Light, Reliable Spanners

Arnold Filtser
Bar-Ilan University, Ramat-Gan, Israel

Yuval Gitlitz
Ben-Gurion University of the Negev, Be'er Sheva, Israel

Ofer Neiman
Ben-Gurion University of the Negev, Be'er Sheva, Israel

Abstract

A *ν-reliable spanner* of a metric space (X, d), is a (dominating) graph H, such that for any possible failure set $B \subseteq X$, there is a set B^+ just slightly larger $|B^+| \leq (1+\nu) \cdot |B|$, and all distances between pairs in $X \setminus B^+$ are (approximately) preserved in $H \setminus B$. Recently, there have been several works on sparse reliable spanners in various settings, but so far, the weight of such spanners has not been analyzed at all. In this work, we initiate the study of *light* reliable spanners, whose weight is proportional to that of the Minimum Spanning Tree (MST) of X.

We first observe that unlike sparsity, the lightness of any deterministic reliable spanner is huge, even for the metric of the simple path graph. Therefore, randomness must be used: an *oblivious* reliable spanner is a distribution over spanners, and the bound on $|B^+|$ holds in expectation.

We devise an oblivious ν-reliable $(2 + \frac{2}{k-1})$-spanner for any k-HST, whose lightness is $\approx \nu^{-2}$. We demonstrate a matching $\Omega(\nu^{-2})$ lower bound on the lightness (for any finite stretch). We also note that any stretch below 2 must incur linear lightness.

For general metrics, doubling metrics, and metrics arising from minor-free graphs, we construct *light* tree covers, in which every tree is a k-HST of low weight. Combining these covers with our results for k-HSTs, we obtain oblivious reliable light spanners for these metric spaces, with nearly optimal parameters. In particular, for doubling metrics we get an oblivious ν-reliable $(1+\varepsilon)$-spanner with lightness $\varepsilon^{-O(\mathrm{ddim})} \cdot \tilde{O}(\nu^{-2} \cdot \log n)$, which is best possible (up to lower order terms).

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

Keywords and phrases light spanner, reliable spanner, HST cover, doubling metric, minor free graphs

Digital Object Identifier 10.4230/LIPIcs.SoCG.2024.56

Related Version *Full Version*: https://arxiv.org/abs/2307.16612 [20]

Funding *Arnold Filtser*: Israel Science Foundation (grant No. 1042/22).
Yuval Gitlitz: Partially supported by the Lynn and William Frankel Center for Computer Sciences and Israel Science Foundation (grant No. 970/21).
Ofer Neiman: Israel Science Foundation (grant No. 970/21).

1 Introduction

Given a metric space (X, d_X), a *t-spanner* is a graph H over X such that for every $x, y \in X$, $d_X(x,y) \leq d_H(x,y) \leq t \cdot d_X(x,y)$, where d_H is the shortest path metric in H. [1] The parameter t is often referred to as the *stretch*. In essence, the purpose of spanners is to represent the distance metric using a sparse graph. Spanners where introduced by Peleg and Schäffer [34], and found numerous applications throughout computer science. For a

[1] Often in the literature, the input metric is the shortest path metric of a graph, and a spanner is required to be a subgraph of the input graph. Here we study metric spanners where there is no such requirement.

more systematical study, we refer to the book [33] and survey [2]. In many cases, the goal is to minimize the total weight of the spanner and not just the number of edges. E.g., when constructing a road network, the cost is better measured by the total length of paved roads, as opposed to their number. This parameter of interest is formalized as the *lightness* of a spanner, which is the ratio between the weight of the spanner (sum of all edge weights), and the weight of the Minimum Spanning Tree (MST) of X: $\frac{w(H)}{w(\mathrm{MST})}$. Note that the MST is the minimal weight of a connected graph, and thus of a spanner with finite stretch. So the lightness is simply a "normalized" notion of weight.

