Labeled Nearest Neighbor Search and Metric Spanners via Locality Sensitive Orderings

Arnold Filtser
Department of Computer Science, Bar-Ilan University, Ramat-Gan, Israel

Abstract

Chan, Har-Peled, and Jones [SICOMP 2020] developed locality-sensitive orderings (LSO) for Euclidean space. A (τ,ρ)-LSO is a collection Σ of orderings such that for every $x,y\in\mathbb{R}^d$ there is an ordering $\sigma\in\Sigma$, where all the points between x and y w.r.t. σ are in the ρ-neighborhood of either x or y. In essence, LSO allow one to reduce problems to the 1-dimensional line. Later, Filtser and Le [STOC 2022] developed LSO's for doubling metrics, general metric spaces, and minor free graphs.

For Euclidean and doubling spaces, the number of orderings in the LSO is exponential in the dimension, which made them mainly useful for the low dimensional regime. In this paper, we develop new LSO's for Euclidean, ℓ_p, and doubling spaces that allow us to trade larger stretch for a much smaller number of orderings. We then use our new LSO's (as well as the previous ones) to construct path reporting low hop spanners, fault tolerant spanners, reliable spanners, and light spanners for different metric spaces.

While many nearest neighbor search (NNS) data structures were constructed for metric spaces with implicit distance representations (where the distance between two metric points can be computed using their names, e.g. Euclidean space), for other spaces almost nothing is known. In this paper we initiate the study of the labeled NNS problem, where one is allowed to artificially assign labels (short names) to metric points. We use LSO's to construct efficient labeled NNS data structures in this model.

2012 ACM Subject Classification Theory of computation → Computational geometry; Theory of computation → Sparsification and spanners

Keywords and phrases Locality sensitive ordering, nearest neighbor search, high dimensional Euclidean space, doubling dimension, planar and minor free graphs, path reporting low hop spanner, fault tolerant spanner, reliable spanner, light spanner

Digital Object Identifier 10.4230/LIPIcs.SoCG.2023.33

Related Version *Full Version*: `https://arxiv.org/abs/2211.11846` [46]

Funding This research was supported by the Israel Science Foundation (grant No. 1042/22).

1 Introduction

1.1 Locality Sensitive Ordering

Chan, Har-Peled, and Jones [30] recently introduce a new and powerful tool into the algorithmist's toolkit, called *locality sensitive ordering* (abbreviated LSO). LSO provides an order over the points of a metric space (X,d_X), this order being very useful, as it helps to store, sort, and search the data (among other manipulations).

▶ **Definition 1** ((τ,ρ)-LSO)**.** *Given a metric space (X,d_X), we say that a collection Σ of orderings is a (τ,ρ)-LSO if $|\Sigma|\le\tau$, and for every $x,y\in X$, there is a linear ordering $\sigma\in\Sigma$ such that (w.l.o.g.[1]) $x\preceq_\sigma y$ 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 $I_y\subseteq B_X(y,\rho\cdot d_X(x,y))$. ρ is called the* stretch *parameter.*

[1] That is either $x\preceq_\sigma y$ or $y\preceq_\sigma x$, and the guarantee holds w.r.t. all the points between x and y in the order σ.

39th International Symposium on Computational Geometry (SoCG 2023).
Editors: Erin W. Chambers and Joachim Gudmundsson; Article No. 33; pp. 33:1–33:18
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

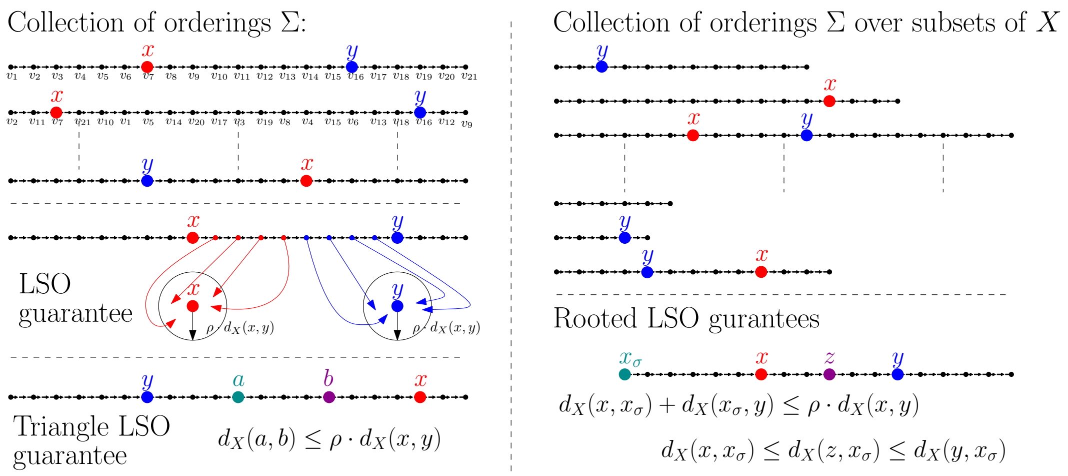

Figure 1 Illustration of different types of LSO.

Morally speaking, given a problem, LSO can reduce it from a general and complicated space to a much simpler space: 1-dimensional line. Chan et al. [30] constructed $\left(O_d(\epsilon^{-d}) \cdot \log \frac{1}{\epsilon}, \epsilon\right)$-LSO for the d-dimensional Euclidean space. They used their LSO to design simple dynamic algorithms for approximate nearest neighbor search, bichromatic closest pair, MST, spanners, and fault-tolerant spanners. Later, Buchin, Har-Peled, and Oláh [27, 28] constructed reliable spanners using LSO, obtaining considerably superior results compared with previous techniques.

Filtser and Le [49] generalized Chan et al. [30] result to doubling spaces,[2] showing that every metric space with doubling dimension d admits a $\left(\epsilon^{-O(d)}, \epsilon\right)$-LSO. Furthermore, they generalized the concept of LSO to other metric spaces, defining the two related notions of *triangle-LSO* (which turn to be useful for general metric spaces), and *left-sided LSO* (which turn to be useful for topologically restricted graphs). Here, instead of presenting the left-sided LSO's of [49], we introduce the closely related notion of *rooted-LSO*, which has some additional structure. All the results and constructions for left-sided LSO in [49] hold for rooted LSO as well. We refer to [49] for a comparison between the different notions, and to Figure 1 for an illustration.

▶ **Definition 2** ((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 $|\Sigma| \le \tau$, and for every $x,y \in X$, there is an ordering $\sigma \in \Sigma$ such that (w.l.o.g.[1]) $x \prec_\sigma y$, and for every $a,b \in X$ such that $x \preceq_\sigma a \preceq_\sigma b \preceq_\sigma y$ it holds that $d_X(a,b) \le \rho \cdot d_X(x,y)$.*

▶ **Definition 3** ((τ,ρ)-rooted-LSO). *Given a metric space (X,d_X), we say that a collection Σ of* orderings *over subsets of X is a (τ,ρ)-rooted-LSO if the following hold:*

- *Each point $x \in X$ belongs to at most τ orderings in Σ.*
- *Each ordering $\sigma \in \Sigma$ is associated with a point $x_\sigma \in X$, which is the first in the order, and such that the ordering is w.r.t. distances from x_σ (i.e. $y \prec_\sigma z \Rightarrow d_X(x_\sigma,y) \le d_X(x_\sigma,z)$).*
- *For every pair of points u,v, there is some $\sigma \in \Sigma$ containing both x,y, and such that $d_G(u,x_\sigma) + d_G(x_\sigma,v) \le \rho \cdot d_G(u,v)$.*

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

Filtser and Le [49] constructed triangle LSO for general metrics, and rooted LSO for the shortest path metrics of trees, treewidth graphs, planar graphs, and graph excluding a fixed minor. They used their LSO's to construct oblivious reliable spanners for the respective metric spaces, considerably improving previous constructions (that used different techniques). All the known results on LSO's are summarized in Table 1.

Table 1 Summary of all known results, on all the different types of locality sensitive orderings (LSO). $k \in \mathbb{N}$, $t > 1$, $\epsilon \in (0,1)$ is an arbitrarily small parameter. $^{(*)}$ O_d hides an arbitrary function of d, the number of orderings in [30] LSO is $O_d(\epsilon^{-d}) \cdot \log \frac{1}{\epsilon} = 2^{O(d)} \cdot d^{\frac{3}{2}d} \cdot \epsilon^{-d} \cdot \log \frac{1}{\epsilon}$.

