. Note that the stretch guarantee of a (τ, ρ) -left-sided LSO also fulfills the stretch requirement of a (τ, ρ) -LSO. To see this, consider a pair of points $x, y \in X$, and let $\sigma \in \Sigma$ be the ordering guaranteed above, where w.l.o.g $x \preceq_\sigma y$. Then for every z such that $x \preceq_\sigma z \preceq_\sigma y$, it holds that $d_X(x, z) \leq \rho \cdot d_X(x, y)$ (as $x \preceq_\sigma z$ and $z \preceq_\sigma y$). In particular $z \in B_X(x, \rho \cdot d_X(x, y))$. However, there is no $f(\cdot)$ such that a (τ, ρ) -left-sided LSO will be guaranteed to be an $(f(\tau), \rho)$ -LSO as the number of orderings in a (τ, ρ) -left-sided LSO, by definition, could be $\Omega(n)$ even if τ is a constant. We show that minor-free metrics admit a left-sided LSO where each point belongs to a small number of orderings.

Theorem 6. *Treewidth- k metrics admit a $(k \log n, 1)$ -left-sided LSO. In particular, tree metrics admit a $(\log n, 1)$ -left-sided LSO.*

Theorem 7. *For every $\epsilon \in (0, \frac{1}{2})$, every graph with $\text{SPDdepth } k$ admits an $(O(\frac{k}{\epsilon} \log n), 1 + \epsilon)$ -left-sided-LSO. In particular, every graph excluding a fixed minor (e.g. a planar graph) admits an $(O(\frac{1}{\epsilon} \log^2 n), 1 + \epsilon)$ -left-sided-LSO.*

The proof of Theorem 6 is deferred to Section 6.1. The definition of SPDdepth and the proof of Theorem 7 are deferred to Section 6.2.

Reliable Spanners from left-sided LSO's. Similar to Theorem 4, we can construct a reliable spanner with a nearly linear number of edges from a left-sided LSO.

Theorem 8 (Reliable Spanner From Left-sided LSO). *Consider an n -point metric space (X, d_X) that admits a (τ, ρ) -left-sided LSO. Then for every $\nu \in (0, 1)$, X admits an oblivious ν -reliable, (2ρ) -spanner with $n \cdot O(\nu^{-1}\tau^2 \log n)$ edges.*

The proof of Theorem 8 is deferred to Section 6.3. By applying Theorem 8 on Theorem 6 and Theorem 7 we have:

Theorem 9. *Treewidth- k graphs admit oblivious ν -reliable 2-spanners with $n \cdot O(\nu^{-1}k^2 \log^3 n)$ edges. In particular, trees admit oblivious ν -reliable 2-spanners with $n \cdot O(\nu^{-1} \log^3 n)$ edges.*

Theorem 10. *For every $\epsilon \in (0, \frac{1}{2})$, every graph with $\text{SPDdepth } k$ admits a ν -reliable $2(1+\epsilon)$ -spanner with $n \cdot O(\nu^{-1} \frac{k^2}{\epsilon^2} \log^3 n)$ edges. In particular, every graph excluding a fixed minor (e.g. a planar graph) admits a ν -reliable $2(1+\epsilon)$ -spanner with $n \cdot O(\nu^{-1}\epsilon^{-2} \log^5 n)$ edges.*

Finally, we observe that the stretch of reliable spanners in Theorem 9 and Theorem 10 is essentially optimal even for tree metrics.

Observation 1. *For every $\nu \in (0, 1)$, and $n \in \mathbb{N}$, there is an n -vertex unweighted tree $T = (V, E)$ such that every oblivious ν -reliable spanner with stretch $t < 2$ has $\Omega(n^2)$ edges.*

Proof. Let T be the star graph. That is, there is a single vertex r of degree $n-1$ and vertices of degree 1, and let H be a ν -reliable spanner of stretch t . Let $B = \{r\}$ be an attack. Every pair of vertices $u, v \notin B^+$ must be adjacent, as every path of length at least 2 has weight $4 > t \cdot d_T(u, v)$. It follows that H must contain $\mathbb{E} \left[\binom{|V \setminus B^+|}{2} \right] = \Omega(n - \mathbb{E}[|B^+|])^2 = \Omega(n^2)$ edges. $\square$

Remark 1. *Note that Observation 1 will also hold for ultrametrics, as the uniform metric on n points is an ultrametric (here the attack B could be $\emptyset$). In particular, the stretch parameter in Corollary 4 is tight.*

6.1 k -separable graphs

Definition 9 (k -separable graph family). *A graph family $\mathcal{G}$ is k -separable if for every graph $G = (V, E, w) \in \mathcal{G}$, and every induced subgraph $G' = G[V']$ for $V' \subseteq V$, there is a set $K \subseteq V'$ of at most k vertices such that each connected component in $G' \setminus K$ contains at most $|V'|/2$ vertices.*

Note that trees belong to a 1-separable graph family and treewidth k graphs belong to a k -separable graph family. Hence Theorem 6 is an immediate corollary of the following lemma.

Lemma 14 (k -Separators to LSO). *Let $\mathcal{G}$ be a k -separable family, then every n -vertex graph $G = (V, E, w) \in \mathcal{G}$ admits a $(k \log n, 1)$ -left-sided LSO.*

Proof. We recursively construct a set of orderings Σ ; initially, $\Sigma = \emptyset$. Let K be a separator containing at most k vertices such that every connected component in $G \setminus K$ has size at most $\frac{n}{2}$. For every $r \in K$, let l_r be an order of V w.r.t. distances from r . That is $l_r = v_1, v_2, \dots, v_n$ where $d_G(v_1, r) \leq d_G(v_2, r) \leq \dots \leq d_G(v_n, r)$. The set of k orderings $\{l_r\}_{r \in K}$ are added to Σ .

Let $\{G_1, \dots, G_s\}$ be the set of connected components of $G \setminus K$. For each G_i , let Σ_i be the set of orderings obtained by recursively applying the construction to G_i ; each ordering in Σ_i is a partial ordering of G . We then update $\Sigma \leftarrow \Sigma \cup (\cup_{i \in [1, s]} \Sigma_i)$. This completes the construction of Σ .

Next, we argue that Σ is a $(k \log n, 1)$ -left-sided LSO. By construction, each vertex belongs to at most $k \log n$ orderings since the depth of the recursion is at most $\log n$ and at each level of the recursion, each vertex belongs to at most k partial orderings.