Light spanners have been thoroughly studied. It is known that general n-point metric spaces admit a $(2k-1)(1+\varepsilon)$ spanner (for $k \in \mathbb{N}$, $\varepsilon \in (0,1)$) with $O(n^{1+1/k})$ edges and lightness $O(\varepsilon^{-1} \cdot n^{1/k})$ [30, 6] (see also [4, 18, 14, 23]). Every n-point metric space with doubling dimension[2] ddim admits a $(1+\varepsilon)$-spanner with $n \cdot \varepsilon^{-O(\mathrm{ddim})}$ edges and lightness $\varepsilon^{-O(\mathrm{ddim})}$ [8] (see also [24, 23]). Finally, the shortest path metric of a graph excluding a fixed minor admits a (sub-graph, which already implies sparsity) $(1+\varepsilon)$-spanner with lightness $\tilde{O}(\varepsilon^{-3})$ [7].

A highly desirable properly of a spanner is the ability to withstand massive node-failures. To this end, Bose *et. al.* [9] introduced the notion of a *reliable spanner.* [3] Here, given a set of failed nodes $B \subseteq X$, 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. For the case of points in d-dimensional Euclidean space, for constant d, Bose *et. al.* [9] constructed $O(1)$ spanner such that $|B^+| \le O(|B|^2)$. Later, Buchin, Har-Peled, and Oláh [10] constructed $1+\varepsilon$ reliable spanner with $n \cdot \varepsilon^{-O(d)} \cdot \nu^{-6} \cdot \tilde{O}(\log n)$ edges, guaranteeing that for every set of failed nodes B, $|B^+| \le (1+\nu) \cdot |B|$. This result was generalized to metric spaces with doubling dimension ddim by Filtser and Le [21].

While reliable spanners for Euclidean and doubling metrics admit sparsity which is comparable to their non-reliable counter-parts, the situation is very different for other metric families. Indeed, Har-Peled *et. al.* [25] showed that every reliable k-spanner of the simple uniform metric (which is also a tree metric) must have $\Omega(n^{1+1/k})$ edges. Nevertheless, it is possible to construct *oblivious* reliable spanner for other metric spaces with good parameters, where the bound on the size of B^+ is only in expectation.

▶ **Definition 1** (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 dominates*[4] *d_X, and for every set $B \subseteq X$ of points, called an* attack set*, there is a set $B^+ \supseteq B$, called a* faulty extension *of B, such that: (1) $|B^+| \le (1+\nu)|B|$. (2) For every $x, y \notin B^+$, $d_{H[X\setminus B]}(x,y) \le t \cdot d_X(x,y)$.*

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 $H \in \mathrm{supp}(\mathcal{D})$, there exist a superset $B^+_H \supseteq B$ such that, for every $x, y \notin B^+_H$, $d_{H[X\setminus B]}(x,y) \le t \cdot d_X(x,y)$, and $\mathbb{E}_{H\sim\mathcal{D}}\left[|B^+_H|\right] \le (1+\nu)|B|$. We say that the oblivious spanner $\mathcal{D}$ has m edges and lightness ϕ if every $H \in \mathrm{supp}(\mathcal{D})$ has at most m edges and lightness at most ϕ.

For general n-point metrics, Filtser and Le [21] (improving over [25]) constructed an oblivious ν-reliable $8k+\varepsilon$-spanner with $\tilde{O}(n^{1+\frac{1}{k}} \cdot \varepsilon^{-2}) \cdot \nu^{-1}$ edges. For the shortest path metric of graph excluding a fixed minor, there is oblivious ν-reliable $(2+\varepsilon)$-spanner with $\varepsilon^{-2} \cdot \nu^{-1} \cdot \tilde{O}(n)$ edges, while every oblivious reliable spanner with stretch $t < 2$ requires $\Omega(n^2)$ edges [21]. For Euclidean and doubling metrics, oblivious ν-reliable $(1+\varepsilon)$-spanners can be constructed with only $n \cdot \varepsilon^{-O(d)} \cdot \tilde{O}(\nu^{-1} \cdot \log^2 \log n)$ edges [11, 21].