LSO type	Metric Space	# of orderings (τ)	Stretch (ρ)	Ref
(Classic) LSO	Euclidean space $\mathbb{R}^d$	$O_d(\epsilon^{-d}) \cdot \log \frac{1}{\epsilon}$ $^{(*)}$	ϵ	[30]
	Doubling dimension d	$\epsilon^{-O(d)}$	ϵ	[49]
Triangle-LSO	General metric	$O(n^{\frac{1}{k}} \cdot \log n \cdot \frac{k^2}{\epsilon} \cdot \log \frac{k}{\epsilon})$	$2k + \epsilon$	[49]
	Euclidean space $\mathbb{R}^d$	$e^{\frac{d}{2t^2} \cdot (1+\frac{2}{t^2})} \cdot \tilde{O}(\frac{d^{1.5}}{\epsilon \cdot t})$	$(1+\epsilon)t$	Thm. 4
	ℓ_p^d for $p \in [1,2]$	$e^{O(d/t^p)} \cdot \tilde{O}(d)$	t	Thm. 5
	ℓ_p^d for $p \in [2,\infty]$	$\tilde{O}(d)$	$d^{1-\frac{1}{p}}$	FullV[46]
	Doubling dimension d	$2^{O(d/t)} \cdot d \cdot \log^2 t$	t	Thm. 6
Rooted LSO	Tree	$\log n$	1	[49]
	Treewidth k	$k \cdot \log n$	1	[49]
	Planar / fixed minor free	$O(\frac{1}{\epsilon} \cdot \log^2 n)$	$1 + \epsilon$	[49]

Previously constructed LSO for the Euclidean space [30], as well as for metric spaces with doubling dimension d [49], have exponential dependency on the dimension in their cardinality, a phenomena often referred to as "the curse of dimensionality". When the dimension is high, it can be a major obstacle. Indeed, the distances induced by n point in an $O(\log n)$-dimensional Euclidean space induce a metric space which is much more structured than a general metric space. Therefore one might expect them to admit better LSO. However, using [30] one can only construct (n, ϵ)-LSO (note that every metric admits $(\lceil \frac{n}{2} \rceil, 0)$-LSO [3]).

Every n point metric space has doubling dimension at most $\log n$. Consider the case where the doubling dimension is somewhat large (e.g. $\sqrt{\log n}$) but not maximal. It is much more structured than general metric, however the only construction we have [49] gives us $\epsilon^{-O(d)}$ orderings, which might be too large. In the small number of orderings regime, could we take advantage of the doubling structure to construct better LSO then for general metrics?

Our Contribution. In this paper we construct new triangle-LSO for high dimensional spaces. We then present many applications for the newly constructed LSO's, as well as for the previously constructed LSO's. Old and new LSO construction are summarized in Table 1.

▶ **Theorem 4.** *For every $t \in [4, 2\sqrt{d}]$, $\delta \in (0,1]$, and $d \geq 1$, the d-dimensional Euclidean space $\mathbb{R}^d$ admits $\left(O\left(\frac{d^{1.5}}{\delta \cdot t} \cdot \log(\frac{2\sqrt{d}}{t}) \cdot \log \frac{d}{\delta} \cdot e^{\frac{d}{2t^2} \cdot (1+\frac{2}{t^2})}\right), (1+\delta)t\right)$-triangle LSO.*

For $t = \frac{2}{3}\sqrt{d}$ and $\delta = \frac{1}{2}$, we obtain $\left(O(d \log d), \sqrt{d}\right)$-triangle LSO. In particular, for every set of n points in ℓ_2, using the Johnson Lindenstrauss dimension reduction [61], for every fixed $t > 1$, we can construct $\left(n^{\frac{1}{t^2}} \cdot \tilde{O}(\frac{\log^{1.5} n}{t}), O(t)\right)$-triangle LSO, or $\left(\tilde{O}(\log n), O(\sqrt{\log n})\right)$-triangle LSO, a quadratic improvement compared with general n-point metric spaces!

[3] This follows from a theorem by Walecki [7] who showed that the edges of the K_n clique graph can be partitioned into $\lceil \frac{n}{2} \rceil$ Hamiltonian paths.

Interestingly, we show that the $\left(O(d\log d), \sqrt{d}\right)$-triangle LSO Σ for ℓ_2, is in the same time also a $\left(O(d\log d), d^{\frac{1}{p}}\right)$-triangle LSO for ℓ_p where $p \in [1,2]$, and $\left(O(d\log d), d^{1-\frac{1}{p}}\right)$-triangle LSO for ℓ_p where $p \in [2,\infty]$. For $p \in [1,2]$, we generalize Theorem 4 to ℓ_p spaces to get the entire #ordering-stretch trade-off. Finally, we generalize Theorem 4 to general metric spaces with doubling dimension d.

▶ **Theorem 5.** *For every $p \in [1,2]$, $t \in [5, d^{\frac{1}{p}}]$ and $d \geq 1$, the d-dimensional ℓ_p space admits $\left(e^{O(\frac{d}{t^p})} \cdot \tilde{O}(d), t\right)$-triangle LSO.*

▶ **Theorem 6.** *Given a metric space (X, d_X) with doubling dimension d, and parameter $t \in [\Omega(1), d]$, X admits $\left(2^{O(d/t)} \cdot d \cdot \log^2 t, t\right)$-triangle LSO.*

For $t = d$, we get $(\tilde{O}(d), d)$-triangle LSO, again much better then general metric spaces!

1.2 Labeled Nearest Neighbor Search

Nearest neighbor search (abbreviated NNS) is a classical and fundamental task used in numerous domains including machine learning, clustering, document retrieval, databases, statistics, data compression, database queries, computational biology, data mining, pattern recognition, and many others. In the NNS problem we are given a set P of points in a metric space (X, d_X). The goal is to construct a succinct data structure that given a query point $q \in X$, quickly returns a point $p \in P$ closest to q (i.e. arg $\min_{p\in P} d_X(p,q)$). In order to keep the size of the data structure, and the query time small, usually approximation is allowed. In the t-approximate nearest neighbor problem (abbreviated t-NNS) the goal is to return a point p at distance at most $t \cdot \min_{p\in P} d_X(p,q)$ from q. The problem was extensively studied in ℓ_p spaces (see the survey [11]), and also in various norm spaces over $\mathbb{R}^d$ (see e.g. [12, 13]). NNS data structures were also constructed beyond normed spaces. Some examples are Earth-Mover distance [60], Edit Distance [79, 11], and Fréchet distance [59, 41, 43, 47]. We observe that a crucial property shared by these examples, is that they have an "implicit distance representation". That is, it is possible to compute the distance between two points using only their names (e.g. the coordinates values in $\mathbb{R}^d$ used as names: $d_{\mathbb{R}^d}\left((x_1,\dots,x_d),(y_1,\dots,y_d)\right) = \|(x_1,\dots,x_d)-(y_1,\dots,y_d)\|_2$).

For general metric spaces, Krauthgamer and Lee [66] introduced the *black box model*. Here one is given access to an exact distance oracle [4] DO that answer distance queries in t_{DO} time. They showed that one can construct an efficient $(1+\epsilon)$-NNS (that is with polynomial space, and polylogarithmic query time), if and only if the doubling dimension of X is at most $O(\log\log n)$.

Indeed, for metric spaces with large doubling dimension, distance queries provide very limited information. Consider for example the case where the input metric is the star graph (inducing uniform metric on the leaves, see illustration below), and the query point attached to one of the leaves with an edge of infinitesimal weight, one must query all the points before finding any finite approximation to the nearest neighbor.

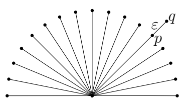

[4] An exact distance oracle D is a data structure that given two points x, y, returns $\mathrm{est}(x,y) = d_X(x,y)$. A distance oracle of stretch t returns a value $\mathrm{est}(x,y)$ in $[d_X(x,y), t \cdot d_X(x,y)]$.

An interesting case studied by Abraham, Chechik, Krauthgamer, and Wieder [3] is that of planar graphs. Here we are given a huge weighted planar graph $G = (V, E, w)$ with N vertices, and a subset of n vertices $X \subseteq V$. The goal is to solve the $(1+\epsilon)$-NNS problem w.r.t. the shortest path metric d_G, input set X and queries from V. Assuming access to an exact distance oracle[4] DO that answer distance queries in t_{DO} time, and given a planar graph G of maximum degree Δ, Abraham et al. [3] constructed a $(1+\epsilon)$-NNS data structure for planar graph of size $n \cdot O(\epsilon^{-1} \cdot \log\log N + \Delta \cdot \log^2 n)$ and query time $O((\epsilon^{-1} \cdot \log\log n + t_{\mathrm{DO}}) \cdot \log\log N + \log n \cdot \Delta \cdot t_{\mathrm{DO}})$.

Linear dependence on the degree is a very limiting requirement, as planar graphs have apriori unbounded degree. Moreover, exact distance computations (even in planar graphs) are time consuming, and if the graph is big enough could be infeasible. Exact distance oracle is a highly non-trivial assumption, it is an expensive data structure, [5] better to be avoided. One might hope to relax either the max degree assumption, or to use the much more reasonable and efficient data structure of approximate distance oracle [84, 64, 70]. Unfortunately, Abraham et al. [3] showed both assumptions to be necessary. Specifically, the dependence on the degree is necessary, as every NNS data structure with space at most $O(\frac{N}{\Delta \log_\Delta n})$ must probe the distance oracle at least $\Omega(\Delta \log_\Delta n)$ times. Furthermore, they show that if one is only given access to a $(1+\epsilon)$-distance oracle, then there is a planar graph (in fact a tree) with maximum degree $O(\log n)$, aspect ratio $O(\frac{\log n}{\epsilon})$, $N \le n^2$, and the NNS data structure is forced to make $\Omega(n)$ queries to the distance oracle.

