

Locality-Sensitive Orderings and Applications to Reliable Spanners*

Arnold Filtser
Bar-Ilan University
Ramat-Gan, Israel
arnold.filtser@biu.ac.il

Hung Le
University of Massachusetts at Amherst
Amherst, Massachusetts, USA
hungle@cs.umass.edu

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.

CCS CONCEPTS

• **Theory of computation** → **Computational geometry; Sparsification and spanners.**

KEYWORDS

Locality-Sensitive Orderings, Reliable Spanners, Doubling Metric, Ultrametric cover, 2-hop spanners, Minor Free graphs

ACM Reference Format:

Arnold Filtser and Hung Le. 2022. Locality-Sensitive Orderings and Applications to Reliable Spanners. In *Proceedings of the 54th Annual ACM SIGACT STOC '22, June 20–24, 2022, Rome, Italy*

*The reader is encouraged to read the full version of the paper [62].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

STOC '22, June 20–24, 2022, Rome, Italy

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9264-8/22/06...\$15.00

<https://doi.org/10.1145/3519935.3520042>

Symposium on Theory of Computing (STOC '22), June 20–24, 2022, Rome, Italy.
ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3519935.3520042>

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 [33, 55] and padded decompositions [12, 56, 74], while many tools apply mainly to Euclidean spaces, such as dimension reduction [71], locality-sensitive hashing [70], well-separated pair decomposition (WSPD) [40], and many others. Recently, Chan, Har-Peled, and Jones [42] developed a new and exciting tool for Euclidean spaces called *Locality-Sensitive Ordering* (LSO).

Definition 1.1 ((τ, ρ) -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. [42] 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.* [42] 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 [36, 37] used the LSO of Chan *et al.* [42] to construct *reliable spanners* (see Definition 1.2) 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.

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.

A highly desirable property of a t -spanner is the ability to withstand extensive vertex failures. Levcopoulos, Narasimhan, and Smid [76] 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.* [32] 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.* [37] relaxed the notion of reliable spanners by allowing the size of B^+ to be bounded only in expectation.

Definition 1.2 (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 , s.t.:

- (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 1.2 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.* [32] 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.* [36] used the LSO of Chan *et al.* [42] 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 [31]). Later, for the *oblivious* case, Buchin *et al.* [37] 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 [69] 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

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[69] did not compute the constant explicitly. Their construction is based on the Ramsey-trees of Mendel-Naor [79], which have stretch $128k$. Using state of the art Ramsey trees [80] of stretch $2ek$ instead (see also [2]), the approach of [69] provides stretch $8ek$.

$n^{1+1/k} \cdot O(\nu^{-1}k \log^2 \Phi \log \frac{n}{\nu})$ 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 3.1) 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.* [69] are based on *sparse covers*. A (τ, ρ) -sparse cover is a collection 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 $C \in C$ 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 C . They then treat each cluster in 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 [37], Har-Peled *et al.* [69] constructed an oblivious ν -reliable 2-spanner for the uniform metric, which is the best stretch possible for subquadratic size spanners (see full version [62]). Indeed, this additional factor 2 appears in the stretch parameter in all the spanners in [69]. Most prominently, for trees they 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 (see full version [62]). A similar phenomenon occurs for planar graphs. An additional drawback in the sparse cover based approach of [69] 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 [69]).

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 [42]; 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),

⁵With an additional effort, [69] 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.

Table 1: Summary 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. \star denotes the full version of this paper [62].

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	[42]
	Doubling dimension d	$O(\epsilon^{-1} \cdot \log n)^{O(d)}$ $\epsilon^{-O(d)}$	ϵ	[42] Cor. 4.4
Triangle-LSO	General metric	$\tilde{O}(n^{1/k} \cdot \epsilon^{-1})$	$2k + \epsilon$	Cor. 5.3
	Ultrametric	1	1	Lem. 5.2
Left-sided LSO	Tree	$\log n$	1	\star
	Treewidth k	$k \cdot \log n$	1	\star
	Planar graph	$\frac{1}{\epsilon} \cdot \log^2 n$	$1 + \epsilon$	\star
	Minor Free	$\frac{1}{\epsilon} \cdot \log^2 n$	$1 + \epsilon$	\star

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.* [36, 37] 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 [67] regarding sub-graph reliable spanners, by providing matching upper and lower bounds for reliable connectivity preservers.

Classic LSO. Chan *et al.* [42] constructed a $(\epsilon^{-1} \cdot \log n)^{O(d)}, \epsilon$ -LSO for metric spaces of doubling dimension d . Applying the (implicit) framework of [36, 37] 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 4.4); 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 4.6).

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 $x \preceq_\sigma z \preceq_\sigma q \preceq_\sigma y$ (or $y \preceq_\sigma z \preceq_\sigma q \preceq_\sigma x$) it holds that $d_X(z, q) \leq \rho \cdot d_X(x, y)$ (Definition 5.1). 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 5.2), and show that general n -point metric spaces admit an $(\tilde{O}(n^{1/k} \cdot \epsilon^{-1}), 2k + \epsilon)$ -triangle-LSO (Corollary 5.3).

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 5.4). This gives oblivious reliable spanners for ultrametrics and general metric spaces. Our spanners for general metrics have significantly smaller

stretch compared to [69] ($8k$ 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 points $x', y' \in X$ satisfying $x \preceq_\sigma x'$ and $y' \preceq_\sigma y$, it holds that $d_X(x', y') \leq \rho \cdot d_X(x, y)$ (see the full version [62] for formal definition and statements). 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 [69].

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 (see full version [62]). 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. 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 $\epsilon \in (0, 1)$. An oblivious ν -reliable $(2 + \epsilon)$ -spanner with $n \cdot O(\nu^{-1} \epsilon^{-2} \log^5 n)$ edges follows. 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 3.2). The first study on tree covers was for Euclidean spaces by Arya *et al.* [9] who constructed the so-called Dumbbell trees. For general metrics, Mendel and Naor [79] (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 [2, 13, 14, 17, 23, 35, 59, 61, 79, 80]. The state of the art covers follow from Naor and Tao [80], and implies a $(2e \cdot k, O(k \cdot n^{1/k}))$ ultrametric cover. For doubling metrics, Bartal, Fandina, and Neiman [15] constructed a $(1 + \epsilon, \epsilon^{-O(d)})$ tree cover. We refer to [15] 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 5.2). Indeed, the main step in our construction of a triangle LSO for general metrics is a construction of a $(\tilde{O}(n^{1/k} \cdot \epsilon^{-1}), 2k + \epsilon)$ -ultrametric cover (Theorem 3.3). 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.