---

[2] A metric space (X, d) has doubling dimension ddim if every ball of radius $2r$ can be covered by 2^{ddim} balls of radius r. The d-dimensional Euclidean space has doubling dimension $\Theta(d)$.

[3] For a comprehensive discussion with the related notion of fault-tolerant spanners, see Section 1.3.

[4] Metric space (X, d_H) dominates metric space (X, d_X) if $\forall u, v \in X$, $d_X(u,v) \le d_H(u,v)$.

Table 1 Our results for constructing light ν-reliable spanners for various metric spaces. All the results in the table (other than the one specified as deterministic) are for oblivious reliable spanners. Stretch $< \infty$ stands for the requirement that all the points in $X \setminus B^+$ belong to the same connected component in $H \setminus B$. $*$ stands for $\mathrm{poly}(\log\log n)$ factors.

Family	Stretch	Lightness	Size	Ref
Doubling ddim	$1+\varepsilon$	$\varepsilon^{-O(\mathrm{ddim})} \cdot \tilde{O}(\nu^{-2} \cdot \log n)$	$n \cdot \varepsilon^{-O(\mathrm{ddim})} \cdot \tilde{O}(\nu^{-2}) \cdot *$	FullV[20]
	ddim	$\tilde{O}(\log n \cdot \nu^{-2}) \cdot \mathrm{ddim}^{O(1)}$	$n \cdot \tilde{O}\left(\nu^{-2}\right) \cdot \mathrm{ddim}^{O(1)} \cdot *$	FullV[20]
General Metric	$12t+\varepsilon$	$n^{1/t} \cdot \tilde{O}(\nu^{-2} \cdot \varepsilon^{-4}) \cdot \log^{O(1)} n$	$\tilde{O}\left(n^{1+1/t} \cdot \nu^{-2} \cdot \varepsilon^{-3}\right)$	FullV[20]
	$O(\log n)$	$\tilde{O}(\nu^{-2} \cdot \log^4 n)$	$n \cdot \tilde{O}\left(\nu^{-2} \cdot \log^3 n\right)$	FullV[20]
Euclidean $\mathbb{R}^d$	$\sqrt{d}$	$\tilde{O}\left(\log n \cdot \nu^{-2} \cdot d^{7.5}\right)$	$n \cdot \tilde{O}\left(\nu^{-2} \cdot d^6\right) \cdot *$	FullV[20]
Minor-Free	$2+\varepsilon$	$\tilde{O}(\nu^{-2} \cdot \varepsilon^{-7} \cdot \log^8 n)$	$\tilde{O}(n \cdot \nu^{-2} \cdot \varepsilon^{-6})$	FullV[20]
Tree	< 2	$\Omega(n)$	$\Omega(n^2)$	[21]
Weighted Path	1	$\nu^{-2} \cdot \tilde{O}(\log n)$	$n \cdot \tilde{O}(\nu^{-1}) \cdot *$	FullV[20]
Unweighted Path	$< \infty$	$\Omega(\nu^{-2} \cdot \log(\nu \cdot n))$	-	FullV[20]
	$< \infty$	$\Omega(n)$ (deterministic)	-	FullV[20]
HST (ultrametric)	$2+\varepsilon$	$\tilde{O}(\varepsilon^{-4} \cdot \nu^{-2}) \cdot *$	$n \cdot \tilde{O}\left(\varepsilon^{-3} \cdot \nu^{-2}\right) \cdot *$	FullV[20]
	$< \infty$	$\Omega(\nu^{-2})$	-	FullV[20]