Finally, consider a pair of vertices $u, v \in V$. Let $P_{u,v}$ be a shortest path from u to v in G . Let G' be the subgraph of G at highest depth in the recursive construction that contains all the vertices of $P_{u,v}$. Let K be the separator of G' of size at most k ; by definition, there is a vertex $r \in K \cap P_{u,v}$. As G' contains all the vertices of $P_{u,v}$ it holds that $d_{G'}(u, v) = d_{G'}(u, r) + d_{G'}(r, v) = d_G(u, r) + d_G(r, v)$. Let l_r be the ordering that is constructed for r . Let u', v' be any vertices such that $u' \preceq_{l_r} u$ and $v' \preceq_{l_r} v$. Then:

$$d_G(u', v') \leq d_G(u', r) + d_G(r, v') \leq d_{G'}(u', r) + d_{G'}(r, v') \leq d_{G'}(u, r) + d_{G'}(r, v) = d_G(u, v) ;$$

the lemma follows. $\square$

6.2 SPD Graphs

This section is devoted to proving Theorem 10. Abraham *et al.* [AFGN18] introduced the family of graphs with a shortest path decomposition (abbreviated SPD) of bounded depth.

Definition 10 (SPD depth). *A graph has an SPD of depth 1 if and only if it is a (weighted) path. A graph G has an SPD of depth $k \geq 2$ if there exists a shortest path P , such that deleting P from the graph G results in a graph whose connected components all have SPD of depth at most $k - 1$.*

It is known that (see [AFGN18]) that minor-free graphs have SPD of depth $k = O(\log n)$. However, the family of bounded SPDdepth graphs is much richer and contains dense graphs with K_r as a minor, for arbitrarily large r . To prove Theorem 10, we will show that graphs with an SPDdepth k admit a $(\frac{k}{\epsilon} \log \frac{n}{\epsilon}, 1 + \epsilon)$ -left-sided LSO, the second assertion in Theorem 10 is an immediate corollary. Our construction will rely on the following lemma, which implicitly appeared in [Kle02, Tho04] (see [BFN19a] for an explicit proof).

Lemma 15. *Consider a weighted graph $G = (V, E, w)$ with parameter $\epsilon \in (0, 1)$, and let P be a shortest path in G . Then one can find for each $v \in V$ a set of vertices, called landmarks, L_v on P of size $|L_v| = O(\frac{1}{\epsilon})$, such that for any $v, u \in V$ whose shortest path between them intersects P , there exists $x \in L_v$ and $y \in L_u$ satisfying $d_G(v, x) + d_P(x, y) + d_G(y, u) \leq (1 + \epsilon) \cdot d_G(v, u)$.*

Proof of Theorem 10. We will assume that $\epsilon > \frac{1}{n^2}$, as otherwise we can simply return $\binom{n}{2}$ orderings, where for every pair u, v , there is an ordering where u, v are the first two vertices. We recursively construct a set of orderings Σ ; initially, $\Sigma = \emptyset$. Let P be a shortest path of G such that every component of $G \setminus P$ has an SPD of depth $k - 1$.

We construct a set of orderings for G from P to add to Σ as follows. Let $\{L_v\}_{v \in V}$ be the set of landmarks provided by Lemma 15 w.r.t. P . For every vertex $v \in V$, denote $l(v) = |L_v|$. We construct an auxiliary tree T that initially contains the shortest path P only. For every vertex $v \in V$, let $L_v = \{x_1, \dots, x_{l(v)}\} \subseteq P$. We will add the vertices $\{v_i\}_{i=1}^{l(v)}$, which are $l(v)$ copies of v , to T , and for every i connect v_i to x_i with an edge of weight $d_G(v, x_i)$. Note that T has $\sum_{v \in V} l(v) = O(\frac{n}{\epsilon})$ vertices. The distances in T dominate the distances in the original graph, that

is for every v_i, u_j copies of v, u respectively, it holds that $d_T(v_i, u_j) \geq d_G(v, u)$. Additionally, by Lemma 15, for every pair of vertices u, v , such that the shortest path between them intersects P , it holds that $\min_{i,j} d_T(v_i, u_j) \leq (1 + \epsilon) \cdot d_G(v, u)$. Let Σ_T be an $(O(\log \frac{n}{\epsilon}), 1)$ -left-sided LSO for T provided by Lemma 14. Let Σ'_T be the collection of orderings Σ_T where we delete duplicated occurrences. That is, for every ordering $\sigma \in \Sigma$ and vertex v , σ might contain multiple copies of v : $v_{i_1}, v_{i_2}, \dots$, we will keep only the leftmost copy of v in σ ; the resulting ordering is denoted by σ' . As $\epsilon > n^{-2}$, each copy v_i of a vertex v appears in at most $O(\log \frac{n}{\epsilon}) = O(\log n)$ orderings in Σ_T . As each vertex has $O(\frac{1}{\epsilon})$ copies, we conclude that each vertex v appears in at most $O(\frac{1}{\epsilon} \cdot \log n)$ orderings in Σ'_T .

Let $\{G_1, \dots, G_s\}$ be the set of connected components of $G \setminus P$. For each G_i , let Σ_i be the set of orderings obtained by recursively applying the construction to G_i ; each ordering in Σ_i is a partial ordering of G . We then set $\Sigma \leftarrow \Sigma'_T \cup (\cup_{i \in [1, s]} \Sigma_i)$. This completes the construction of Σ .

Next, we argue that Σ is an $(O(\frac{k}{\epsilon} \cdot \log n), 1)$ -left-sided LSO. Since the recursion depth is k , and each vertex belongs to at most $O(\frac{1}{\epsilon} \cdot \log n)$ partial orderings in each recursive level, it follows that each vertex belongs to at most $O(\frac{k}{\epsilon} \cdot \log n)$ partial orderings in Σ .