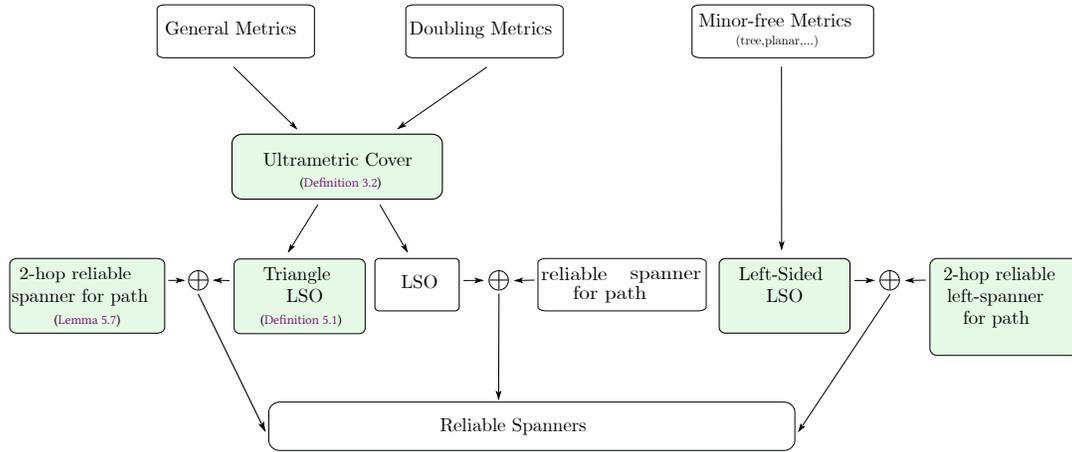


Figure 1: Relationships between different concepts; new concepts introduced in this papers are green-shaded.

Table 2: Comparison between previous and new constructions of reliable spanners. \star denotes the full version of this paper [62]. All spanners (except [32]) 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. Both lower bounds hold for the uniform metric. For all other metric spaces, we improve both stretch and sparsity, and remove the undesirable dependence on the aspect ratio Φ . For trees and planar graphs, the stretch was improved from $3 + \epsilon$, to the best possible stretch 2. For general graphs, our spanner has stretch $8k$, considerably improving the constant hiding in [69]. This constant governs the parameter in the power of n .

Family	stretch	guarantee	size	ref
Euclidean ($\mathbb{R}^d, \ \cdot\ _2$)	$O(1)$	Deterministic	$\Omega(n \log n)$	[32]
	$1 + \epsilon$	Deterministic	$n \cdot \tilde{O}(\epsilon)^{-7d} \nu^{-6} \cdot \tilde{O}(\log n)$	[36]
	$1 + \epsilon$	Oblivious	$n \cdot \tilde{O}(\epsilon)^{-2d} \cdot \tilde{O}(\nu^{-1} (\log \log n)^2)$	[37]
Doubling dimension d	$1 + \epsilon$	Deterministic	$n \cdot \epsilon^{-O(d)} \nu^{-6} \cdot \tilde{O}(\log n)$	Corollary 4.6
	$1 + \epsilon$	Oblivious	$n \cdot \epsilon^{-O(d)} \nu^{-1} \log \nu^{-1} \cdot \tilde{O}(\log \log n)^2$	Corollary 4.6
General metric	$512 \cdot k^4$	Oblivious	$\tilde{O}(n^{1+1/k} \cdot \nu^{-1}) \cdot \log^2 \Phi$	[69]
	$8k + \epsilon$	Oblivious	$\tilde{O}(n^{1+1/k} \cdot \epsilon^{-2}) \cdot \nu^{-1}$	Theorem 5.5
Ultrametric, Tree, planar	k	Deterministic	$\Omega(n^{1+1/k})$	[69]
	$k < 2$	Oblivious	$\Omega(n^2)$	\star
Ultrametric	$2 + \epsilon$	Oblivious	$n \cdot \tilde{O}(\nu^{-1} \epsilon^{-2} \log^2 \Phi)$	[69]
	2 (tight)	Oblivious	$n \cdot \tilde{O}(\log^2 n + \nu^{-1} \log n)$	Corollary 5.6
Tree	$3 + \epsilon$	Oblivious	$n \cdot \tilde{O}(\nu^{-1} \epsilon^{-2} \log^2 n \log^2 \Phi)$	[69]
	2 (tight)	Oblivious	$n \cdot O(\nu^{-1} \log^3 n)$	\star
Treewidth k	2 (tight)	Oblivious	$n \cdot O(\nu^{-1} k^2 \log^3 n)$	\star
Planar	$3 + \epsilon$	Oblivious	$n \cdot \tilde{O}(\nu^{-1} \epsilon^{-4} \log^2 n \log^2 \Phi)$	[69]
	$2 + \epsilon$ (tight)	Oblivious	$n \cdot O(\nu^{-1} \epsilon^{-2} \log^5 n)$	\star
Minor-free	$2 + \epsilon$ (tight)	Oblivious	$n \cdot O(\nu^{-1} \epsilon^{-2} \log^5 n)$	\star

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 3.1). We show that every $\Omega(\frac{1}{\epsilon})$ -HST of degree bounded by δ admits a (classic) $(\frac{\delta}{2}, \epsilon)$ -LSO (Lemma 4.2). 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 4.3). The trees in the tree cover for doubling metrics of [15] 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 3.4), which implies the respective LSO. Interestingly, having such an ultrametric cover is a characterizing property for metric spaces of bounded doubling dimension (Theorem 3.4). See Table 3 for a summary.