But what about lightness? no previous work attempted to construct reliable spanners of low total weight even though it is clearly desirable to construct reliable networks of low total cost. The single most studied metric in the context of reliable spanners is the unweighted path P_n. Indeed, most of the previous work [10, 11, 21, 19] focused on constructing various reliable 1-spanners for the path graph, and then generalized it other metric spaces using *locality sensitive orderings* [5]. A reliable spanner should have many edges between every two large enough sets, so that they could not be easily disconnected. Consider an attack B consisting of the middle $\frac{n}{2}$ vertices on P_n. If there are less than $\frac{n}{8}$ crossing edges from left to right, then an attack $B' \supseteq B$ that contains also one endpoint per crossing edge, will disconnect two sets of size $\frac{n}{8}$. Therefore a linear number of vertices should be added to B'^+. We conclude that every deterministic reliable spanner (for any finite stretch) must have lightness $\Omega(n)$ (see full version [20] for a formal proof). Thus, all hope lies in oblivious reliable spanners. However, even here any two large sets must be well connected. Previous oblivious reliable spanners for P_n all had unacceptable polynomial lightness.

As reliable spanners for P_n are the main building blocks for reliable spanners for other metric spaces, all previous constructions have inherent polynomial lightness.[6]

1.1 Our Results

The results of this paper are summarized in Table 1. Our results on light reliable spanners for various metric families are based on constructing such spanners for k-HSTs, this lies in contrast to previous results on sparse reliable spanners, which were mostly based on reliable spanners for the path graph.

[5] Locality sensitive ordering is a generic tool that "reduces" metric spaces into the line, by devising a collection of orderings such that every two points are "nearby" in one of the orderings, see [13, 21].

[6] The only previous work that did not reduced to P_n is by Har-Peled *et. al.* [25] who reduced to uniform metrics. Nevertheless, their approach on P_n will have stretch 3, and lightness $\Omega(n)$.

Roughly speaking, previous works on reliable spanners show us that the "cost" of making a spanner ν-reliable, is often a ν^{-1} factor in its size. Our results in this paper offer a similar view for light spanners: here the "cost" of reliability is a factor of ν^{-2} in the lightness. That is, an $\Omega(\nu^{-2})$ factor must be paid in the most basic cases (path graph, HST), while in more interesting and complicated metric families, we essentially match the best non-reliable light spanner constructions, up to this ν^{-2} factor (and in some cases, such as minor-free graphs, an unavoidable constant increase in the stretch). For brevity, in the discussion that follows we omit the bounds on the size of our spanners (which can be found in Table 1).

k-HSTs. We devise an oblivious ν-reliable $2+\frac{O(1)}{k}$-spanner for any k-HST , whose lightness is $\tilde{O}(\nu^{-1}\cdot\log\log n)^2$ (see Theorem 2). It is implicitly shown in [21, Observation 1] that with stretch smaller than 2, the lightness must be $\Omega(n)$. So when k is large, our stretch bound is nearly optimal.[7] We also show that the lightness must be at least $\Omega(\nu^{-2})$, regardless of the stretch, thus nearly matching our upper bound.

Light k-HST Covers. To obtain additional results for other metric families, following [21], we use the notion of *tree covers*, in which every tree is a k-HST . We design these covers for metrics admitting a pairwise partition cover scheme (see definition in the full version [20]), such that each k-HST in the cover has lightness $O(k\cdot\log n)$.

General Metrics. For any metric space, by building a light k-HST cover, and applying our oblivious reliable spanner for every k-HST in the cover, we obtain an oblivious ν-reliable $O(k)$-spanner with lightness $\tilde{O}(\nu^{-2}\cdot n^{1/k})$. Note that up to a constant in the stretch (and lower order terms), this result is optimal, even omitting the reliability requirement.