To conclude this discussion, exact distance oracle (assumed both by the black box model [66] and [3]) is an expensive data structure, which enables us to construct efficient NNS only under very limiting assumptions (small doubling dimension / constant maximum degree in planar graphs). On the other hand in many metric spaces with "implicit distance representation" efficient NNS were constructed. The crux is that the information stored in the name (e.g. coordinate values) used to preform various manipulations on the data, in addition to distance computation. What if in planar graphs, or even in completely general metric spaces, we could choose the names of the metric points, or alternatively assigning each point a short label, would it be possible to construct efficient NNS data structures?

To answer this question, we introduce the *labeled t-NNS* problem.

▶ **Definition 7** (Labeled t-NNS). *Consider an N-point metric space (X, d_X), where one can assign to each point $x \in X$ an arbitrary short label l_x. Given a subset $P \subseteq X$ of size n (unknown in advance) together with their labels $\{l_x\}_{x \in P}$ (but without access to (X, d_X) or any additional information) the goal is to construct a NNS search data structure as follows: given a query $q \in X$ together with its assigned label ℓ_q, the data structure will return a t-approximate nearest neighbor $p \in P$: $d_X(p,q) \le t \cdot \min_{x \in P} d_X(x,q)$. The parameters of study are: label size, data structure size, query time, and approximation factor t.*

We also consider the scenario where the set P is changing dynamically: points are added and removed from P. Here we are required to maintain a data structure for P, while minimizing the update time (as well as all the other parameters).[6]

In the labeled NNS model we get to assign a short label (alternatively choose a name) for each point in a big metric space (X, d_X). These labels try to imitate the natural hint provided by the name of the points themselves in metric spaces with implicit distance

[5] After a long line of work, the state of the art (by Long and Pettie [73]) requires either super-linear space $N^{1+o(1)}$, or very large query time $N^{o(1)}$, both quite undesirable.

[6] For example, consider a NNS data structure for a set P. Dynamic NNS, should be able to efficiently update the data structure to work w.r.t. a slightly updated set $P' = P \cup \{x\} \setminus \{y\}$ instead of P.

representation. The main object of study here is the trade-off between label size, and the approximation of the resulting NNS. A trivial choice of label for each point x will be simply to store distances to all other points. However the label size $\Omega(N)$ is infeasible. A more sophisticated solution is the following: fix constants $k, t \in \mathbb{N}$, and embed all the points in (X, d_X) into $d = \tilde{O}(N^{\frac{1}{k}})$-dimensional ℓ_∞ [75, 1]. That is we assign each point x a vector $v_x \in \mathbb{R}^d$ such that $\forall x, y \in X$, $d_X(x, y) \le \|v_x - v_y\|_\infty \le (2k-1) \cdot d_X(x, y)$, and use the vectors as labels. Given an n point subset $P \subseteq X$ with its respective labels (vectors), use Indyk's NNS [58] over $\{v_x\}_{x\in P}$ to construct a NNS data structure $\mathbb{D}_{\mathrm{Ind}}$ with approximation factor $O(\log_{1+\frac{1}{t}} \log d) = O(t \cdot \log\log N)$ w.r.t. the ℓ_∞ vectors, space $\tilde{O}(d \cdot n^{1+\frac{1}{t}}) = \tilde{O}(N^{\frac{1}{k}} \cdot n^{1+\frac{1}{t}})$ and query time $\tilde{O}(n^{1+\frac{1}{k}})$. Given a query q, we will simply query $\mathbb{D}_{\mathrm{Ind}}$ on the vector v_q, and on answer v_p will return p. Note that the query time and space are the same as above, while the approximation factor will be $O(k \cdot t \cdot \log\log N)$.

Our Contribution. Our results for the labeled t-NNS are summarized in Table 2. We begin by proving meta theorem showing that (τ, ρ)-rooted LSO implies a labeled ρ-NNS with label size $O(\tau)$, space $O(n \cdot \tau)$, query time $O(\tau)$, and update time $O(\tau \cdot \log\log N)$. As a result we conclude efficient labeled $(1+\epsilon)$-NNS data structures for fixed minor free graphs (and planar), and exact labeled NNS for treewidth graphs. Another interesting corollary is an efficient labeled NNS for metrics with small *correlation dimension* (a generalization of doubling, see [29]).

Table 2 Labeled NNS data structures for different families. The $*$ sign is replacing $O(\log\log N)$. The second to last line is a lower bound. Space is measured in machine words. The label size and query time always equal. The space in all the cases above equals n times the label size.

Family	stretch	label	query time	update time	Ref
Minor free	$1+\epsilon$	$O(\frac{1}{\epsilon}\log^2 N)$	$O(\frac{1}{\epsilon}\log^2 N)$	$\frac{1}{\epsilon} \cdot \tilde{O}(\log^2 N)$	FullV[46]
Treewidth k	1	$O(k \log N)$	$O(k \log N)$	$k \cdot \tilde{O}(\log N)$	FullV[46]
Correlation k	$1+\epsilon$	$\tilde{O}_{k,\epsilon}(\sqrt{N})$	$\tilde{O}_{k,\epsilon}(\sqrt{N})$	$\tilde{O}_{k,\epsilon}(\sqrt{N})$	FullV[46]
Ultrametric	1	$O(\log N)$	$*$	$*$	FullV[46]
General Metric	$8(1+\epsilon)k$	$O(\frac{k}{\epsilon} N^{\frac{1}{k}} \cdot \log N)$	$O(\frac{1}{\epsilon} \cdot *)$	$O(\frac{k}{\epsilon} N^{\frac{1}{k}} \cdot *)$	FullV[46]
	$t < 2k+1$	$\tilde{\Omega}(N^{\frac{1}{k}})$	arbitrary	arbitrary	FullV[46]
Doubling d	t	$2^{O(d/t)} \cdot \tilde{O}(d) \cdot \log N$	$2^{O(d/t)} \cdot \tilde{O}(d) \cdot *$	$2^{O(d/t)} \cdot \tilde{O}(d) \cdot *$	FullV[46]

Next, we prove a meta theorem, showing that (τ, ρ)-triangle LSO implies a labeled 2ρ-NNS with label size $O(\tau \cdot \log N)$, space $O(n \cdot \tau \cdot \log N)$, and query and update time $O(\tau \cdot \log\log N)$. We conclude an efficient labeled NNS for graphs with large doubling dimension. For the high-dimensional Euclidean space, approximate nearest neighbor search was extensively studied (see the survey [11], and additional discussion in the full version [46]). However, for the case of doubling metrics, NNS never went beyond $1+\epsilon$ approximation. In particular, in all existing solutions the query time and space have exponential dependence on the dimension (see references in the full version [46]). Thus ours are the first results in this regime, removing "the curse of dimensionality".

As an additional corollary of the triangle LSO to labeled NNS meta theorem one can derive a NNS of for general metric spaces which considerably improved upon the labeled NNS based on [75]+[58] discussed above. However, the query time turns out to be somewhat large. We provide direct constructions for labeled NNS for general metrics, getting label size $\tilde{O}(\epsilon^{-1} \cdot N^{\frac{1}{k}})$, stretch $8(1+\epsilon)k$ and very small query time: $O(\epsilon^{-1} \cdot \log\log N)$. We show that the standard information theoretic bound applies for the labeled NNS as well, specifically, for stretch $t < 2k+1$, the label size must be $\tilde{\Omega}(n^{\frac{1}{k}})$ (regardless of query time). Finally, we put special focus on the regime where the stretch is $O(\log N)$. We obtain labeled NNS scheme with very short label and small query time. Most notably, assuming polynomial aspect ratio, and allowing the bound on the label to be only in expectation, we can obtain $O(1)$ label size, and $O(\log\log N)$ query time.

1.3 Spanners

Given a metric space (X, d_X), a metric *spanner* is a graph H over X points, such that that the shortest path metric d_H in H, closely resembles the metric d_X. Formally, a t*-spanner* for X is a weighted graph $H(X, E, w)$ that has $w(u, v) = d_X(u, v)$ for every edge $(u, v) \in E$ and $d_H(x, y) \le t \cdot d_X(x, y)$ for every pair of points $x, y \in X$. [7] The classic parameter of study is the trade-off between stretch and sparsity (number of edges). Althöfer et al. [8] showed that every n point metric space admits a $2k-1$ spanner with $O(n^{1+\frac{1}{k}})$ edges, while every set of n points in $\mathbb{R}^d$, or more generally metric space of doubling dimension d, admits a $(1+\epsilon)$-spanner with $n \cdot \epsilon^{-O(d)}$ edges [38, 52]. We refer to the book [77], and the survey [4] for an overview.

Path Reporting Low Hop Spanners. Recently, Kahalon, Le, Milenkovic, and Solomon [62] studied *path reporting low-hop spanners*. While a t-spanner guarantees that a "short" path exists between every two points, such a path might be very long, and finding it is a time consuming operation. A path reporting t-spanner, is a spanner accompanied with a data structure that given a query pair $\{x, y\}$, efficiently retrieves a path between x and y (of total weight $\le t \cdot d_X(x, y)$). A path P with h edges is called an h-hop path. H is an h-hop t-spanner of X if for every $x, y \in X$, there is an h-hop path P from x to y in H, such that $w(P) \le t \cdot d_X(x, y)$. Clearly, the time required to report a path is at least as large as the number of edges along the path, thus we wish to minimize the number of hops.