Finally, consider a pair of vertices $u, v \in V$. Let $P_{u,v}$ be a shortest path from u to v in G . Let G' be the subgraph of G at highest depth in the recursive construction that contains all the vertices of $P_{u,v}$. In particular $d_{G'}(v, u) = d_G(v, u)$. Let P be the shortest path deleted from G' . Let T be the tree constructed for P, G' . By maximality $P \cap P_{u,v} \neq \emptyset$, thus there are two copies v_i, u_j of v, u such that $d_T(v_i, u_j) \leq (1 + \epsilon) \cdot d_{G'}(v, u) = (1 + \epsilon) \cdot d_G(v, u)$. By Lemma 14, Σ_T contains an ordering σ such that $v_i, u_j \in \sigma$, and for every $x \preceq_\sigma v_i$, $y \preceq_\sigma u_j$ it holds that $d_T(x, y) \leq d_T(v_i, u_j)$. The collection of orderings Σ'_T contains an ordering σ' , which is simply σ with removed duplicates. In particular, σ' contains only original vertices from V . Let $v_{i'}, u_{j'}$ be the leftmost duplicates of v, u (respectively) in σ . For every $x, y \in V$ such that $x \preceq_{\sigma'} v$, $y \preceq_{\sigma'} u$, there are copies x_p, y_q of x, y such that $x_p \preceq_\sigma v_{i'} \preceq_\sigma v_i$ and $y_q \preceq_\sigma u_{j'} \preceq_\sigma u_j$. It thus holds that

$$d_G(x, y) \leq d_T(x_p, y_q) \leq d_T(v_i, u_j) \leq (1 + \epsilon) \cdot d_G(v, u) .$$

This completes the proof of Theorem 10. □

6.3 Reliable Spanners from Left-sided LSO

To create a reliable spanner from a left-sided LSO, we will use *left spanners*:

Definition 11 (Reliable Left Spanner). *Given a path graph P_n , a 2-hop left spanner H is a graph such that for every $a < b$ there is $c \leq a, b$, such that $\{v_c, v_a\}, \{v_c, v_b\} \in H$. A distribution over 2-hop left spanners H is ν -reliable if for every attack $B \subseteq X$, there is a set $B \subseteq B^+$ such that $\mathbb{E}[|B^+|] \leq (1 + \nu)|B|$, and for every $a < b$ such that $v_a, v_b \notin B^+$, there is $c \leq a, b$, $c \notin B$ such that $\{v_c, v_a\}, \{v_c, v_b\} \in E(H)$. An oblivious spanner $\mathcal{D}$ is said to have m edges if every spanner in the support $\mathcal{D}$ has at most m edges (see Lemma 2).*

In the following lemma, we construct a reliable left spanner. Interestingly, it is sparser than the reliable 2-hop spanner for the path graph constructed in Lemma 12. The proof is deferred to Section 6.3.1.

Lemma 16 (2-hop-left-spanner). *For every $\nu \in (0, 1)$, the n -vertex path P_n admits an oblivious ν -reliable 2-hop left spanner with $n \cdot O(\nu^{-1} \log n)$ edges.*

Using Lemma 16, we can construct a reliable spanner using a left-sided LSO, and thus prove Theorem 8. The meta theorem is restated bellow for convenience.

Theorem 8 (Reliable Spanner From Left-sided LSO). *Consider an n -point metric space (X, d_X) that admits a (τ, ρ) -left-sided LSO. Then for every $\nu \in (0, 1)$, X admits an oblivious ν -reliable, (2ρ) -spanner with $n \cdot O(\nu^{-1}\tau^2 \log n)$ edges.*

Proof. Let Σ be a (τ, ρ) -left-sided LSO; Σ is a collection of partial orderings. Set $\nu' = \frac{\nu}{\tau}$. For every ordering $\sigma \in \Sigma$ which contains n_σ vertices, using Lemma 16, we construct a ν' -reliable 2-hop left spanner H_σ with $n_\sigma \cdot O(\nu'^{-1} \log n_\sigma)$ edges. Set $H = \cup_\sigma H_\sigma$. The total number of edges is thus bounded by

$$|H| = \sum_\sigma |H_\sigma| = \sum_\sigma n_\sigma \cdot O(\nu'^{-1} \log n_\sigma) = O(\nu'^{-1} \log n) \cdot \sum_\sigma n_\sigma = n \cdot O(\nu^{-1}\tau^2 \log n) .$$

For an attack B , let B_σ^+ be the faulty extension of B w.r.t. H_σ , and let $B^+ = \cup_\sigma B_\sigma^+$. For every pair of points $x, y \notin B^+$, let σ be the promised left-sided partial ordering from Definition 8. As $x, y \notin B_\sigma^+$, there is some vertex $z \notin B$ such that $z \leq_\sigma x, y$ and $(x, z), (y, z) \in E(H_\sigma)$. We conclude that

$$d_H(x, y) \leq d_X(x, z) + d_X(z, y) \leq 2\rho \cdot d_X(x, y) .$$

To bound the size of B^+ , we observe that

$$\mathbb{E}[|B^+|] \leq |B| + \sum_\sigma \mathbb{E}[|B_\sigma^+ \setminus B|] \leq |B| + \tau\nu' \cdot |B| \leq (1 + \nu) \cdot |B| .$$

The theorem now follows. □

6.3.1 Left Spanner: proof of Lemma 16

We restate the lemma for convenience.

Lemma 16 (2-hop-left-spanner). *For every $\nu \in (0, 1)$, the n -vertex path P_n admits an oblivious ν -reliable 2-hop left spanner with $n \cdot O(\nu^{-1} \log n)$ edges.*

Proof. We will construct a spanner with reliability parameter $O(\nu)$; afterward, the constants can be readjusted accordingly. Let h be a universal constant such that $H_s - H_t = \frac{1}{t+1} + \frac{1}{t+2} + \dots + \frac{1}{s} > h \cdot \ln \frac{s}{t}$. We create a set N of centers as follows: sample each vertex $a \in [n]$ to N with probability $\frac{1}{a} \cdot \frac{c}{\nu}$ (or 1 for small enough a) where $c = \max\{\frac{2}{h}, 1\}$. The spanner H is then defined as follows: for each vertex $a \in V$ set $N_a = [1, a] \cap N$. We add edges between a and all the vertices in N_a (alternatively, each vertex add edges to all the centers with smaller indices). Clearly, the expected number of edges is bounded by

$$\sum_{a=1}^n \frac{1}{a} \cdot \frac{c}{\nu} \cdot (n - a) \leq n \cdot \frac{c}{\nu} \cdot \sum_{a=1}^n \frac{1}{a} = n \cdot O(\nu^{-1} \log n) .$$

By Lemma 2, we can obtain the same guarantee on the number of edges also in the worst case.

Given an attack B , we add a vertex a to B^+ iff $N_a \subseteq B$. The stretch bound easily follows: for a pair of vertices $a < b \notin B^+$ there is a center $x \in N_a \subseteq N_b$, thus $\{x, a\}, \{x, b\} \in H$ as required.