Doubling Metrics. For any metric with doubling dimension ddim,[2] and $\varepsilon\in(0,1)$, we devise an oblivious ν-reliable $(1+\varepsilon)$-spanner with lightness $\varepsilon^{-O(\mathrm{ddim})}\cdot\tilde{O}\left(\nu^{-2}\cdot\log n\right)$. This result is tight up to second order terms. Indeed, every $(1+\varepsilon)$-spanner for doubling metrics must have lightness $\varepsilon^{-\Omega(\mathrm{ddim})}$ (see e.g., [8]). In the full version [20], we show that every oblivious ν-reliable spanner (for any finite stretch) for the shortest path metric of the unweighted path graph (which has ddim 1) must have lightness $\Omega(\nu^{-2}\cdot\log(\nu n))$. This dependence on n in the lower bound is somewhat surprising, and does not appear in the closely related fault-tolerant spanners for doubling metrics (see Section 1.3 for further details).

In our doubling reliable spanner construction, we adapt the framework used for general metrics. Note that general k-HSTs must suffer stretch at least 2. Fortunately, the k-HSTs in the cover for doubling metrics have bounded maximum degree. For such HSTs we construct oblivious reliable $1+O(\frac{1}{k})$-spanner with lightness $\tilde{O}(\nu^{-1}\cdot\log\log n)^2$. Whenever $k\geq\frac{1}{\varepsilon}$, this is $1+\varepsilon$ stretch.

High Dimensional Euclidean and Doubling Metrics. Given n points in high dimensional Euclidean (or doubling) space, our previous construction has exponential dependence on the dimension, which might be too large. Following our approach for general metrics, we construct a t-spanner with lightness $2^{O(\frac{\mathrm{ddim}}{t})}\cdot\mathrm{ddim}^{O(1)}\cdot\tilde{O}(\nu^{-2}\cdot\log n)$, which can be further improved to $2^{O(\frac{d}{t^2})}\cdot d^{O(1)}\cdot\tilde{O}\left(\nu^{-2}\cdot\log n\right)$ for the case of d-dimensional Euclidean space.

[7] We also have a similar result for every $k\geq 1$, with stretch $2+\varepsilon$ and lightness $\tilde{O}(\varepsilon^2\cdot\nu^{-1}\cdot\log\log n)^2$.

Metrics of Minor-free Graphs. Consider a metric (X, d) arising from shortest paths of a graph G that excludes a fixed minor. In the full version [20]we show that X admits "good" pairwise partition cover with stretch 2, and thus by using the framework mentioned above as a black-box, we can get oblivious ν-reliable $(4 + \varepsilon)$-spanner. However, the lower bound on the stretch is the same as for k-HST, which is only 2 (whenever the lightness is sub-linear). To obtain near optimal results, we exploit a certain property of our pairwise partition cover for these metrics, and achieve (in a non-black-box manner) the nearly optimal oblivious ν-reliable $(2 + \varepsilon)$-spanner with lightness $\nu^{-2} \cdot \text{poly}(\log n, 1/\varepsilon)$.

The path graph. We conclude our journey on light reliable spanners by constructing an oblivious ν-reliable 1-spanner for the weighted path graph P_n, whose lightness is $\tilde{O}(\nu^{-2} \cdot \log n)$. As mentioned above, we prove that this bound on the lightness is optimal (up to lower order terms), for any finite stretch. A useful[8] property of our spanner is that it is hop-bounded, that is, every pair outside B^+ admits a shortest path with at most $\log n$ edges.[9]

1.2 Technical Overview

From a high level, our construction of light reliable spanners for various graph families has the following structure.

- We first devise light reliable spanners for k-HSTs.
- We construct *light* tree covers for the relevant family, where all the trees in the cover are k-HSTs.
- The final step is to sample a reliable spanner for each tree in the cover, and take as a final spanner the union of these spanners.

In what follows we elaborate more on the main ideas and techniques for each of those steps.

1.2.1 Reliable Light Spanner for k-HSTs

Let T be the tree representing the k-HST . Our construction consists of a collection of randomly chosen bi-cliques: For every node $x \in T$ we choose at random a set Z_x of $\ell \approx \nu^{-1}$ vertices from the leaves of the subtree rooted at x (denoted $L(x)$). Then, for every $x \in T$ with children $x_1, \dots, x_t$, add to the spanner H all edges in $Z_x \times Z_{x_j}$ for every $j = 1, \dots, t$.