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.* [37] constructed an oblivious ν -reliable 1-spanner

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

Space	type	stretch	# of trees	ref
Euclidean \mathbb{R}^d	tree	$1 + \epsilon$	$O((\frac{d}{\epsilon})^d \log \frac{d}{\epsilon})$	[9]
Doubling dimension d	ultrametric	$O(d^2)$	$O(d \log d)$	[41]
	tree	$1 + \epsilon$	$\epsilon^{-O(d)}$	[16]
	ultrametric	$1 + \epsilon$	$\epsilon^{-O(d)}$	Thm. 3.4
General metric	ultrametric	$2e \cdot k$	$O(k \cdot n^{1/k})$	[79, 80]
	ultrametric	$2k + \epsilon$	$\tilde{O}(n^{\frac{1}{k}} \cdot \epsilon^{-1})$	Thm. 3.3

with $n \cdot \tilde{O}(v^{-1}(\log \log n)^2)$ 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 v -reliable, 2-hop 1-spanner with $n \cdot O(\log^2 n + v^{-1} \log n \cdot \log \log n)$ edges for the path graph (Lemma 5.7). This spanner is later used in our meta Theorem 5.4 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* (see the full version [62] for formal definition and statements); we then construct an oblivious v -reliable 2-hop left-spanner with $n \cdot O(v^{-1} \log n)$ edges. The left-spanner is later used 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 [67] 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 v -reliable spanner defined in Definition 1.2 is a *g-reliable* (non-subgraph) spanner for the linear function $g(x) = (1+v)x$. 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 $g = O(x^k)$. Taking $k = 1$ gives a lower bound $\Omega(n^2)$ on the number of edges of *subgraph* v -reliable spanners. On the positive side, we provide a construction of a *deterministic* connectivity preserver matching the lower bound (see the full version [62]).

1.2 Related Work

The tradeoff between stretch and sparsity (number of edges) of (regular) t -spanners has been extensively studied [8, 43, 63, 68, 75, 82]; see the recent survey of Ahmed *et al.* [4], and the book [81] 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 [43], 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 [8]. Similarly, n -point metric spaces with doubling dimension d admit $(1 + \epsilon)$ -spanners with $n \cdot \epsilon^{-O(d)}$ edges [41].

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 [76, 77, 84]. For general graphs, after a long line of works [25, 26, 28, 44, 47, 48], 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 [54]. 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 [24].

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 [27] 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 $\Omega(f^{\frac{1}{2} - \frac{1}{2k}} n^{1+1/k})$ [25].

Abam *et al.* [1] 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.* [10]. State of the art is a $(1 + \epsilon)$ -spanner constructible in $O(n \log n)$ time by Solomon [83] 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 [7, 29, 30, 43, 45, 46, 51, 52, 64, 65, 75]. Sparse (and light) spanners were constructed efficiently in different computational models such as LOCAL [72], CONGEST [50], streaming [5, 18, 20, 49, 60, 73], massive parallel computation (MPC) [22] and dynamic graph algorithms [19, 21].

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 $[a : b] = \{a, a + 1, \dots, b\}$.

⁶ $\alpha_k(n)$ is the inverse function of a very fast growing function at level k of the primitive recursive hierarchy; see [83] for a more formal description.

Table 4: Construction of reliable spanners for the line. [36] and [37] 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.

Type	guarantee	size	hops	ref
1 spanner	Deterministic	$n \cdot O(\log n \cdot v^{-6})$	$O(\log n)$	[36]
	Oblivious	$n \cdot O(v^{-1} \cdot \log v^{-1})$	$O(\log n)$	[37]
	Oblivious	$n \cdot O(\log^2 n + v^{-1} \log n \cdot \log \log n)$	2	Lemma 5.7
Left spanner	Oblivious	$n \cdot O(v^{-1} \log n)$	2	Full version [62]

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., [66]).

LEMMA 2.1 (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 2.2. *Consider an n -vertex graph $G = (V, E, w)$ that admits an oblivious v -reliable t -spanner with m edges in expectation. Then G admits an oblivious $2v$ -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 $B \subseteq V$, $\mathbb{E}[|B^+ \setminus B|] \leq v|B|$, 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

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

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 3.1 (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 [17]).

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 [15] 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.* [15] observed that the previous constructions of Ramsey trees⁷ [2, 79, 80] 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

⁷Ramsey trees have additional de'sired 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.

doubling dimension d . Here we initiate the study of ultrametric covers.

Definition 3.2 (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 $\sqrt{\log \log n}$ [34, 78]. Later, we will show how to use ultrametric covers to construct locality-sensitive orderings (see Lemma 4.3 and Lemma 5.2).

The first main result of this section is Theorem 3.3 where we construct an ultrametric cover for general metrics.

THEOREM 3.3 (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 [15] for general metrics actually consists of ultrametrics; in other words, Bartal *et al.* [15] 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 $\Omega(n^{\frac{\epsilon}{k+1}})$ ultrametrics.

Next, in Theorem 3.4 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 3.4 is provided in Section 3.3.

THEOREM 3.4 (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 3.3 and 3.4 is *pairwise partition cover*, a newly introduced notion, which is closely related to the previously introduced stochastic/padded decompositions and sparse covers [3, 11, 12, 56, 66, 74]. 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 3.5 (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 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 3.6. 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 3.7. 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 3.3:

PROOF OF THEOREM 3.3. Let (X, d_X) be an n -point metric space, and fix $\delta = \frac{\epsilon}{8}$. By Lemma 3.7, X admits an $(O(n^{\frac{1}{k}} \log n), 2k + \delta, \frac{\delta}{4k(2k+\delta)})$ -pairwise partition cover. By Lemma 3.6, 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 3.6

Lemma 3.6 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 [58], stochastic Steiner point removal [53], universal Steiner tree [38], 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 3.6. 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^{i-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 $\tilde{\mathcal{P}}_j^i$. 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 $\epsilon < 1$.

Next we argue that $\tilde{\mathcal{P}}_i = \{\tilde{\mathcal{P}}_1^i, \dots, \tilde{\mathcal{P}}_\tau^i\}$ 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 $\tilde{C}_u, \tilde{C}_v \in \tilde{\mathcal{P}}_j^{i-1}$ be the clusters containing u, v 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 $\tilde{C}_u \subseteq B_X(v, \epsilon\Delta_i) \subseteq C_i$, and similarly $\tilde{C}_v \subseteq C_i$. 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 to 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 $l \in [0, \lfloor \log_{1+\epsilon} \frac{4\rho}{\epsilon} \rfloor]$. 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 $l \in [0, \lfloor \log_{1+\epsilon} \frac{4\rho}{\epsilon} \rfloor]$, 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 descend 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 3.7

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 3.7 will follow by taking the union of many independently sampled such partitions.