Finally we bound the expected size of B^+ for an attack B . We will begin by analyzing the specific attack $B = [1, x]$ for some $x \in [n]$. Later, we will argue that our analysis holds also for an

arbitrary attack B . The vertex $x+i$ joins B^+ only if none of the vertices $x+1, \dots, x+i$ becomes a center in N . The probability of this event is

$$\Pr[x+i \in B^+] = \prod_{j=1}^i \left(1 - \frac{1}{x+j} \cdot \frac{c}{\nu}\right) < e^{-\frac{c}{\nu} \sum_{j=1}^i \frac{1}{x+j}} < e^{-\frac{c}{\nu} \cdot h(\ln \frac{x+i}{x})} \leq \left(1 + \frac{i}{x}\right)^{-\frac{2}{\nu}}.$$

The expect size of $B^+ \setminus B$ is thus bounded by

$$\mathbb{E}[|B^+ \setminus B|] < \sum_{i=1}^{n-x} \left(1 + \frac{i}{x}\right)^{-\frac{2}{\nu}} < \sum_{i \geq 1} \left(1 + \frac{i}{x}\right)^{-\frac{2}{\nu}}$$

For every $s \geq 0$, and $i \in [s\nu x, (s+1)\nu x)$, we have that $\left(1 + \frac{i}{x}\right)^{-\frac{2}{\nu}} < (1 + s\nu)^{-\frac{2}{\nu}} < e^{-\frac{s\nu}{2} \cdot \frac{2}{\nu}} = e^{-s}$. We conclude:

$$\mathbb{E}[|B^+ \setminus B|] \leq \sum_{s \geq 0} \sum_{i \in [s\nu x, (s+1)\nu x)} \left(1 + \frac{i}{x}\right)^{-\frac{2}{\nu}} < \nu x \sum_{s \geq 0} e^{-s} = O(\nu) \cdot |B|.$$

It follows that $\mathbb{E}[|B^+|] = (1 + O(\nu))|B|$ as required. Finally, consider an arbitrary attack $B \subseteq [n]$ of size x . Let $a_1, a_2, \dots, a_{n-x}$ the vertices not in B , ordered from left to right. Note that a_i has i vertices to the left (or equal) of it not in B , all from the interval $[1, x+i]$. To bound $\Pr[a_i \in B^+]$, we observe that this probability is maximum when $B = [1, x]$ since the probability that vertices are sampled to N decreases monotonically. Thus, $\Pr[a_i \in B^+] \leq \prod_{j=1}^i \left(1 - \frac{1}{x+j} \cdot \frac{c}{\nu}\right)$. The rest of the analysis of $|B^+|$ follows exactly the same argument. $\square$

7 Subgraph Reliable Connectivity Preserves

Classically, throughout the literature, given a graph $G = (V, E, w)$ a spanner H of G is required to be a subgraph of G . A natural question in the context of this paper is: Is it possible to construct a reliable spanner with a subquadratic number of edges that only uses edges of the input graph? Since removing a single vertex could disconnect the graph into two equal parts, a necessary relaxation is to require that distances are preserved only w.r.t. the induced graph $G[V \setminus B]$ (similarly to the case of fault-tolerant spanners). Here we show that even for the less ambitious task of constructing a ν -reliable connectivity preserver, it is not possible without using $\Omega(n^2)$ edges. To this end, we introduce a generalized definition of reliability, and present an oblivious lower bound and a matching deterministic upper bound.

Definition 12 (Reliable Connectivity Preserver). *Given a graph $G = (V, E)$, and a monotonically non-decreasing function $g : \mathbb{N} \rightarrow \mathbb{N}$, a subgraph H of G is a deterministic g -reliable connectivity preserver if for every attack $B \subseteq V$, there is a superset $B^+ \subseteq B^+$ of size at most $g(|B|)$, such that for every $u, v \in V \setminus B^+$, if u and v are connected in $G \setminus B$, then they are also connected in $H \setminus B$.*

An oblivious g -reliable connectivity preserver is a distribution $\mathcal{D}$ over subgraphs H of G , such that for every attack B and $H \in \text{supp}(\mathcal{D})$, there is an algorithm producing a superset B^+ of B such that, for every $u, v \notin B^+$, if u and v are connected in $G \setminus B$, then they are also connected in $H \setminus B$, and furthermore, $\mathbb{E}_{H \sim \mathcal{D}}[|B^+|] \leq g(|B|)$. We say that an oblivious g -reliable connectivity preserver $\mathcal{D}$ has m edges if every graph H in the support of $\text{supp}(\mathcal{D})$ has at most m edges.

Note that using the notation of Definition 12, the rest of the paper is concerned with constructing oblivious g -reliable (non-subgraph) spanners for the linear function $g(x) = (1 + \nu)x$. We begin with the upper bound.

Theorem 11. *Let $k > 1$ be an integer and $g_k(x) = \Omega(x^k)$. Then every n -vertex graph admits a deterministic g_k -reliable connectivity preserver with $O(n^{1+1/k})$ edges.*

Proof. The construction is by reduction to f -fault-tolerant spanners. Recall that a subgraph H of G is an f -vertex-fault-tolerant t -spanner if for every subset B of at most f vertices, for every $u, v \in V \setminus B$, $d_{H \setminus B}(u, v) \leq t \cdot d_{G \setminus B}(u, v)$. Bodwin and Patel [BP19] constructed f -vertex-fault-tolerant $2t - 1$ -spanners with $O(n^{1+\frac{1}{t}} \cdot f^{1-\frac{1}{t}})$ edges (later Bodwin, Dinitz, and Robelle [BDR21] showed how to construct such spanners efficiently). Assume that $g(x) \geq cx^k$ for some constant c . Set $f = (n/c)^{\frac{1}{k}}$, and construct an f -vertex-fault-tolerant $O(\log n)$ -spanner H with $O(n \cdot f) = O(n^{1+1/k})$ edges. We argue that H is a deterministic g_k -reliable connectivity preserver. Clearly H has $O(n^{1+1/k})$ edges. Consider an attack $B \subseteq V$,

- If $|B| \leq f$, set $B^+ = B$. Clearly $|B^+| \leq g(|B|)$, and for every $u, v \in V \setminus B^+$, if u, v are connected in $G \setminus B$, then $d_{H \setminus B}(u, v) \leq t \cdot d_{G \setminus B}(u, v) < \infty$, and in particular u and v are connected in $H \setminus B$.
- Else, $|B| > f$, set $B^+ = V$. Then $|B^+| = n = c((n/c)^{\frac{1}{k}})^k \leq g(|B|)$. The connectivity preservation holds trivially.

□

Note that in Theorem 11 we actually constructed a g_k -reliable deterministic $O(\log n)$ spanner. One can also observe that by the same construction, we get the following corollary.