Fix a pair of leaves $u, v \in T$, let $x = \text{lca}(u, v)$, and let x_i (resp., x_j) be the child of x whose subtree contains u (resp., v). The idea behind finding a spanner path between u, v is as follows. We will connect both u, v to a certain chosen leaf $x' \in Z_x$. To this end, we first connect recursively u to a $u' \in Z_{x_i}$ and v to $v' \in Z_{x_j}$. Now, if x, x_i, and x_j have all chosen such leaves x', u', v' to the sets Z_x, Z_{x_i}, Z_{x_j} respectively, that survive the attack B, and also we managed the $u - u'$ and $v - v'$ connections recursively, then we can complete the $u - v$ path. That path will consists of the two "long" bi-clique edges $\{u', x'\}, \{x', v'\}$, and the recursive $u - u'$ and $v - v'$ paths. Note that since $u, u' \in L(x_i)$, $d_T(u, u') \le d_T(u, v)/k$ (and similarly $d_T(v, v') \le d_T(u, v)/k$), so we can show inductively that the total distance taken by these recursive paths is only $O(d_T(u, v)/k)$. See Figure 1 for an illustration of a path in H between two vertices u, v.

[8] Buchin *et. al.* 's [11] oblivious reliable spanner for the path is $O(\log n)$ hop-bounded. This property was crucial for the construction of sparse oblivious reliable spanners for Euclidean and doubling metrics [11, 21]. Filtser and Le [21] constructed 2 hop-bounded oblivious reliable spanners for the path (and used it in their construction of oblivious reliable spanner for general metrics).

[9] Additionally, for any $h \ge 1$, we can also devise a h-hop-bounded reliable spanner, while achieving lightness $\approx \nu^{-2} \cdot h \cdot n^{1/h}$.

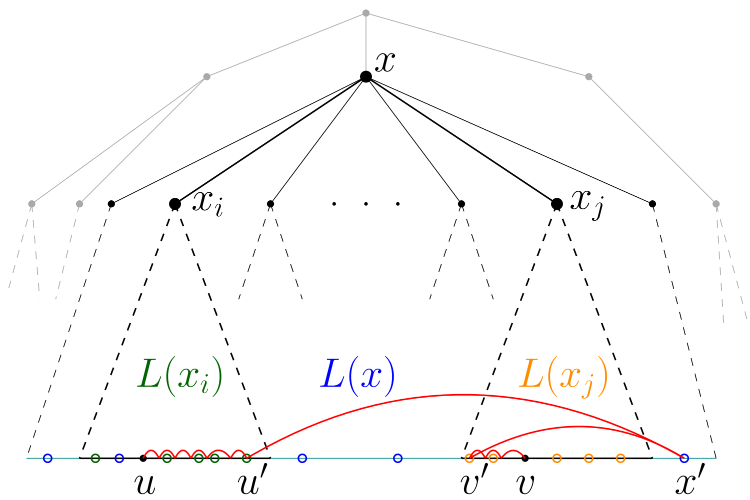

Figure 1 Illustration of the construction of spanner for a k-HST. For each internal node x we sample a subset Z_x of leaves from $L(x)$, and connect all of Z_x to $Z_{x'}$ for every child x' of x. The path from u to v will first go from u to a surviving vertex in Z_{x_i} (using recursion), from there to a surviving vertices in Z_x and Z_{x_j}, and finally to v (again by recursion).

Having established what is needed for finding a spanner path, we say that a leaf is *safe* if all its ancestors x in T have that Z_x is not fully included in B. The failure set B^+ consists of B and all leaves that are not safe.