Low number of hops is a highly desirable property in network design, as each transmission causes delays, which are non-negligible when the number of transmissions is large [5, 23]. Low hop networks are also known to be more reliable [23, 87, 82], and used in electricity and telecommunications [23], and many other (practical) network design problems [71, 16, 55, 54, 81]. Hop-constrained network approximation is often used in parallel computing [36, 14], as the number of hops governs the number of required parallel rounds (e.g. in Dijkstra).

Kahalon et al. [62] constructed path reporting low-hop spanners for many spaces, such as path reporting 2-hop $O(k)$-spanners with $O(n^{1+\frac{1}{k}} \cdot k \cdot \log n)$ edges, and $O(1)$ query time for general metrics, and path reporting 2-hop $(1+\epsilon)$-spanners with $O(\frac{n}{\epsilon^2} \cdot \log^3 n)$ edges and $O(\epsilon^{-2} \cdot \log^2 n)$ query time for planar graphs. They showed a plethora of applications for their spanners: compact routing schemes, fault tolerant routing, spanner sparsification, approximate shortest path trees, minimum weight trees (MST), and online MST verification.

Our Contribution. Kahalon et al. [62] first constructed path reporting low hop spanners for trees, and then reduced each type of metric to the case of trees. We observe that it is actually enough to reduce to the even simpler case of paths, and obtain a host of such spanners using LSO's. We then manually improve some of the resulting spanners, most notably we create

[7] Frequently the literature is concerned with graph spanners, where given a graph $G = (V, E, w)$ the goal is to find a subgraph H preserving distances. Here we study metric spanners, where there is no underlying graph.

Table 3 Summary of old and new results on path reporting low hop spanners. The spanners are for n point metrics, and all report paths with hop bound 2. Here $\epsilon \in (0,1)$, $k, d \geq 1$ are integers. The space required for the path reporting data structure is asymptotically equal to the sparsity of the spanner in all the cases other than Euclidean space, where there is an additional additive factor of $O_d(\epsilon^{-2d}) \log \frac{1}{\epsilon}$.

Metric family	stretch	sparsity	query time	Ref
General Metric	$O(k)$	$O\left(n^{1+\frac{1}{k}} \cdot k \cdot \log n\right)$	$O(1)$	[62]
	$2k-1$	$O(n^{1+\frac{1}{k}} \cdot k)$	$O(k)$	FullV[46], [85]
	$(1+\epsilon)(4k-2)$	$O(n^{1+\frac{1}{k}} \cdot \epsilon^{-1} \cdot k \cdot \log \Phi)$	$O(\epsilon^{-1} \cdot \log 2k)$	FullV[46]
Doubling Dimension	$1+\epsilon$	$\epsilon^{-O(d)} \cdot n \cdot \log n$	$\epsilon^{-O(d)}$	[62]
	t	$2^{-O(d/t)} \cdot \tilde{O}(n)$	$2^{-O(d/t)} \cdot d \cdot \log^2 t$	FullV[46]
Euclidean $\mathbb{R}^d$	$1+\epsilon$	$O_d(\epsilon^{-d}) \cdot \log \frac{1}{\epsilon} \cdot n \cdot \log n$	$O_d(1)$	FullV[46]
	$(1+\epsilon)t$	$\tilde{O}(\frac{d^{1.5}}{\epsilon \cdot t}) \cdot e^{\frac{2d}{t^2} \cdot (1+\frac{8}{t^2})} \cdot n \log n$	$\tilde{O}(\frac{d^{1.5}}{\epsilon \cdot t}) \cdot e^{\frac{2d}{t^2} \cdot (1+\frac{8}{t^2})}$	FullV[46]
ℓ_p^d, $p \in [1,2]$	t	$\tilde{O}(d) \cdot e^{O(\frac{d}{t^p})} \cdot n \log n$	$\tilde{O}(d) \cdot e^{O(\frac{d}{t^p})}$	FullV[46]
ℓ_p^d, $p \in [2,\infty]$	$2 \cdot d^{1-\frac{1}{p}}$	$\tilde{O}(d) \cdot n \log n$	$\tilde{O}(d)$	FullV[46]
Tree	1	$O\left(n \cdot \log n\right)$	$O(1)$	[62]
Fixed Minor Free	$1+\epsilon$	$O\left(n \cdot \epsilon^{-2} \cdot \log^3 n\right)$	$O(\epsilon^{-2} \cdot \log^2 n)$	[62]
	$1+\epsilon$	$O(n \cdot \epsilon^{-1} \cdot \log^2 n)$	$O(\epsilon^{-1} \cdot \log n)$	FullV[46]
Planar	$1+\epsilon$	$O(n \cdot \epsilon^{-1} \cdot \log^2 n)$	$O(\epsilon^{-1})$	FullV[46]
Treewidth k	1	$O(n \cdot k \cdot \log n)$	$O(k)$	FullV[46]

path reporting 2-hop $(1+\epsilon)$-spanner for planar graph with $O(\frac{n}{\epsilon}\log^2 n)$ edges and $O(\frac{1}{\epsilon})$-query time, and a path reporting 2-hop $(1+\epsilon)$-spanner for points in d-dimensional Euclidean space with $O_d(\epsilon^{-d}) \cdot \log \frac{1}{\epsilon} \cdot n \log n$ edges and $O_d(1)$-query time. See Table 3 for a summary of old and new results.

Fault tolerant spanners. Levcopoulos, Narasimhan, and Smid [72] introduced the notion of a fault-tolerant spanner. A graph $H = (X, E_H, w)$ is an f-*vertex-fault-tolerant* t-spanner of a metric space (X, d_X), if for every set $F \subset X$ of at most f vertices, it holds that $\forall u, v \notin F$, $d_{H \setminus F}(u,v) \leq t \cdot d_X(u,v)$. For general metrics, after a long line of work [34, 39, 20, 22, 40, 21, 80], 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 Erdös' Girth Conjecture [44]. For n-points in d dimensional Euclidean space, or more generally in a space of doubling dimension d, f-*vertex fault tolerant* $(1+\epsilon)$-spanner were constructed with $\epsilon^{-O(d)} \cdot f \cdot n$ edges [72, 74, 83].

Kahalon et al. [62] initiated the study of low-hop fault tolerant spanners (previous constructions had $\Omega(\log n)$ hops). An h-hop f-fault tolerant t-spanner H of a metric (X, d_x) is a graph over X such that for every set $F \subseteq X$ of at most f vertices, for every $x, y \notin F$, the spanner without F: $H[X \setminus F]$ contains an h-hop path between x to y of weight at most $t \cdot d_X(x,y)$. The advantages of such a spanner are straightforward, we refer to [62] for a discussion. Kahalon et al. constructed a 2-hop f-fault tolerant spanner for doubling spaces with $n \cdot f^2 \cdot \epsilon^{-O(d)} \cdot \log n$ edges. Note that a linear dependence on f is necessary (as if a point has degree $\leq f$ in H, we can delete all it's neighbors and get distortion ∞). It is natural to ask whether it is possible to construct such a spanner with only a linear dependence, and not quadratic as in [62].

Our Contribution. One can easily construct f-fault tolerant 1-spanner for the path graph with $O(nf)$ edges. We observe that using $O(nf \log n)$ edges, it is possible to obtain f-fault tolerant 2-hop 1-spanner for the path graph (note that $O(n \log n)$ edges are necessary for

Table 4 Summary of old and new results on 2-hop f-fault tolerant spanners. The spanners are for n point metrics, and all report paths with hop bound 2. Here $\epsilon \in (0,1)$, $k,d \geq 1$ are integers.

Family	Stretch	Edges	Ref
Doubling dimension d	$1+\epsilon$	$\epsilon^{-O(d)} \cdot f^2 \cdot n \cdot \log n$	[62]
	$1+\epsilon$	$\epsilon^{-O(d)} \cdot f \cdot n \cdot \log n$	FullV[46]
	t	$2^{-O(d/t)} \cdot f \cdot \tilde{O}(n)$	FullV[46]
General Metric	$4k+\epsilon$	$\tilde{O}(n^{1+\frac{1}{k}} \cdot f \cdot \epsilon^{-1})$	FullV[46]
Euclidean $\mathbb{R}^d$	$1+\epsilon$	$O_d(\epsilon^{-d}) \log \frac{1}{\epsilon} \cdot f \cdot n \cdot \log n$	FullV[46]
	$(1+\epsilon)k$	$e^{\frac{2d}{k^2} \cdot (1+\frac{8}{k^2})} \cdot \tilde{O}(\frac{d^{1.5}}{\epsilon \cdot k}) \cdot f \cdot n \cdot \log n$	FullV[46]
ℓ_p^d, $p \in [1,2]$	k	$e^{O(\frac{d}{k^p})} \cdot \tilde{O}(d) \cdot f \cdot n \cdot \log n$	FullV[46]
ℓ_p^d, $p \in [2,\infty]$	$2 \cdot d^{1-\frac{1}{p}}$	$\tilde{O}(d) \cdot f \cdot n \cdot \log n$	FullV[46]
Treewidth k	2	$O(n \cdot k \cdot f \cdot \log n)$	FullV[46]
Fixed Minor Free	$2+\epsilon$	$O(\frac{n}{\epsilon} \cdot f \cdot \log^2 n)$	FullV[46]

every 2-hop spanner [6, 68]). Using the various old and new LSO's, we obtain a host of f-fault tolerant 2-hop spanners for various metric spaces. Most notably, for metrics with doubling dimension d, we obtain an f-fault tolerant 2-hop $(1+\epsilon)$-spanner with $\epsilon^{-O(d)} \cdot f \cdot n \cdot \log n$ edges, getting the desired linear dependence on f. See Table 4 for a summary of results.