LEMMA 3.8. *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 [57].⁸ Our proof here follows the steps of [57] (which is based on the partition in [39]).

PROOF OF LEMMA 3.8. 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

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

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

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

⁸This is the full version, see also the conference version [56].

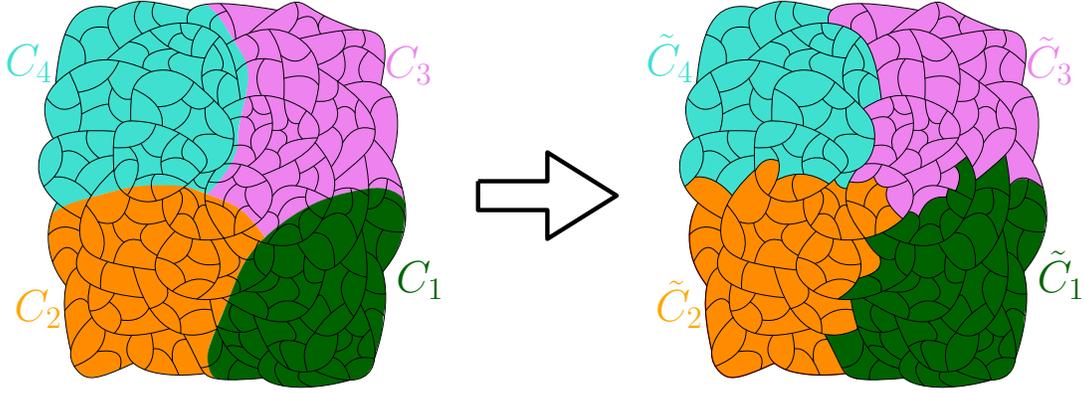


Figure 2: Illustration of the construction of the partition $\tilde{\mathcal{P}}_j^i$ given \mathcal{P}_j^i and $\tilde{\mathcal{P}}_j^{i-1}$. The black lines in both the left and right parts of the figure border clusters in $\tilde{\mathcal{P}}_j^{i-1}$. On the left illustrated the partition $\mathcal{P}_j^i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

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

We now ready to prove [Lemma 3.7](#) (restated for convenience).

LEMMA 3.7. *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 $k \geq 1$.*

PROOF. Fix Δ . Sample $s = n^{\frac{1}{k}} \cdot 2 \ln n$ i.i.d. partitions using [Lemma 3.8](#). 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 3.4](#). We begin with the following combinatorial lemma.

LEMMA 3.9. *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 red neighbor of v is matched by M .
- 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 3.10. *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 $(u, v) \in E_r$ 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 2.1](#)), every vertex in G has blue degree $\epsilon^{-O(d)}$ and

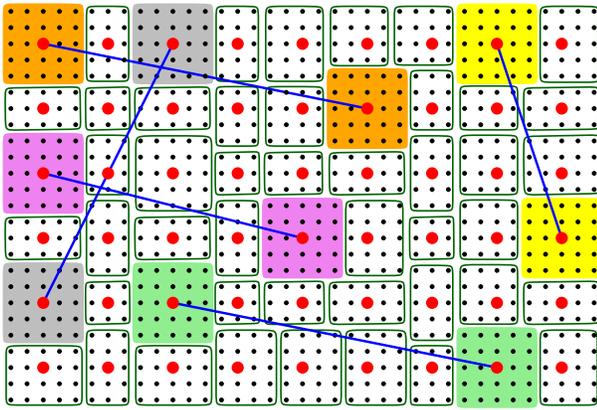


Figure 3: Illustration of the partition \mathcal{P}_M . The black points represent metric points, while the red points represent the net points N . The blue edges are the matching M . For each edge $\{u, v\} \in M$, the cluster $B_X(u, 2\epsilon\Delta) \cup B_X(v, 2\epsilon\Delta)$ is added to \mathcal{P}_M . 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.

red degree $2^{O(d)}$. Let \mathcal{M} be the set of matching of $G(N, E_b \cup E_r)$ guaranteed by Lemma 3.9; $|\mathcal{M}| = O(\epsilon^{-O(d)} 2^{O(d)}) = \epsilon^{-O(d)}$.

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 3.9. 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 3.4 that we restate below for convenience.

THEOREM 3.4 (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 3.6 and Lemma 3.10, 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 3.6 w.r.t. the pairwise partition cover scheme constructed in Lemma 3.10 has bounded degree.

We will use the terminology of Lemmas 3.6 and 3.10. 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. [42] introduced and studied the notion of locality-sensitive ordering (Definition 1.1). In the same paper, Chan et al. [42] 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 [6].

LEMMA 4.1. 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 4.2. *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 4.1](#), 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_δ . 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 into 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 4.3. *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 4.2](#). 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 4.2](#), 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 3.4](#) and [Lemma 4.3](#) with $\rho = (1+\epsilon)$ and $k = O(\frac{1}{\epsilon})$, we conclude

COROLLARY 4.4 (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.* [[36](#), [37](#)] constructed reliable $(1+\epsilon)$ spanners for Euclidean metrics using $(\tau(\epsilon), \epsilon)$ -LSO by Chan *et al.* [[42](#)]. Specifically, their spanner for the deterministic case has $n \cdot O(\epsilon)^{-7d} \log^7(\frac{1}{\epsilon}) \cdot v^{-6} \log n (\log \log n)^6$ edges, while for the oblivious case, they constructed a spanner with an almost linear number of edges: $n \cdot O(\epsilon)^{-2d} \log^3 \frac{1}{\epsilon} \cdot v^{-1} \log v^{-1} (\log \log n)^2 \log \log n$. 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 4.5 ([[36](#), [37](#)], 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 $v \in (0, 1)$ and $\epsilon \in (0, 1)$:*

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

By [Theorem 4.5](#) and [Corollary 4.4](#), we have:

COROLLARY 4.6. *Consider a metric space (X, d_X) with doubling dimension d . Then for every $v \in (0, 1)$ and $\epsilon \in (0, 1)$, (X, d_X) admits a deterministic v -reliable $(1+\epsilon)$ -spanner with $n \cdot \epsilon^{-O(d)} v^{-6} \log n (\log \log n)^6$ edges, and an oblivious v -reliable $(1+\epsilon)$ -spanner with*

$$n \cdot \epsilon^{-O(d)} \cdot v^{-1} (\log \log n)^2 \cdot \log\left(\frac{\log \log n}{\epsilon v}\right) = n \cdot \epsilon^{-O(d)} \cdot \tilde{O}(v^{-1} (\log \log n)^2) \text{ edges.}$$

5 TRIANGLE LOCALITY SENSITIVE ORDERING

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

Definition 5.1 ((τ, ρ) -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 4.3](#), we show that a metric space admitting an ultrametric cover has a triangle-LSO with a small number of orderings. The proof of the following lemma is deferred to the full version [[62](#)].