A subtle issue is that a vertex may have a linear number of ancestors, and we will need ℓ to be at least logarithmic to ensure good probability for success in all of them. To avoid this, we use the following approach. For any node x that has a "heavy"child y (that is, $L(y)$ is almost as large as $L(x)$), we use the sample Z_y for x, instead of sampling Z_x. This way, any leaf will have only logarithmically many ancestors that are not heavy parents, which reduce dramatically the sample size needed for success in all ancestors.

For the reliability analysis, we distinguish between leaves that have an ancestor x with a very large $1 - \nu$ fraction of vertices in $L(x)$ that fall in the attack set B. These leaves are immediately taken as failed, but there can be only $\approx \nu|B|$ such leaves. For the other leaves, a delicate technical analysis follows to show that only a small fraction $\approx \nu \cdot |B|$ new vertices are expected to join B^+. Note that if some node has a heavy child, we take the child's sample, so some care is needed in the analysis to account for this – roughly speaking, the definition of "heavy" must depend on the reliability parameter ν, in order to ensure sufficiently small failure probability.

Improved stretch for bounded degree HSTs. In the case the k-HST has bounded degree δ, we can alter the construction slightly, and for every x with children $x_1, \ldots, x_s$, also add all edges in $Z_{x_i} \times Z_{x_j}$ for every $1 \le i < j \le s$. While this alternative increases the lightness and size by a factor of δ, the stretch improves to $1 + \frac{O(1)}{k}$, since we only use one long edge. This variation will be useful for the class of doubling metrics.

1.2.2 Reliable Spanners via Light k-HST Covers

A (τ, ρ)-tree cover of a metric space (X, d), is a collection of τ dominating trees, such that for every pair $u, v \in X$, there exists a tree T in the cover with $d_T(u, v) \le \rho \cdot d(u, v)$. Let (X, d) be any metric that admits a (τ, ρ)-tree cover in which all trees are k-HSTs of weight at most $O(l \cdot w(MST(X))$, then we can devise an oblivious reliable spanner for X as follows. Sample an oblivious light ν/τ-reliable spanner H_T for each tree T, and define $H = \bigcup_T H_T$ as their union. We define B^+ as the union of all the failure sets B_T^+ over all tree spanners.

Since in every ν/τ-reliable spanner of a tree only $\nu/\tau \cdot |B|$ additional vertices fail in expectation, the total expected number of additional failures is at most $\nu \cdot |B|$, as required. Now, if a pair u, v did not fail, there is a k-HST T in which $d_T(u,v) \leq \rho \cdot d(u,v)$, and thus H has stretch at most $\rho \cdot (2 + \frac{O(1)}{k})$ for such a pair.

Light k-HST Covers using Pairwise Partition Cover Scheme. A $(\tau, \rho, \varepsilon, \Delta)$-Pairwise Partition Cover for a metric space (X, d) is a collection of τ partitions, each cluster in each partition has diameter at most Δ, and every pair $u, v \in X$ with $\frac{\Delta}{2\rho} \leq d(u,v) \leq \frac{\Delta}{\rho}$ is *padded* in at least one cluster C of a partition. This means that the cluster C contains u, v, and also the balls of radius $\varepsilon\Delta$ around them . If (X, d) admits such a cover for every Δ, we say it has a Pairwise Partition Cover Scheme (PPCS). In [21], PPCS were shown for general metrics and doubling metrics. In this paper, for any parameter $0 < \varepsilon < 1/6$, we devise a $\left(\frac{\log n}{\varepsilon}, \frac{2}{1-6\varepsilon}, \varepsilon\right)$-PPCS for minor-free graphs.