Corollary 5. *Let $k > 1$ be an integer and $g_k(x) = \Omega(x^k)$. Then every graph n -vertex graph $G = (V, E, w)$ admits a deterministic g_k -reliable $(2t - 1)$ -spanner with $O(n^{1+\frac{t+k-1}{t \cdot k}})$ edges that only uses edges of G .*

Next, we provide a matching lower bound for the connectivity preserver; our lower bound is actually stronger as it holds in the oblivious case.

Theorem 12. *Let $k > 1$ be an integer and $g_k(x) = O(x^k)$. For every large enough n , there is an n -vertex graph G such that every oblivious g_k -reliable connectivity preserver has $\Omega(n^{1+1/k})$ edges.*

Proof. For simplicity, we show the lower bound for $g_k(x) = x^k$; the lower bound could be extended to any function $g_k(x) \leq c_0 \cdot x^k$ for some constant c_0 by adjusting the constants in our proof. We define a graph G to be a “thick” cycle, specifically, for $c = 12$ we will have $s = cn^{1-\frac{1}{k}}$ sets $A_0, A_1, \dots, A_{s-1}$, each containing $\frac{1}{c}n^{1/k}$ vertices. There will be an edge between $v \in A_i$ to $u \in A_j$ if and only if $i - j = 1 \pmod{s}$. All the index computation in the rest of the proof will be done modulo s , and we will omit it in the notation. We will also ignore rounding issues (which can be easily fixed). Note that $|V| = n$, while $|E| = cn^{1-\frac{1}{k}} \cdot \frac{1}{c}n^{\frac{2}{k}} = \frac{1}{c}n^{1+1/k}$. See Figure 4 for illustration.

We begin by proving that every deterministic reliable connectivity preserver has $\Omega(n^{1+1/k})$ edges. Later, we will generalize to the oblivious case. We argue that only the graph G itself is a deterministic g_k -reliable connectivity preserver. Suppose for contradiction that there is a g_k -reliable connectivity preserver H missing the edge $\{u, v\}$, where $u \in A_i$ and $v \in A_{i+1}$. Consider an attack $B = V_i \cup V_{i+1} \cup V_{i+\frac{s}{2}} \setminus \{u, v\}$. Note that $|B| < \frac{3}{c} \cdot n^{1/k}$. Note that the graph $G \setminus B$ is connected. The graph G' obtained by removing the edge $\{u, v\}$ from $G \setminus B$ has two connected components




Figure 4: On the left illustrated the thick cycle consisting of 12 sets $A_0, \dots, A_{11}$, where there is complete bipartite graph between A_i to A_{i+1} . On the right, illustrated the graph G' lacking an edge $\{u, v\}$ colored in blue. The attack $B = V_i \cup V_{i+1} \cup V_{i+\frac{s}{2}} \setminus \{u, v\}$ (encircled in red) disconnects the graph G' into two nearly equal halves, while keeping G connected.

$C_1 = \{v\} \cup A_{i+2} \cup \dots \cup V_{i+\frac{s}{2}-1}$, and $C_2 = V_{i+\frac{s}{2}-1} \cup \dots \cup V_{i-1} \cup \{u\}$, see Figure 4 for illustration. The set B^+ must contain all the vertices in either C_1 or C_2 , as otherwise there will be vertices connected in $G \setminus B$, but not connected in $H \setminus B \subseteq G' \setminus B$. It follows that $|B^+| \geq \min\{|C_1 \cup B|, |C_2 \cup B|\} \geq \frac{n}{2}$. But $g_k(|B|) = \left(\frac{3}{c} \cdot n^{1/k}\right)^k = \left(\frac{3}{c}\right)^k \cdot n < \frac{n}{4} \leq |B^+|$, a contradiction. It follows that every deterministic g_k -reliable connectivity preserver has $\Omega(n^{1+1/k})$ edges.

Next we generalize the lower bound to the oblivious case. Let $\mathcal{D}$ be a distribution over connectivity preservers, and suppose that there is an edge $\{u, v\}$ as above such that $\Pr_{H \sim \mathcal{D}}[e \in H] < \frac{1}{2}$. Then using the same argument as above, with the same attack B (w.r.t. $\{u, v\}$) we get that in all the preservers $H \in \text{supp}(\mathcal{D})$ not containing $\{u, v\}$, it holds that $|B^+| \geq \frac{n}{2}$. We conclude

$$\mathbb{E}_{H \sim \mathcal{D}}[|B^+|] \geq \Pr[e \notin H] \cdot \mathbb{E}_{H \sim \mathcal{D}}[|B^+| \mid e \notin H] \geq \frac{1}{2} \cdot \frac{n}{2},$$

As $g_k(|B|) < \frac{n}{4}$, it follows that the probability that H contains an arbitrary edge $\{u, v\}$ is at least $\frac{1}{2}$. In particular $\mathbb{E}_{H \sim \mathcal{D}}[|H|] \geq \frac{1}{2}|G| = \Omega(n^{1+1/k})$. The theorem now follows. $\square$

By setting $g(x) = (1 + \nu)x$, we obtain a lower bound $\Omega(n^2)$ on the number of edges of any oblivious ν -reliable *subgraph* connectivity preserver, and hence any oblivious ν -reliable *subgraph* t -spanner for any finite t .

8 Conclusions

In this paper, we have presented different types of locality-sensitive orderings and used them to construct reliable spanners. For the construction of the LSO's, we introduced and constructed ultrametric covers. Finally, in order to use the LSO's to construct reliable spanners, we construct 2-hop spanners and left spanners for the path graph. Several open questions naturally arise from our work:

1. Can we construct a ν -reliable 2-hop 1-spanner for the path graph P_n with $O(n \log n)$ edges for constant ν ? Note that a 2-hop spanner for the path graph P_n even without any reliability guarantee must contain $\Omega(n \log n)$ edges (see Exercise 12.10 [NS07]).
2. 1
3. Can we construct a reliable spanner of stretch 2 (as opposed to stretch $2 + \epsilon$ presented in this paper) for planar metrics with a nearly linear number of edges?