Reliable spanners. 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. [26] introduced the notion of a *reliable spanner*. A ν-reliable spanner is a graph such that for every 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 superset of cardinality at most $(1+\nu) \cdot |B|$. An oblivious ν-reliable t-spanner is a distribution $\mathcal{D}$ over spanners, such that for every failure set B, $H \setminus B$ is a t-spanner for $X \setminus B_H^+$, where the superset B_H^+ depends on both B and the sampled spanner H. The guarantee is that the cardinality of B_H^+ is bounded by $(1+\nu) \cdot |B|$ in expectation.

ν-Reliable spanners were constructed for d dimensional Euclidean and doubling spaces with $n \cdot \epsilon^{-O(d)} \cdot \tilde{O}(\log n)$ edges [27, 28, 49] by a reduction from (classic) LSO's. Oblivious reliable spanners were constructed also for planar, minor free, treewidth graphs, and general metrics [49] by reductions from triangle, and rooted LSO's (as well as from sparse covers [57]).

Our Contribution. Our newly constructed triangle LSO's for high dimensional Euclidean, ℓ_p spaces, and doubling spaces, directly imply reliable spanners for these spaces, obtaining the first results without exponential dependence on the dimension. See Table 5 for a summary.

Light spanners. An extensively studied parameter is the *lightness* of a spanner, defined as the ratio $w(H)/w(MST(X))$, where $w(H)$ resp. $w(MST(X))$ is the total weight of edges in H resp. a minimum spanning tree (MST) of X. Obtaining spanners with small lightness (and thus total weight) is motivated by applications where edge weights denote e.g. establishing cost. The best possible total weight that can be achieved in order to ensure finite stretch is the weight of an MST, thus making the definition of lightness very natural.

Table 5 Summary of previous and new constructions of ν-reliable spanners.

Family	stretch	guarantee	size	ref
Euclidean $(\mathbb{R}^d, \|\cdot\|_2)$	$1+\epsilon$	Deterministic	$n \cdot \tilde{O}_d(\epsilon^{-7d})\nu^{-6} \cdot \tilde{O}(\log n)$	[27]
	$1+\epsilon$	Oblivious	$n \cdot \tilde{O}_d(\epsilon^{-2d}) \cdot \tilde{O}(\nu^{-1}(\log\log n)^2)$	[28]
	$(1+\epsilon)t$	Oblivious	$\nu^{-1} \cdot e^{\frac{4d}{t^2}\cdot(1+\frac{8}{t^2})} \cdot \tilde{O}(n \cdot \frac{d^3}{\epsilon^2 \cdot t^2})$	FullV[46]
ℓ_p^d for $p \in [1,2]$	t	Oblivious	$\nu^{-1} \cdot e^{O(\frac{d}{t^p})} \cdot \tilde{O}(n \cdot d^2)$	FullV[46]
ℓ_p^d for $p \in [2,\infty]$	$2 \cdot d^{1-\frac{1}{p}}$	Oblivious	$\nu^{-1} \cdot \tilde{O}\left(n \cdot d^2\right)$	FullV[46]
Doubling dimension d	$1+\epsilon$	Deterministic	$n \cdot \epsilon^{-O(d)}\nu^{-6} \cdot \tilde{O}(\log n)$	[49]
	$1+\epsilon$	Oblivious	$n \cdot \epsilon^{-O(d)}\nu^{-1}\log\nu^{-1} \cdot \tilde{O}(\log\log n)^2$	[49]
	t	Oblivious	$\tilde{O}(n) \cdot \nu^{-1} \cdot 2^{O(d/t)}$	FullV[46]
General metric	$8t+\epsilon$	Oblivious	$\tilde{O}(n^{1+1/t} \cdot \epsilon^{-2}) \cdot \nu^{-1}$	[49]
Tree	2	Oblivious	$n \cdot O(\nu^{-1}\log^3 n)$	[49]
Treewidth k	2	Oblivious	$n \cdot O(\nu^{-1}k^2\log^3 n)$	[49]
Planar/Minor-free	$2+\epsilon$	Oblivious	$n \cdot O(\nu^{-1}\epsilon^{-2}\log^5 n)$	[49]

Table 6 Summary of previews and new results of light spanners for high dimensional metric spaces. Interestingly, for $p \in [1,2]$ [49] obtain lightness $O(\frac{t^{1+p}}{\log^2 t} \cdot n^{O(\frac{\log^2 t}{t^p})} \cdot \log n)$ regardless of dimension, which is superior to ours for $d \gg \log n$.

Metric space	**Stretch**	**Lightness**	**Ref**
Euclidean space	$O(t)$	$O(n^{\frac{1}{t^2}} \cdot \log n \cdot t)$	[69]
	$O(t)$	$O(e^{\frac{d}{t^2}} \cdot \log^2 n \cdot t)$	[50]
	$(1+\epsilon)2t$	$e^{\frac{d}{2t^2}\cdot(1+\frac{2}{t^2})} \cdot \tilde{O}(\frac{d^{1.5}}{\epsilon^2}) \cdot \log n$	FullV[46]
	$(1+\epsilon)4t$	$e^{\frac{d}{2t^2}\cdot(1+\frac{2}{t^2})} \cdot \tilde{O}(\frac{d^{1.5}}{\epsilon^2}) \cdot \log^* n$	FullV[46]
Doubling dimension d	$O(t)$	$O(2^{\frac{d}{t}} \cdot t \cdot \log^2 n)$	[50]
	$O(t)$	$2^{O(d/t)} \cdot d \cdot \log^2 t \cdot \log^* n$	FullV[46]
	d	$O(d \cdot \log^2 n)$	[50]
	d	$O(d \cdot \log^2 d \cdot \log^* n)$	FullV[46]
ℓ_p^d for $p \in [1,2]$	t	$O(\frac{t^{1+p}}{\log^2 t} \cdot n^{O(\frac{\log^2 t}{t^p})} \cdot \log n)$	[50]
	t	$e^{O(\frac{d}{t^p})} \cdot \tilde{O}(d \cdot t) \cdot \log^* n$	FullV[46]
ℓ_p^d for $p \in [2,\infty]$	$4 \cdot d^{1-\frac{1}{p}}$	$\tilde{O}(d^{2-\frac{1}{p}}) \cdot \log^* n$	FullV[46]

Obtaining light spanners for general graphs has been the subject of an active line of work [8, 31, 42, 18, 35, 51], where the state of the are by Le and Solomon [69] who obtained $(1+\epsilon)(2k-1)$ spanner with lightness $O(\epsilon^{-1} \cdot n^{\frac{1}{k}})$. Light spanners were also studied extensively in Euclidean spaces (see the book [77]), doubling spaces [53, 51, 25], planar and minor free graphs [63, 65, 24, 67, 37], and high dimensional Euclidean and doubling spaces [56, 50, 69].

Our Contribution. Recently Le and Solomon [69] obtain a general framework for constructing light spanners from spanner oracles. We construct new spanner oracles using LSO's. As a result we derive new light spanners, that improve the state of the art for high dimensional spaces (and match the state of the art for low dimensional doubling spaces). See Table 6 for a summary of results.

1.4 Technical ideas

Triangle LSO for high dimensional Euclidean space. Our construction is very natural: partition the space randomly in every distance scale ξ_i (for some large ξ) into clusters of diameter ξ_i, such that close-by points are likely to be clustered together. In the created ordering σ, points in each cluster will be ordered consecutively and recursively. In particular, the ordering σ will correspond to a laminar partition obtained by the clustering in all possible scales. For a pair of points $x, y \in \mathbb{R}^d$ to be satisfied in the resulting ordering σ, they have to be clustered together in all the distance scales $\xi^i \geq t \cdot \|x - y\|_2$.

Our space partition in each scale is done using ball carving (ala [10]): pick a uniformly random series of centers $z_1, z_2, \dots,$. Each points is assigned to the cluster of the first center at distance at most $R = \frac{1}{2} \cdot \xi^i$. We show that a finite random seed of size $d^{O(d)}$ is enough to sample such a clustering (in all possible distance scales, simultaneously). The probability that two points x, y are clustered together is then equal to the ratio between the volumes of intersection and union of balls: $\Pr[x, y \text{ clustered together}] = \frac{\mathrm{Vol}_d(B(x,R) \cap B(y,R))}{\mathrm{Vol}_d(B(x,R) \cup B(y,R))} \geq \Omega(\frac{1}{\sqrt{d}}) \cdot \left(1 - (\frac{\|x-y\|_2}{R})^2\right)^{d/2}$. We bound this ratio for the case $\|x - y\|_2 \leq \frac{R}{\sqrt{d}}$ using a lemma from [33]. For the general case, we prove that the ratio between these volumes is at least $\Omega(\frac{R}{\sqrt{d} \cdot \|p-q\|_2}) \cdot (1 - (\frac{\|p-q\|_2}{R})^2)^{\frac{d}{2}}$, slightly improving a similar fact from [9], by a $\frac{R}{\|p-q\|_2}$ factor. This ratio eventually governs our success probability (when replacing $R/\|p-q\|_2$ by twice the stretch $2t$). The improved analysis of the volumes ratio is significant for the $O(\sqrt{d})$-stretch regime, improving the number of orderings to $\tilde{O}(d)$ (compered with $\tilde{O}(d^{1.5})$ orderings if we were using [9]).