LEMMA 5.2. *If a metric (X, d_X) admits a (τ, ρ) -ultrametric cover \mathcal{U} , then it has a (τ, ρ) -triangle-LSO.*

Using [Theorem 3.3](#) and [Lemma 5.2](#), we conclude.

COROLLARY 5.3. *For every $k \in \mathbb{N}$, and $\epsilon \in (0, 1)$, every n -point metric space admits an $\left(O(n^{\frac{1}{k}} \cdot \log n \cdot \frac{k^2}{\epsilon} \cdot \log \frac{k}{\epsilon}), 2k + \epsilon\right)$ -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 5.4. *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 $n\tau \cdot O\left(\log^2 n + \nu^{-1}\tau \log n \cdot \log \log n\right)$ edges.*

The proof of [Theorem 5.4](#) is deferred to the end of the section. Using [Corollary 5.3](#) with parameters $2k$ and $\frac{\epsilon}{2}$, and [Theorem 5.4](#) we conclude:

THEOREM 5.5 (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 $n^{1+\frac{1}{k}} \cdot \nu^{-1} \cdot \log^3 n \cdot \log \log n \cdot \frac{k^4}{\epsilon^2} O(\log \frac{k}{\epsilon})^2 = n^{1+\frac{1}{k}} \cdot \nu^{-1} \cdot \tilde{O}\left(\log^3 n \cdot \frac{k^4}{\epsilon^2}\right)$ edges.*

By [Lemma 5.2](#) and [Theorem 5.4](#), we obtain:

COROLLARY 5.6. *For every parameter $\nu \in (0, 1)$, every n -point ultrametric space (X, d) admits an oblivious ν -reliable, 2-spanner with $n \cdot O\left(\log^2 n + \nu^{-1} \log n \cdot \log \log n\right)$ edges.*

The stretch parameter in [Corollary 5.6](#) is tight (see the full version [62]).

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.* [36, 37], 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.* [37], 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.* [37] 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 [81]). Our result is summarized in the following lemma whose proof is deferred to the full version [62].

LEMMA 5.7 (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 $n \cdot O\left(\log^2 n + \nu^{-1} \log n \cdot \log \log n\right)$ edges.*

Using [Lemma 5.7](#), we can construct a reliable spanner for metric spaces admitting a (τ, ρ) -triangle LSO as claimed by [Theorem 5.4](#).

PROOF OF THEOREM 5.4. 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 $n \cdot O\left(\log^2 n + \nu'^{-1} \log n \cdot \log \log n\right)$ edges by [Lemma 5.7](#). 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

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

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. \square

6 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 [81]).
- (2) A major open question is the construction of deterministic reliable spanners general metric spaces. [69] constructed deterministic reliable $O(t^2)$ -spanners for general metrics with $\tilde{O}(n^{1+\frac{1}{t}})$ edges, and deterministic reliable $O(t)$ -spanners for

trees and planar graphs with $\tilde{O}(n^{1+\frac{1}{t}})$ edges, while showing an $\Omega(n^{1+\frac{1}{t}})$ lower bound on the number of edges in a deterministic reliable t -spanners for the uniform metric. Using the new LSO's constructed in this paper, it could be possible to improve the stretch parameters by a constant factor and remove the dependency on aspect ratio from the sparsity. However, the lower bound for uniform metrics applies to trees and planar graphs as well; thus using $\tilde{O}(n^{1+\frac{1}{t}})$ edges to obtain stretch t is necessary. On the other hand, general metrics are far from understood. Closing the gap between the current $O(t^2)$ upper bound to the t lower bound is a fascinating open question.