References

- [ABS⁺20] A. R. Ahmed, G. Bodwin, F. D. Sahneh, K. Hamm, M. J. L. Jebelli, S. G. Kobourov, and R. Spence. Graph spanners: A tutorial review. *Comput. Sci. Rev.*, 37:100253, 2020. doi:10.1016/j.cosrev.2020.100253.
- [ACE⁺20] I. Abraham, S. Chechik, M. Elkin, A. Filtser, and O. Neiman. Ramsey spanning trees and their applications. *ACM Trans. Algorithms*, 16(2):19:1–19:21, 2020. preliminary version published in SODA 2018. doi:10.1145/3371039.
- [AdBFG09] M. A. Abam, M. de Berg, M. Farshi, and J. Gudmundsson. Region-fault tolerant geometric spanners. *Discret. Comput. Geom.*, 41(4):556–582, 2009. Preliminary version published in SODA 2007. doi:10.1007/s00454-009-9137-7.
- [ADD⁺93] I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete Computational Geometry*, 9(1):81–100, 1993. doi:10.1007/BF02189308.
- [ADF⁺19] S. Alstrup, S. Dahlgaard, A. Filtser, M. Stöckel, and C. Wulff-Nilsen. Constructing light spanners deterministically in near-linear time. In *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*, pages 4:1–4:15, 2019. doi:10.4230/LIPIcs.ESA.2019.4.
- [ADM⁺95] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid. Euclidean spanners: Short, thin, and lanky. In *Proceedings of the Twenty-seventh Annual ACM Symposium on Theory of Computing, STOC '95*, pages 489–498, 1995.
- [AFGN18] I. Abraham, A. Filtser, A. Gupta, and O. Neiman. Metric embedding via shortest path decompositions. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 952–963, 2018. full version: <https://arxiv.org/abs/1708.04073>. doi:10.1145/3188745.3188808.
- [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. preliminary version published in STOC 2014. doi:10.1137/17M1112406.

- [AGM12] K. J. Ahn, S. Guha, and A. McGregor. Graph sketches: sparsification, spanners, and subgraphs. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2012, Scottsdale, AZ, USA, May 20-24, 2012*, pages 5–14, 2012. doi:10.1145/2213556.2213560.
- [Als08] B. Alspach. The wonderful Walecki construction. *Bull. Inst. Combin. Appl*, 52:7–20, 2008. see here.
- [AMS94] S. Arya, D. M. Mount, and M. H. M. Smid. Randomized and deterministic algorithms for geometric spanners of small diameter. In *Proc. of 35th FOCS*, pages 703–712, 1994.
- [AP90] B. Awerbuch and D. Peleg. Sparse partitions. In *Proceedings of the 31st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 503–513, 1990. doi:10.1109/FSCS.1990.89571.
- [AS87] N. Alon and B. Schieber. Optimal preprocessing for answering on-line product queries, 1987. Manuscript, see <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.79.1554>.
- [Bar96] Y. Bartal. Probabilistic approximations of metric spaces and its algorithmic applications. In *37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, Vermont, USA, 14-16 October, 1996*, pages 184–193, 1996. doi:10.1109/SFCS.1996.548477.
- [Bar21] Y. Bartal. Advances in metric ramsey theory and its applications. *CoRR*, abs/2104.03484, 2021. arXiv:2104.03484.
- [Bas08] S. Baswana. Streaming algorithm for graph spanners - single pass and constant processing time per edge. *Inf. Process. Lett.*, 106(3):110–114, 2008. doi:10.1016/j.ipl.2007.11.001.
- [BBM06] Y. Bartal, B. Bollobás, and M. Mendel. Ramsey-type theorems for metric spaces with applications to online problems. *J. Comput. Syst. Sci.*, 72(5):890–921, 2006. Special Issue on FOCS 2001. doi:10.1016/j.jcss.2005.05.008.
- [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.
- [BDG⁺20] A. S. Biswas, M. Dory, M. Ghaffari, S. Mitrovic, and Y. Nazari. Massively parallel algorithms for distance approximation and spanners. *CoRR*, abs/2003.01254, 2020. arXiv:2003.01254.
- [BDMS13] P. Bose, V. Dujmovic, P. Morin, and M. H. M. Smid. Robust geometric spanners. *SIAM J. Comput.*, 42(4):1720–1736, 2013. preliminary version published in SOCG 2013. doi:10.1137/120874473.
- [BDN22] G. Bodwin, M. Dinitz, and Y. Nazari. Vertex Fault-Tolerant Emulators. In M. Braverman, editor, *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:22, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ITCS.2022.25.

- [BDPW18] G. Bodwin, M. Dinitz, M. Parter, and V. V. Williams. Optimal vertex fault tolerant spanners (for fixed stretch). In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1884–1900, 2018. doi:10.1137/1.9781611975031.123.
- [BDR⁺12] C. Busch, C. Dutta, J. Radhakrishnan, R. Rajaraman, and S. Srivathsan. Split and join: Strong partitions and universal steiner trees for graphs. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 81–90. IEEE Computer Society, 2012. doi:10.1109/FOCS.2012.45.
- [BDR21] G. Bodwin, M. Dinitz, and C. Robelle. Optimal vertex fault-tolerant spanners in polynomial time. In D. Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 2924–2938. SIAM, 2021. doi:10.1137/1.9781611976465.174.
- [BDR22] G. Bodwin, M. Dinitz, and C. Robelle. Partially optimal edge fault-tolerant spanners. In J. S. Naor and N. Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference, January 9 - 12, 2022*, pages 3272–3286. SIAM, 2022. doi:10.1137/1.9781611977073.129.
- [BFH19] A. Bernstein, S. Forster, and M. Henzinger. A deamortization approach for dynamic spanner and dynamic maximal matching. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1899–1918, 2019. doi:10.1137/1.9781611975482.115.
- [BFM86] J. Bourgain, T. Figiel, and V. Milman. On Hilbertian subsets of finite metric spaces. *Israel J. Math.*, 55(2):147–152, 1986. doi:10.1007/BF02801990.
- [BFN19a] Y. Bartal, N. Fandina, and O. Neiman. Covering metric spaces by few trees. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, pages 20:1–20:16, 2019. doi:10.4230/LIPIcs.ICALP.2019.20.
- [BFN19b] Y. Bartal, A. Filtser, and O. Neiman. On notions of distortion and an almost minimum spanning tree with constant average distortion. *J. Comput. Syst. Sci.*, 105:116–129, 2019. preliminary version published in SODA 2016. doi:10.1016/j.jcss.2019.04.006.
- [BGS16] G. E. Blelloch, Y. Gu, and Y. Sun. A new efficient construction on probabilistic tree embeddings. *CoRR*, abs/1605.04651, 2016. <https://arxiv.org/abs/1605.04651>. arXiv:1605.04651.
- [BHO19] K. Buchin, S. Har-Peled, and D. Oláh. A spanner for the day after. In *35th International Symposium on Computational Geometry, SoCG 2019, June 18-21, 2019, Portland, Oregon, USA*, pages 19:1–19:15, 2019. doi:10.4230/LIPIcs.SocG.2019.19.
- [BHO20] K. Buchin, S. Har-Peled, and D. Oláh. Sometimes reliable spanners of almost linear size. In *28th Annual European Symposium on Algorithms, ESA 2020, September 7-9,*