To generalize this construction to ℓ_p spaces, we use the exact same construction, replacing ℓ_2 balls with ℓ_p balls. The volume ratio lemma from [32] for close-by points is replaced by a crude observation without any significant consequences to the resulting number of orderings. For the general case, we directly analyze the ratio of volumes for ℓ_p-balls (our computation is similar to [78]). The rest of the analysis is the same.

Triangle LSO for doubling spaces. Ultrametrics are trees with additional structure, where each ultrametric admits a $(1, 1)$-triangle LSO. (τ, ρ)-ultrametric cover of a metric space (X, d_X) is a collection $\mathcal{U}$ of τ ultrametrics such that every pair $x, y \in X$ is well approximated by the ultrametrics: $d_X(x, y) \leq \min_{U \in \mathcal{U}} d_U(x, y) \leq \rho \cdot d_X(x, y)$. Filtser and Le [49] showed that (τ, ρ)-ultrametric cover implies (τ, ρ)-triangle LSO. We construct $\left(2^{O(d/t)} \cdot d \cdot \log^2 t, t\right)$-ultrametric cover for spaces with doubling dimension d, implying Theorem 6.

Our starting point for constructing the ultrametric cover is Filtser's [45] padded partition cover, which is a collection of $\approx 2^{O(d/t)}$ space partitions where all clusters are of diameter at most Δ, and every ball of radius $\frac{\Delta}{t}$ is fully contained in a single cluster in one of the partitions. We take a single partition from each distance scales, where the gap between the distance scales is somewhat large: $O(\frac{t}{\epsilon})$. Initially these partitions are unrelated, and we "force" them to be laminar, while keeping the padding property. Each such laminar partition induces an ultrametric, and their union is the desired ultrametric cover.

Labeled NNS. Morally, given a (τ, ρ) LSO (or triangle LSO), the NNS label of every point is simply its position in each ordering. Given a query q, we simply find its successor and predecessor in each one of the orderings, one of them is guaranteed to be an approximate nearest neighbor (abbreviated ANN). We can find the successor and predecessor in each ordering in $O(\log \log N)$ time using Y-fast trie [86], it only remains to choose one of the 2τ candidates to be the ANN. To solve this problem we again deploy the LSO structure, and construct a 2-hop 1-spanner for the implicit path graph induced by each ordering. Specifically, each point will be associated with $O(\log N)$ edges (the name and weight of which will be added to the NNS label), where given two points $x \prec_\sigma y$, in $O(1)$ time we will be able to find a point z such that $x \preceq_\sigma z \preceq_\sigma y$ and x and y stored $\{x,z\},\{y,z\}$ respectively. Then $d_X(x,z)+d_X(z,y)$ will provide us the desired estimate of $d_X(x,y)$, which will be used to choose the ANN.

The case of rooted LSO is simpler- the label of each point z will consist of its position in all the orderings σ it belongs to, and the distance to the first point x_σ (w.r.t. d_X). Given a query q, for each ordering σ containing q, the leftmost point $y_\sigma \in P$ in the ordering will be a candidate ANN. We will estimate the distance from q to y_σ by $d_X(q,x_\sigma)+d_X(x_\sigma,y_\sigma)$, and return the point with minimum estimation.

For general metrics, the number of orderings is polynomial, $N^{\frac{1}{k}}$ which results in similar NNS label size, and query time (following the approach above). While the NNS label essentially cannot be improved, we can significantly reduce the query time. Our solution is to use Ramsey trees [19, 76, 17, 2], which are a collection of embeddings into ultrametrics $\mathcal{U}$ such that each point x has a single home ultrametric $U_x \in \mathcal{U}$ which well approximate all the distances to x. We thus reduce the labeled NNS problem to ultrametrics, where it can be efficiently solved. For the case of approximation factor $O(\log N)$ the required number of ultrametrics is $O(\log N)$, which leads us to label size $O(\log^2 N)$. To reduce it even farther, we use the novel clan embedding [48], where instead of embedding the space X into a collection of ultrametrics, we embed it into a single ultrametric (but where each point might have several copies). This allows us to reduce the label size to $O(\log N)$ (in expectation), and with one additional easing assumption (either polynomial aspect ratio or small failure probability) to even $O(1)$ label size.

Path reporting low hop spanners. A (τ,ρ)-tree cover is similar to ultrametric cover discussed above, where the ultrametrics are replaced by trees. Kahalon et al. [62] first constructed path reporting low hop spanner for a tree metric, and then for each metric space of interest, they considered it's tree cover, and constructed a path reporting low hop spanner for each tree in the cover. The spanner for the global metric is obtained by taking the union of all these spanners constructed for the trees in the cover. To report a queried distance, they simply computed the paths in all the trees, and returned the shortest observed path.

Thus Kahalon et al. idea is to reduce the problem to the fairly simple case of tree metrics. We reduce each metric space into the even simpler case of paths using LSO. Given an LSO (or triangle LSO) we simply construct a path reporting 2-hop path for each path associated with an ordering, and similarly to [62], check all the path spanners and return the shortest observed path. The resulting query time has linear dependence on the number of orderings. The case of rooted LSO is simpler, where it is enough to add a single edge per ordering, to the leftmost point in the ordering.

Next we present some improvements to the query. First, for the case of Euclidean space (low dimensional), we observe that given two points x,y, the ordering satisfying them could be computed in $O_d(1)$ time, implying that we don't need to check all the orderings, and return a 2 hop path in $O_d(1)$ time. Next, for the case of planar graphs, using the structure of cycle separators (which are used to construct the rooted LSO), in $O(1)$ time one can narrow the number of potential orderings to $O(\epsilon^{-1})$, implying $O(\epsilon^{-1})$ query time. For general graphs we observe that the celebrated Thorup Zwick distance oracle [85] can be used to produce a path reporting 2-hop $(2k-1)$-spanner with $O(n^{1+\frac{1}{k}}\cdot k)$ edges and $O(k)$ query time. Finally, we use sparse covers [15] to obtain an exponential improvement in the query time, while incurring a factor 2 increase in the stretch.

Fault tolerant spanners. The 2-hop f-fault tolerant spanner for doubling metrics by Kahalon et al. [62] is based on a quite sophisticated tool of robust tree cover. We have a superior, and an extremely simple construction. First we observe that the path graph has a 2-hop f-fault tolerant 1-spanner with $O(nf\log n)$ edges. Indeed, add edges from all the vertices to the middle $f+1$ vertices, delete the middle vertices and recurse on each side. We then apply this construction on each of the path graphs induced by the LSO (or triangle LSO) to obtain our results. The case of rooted LSO is even simpler: for every path it is enough to add all the edges to the first $f+1$ points.

References

1 Ittai Abraham, Yair Bartal, and Ofer Neiman. Advances in metric embedding theory. *Advances in Mathematics*, 228(6):3026–3126, 2011. doi:10.1016/j.aim.2011.08.003.

2 Ittai Abraham, Shiri Chechik, Michael Elkin, Arnold Filtser, and Ofer Neiman. Ramsey spanning trees and their applications. *ACM Trans. Algorithms*, 16(2):19:1–19:21, 2020. preliminary version published in SODA 2018. doi:10.1145/3371039.

3 Ittai Abraham, Shiri Chechik, Robert Krauthgamer, and Udi Wieder. Approximate nearest neighbor search in metrics of planar graphs. In Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015, August 24-26, 2015, Princeton, NJ, USA*, volume 40 of *LIPIcs*, pages 20–42. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.APPROX-RANDOM.2015.20.

4 Abu Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen G. Kobourov, and Richard Spence. Graph spanners: A tutorial review. *Comput. Sci. Rev.*, 37:100253, 2020. doi:10.1016/j.cosrev.2020.100253.

5 Ibrahim Akgün and Barbaros Ç. Tansel. New formulations of the hop-constrained minimum spanning tree problem via miller-tucker-zemlin constraints. *Eur. J. Oper. Res.*, 212(2):263–276, 2011. doi:10.1016/j.ejor.2011.01.051.

6 N. Alon and B. Schieber. Optimal preprocessing for answering on-line product queries. Technical report, Tel-Aviv University, 1987.

7 B. Alspach. The wonderful Walecki construction. *Bull. Inst. Combin. Appl*, 52:7–20, 2008. see here.

8 Ingo Althöfer, Gautam Das, David P. Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discret. Comput. Geom.*, 9:81–100, 1993. doi:10.1007/BF02189308.