- (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

- [1] Mohammad Ali Abam, Mark de Berg, Mohammad Farshi, and Joachim Gudmundsson. 2009. Region-Fault Tolerant Geometric Spanners. *Discret. Comput. Geom.* 41, 4 (2009), 556–582. <https://doi.org/10.1007/s00454-009-9137-7>
- [2] Ittai Abraham, Shiri Chechik, Michael Elkin, Arnold Filtser, and Ofer Neiman. 2020. Ramsey Spanning Trees and Their Applications. *ACM Trans. Algorithms* 16, 2 (2020), 19:1–19:21. <https://doi.org/10.1145/3371039>
- [3] Ittai Abraham, Cyril Gavoille, Anupam Gupta, Ofer Neiman, and Kunal Talwar. 2019. Cops, Robbers, and Threatening Skeletons: Padded Decomposition for Minor-Free Graphs. *SIAM J. Comput.* 48, 3 (2019), 1120–1145. <https://doi.org/10.1137/17M1112406>
- [4] Abu Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen G. Kobourov, and Richard Spence. 2020. Graph spanners: A tutorial review. *Comput. Sci. Rev.* 37 (2020), 100253. <https://doi.org/10.1016/j.cosrev.2020.100253>
- [5] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. 2012. 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*, 5–14. <https://doi.org/10.1145/2213556.2213560>
- [6] B. Alspach. 2008. The wonderful Walecki construction. *Bull. Inst. Combin. Appl.* 52 (2008), 7–20. <https://doi.org/10.1137/1.9781611975031.123>
- [7] Stephen Alstrup, Søren Dahlgaard, Arnold Filtser, Morten Stöckel, and Christian Wulff-Nilsen. 2019. Constructing Light Spanners Deterministically in Near-Linear Time. In *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*, 4:1–4:15. <https://doi.org/10.4230/LIPIcs.ESA.2019.4>
- [8] I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. 1993. On Sparse Spanners of Weighted Graphs. *Discrete Computational Geometry* 9, 1 (1993), 81–100. <https://doi.org/10.1007/BF02189308>
- [9] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid. 1995. Euclidean Spanners: Short, Thin, and Lanky. In *Proceedings of the Twenty-seventh Annual ACM Symposium on Theory of Computing (STOC '95)*, 489–498.
- [10] S. Arya, D. M. Mount, and M. H. M. Smid. 1994. Randomized and deterministic algorithms for geometric spanners of small diameter. In *Proc. of 35th FOCS*, 703–712.
- [11] Baruch Awerbuch and David Peleg. 1990. Sparse partitions. In *Proceedings of the 31st IEEE Symposium on Foundations of Computer Science (FOCS)*, 503–513. <https://doi.org/10.1109/FSCS.1990.89571>
- [12] Yair Bartal. 1996. 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*, 184–193. <https://doi.org/10.1109/SFCS.1996.548477>
- [13] Yair Bartal. 2021. Advances in Metric Ramsey Theory and its Applications. *CoRR abs/2104.03484* (2021). [arXiv:2104.03484](https://arxiv.org/abs/2104.03484) <https://arxiv.org/abs/2104.03484>
- [14] Yair Bartal, Béla Bollobás, and Manor Mendel. 2006. Ramsey-type theorems for metric spaces with applications to online problems. *J. Comput. Syst. Sci.* 72, 5 (2006), 890–921. <https://doi.org/10.1016/j.jcss.2005.05.008> Special Issue on FOCS 2001.
- [15] Yair Bartal, Nova Fandina, and Ofer Neiman. 2019. Covering Metric Spaces by Few Trees. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, 20:1–20:16. <https://doi.org/10.4230/LIPIcs.ICALP.2019.20>
- [16] Yair Bartal, Arnold Filtser, and Ofer Neiman. 2019. On notions of distortion and an almost minimum spanning tree with constant average distortion. *J. Comput. Syst. Sci.* 105 (2019), 116–129. <https://doi.org/10.1016/j.jcss.2019.04.006>
- [17] Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. 2005. Some Low Distortion Metric Ramsey Problems. *Discret. Comput. Geom.* 33, 1 (2005), 27–41. <https://doi.org/10.1007/s00454-004-1100-z>
- [18] Surender Baswana. 2008. Streaming algorithm for graph spanners - single pass and constant processing time per edge. *Inf. Process. Lett.* 106, 3 (2008), 110–114. <https://doi.org/10.1016/j.ipl.2007.11.001>
- [19] Surender Baswana, Sumeet Khurana, and Soumojit Sarkar. 2012. Fully dynamic randomized algorithms for graph spanners. *ACM Trans. Algorithms* 8, 4 (2012), 35:1–35:51. <https://doi.org/10.1145/2344422.2344425>
- [20] Surender Baswana and Sandeep Sen. 2007. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Struct. Algorithms* 30, 4 (2007), 532–563. <https://doi.org/10.1002/rsa.20130>
- [21] Aaron Bernstein, Sebastian Forster, and Monika Henzinger. 2019. 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*, 1899–1918. <https://doi.org/10.1137/1.9781611975482.115>
- [22] Amartya Shankha Biswas, Michal Dory, Mohsen Ghaffari, Slobodan Mitrovic, and Yasamin Nazari. 2020. Massively Parallel Algorithms for Distance Approximation and Spanners. *CoRR abs/2003.01254* (2020). [arXiv:2003.01254](https://arxiv.org/abs/2003.01254) <https://arxiv.org/abs/2003.01254>
- [23] Guy E. Blelloch, Yan Gu, and Yihan Sun. 2016. A New Efficient Construction on Probabilistic Tree Embeddings. *CoRR abs/1605.04651* (2016). [arXiv:1605.04651](https://arxiv.org/abs/1605.04651) <https://arxiv.org/abs/1605.04651>
- [24] Greg Bodwin, Michael Dinitz, and Yasamin Nazari. 2022. Vertex Fault-Tolerant Emulators. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 215)*, Mark Braverman (Ed.), Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 25:1–25:22. <https://doi.org/10.4230/LIPIcs.ITCS.2022.25>
- [25] Greg Bodwin, Michael Dinitz, Merav Parter, and Virginia Vassilevska Williams. 2018. 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*, 1884–1900. <https://doi.org/10.1137/1.9781611975031.123>
- [26] Greg Bodwin, Michael Dinitz, and Caleb Robelle. 2021. Optimal Vertex Fault-Tolerant Spanners in Polynomial Time. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, Daniel Marx (Ed.), SIAM, 2924–2938. <https://doi.org/10.1137/1.9781611976465.174>
- [27] Greg Bodwin, Michael Dinitz, and Caleb Robelle. 2022. Partially Optimal Edge Fault-Tolerant Spanners. In *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference, January 9 - 12, 2022*, Joseph (Seffi) Naor and Niv Buchbinder (Eds.), SIAM, 3272–3286. <https://doi.org/10.1137/1.9781611977073.129>
- [28] Greg Bodwin and Shyamal Patel. 2019. 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*, 541–543. <https://doi.org/10.1145/3293611.3331588>
- [29] G. Borradaile, H. Le, and C. Wulff-Nilsen. 2017. Minor-free graphs have light spanners. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS '17)*, 767–778. <https://doi.org/10.1109/FOCS.2017.76>
- [30] G. Borradaile, H. Le, and C. Wulff-Nilsen. 2019. Greedy spanners are optimal in doubling metrics. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '19)*, 2371–2379. <https://doi.org/10.1137/1.9781611975482.145>
- [31] Prosenjit Bose, Paz Carmi, Vida Dujmovic, and Pat Morin. 2018. Near-Optimal $O(k)$ -Robust Geometric Spanners. *CoRR abs/1812.09913* (2018). [arXiv:1812.09913](https://arxiv.org/abs/1812.09913) <https://arxiv.org/abs/1812.09913>
- [32] Prosenjit Bose, Vida Dujmovic, Pat Morin, and Michiel H. M. Smid. 2013. Robust Geometric Spanners. *SIAM J. Comput.* 42, 4 (2013), 1720–1736. <https://doi.org/10.1137/120874473>
- [33] J. Bourgain. 1985. On lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics* 52, 1-2 (1985), 46–52. <https://doi.org/10.1007/BF02776078>
- [34] Jean Bourgain. 1986. The metrical interpretation of superreflexivity in banach spaces. *Israel Journal of Mathematics* 56 (1986), 222–230. <https://doi.org/10.1007/BF02766125>
- [35] J. Bourgain, T. Figiel, and V. Milman. 1986. On Hilbertian subsets of finite metric spaces. *Israel J. Math.* 55, 2 (1986), 147–152. <https://doi.org/10.1007/BF02801990>
- [36] Kevin Buchin, Sarel Har-Peled, and Dániel Oláh. 2019. A Spanner for the Day After. In *35th International Symposium on Computational Geometry, SoCG 2019, June 18-21, 2019, Portland, Oregon, USA*, 19:1–19:15. <https://doi.org/10.4230/LIPIcs.SoCG.2019.19>
- [37] Kevin Buchin, Sarel Har-Peled, and Dániel Oláh. 2020. Sometimes Reliable Spanners of Almost Linear Size. In *28th Annual European Symposium on Algorithms, ESA 2020, September 7-9, 2020, Pisa, Italy (Virtual Conference)*, 27:1–27:15. <https://doi.org/10.4230/LIPIcs.ESA.2020.27>