- 2020, Pisa, Italy (Virtual Conference), pages 27:1–27:15, 2020. doi:10.4230/LIPIcs.ESA.2020.27.
- [BKS12] S. Baswana, S. Khurana, and S. Sarkar. Fully dynamic randomized algorithms for graph spanners. *ACM Trans. Algorithms*, 8(4):35:1–35:51, 2012. doi:10.1145/2344422.2344425.
- [BLMN05] Y. Bartal, N. Linial, M. Mendel, and A. Naor. Some low distortion metric ramsey problems. *Discret. Comput. Geom.*, 33(1):27–41, 2005. doi:10.1007/s00454-004-1100-z.
- [BLW17] G. Borradaile, H. Le, and C. Wulff-Nilsen. Minor-free graphs have light spanners. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science, FOCS '17*, pages 767–778, 2017. doi:10.1109/FOCS.2017.76.
- [BLW19] G. Borradaile, H. Le, and C. Wulff-Nilsen. Greedy spanners are optimal in doubling metrics. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '19*, pages 2371–2379, 2019. doi:10.1137/1.9781611975482.145.
- [Bou85] J. Bourgain. On lipschitz embedding of finite metric spaces in hilbert space. *Israel Journal of Mathematics*, 52(1-2):46–52, 1985. doi:10.1007/BF02776078.
- [Bou86] J. Bourgain. The metrical interpretation of superreflexivity in banach spaces. *Israel Journal of Mathematics*, 56:222–230, 1986. doi:10.1007/BF02766125.
- [BP19] G. Bodwin and S. Patel. A trivial yet optimal solution to vertex fault tolerant spanners. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 541–543, 2019. doi:10.1145/3293611.3331588.
- [BS07] S. Baswana and S. Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Struct. Algorithms*, 30(4):532–563, 2007. preliminary version published in ICALP 2003. doi:10.1002/rsa.20130.
- [BTS94] H. L. Bodlaender, G. Tel, and N. Santoro. Trade-offs in non-reversing diameter. *Nord. J. Comput.*, 1(1):111–134, 1994. see here.
- [CDNS95] B. Chandra, G. Das, G. Narasimhan, and J. Soares. New sparseness results on graph spanners. *Int. J. Comput. Geom. Appl.*, 5:125–144, 1995. preliminary version published in SOCG 1992. doi:10.1142/S0218195995000088.
- [CFKL20] V. Cohen-Addad, A. Filtser, P. N. Klein, and H. Le. On light spanners, low-treewidth embeddings and efficient traversing in minor-free graphs. In *61th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, pages 589–600, 2020. See: conference version, arXiv version,.
- [CGMZ16] T. H. Chan, A. Gupta, B. M. Maggs, and S. Zhou. On hierarchical routing in doubling metrics. *ACM Trans. Algorithms*, 12(4):55:1–55:22, 2016. Preliminary version appeared in SODA 2005.

- [CHJ20] T. M. Chan, S. Har-Peled, and M. Jones. On locality-sensitive orderings and their applications. *SIAM J. Comput.*, 49(3):583–600, 2020. preliminary version published in ITCS 2019. doi:10.1137/19M1246493.
- [CK95] P. B. Callahan and S. R. Kosaraju. A decomposition of multidimensional point sets with applications to k -nearest-neighbors and n -body potential fields. *Journal of the ACM*, 42(1):67–90, 1995. doi:10.1145/200836.200853.
- [CKR04] G. Călinescu, H. J. Karloff, and Y. Rabani. Approximation algorithms for the 0-extension problem. *SIAM Journal on Computing*, 34(2):358–372, 2004.
- [Cla87] K. Clarkson. Approximation algorithms for shortest path motion planning. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, STOC ‘87, pages 56–65, 1987.
- [CLPR10] S. Chechik, M. Langberg, D. Peleg, and L. Roditty. Fault tolerant spanners for general graphs. *SIAM J. Comput.*, 39(7):3403–3423, 2010. preliminary version published in STOC 2009. doi:10.1137/090758039.
- [CW18] S. Chechik and C. Wulff-Nilsen. Near-optimal light spanners. *ACM Trans. Algorithms*, 14(3):33:1–33:15, 2018. preliminary version published in SODA 2016. doi:10.1145/3199607.
- [DK11] M. Dinitz and R. Krauthgamer. Fault-tolerant spanners: better and simpler. In *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011, San Jose, CA, USA, June 6-8, 2011*, pages 169–178, 2011. doi:10.1145/1993806.1993830.
- [DR20] M. Dinitz and C. Robelle. Efficient and simple algorithms for fault-tolerant spanners. In *PODC ’20: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, August 3-7, 2020*, pages 493–500, 2020. doi:10.1145/3382734.3405735.
- [EFN20] M. Elkin, A. Filtser, and O. Neiman. Distributed construction of light networks. In *Proceedings of the 39th Symposium on Principles of Distributed Computing, PODC’20*, pages 483–492, 2020. doi:10.1145/3382734.3405701.
- [EGK⁺14] M. Englert, A. Gupta, R. Krauthgamer, H. Räcke, I. Talgam-Cohen, and K. Talwar. Vertex sparsifiers: New results from old techniques. *SIAM J. Comput.*, 43(4):1239–1262, 2014. doi:10.1137/130908440.
- [Elk11] M. Elkin. Streaming and fully dynamic centralized algorithms for constructing and maintaining sparse spanners. *ACM Trans. Algorithms*, 7(2):20:1–20:17, 2011. doi:10.1145/1921659.1921666.
- [ENS15] M. Elkin, O. Neiman, and S. Solomon. Light spanners. *SIAM J. Discret. Math.*, 29(3):1312–1321, 2015. doi:10.1137/140979538.
- [Erd64] P. Erdős. Extremal problems in graph theory. *Theory of Graphs and Its Applications (Proc. Sympos. Smolenice)*, pages 29–36, 1964.