9 Alexandr Andoni. *Nearest neighbor search: the old, the new, and the impossible*. PhD thesis, Massachusetts Institute of Technology, 2009. see here.

10 Alexandr Andoni and Piotr Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Commun. ACM*, 51(1):117–122, 2008. Preliminary version published in FOCS 2006. doi:10.1145/1327452.1327494.

11 Alexandr Andoni, Piotr Indyk, and Ilya P. Razenshteyn. Approximate nearest neighbor search in high dimensions. *CoRR*, abs/1806.09823, 2018. arXiv:1806.09823.

12 Alexandr Andoni, Huy L. Nguyen, Aleksandar Nikolov, Ilya P. Razenshteyn, and Erik Waingarten. Approximate near neighbors for general symmetric norms. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 902–913. ACM, 2017. doi:10.1145/3055399.3055418.

13 Alexandr Andoni, Aleksandar Nikolov, Ilya P. Razenshteyn, and Erik Waingarten. Approximate nearest neighbors beyond space partitions. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 1171–1190. SIAM, 2021. doi:10.1137/1.9781611976465.72.

14 Alexandr Andoni, Clifford Stein, and Peilin Zhong. Parallel approximate undirected shortest paths via low hop emulators. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani,

Gautam Kamath, and Julia Chuzhoy, editors, *Proccedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 322–335. ACM, 2020. doi:10.1145/3357713.3384321.

15 Baruch Awerbuch and David Peleg. Sparse partitions. In *Proceedings of the 31st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 503–513, 1990. doi:10.1109/FSCS.1990.89571.

16 Anantaram Balakrishnan and Kemal Altinkemer. Using a hop-constrained model to generate alternative communication network design. *INFORMS J. Comput.*, 4(2):192–205, 1992. doi:10.1287/ijoc.4.2.192.

17 Yair Bartal. Advances in metric ramsey theory and its applications. *CoRR*, abs/2104.03484, 2021. arXiv:2104.03484.

18 Yair Bartal, Arnold Filtser, and Ofer Neiman. On notions of distortion and an almost minimum spanning tree with constant average distortion. *J. Comput. Syst. Sci.*, 105:116–129, 2019. preliminary version published in SODA 2016. doi:10.1016/j.jcss.2019.04.006.

19 Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. Some low distortion metric ramsey problems. *Discret. Comput. Geom.*, 33(1):27–41, 2005. doi:10.1007/s00454-004-1100-z.

20 Greg Bodwin, Michael Dinitz, Merav Parter, and Virginia Vassilevska Williams. Optimal vertex fault tolerant spanners (for fixed stretch). In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1884–1900, 2018. doi:10.1137/1.9781611975031.123.

21 Greg Bodwin, Michael Dinitz, and Caleb Robelle. Optimal vertex fault-tolerant spanners in polynomial time. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 2924–2938. SIAM, 2021. doi:10.1137/1.9781611976465.174.

22 Greg Bodwin and Shyamal Patel. A trivial yet optimal solution to vertex fault tolerant spanners. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing, PODC 2019, Toronto, ON, Canada, July 29 - August 2, 2019*, pages 541–543, 2019. doi:10.1145/3293611.3331588.

23 Jérôme De Boeck and Bernard Fortz. Extended formulation for hop constrained distribution network configuration problems. *Eur. J. Oper. Res.*, 265(2):488–502, 2018. doi:10.1016/j.ejor.2017.08.017.

24 G. Borradaile, H. Le, and C. Wulff-Nilsen. Minor-free graphs have light spanners. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science*, FOCS '17, pages 767–778, 2017. doi:10.1109/FOCS.2017.76.

25 G. Borradaile, H. Le, and C. Wulff-Nilsen. Greedy spanners are optimal in doubling metrics. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '19, pages 2371–2379, 2019. doi:10.1137/1.9781611975482.145.

26 Prosenjit Bose, Vida Dujmovic, Pat Morin, and Michiel H. M. Smid. Robust geometric spanners. *SIAM J. Comput.*, 42(4):1720–1736, 2013. preliminary version published in SOCG 2013. doi:10.1137/120874473.

27 Kevin Buchin, Sariel Har-Peled, and Dániel Oláh. A spanner for the day after. In *35th International Symposium on Computational Geometry, SoCG 2019, June 18-21, 2019, Portland, Oregon, USA*, pages 19:1–19:15, 2019. doi:10.4230/LIPIcs.SoCG.2019.19.

28 Kevin Buchin, Sariel Har-Peled, and Dániel Oláh. Sometimes reliable spanners of almost linear size. In *28th Annual European Symposium on Algorithms, ESA 2020, September 7-9, 2020, Pisa, Italy (Virtual Conference)*, pages 27:1–27:15, 2020. doi:10.4230/LIPIcs.ESA.2020.27.

29 T.-H. Hubert Chan and Anupam Gupta. Approximating TSP on metrics with bounded global growth. *SIAM J. Comput.*, 41(3):587–617, 2012. preliminary version published in SODA 2008. doi:10.1137/090749396.

30 Timothy M. Chan, Sariel Har-Peled, and Mitchell Jones. On locality-sensitive orderings and their applications. *SIAM J. Comput.*, 49(3):583–600, 2020. preliminary version published in ITCS 2019. doi:10.1137/19M1246493.

31 Barun Chandra, Gautam Das, Giri Narasimhan, and José Soares. New sparseness results on graph spanners. *Int. J. Comput. Geom. Appl.*, 5:125–144, 1995. preliminary version published in SOCG 1992. doi:10.1142/S0218195995000088.

32 Moses Charikar, Chandra Chekuri, Ashish Goel, and Sudipto Guha. Rounding via trees: deterministic approximation algorithms for group steiner trees and k-median. In *STOC '98: Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 114–123, New York, NY, USA, 1998. ACM Press. doi:10.1145/276698.276719.

33 Moses Charikar, Chandra Chekuri, Ashish Goel, Sudipto Guha, and Serge A. Plotkin. Approximating a finite metric by a small number of tree metrics. In *39th Annual Symposium on Foundations of Computer Science, FOCS '98, November 8-11, 1998, Palo Alto, California, USA*, pages 379–388. IEEE Computer Society, 1998. doi:10.1109/SFCS.1998.743488.

34 Shiri Chechik, Michael Langberg, David Peleg, and Liam Roditty. Fault tolerant spanners for general graphs. *SIAM J. Comput.*, 39(7):3403–3423, 2010. preliminary version published in STOC 2009. doi:10.1137/090758039.

35 Shiri Chechik and Christian Wulff-Nilsen. Near-optimal light spanners. *ACM Trans. Algorithms*, 14(3):33:1–33:15, 2018. preliminary version published in SODA 2016. doi:10.1145/3199607.

36 Edith Cohen. Polylog-time and near-linear work approximation scheme for undirected shortest paths. *J. ACM*, 47(1):132–166, 2000. preliminary version published in STOC 1994. doi:10.1145/331605.331610.

37 Vincent Cohen-Addad, Arnold Filtser, Philip N. Klein, and Hung Le. On light spanners, low-treewidth embeddings and efficient traversing in minor-free graphs. *CoRR*, abs/2009.05039, 2020. To appear in FOCS 2020,https://arxiv.org/abs/2009.05039. arXiv:2009.05039.

38 Gautam Das, Paul J. Heffernan, and Giri Narasimhan. Optimally sparse spanners in 3-dimensional euclidean space. In Chee Yap, editor, *Proceedings of the Ninth Annual Symposium on Computational GeometrySan Diego, CA, USA, May 19-21, 1993*, pages 53–62. ACM, 1993. doi:10.1145/160985.160998.

39 Michael Dinitz and Robert Krauthgamer. Fault-tolerant spanners: better and simpler. In *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011, San Jose, CA, USA, June 6-8, 2011*, pages 169–178, 2011. doi:10.1145/1993806.1993830.

40 Michael Dinitz and Caleb Robelle. Efficient and simple algorithms for fault-tolerant spanners. In *PODC '20: ACM Symposium on Principles of Distributed Computing, Virtual Event, Italy, August 3-7, 2020*, pages 493–500, 2020. doi:10.1145/3382734.3405735.

41 Anne Driemel and Francesco Silvestri. Locality-Sensitive Hashing of Curves. In *Proceedings of the 33rd International Symposium on Computational Geometry*, volume 77, pages 37:1–37:16, Brisbane, Australia, July 2017. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2017.37.

42 Michael Elkin, Ofer Neiman, and Shay Solomon. Light spanners. *SIAM J. Discret. Math.*, 29(3):1312–1321, 2015. doi:10.1137/140979538.

43 Ioannis Z. Emiris and Ioannis Psarros. Products of euclidean metrics and applications to proximity questions among curves. In Bettina Speckmann and Csaba D. Tóth, editors, *34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary*, volume 99 of *LIPIcs*, pages 37:1–37:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.SoCG.2018.37.

44 P. Erdős. Extremal problems in graph theory. *Theory of Graphs and Its Applications (Proc. Sympos. Smolenice)*, pages 29–36, 1964. see here.

45 Arnold Filtser. On strong diameter padded decompositions. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2019, September 20-22, 2019, Massachusetts Institute of Technology, Cambridge, MA, USA*, pages 6:1–6:21, 2019. doi:10.4230/LIPIcs.APPROX-RANDOM.2019.6.