- [38] Costas Busch, Chinmoy Dutta, Jaikumar Radhakrishnan, Rajmohan Rajaraman, and Srinivasagopalan Srivathsan. 2012. 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*. IEEE Computer Society, 81–90. <https://doi.org/10.1109/FOCS.2012.45>
- [39] Gruiă Călinescu, Howard J. Karloff, and Yuval Rabani. 2004. Approximation Algorithms for the 0-Extension Problem. *SIAM J. Comput.* 34, 2 (2004), 358–372.
- [40] P. B. Callahan and S. R. Kosaraju. 1995. A decomposition of multidimensional point sets with applications to k -nearest-neighbors and n -body potential fields. *Journal of the ACM* 42, 1 (1995), 67–90. <https://doi.org/10.1145/200836.200853>
- [41] T.-H. Hubert Chan, Anupam Gupta, Bruce M. Maggs, and Shuheng Zhou. 2016. On Hierarchical Routing in Doubling Metrics. *ACM Trans. Algorithms* 12, 4 (2016), 55:1–55:22.
- [42] Timothy M. Chan, Sarel Har-Peled, and Mitchell Jones. 2020. On Locality-Sensitive Orderings and Their Applications. *SIAM J. Comput.* 49, 3 (2020), 583–600. <https://doi.org/10.1137/19M1246493>
- [43] Barun Chandra, Gautam Das, Giri Narasimhan, and José Soares. 1995. New sparseness results on graph spanners. *Int. J. Comput. Geom. Appl.* 5 (1995), 125–144. <https://doi.org/10.1142/S0218195995000088>
- [44] Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. 2010. Fault Tolerant Spanners for General Graphs. *SIAM J. Comput.* 39, 7 (2010), 3403–3423. <https://doi.org/10.1137/090758039>
- [45] Shiri Chechik and Christian Wulff-Nilsen. 2018. Near-Optimal Light Spanners. *ACM Trans. Algorithms* 14, 3 (2018), 33:1–33:15. <https://doi.org/10.1145/3199607>
- [46] Vincent Cohen-Addad, Arnold Filtser, Philip N. Klein, and Hung Le. 2020. On Light Spanners, Low-treewidth Embeddings and Efficient Traversing in Minor-free Graphs. In *61th Annual IEEE Symposium on Foundations of Computer Science, FOCS, 589–600*. See: [conference version](#), [arXiv version](#).
- [47] Michael Dinitz and Robert Krauthgamer. 2011. 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*. 169–178. <https://doi.org/10.1145/1993806.1993830>
- [48] Michael Dinitz and Caleb Robelle. 2020. 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*. 493–500. <https://doi.org/10.1145/3382734.3405735>
- [49] Michael Elkin. 2011. Streaming and fully dynamic centralized algorithms for constructing and maintaining sparse spanners. *ACM Trans. Algorithms* 7, 2 (2011), 20:1–20:17. <https://doi.org/10.1145/1921659.1921666>
- [50] M. Elkin, A. Filtser, and O. Neiman. 2020. Distributed Construction of Light Networks. In *Proceedings of the 39th Symposium on Principles of Distributed Computing (PODC'20)*, 483–492. <https://doi.org/10.1145/3382734.3405701>
- [51] Michael Elkin, Ofer Neiman, and Shay Solomon. 2015. Light Spanners. *SIAM J. Discret. Math.* 29, 3 (2015), 1312–1321. <https://doi.org/10.1137/140979538>
- [52] Michael Elkin and Shay Solomon. 2013. 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*. 804–809. <https://doi.org/10.1137/1.9781611973105.37>
- [53] Matthias Englert, Anupam Gupta, Robert Krauthgamer, Harald Räcke, Inbal Talgam-Cohen, and Kunal Talwar. 2014. Vertex Sparsifiers: New Results from Old Techniques. *SIAM J. Comput.* 43, 4 (2014), 1239–1262. <https://doi.org/10.1137/130908440>
- [54] P. Erdős. 1964. Extremal problems in graph theory. *Theory of Graphs and Its Applications (Proc. Sympos. Smolenice) (1964)*, 29–36.
- [55] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. 2004. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.* 69, 3 (Nov. 2004), 485–497. <https://doi.org/10.1016/j.jcss.2004.04.011>
- [56] Arnold Filtser. 2019. 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*. 6:1–6:21. <https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2019.6>
- [57] Arnold Filtser. 2019. On Strong Diameter Padded Decompositions. *CoRR abs/1906.09783 (2019)*. arXiv:1906.09783 <http://arxiv.org/abs/1906.09783>
- [58] Arnold Filtser. 2020. Scattering and Sparse Partitions, and Their Applications. In *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8–11, 2020, Saarbrücken, Germany (Virtual Conference) (LIPIcs, Vol. 168)*, Artur Czumaj, Anuj Dawar, and Emanuela Merelli (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 47:1–47:20. <https://doi.org/10.4230/LIPIcs.ICALP.2020.47>
- [59] Arnold Filtser. 2021. 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*. IEEE, 492–503. <https://doi.org/10.1109/FOCS52979.2021.00056>
- [60] Arnold Filtser, Michael Kapralov, and Navid Nouri. 2020. Graph Spanners by Sketching in Dynamic Streams and the Simultaneous Communication Model. *CoRR abs/2007.14204 (2020)*. arXiv:2007.14204 <https://arxiv.org/abs/2007.14204> To appear in SODA 2021.