- [ES13] M. Elkin and S. Solomon. Fast constructions of light-weight spanners for general graphs. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 513–525, 2013. doi:10.1137/1.9781611973105.37.
- [Fil19a] A. Filtser. On strong diameter padded decompositions. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2019, September 20-22, 2019, Massachusetts Institute of Technology, Cambridge, MA, USA*, pages 6:1–6:21, 2019. doi:10.4230/LIPIcs.APPROX-RANDOM.2019.6.
- [Fil19b] A. Filtser. On strong diameter padded decompositions. *CoRR*, abs/1906.09783, 2019. arXiv:1906.09783.
- [Fil20] A. Filtser. Scattering and sparse partitions, and their applications. In A. Czumaj, A. Dawar, and E. Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, pages 47:1–47:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ICALP.2020.47.
- [Fil21] A. Filtser. Hop-constrained metric embeddings and their applications. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 492–503. IEEE, 2021. doi:10.1109/FOCS52979.2021.00056.
- [FKN20] A. Filtser, M. Kapralov, and N. Nouri. Graph spanners by sketching in dynamic streams and the simultaneous communication model. *CoRR*, abs/2007.14204, 2020. To appear in SODA 2021. arXiv:2007.14204.
- [FL21] A. Filtser and H. Le. Clan embeddings into trees, and low treewidth graphs. In S. Khuller and V. V. Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 342–355. ACM, 2021. Full version at <https://arxiv.org/abs/2101.01146>. doi:10.1145/3406325.3451043.
- [FN18] A. Filtser and O. Neiman. Light spanners for high dimensional norms via stochastic decompositions. In *26th Annual European Symposium on Algorithms, ESA 2018, August 20-22, 2018, Helsinki, Finland*, pages 29:1–29:15, 2018. doi:10.4230/LIPIcs.ESA.2018.29.
- [FRT04] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, November 2004. preliminary version published in STOC 2003. doi:10.1016/j.jcss.2004.04.011.
- [FS20] A. Filtser and S. Solomon. The greedy spanner is existentially optimal. *SIAM J. Comput.*, 49(2):429–447, 2020. preliminary version published in PODC 2016. doi:10.1137/18M1210678.

- [GKL03] A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, pages 534–543, 2003. doi:10.1109/SFCS.2003.1238226.
- [Got15] L. Gottlieb. A light metric spanner. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 759–772, 2015. doi:10.1109/FOCS.2015.52.
- [Har20] S. Har-Peled. Theory seminar talk, CS, UIUC, 10/12/20: Reliable spanners for metric spaces, 2020. <https://youtu.be/wo0unW9HCtg?t=3630>.
- [HIS13] S. Har-Peled, P. Indyk, and A. Sidiropoulos. Euclidean spanners in high dimensions. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 804–809, 2013. doi:10.1137/1.9781611973105.57.
- [HMO21] S. Har-Peled, M. Mendel, and D. Oláh. Reliable Spanners for Metric Spaces. In *37th International Symposium on Computational Geometry, SoCG’21*, pages 43:1–43:13, 2021. Full version at <https://arxiv.org/abs/2007.08738>. doi:10.4230/LIPIcs.SocG.2021.43.
- [IM98] P. Indyk and R. Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. In J. S. Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 604–613. ACM, 1998. doi:10.1145/276698.276876.
- [JL84] W. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary Mathematics*, 26:189–206, 1984.
- [Kei88] J. M. Keil. Approximating the complete euclidean graph. In *Proceedings of the first Scandinavian Workshop on Algorithm Theory, SWAT ’88*, pages 208–213, 1988.
- [Kle02] P. N. Klein. Preprocessing an undirected planar network to enable fast approximate distance queries. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 6-8, 2002, San Francisco, CA, USA*, pages 820–827, 2002. see here.
- [KPR93] P. N. Klein, S. A. Plotkin, and S. Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing, STOC ’93*, page 682–690, 1993. doi:10.1145/167088.167261.
- [KPX08] I. A. Kanj, L. Perkovic, and G. Xia. Computing lightweight spanners locally. In *Distributed Computing, 22nd International Symposium, DISC 2008, Arcachon, France, September 22-24, 2008. Proceedings*, pages 365–378, 2008. doi:10.1007/978-3-540-87779-0_25.
- [KW14] M. Kapralov and D. P. Woodruff. Spanners and sparsifiers in dynamic streams. In *ACM Symposium on Principles of Distributed Computing, PODC ’14, Paris, France, July 15-18, 2014*, pages 272–281, 2014. doi:10.1145/2611462.2611497.

- [LNS02] C. Levcopoulos, G. Narasimhan, and M. H. M. Smid. Improved algorithms for constructing fault-tolerant spanners. *Algorithmica*, 32(1):144–156, 2002. preliminary version published in STOC 1998. doi:10.1007/s00453-001-0075-x.
- [LS19] H. Le and S. Solomon. Truly optimal euclidean spanners. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 1078–1100, 2019. doi:10.1109/FOCS.2019.00069.
- [Luk99] T. Lukovszki. New results on fault tolerant geometric spanners. In F. Dehne, J.-R. Sack, A. Gupta, and R. Tamassia, editors, *Algorithms and Data Structures*, pages 193–204, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg. doi:10.1007/3-540-48447-7_20.
- [Mat99] J. Matoušek. On embedding trees into uniformly convex banach spaces. *Israel Journal of Mathematics*, 114(1):221–237, 1999. doi:10.1007/BF02785579.
- [MN07] M. Mendel and A. Naor. Ramsey partitions and proximity data structures. *Journal of the European Mathematical Society*, 9(2):253–275, 2007. preliminary version published in FOCS 2006. doi:10.4171/JEMS/79.
- [NS07] G. Narasimhan and M. H. M. Smid. *Geometric spanner networks*. Cambridge University Press, 2007. doi:10.1017/CB09780511546884.
- [NT12] A. Naor and T. Tao. Scale-oblivious metric fragmentation and the nonlinear dvoretzky theorem. *Israel Journal of Mathematics*, 192(1):489–504, 2012. doi:10.1007/s11856-012-0039-7.
- [PS89] D. Peleg and A. A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989. doi:10.1002/jgt.3190130114.
- [Sol13] S. Solomon. Sparse euclidean spanners with tiny diameter. *ACM Trans. Algorithms*, 9(3):28:1–28:33, 2013. doi:10.1145/2483699.2483708.
- [Sol14] S. Solomon. Euclidean steiner shallow-light trees. In *Proceedings of the Thirtieth Annual Symposium on Computational Geometry, SoCG '14*, pages 454:454–454:463, 2014.
- [Tho04] M. Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *Journal of the ACM*, 51(6):993–1024, 2004. doi:10.1145/1039488.1039493.
- [Yao82] A. C. Yao. On constructing minimum spanning trees in k -dimensional spaces and related problems. *SIAM Journal on Computing*, 11(4):721–736, 1982.