46 Arnold Filtser. Labeled nearest neighbor search and metric spanners via locality sensitive orderings. *CoRR*, abs/2211.11846, 2022. doi:10.48550/arXiv.2211.11846.

47 Arnold Filtser, Omrit Filtser, and Matthew J. Katz. Approximate nearest neighbor for curves - simple, efficient, and deterministic. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, pages 48:1–48:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.ICALP.2020.48.

48 Arnold Filtser and Hung Le. Clan embeddings into trees, and low treewidth graphs. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 342–355. ACM, 2021. doi:10.1145/3406325.3451043.

49 Arnold Filtser and Hung Le. Locality-sensitive orderings and applications to reliable spanners. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 1066–1079. ACM, 2022. doi:10.1145/3519935.3520042.

50 Arnold Filtser and Ofer Neiman. Light spanners for high dimensional norms via stochastic decompositions. *Algorithmica*, 2022. doi:10.1007/s00453-022-00994-0.

51 Arnold Filtser and Shay Solomon. The greedy spanner is existentially optimal. *SIAM J. Comput.*, 49(2):429–447, 2020. preliminary version published in PODC 2016. doi:10.1137/18M1210678.

52 J. Gao, L. J. Guibas, and A. Nguyen. Deformable spanners and applications. *Computational Geometry*, 35(1):2–19, 2006. doi:10.1016/j.comgeo.2005.10.001.

53 Lee-Ad Gottlieb. A light metric spanner. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 759–772, 2015. doi:10.1109/FOCS.2015.52.

54 Luis Eduardo Neves Gouveia and Thomas L. Magnanti. Network flow models for designing diameter-constrained minimum-spanning and steiner trees. *Networks*, 41(3):159–173, 2003. doi:10.1002/net.10069.

55 Luis Eduardo Neves Gouveia, Pedro Patrício, Amaro de Sousa, and Rui Valadas. MPLS over WDM network design with packet level qos constraints based on ILP models. In *Proceedings IEEE INFOCOM 2003, The 22nd Annual Joint Conference of the IEEE Computer and Communications Societies, San Francisco, CA, USA, March 30 - April 3, 2003*, pages 576–586, 2003. doi:10.1109/INFCOM.2003.1208708.

56 Sariel Har-Peled, Piotr Indyk, and Anastasios Sidiropoulos. Euclidean spanners in high dimensions. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 804–809, 2013. doi:10.1137/1.9781611973105.57.

57 Sariel Har-Peled, Manor Mendel, and Dániel Oláh. Reliable Spanners for Metric Spaces. In *37th International Symposium on Computational Geometry*, SoCG'21, pages 43:1–43:13, 2021. Full version at https://arxiv.org/abs/2007.08738. doi:10.4230/LIPIcs.SoCG.2021.43.

58 Piotr Indyk. On approximate nearest neighbors under ℓ_∞ norm. *J. Comput. Syst. Sci.*, 63(4):627–638, 2001. doi:10.1006/jcss.2001.1781.

59 Piotr Indyk. Approximate nearest neighbor algorithms for Fréchet distance via product metrics. In *Proceedings of the 8th Symposium on Computational Geometry*, pages 102–106, Barcelona, Spain, June 2002. ACM Press. doi:10.1145/513400.513414.

60 Piotr Indyk and Nitin Thaper. Fast image retrieval via embeddings. In *3rd international workshop on statistical and computational theories of vision*, volume 2(3), page 5. Nice, France, 2003. see here.

61 William Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary Mathematics*, 26:189–206, 1984. see here.

62 Omri Kahalon, Hung Le, Lazar Milenkovic, and Shay Solomon. Can't see the forest for the trees: Navigating metric spaces by bounded hop-diameter spanners. In Alessia Milani and

Philipp Woelfel, editors, *PODC '22: ACM Symposium on Principles of Distributed Computing, Salerno, Italy, July 25 - 29, 2022*, pages 151–162. ACM, 2022. doi:10.1145/3519270.3538414.

63 P. N. Klein. Subset spanner for planar graphs, with application to subset TSP. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, STOC '06, pages 749–756, 2006. doi:10.1145/1132516.1132620.

64 Philip N. Klein. Preprocessing an undirected planar network to enable fast approximate distance queries. In *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 6-8, 2002, San Francisco, CA, USA*, pages 820–827, 2002. see here. URL: http://dl.acm.org/citation.cfm?id=545381.545488.

65 Philip N. Klein. A linear-time approximation scheme for TSP in undirected planar graphs with edge-weights. *SIAM J. Comput.*, 37(6):1926–1952, 2008. doi:10.1137/060649562.

66 Robert Krauthgamer and James R. Lee. The black-box complexity of nearest-neighbor search. *Theor. Comput. Sci.*, 348(2-3):262–276, 2005. doi:10.1016/j.tcs.2005.09.017.

67 Hung Le. A PTAS for subset TSP in minor-free graphs. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2279–2298, 2020. doi:10.1137/1.9781611975994.140.

68 Hung Le, Lazar Milenkovic, and Shay Solomon. Sparse euclidean spanners with tiny diameter: A tight lower bound. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry, SoCG 2022, June 7-10, 2022, Berlin, Germany*, volume 224 of *LIPIcs*, pages 54:1–54:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.SoCG.2022.54.

69 Hung Le and Shay Solomon. A unified framework of light spanners II: fine-grained optimality. *CoRR*, abs/2111.13748, 2021. arXiv:2111.13748.

70 Hung Le and Christian Wulff-Nilsen. Optimal approximate distance oracle for planar graphs. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 363–374. IEEE, 2021. doi:10.1109/FOCS52979.2021.00044.

71 Larry J. LeBlanc, Jerome Chifflet, and Philippe Mahey. Packet routing in telecommunication networks with path and flow restrictions. *INFORMS J. Comput.*, 11(2):188–197, 1999. doi:10.1287/ijoc.11.2.188.

72 Christos Levcopoulos, Giri Narasimhan, and Michiel H. M. Smid. Improved algorithms for constructing fault-tolerant spanners. *Algorithmica*, 32(1):144–156, 2002. preliminary version published in STOC 1998. doi:10.1007/s00453-001-0075-x.

73 Yaowei Long and Seth Pettie. Planar distance oracles with better time-space tradeoffs. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021*, pages 2517–2537. SIAM, 2021. doi:10.1137/1.9781611976465.149.

74 Tamás Lukovszki. New results on fault tolerant geometric spanners. In Frank Dehne, Jörg-Rüdiger Sack, Arvind Gupta, and Roberto Tamassia, editors, *Algorithms and Data Structures*, pages 193–204, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg. doi:10.1007/3-540-48447-7_20.

75 Jiří Matoušek. On the distortion required for embedding finite metric spaces into normed spaces. *Israel Journal of Mathematics*, 93(1):333–344, 1996. doi:10.1007/BF02761110.

76 Manor Mendel and Assaf Naor. Ramsey partitions and proximity data structures. *Journal of the European Mathematical Society*, 9(2):253–275, 2007.

77 Giri Narasimhan and Michiel H. M. Smid. *Geometric spanner networks*. Cambridge University Press, 2007. doi:10.1017/CBO9780511546884.

78 Huy L. Nguyen. Approximate nearest neighbor search in ℓ_p. *CoRR*, abs/1306.3601, 2013. arXiv:1306.3601.

79 Rafail Ostrovsky and Yuval Rabani. Low distortion embeddings for edit distance. *J. ACM*, 54(5):23, 2007. doi:10.1145/1284320.1284322.

80 Merav Parter. Nearly optimal vertex fault-tolerant spanners in optimal time: sequential, distributed, and parallel. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th*

Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022, pages 1080–1092. ACM, 2022. doi:10.1145/3519935.3520047.

81 Hasan Pirkul and Samit Soni. New formulations and solution procedures for the hop constrained network design problem. *Eur. J. Oper. Res.*, 148(1):126–140, 2003. doi:10.1016/S0377-2217(02)00366-1.

82 André Rossi, Alexis Aubry, and Mireille Jacomino. Connectivity-and-hop-constrained design of electricity distribution networks. *Eur. J. Oper. Res.*, 218(1):48–57, 2012. doi:10.1016/j.ejor.2011.10.006.

83 Shay Solomon. From hierarchical partitions to hierarchical covers: optimal fault-tolerant spanners for doubling metrics. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 363–372, 2014. doi:10.1145/2591796.2591864.

84 M. Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *Journal of the ACM*, 51(6):993–1024, 2004. doi:10.1145/1039488.1039493.

85 Mikkel Thorup and Uri Zwick. Approximate distance oracles. *J. ACM*, 52(1):1–24, 2005. preliminary version published in STOC 2001. doi:10.1145/1044731.1044732.

86 Dan E. Willard. Log-logarithmic worst-case range queries are possible in space theta(n). *Inf. Process. Lett.*, 17(2):81–84, 1983. doi:10.1016/0020-0190(83)90075-3.

87 Kathleen A. Woolston and Susan L. Albin. The design of centralized networks with reliability and availability constraints. *Comput. Oper. Res.*, 15(3):207–217, 1988. doi:10.1016/0305-0548(88)90033-0.