- [61] Arnold Filtser and Hung Le. 2021. Clan embeddings into trees, and low treewidth graphs. In *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21–25, 2021*, Samir Khuller and Virginia Vassilevska Williams (Eds.). ACM, 342–355. <https://doi.org/10.1145/3406325.3451043> Full version at <https://arxiv.org/abs/2101.01146>.
- [62] Arnold Filtser and Hung Le. 2021. Locality-Sensitive Orderings and Applications to Reliable Spanners. *CoRR abs/2101.07428 (2021)*. arXiv:2101.07428 <https://arxiv.org/abs/2101.07428>
- [63] Arnold Filtser and Ofer Neiman. 2018. Light Spanners for High Dimensional Norms via Stochastic Decompositions. In *26th Annual European Symposium on Algorithms, ESA 2018, August 20–22, 2018, Helsinki, Finland*. 29:1–29:15. <https://doi.org/10.4230/LIPIcs.ESA.2018.29>
- [64] Arnold Filtser and Shay Solomon. 2020. The Greedy Spanner Is Existentially Optimal. *SIAM J. Comput.* 49, 2 (2020), 429–447. <https://doi.org/10.1137/18M1210678>
- [65] Lee-Ad Gottlieb. 2015. A Light Metric Spanner. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17–20 October, 2015*. 759–772. <https://doi.org/10.1109/FOCS.2015.52>
- [66] Anupam Gupta, Robert Krauthgamer, and James R. Lee. 2003. 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, 534–543. <https://doi.org/10.1109/SFCS.2003.1238226>
- [67] S. Har-Peled. 2020. *Theory seminar talk, CS, UIUC, 10/12/20: Reliable Spanners for Metric Spaces*. Youtube. <https://youtu.be/w00unW9HCtg?t=3630>
- [68] Sarel Har-Peled, Piotr Indyk, and Anastasios Sidiropoulos. 2013. 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*. 804–809. <https://doi.org/10.1137/1.9781611973105.57>
- [69] Sarel Har-Peled, Manor Mendel, and Dániel Oláh. 2021. Reliable Spanners for Metric Spaces. In *37th International Symposium on Computational Geometry (SoCG'21)*. 43:1–43:13. <https://doi.org/10.4230/LIPIcs.SoCG.2021.43> Full version at <https://arxiv.org/abs/2007.08738>.
- [70] Piotr Indyk and Rajeev Motwani. 1998. Approximate Nearest Neighbors: Towards Removing the Curse of Dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23–26, 1998*, Jeffrey Scott Vitter (Ed.). ACM, 604–613. <https://doi.org/10.1145/276698.276876>
- [71] William Johnson and Joram Lindenstrauss. 1984. Extensions of Lipschitz mappings into a Hilbert space. *Contemp. Math.* 26 (1984), 189–206.
- [72] Iyad A. Kanj, Ljubomir Perkovic, and Ge Xia. 2008. Computing Lightweight Spanners Locally. In *Distributed Computing, 22nd International Symposium, DISC 2008, Arcachon, France, September 22–24, 2008. Proceedings*. 365–378. https://doi.org/10.1007/978-3-540-87779-0_25
- [73] Michael Kapralov and David P. Woodruff. 2014. Spanners and sparsifiers in dynamic streams. In *ACM Symposium on Principles of Distributed Computing, PODC '14, Paris, France, July 15–18, 2014*. 272–281. <https://doi.org/10.1145/2611462.2611497>
- [74] P. N. Klein, S. A. Plotkin, and S. Rao. 1993. Excluded Minors, Network Decomposition, and Multicommodity Flow. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC '93)*. 682–690. <https://doi.org/10.1145/167088.167261>
- [75] Hung Le and Shay Solomon. 2019. Truly Optimal Euclidean Spanners. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9–12, 2019*. 1078–1100. <https://doi.org/10.1109/FOCS.2019.00069>
- [76] Christos Levcopoulos, Giri Narasimhan, and Michiel H. M. Smid. 2002. Improved Algorithms for Constructing Fault-Tolerant Spanners. *Algorithmica* 32, 1 (2002), 144–156. <https://doi.org/10.1007/s00453-001-0075-x>
- [77] Tamás Lukovszki. 1999. New Results on Fault Tolerant Geometric Spanners. In *Algorithms and Data Structures, Frank Dehne, Jörg-Rüdiger Sack, Arvind Gupta, and Roberto Tamassia (Eds.)*. Springer Berlin Heidelberg, Berlin, Heidelberg, 193–204. https://doi.org/10.1007/3-540-48447-7_20
- [78] Jiří Matoušek. 1999. On embedding trees into uniformly convex Banach spaces. *Israel Journal of Mathematics* 114, 1 (1999), 221–237. <https://doi.org/10.1007/BF02785579>
- [79] Manor Mendel and Assaf Naor. 2007. Ramsey partitions and proximity data structures. *Journal of the European Mathematical Society* 9, 2 (2007), 253–275. <https://doi.org/10.4171/JEMS/79>
- [80] Assaf Naor and Terence Tao. 2012. Scale-oblivious metric fragmentation and the nonlinear Dvoretzky theorem. *Israel Journal of Mathematics* 192, 1 (2012), 489–504. <https://doi.org/10.1007/s11856-012-0039-7>
- [81] Giri Narasimhan and Michiel H. M. Smid. 2007. *Geometric spanner networks*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511546884>
- [82] D. Peleg and A. A. Schäffer. 1989. Graph spanners. *Journal of Graph Theory* 13, 1 (1989), 99–116. <https://doi.org/10.1002/jgt.13190130114>
- [83] Shay Solomon. 2013. Sparse Euclidean Spanners with Tiny Diameter. *ACM Trans. Algorithms* 9, 3 (2013), 28:1–28:33. <https://doi.org/10.1145/2483699.2483708>
- [84] S. Solomon. 2014. Euclidean Steiner Shallow-Light Trees. In *Proceedings of the Thirtieth Annual Symposium on Computational Geometry (SoCG '14)*. 454:454–454:463.