In [21] it was shown that one can obtain a k-HST cover from a PPCS, in such a way that every cluster of diameter Δ in the PPCS corresponds to an internal node x of one of the k-HSTs, with label $\Gamma_x = \Delta$. For our purposes, we want every k-HST in the cover to be light. To this end, we augment the reduction of [21] by a feature that allows us to bound the lightness of the resulting k-HST. The idea is to use *nets* . A basic observation for a Δ-net $\mathcal{N}$ of a metric space (X, d), is that $w(MST(X)) \geq \Omega(|\mathcal{N}| \cdot \Delta)$. On the other hand, the weight of a k-HST T is roughly $\sum_{x \in T} k \cdot \Gamma_x$ (every node pays for the edge to its parent in T). So as long as the number of internal nodes with label Δ is bounded by $|\mathcal{N}|$, the k-HST will be rather light.

Now, given some partition with diameter bound Δ, we take a $\approx \varepsilon\Delta$-net $\mathcal{N}$, and break all clusters that do not contain a net point. Then the points in the broken clusters are joined to a nearby remaining cluster. Since the net is dense enough, each cluster that was used for padding remains intact, while the number of clusters is bounded by $|\mathcal{N}|$. This enables us to bound the weight of the k-HST accordingly.

1.2.3 Reliable Light Spanner for Minor-free Graphs with $2 + \varepsilon$ stretch

In the special case of minor-free graphs, the framework described above will lose a factor of 2 in the stretch in two places. The first is due to the padding of the PPCS, and the second in the reliable spanners for the k-HSTs. While each of these losses is unavoidable,[10] we can still exploit a certain property of our PPCS for minor-free graphs, to improve the stretch to near optimal $2 + \varepsilon$.

In our previous approach, suppose vertices u, v are padded in some cluster C of the PPCS, with diameter at most Δ. Then in the k-HST cover, we will have some tree with an internal node x corresponding to C, whose label is $\Gamma_x = \Delta$. The way we construct the spanner path between u, v is via some chosen leaf z in $L(x)$, and as both $d(u,z)$, $d(v,z)$ can be as large as Δ, we loose a factor of 2 here.

[10] Stretch 2 for HST is necessary: Consider the uniform metric, every spanner with less than $\binom{n}{2}$ edges has stretch 2. Every PPCS for minor free graphs must have either $\rho \geq 2$ or $\tau = \Omega(n)$: Fix $\rho < 2$, and consider the unweighted star graph. There are $n - 1$ leaf-center pairs, while a single partition can satisfy at most a single pair.

The main observation behind overcoming this loss, is that in our PPCS for minor-free graphs, each cluster C is a ball around some center x, and whenever a pair u, v is padded, then x is very close to the shortest $u-v$ path, meaning that $d(u,x)+d(v,x) \le (1+\varepsilon)\cdot d(u,v)$. While we cannot guarantee that x, or a vertex close to x, will survive the attack B, we can still use this to improve the stretch guarantee. Suppose that Z_x contains a surviving leaf z which is closer to x than both u, v, then

$$d(u,z)+d(z,v) \le (d(u,x)+d(x,z))+d(z,x)+d(x,v)) \le 2(d(u,x)+d(x,v)) \le 2(1+\varepsilon)\cdot d(u,v) .$$

So, instead of sampling a set Z_x of leaves at random from $L(x)$, we create a bias towards vertices closer to the center x. Concretely, order the leaves of $L(x)$ by their distance to x, and we would like that the probability of the j-th leaf in $L(x)$ to join Z_x will be $\approx \frac{1}{j}$. This way, the expected size of Z_x is still small, and if not too many vertices in the appropriate prefix of $L(x)$ are in B, then there is a good probability that such a $z \in Z_x$ exists. However, as it turns out, this requirement it too strict, since every internal node x will force us to move vertices to B^+ that fail due many vertices in B in its induced ordering.

To avoid this hurdle, we use a *global* ordering for all internal nodes – a carefully chosen preorder of T – and prove that the induced order on $L(x)$ is a good enough approximation of distances to x (specifically, up to an additive factor of $\approx \Gamma_x/k$).

1.2.4 Reliable Light Spanner for the Path Graph

There were several construction of a reliable spanner for P_n in previous works [10, 11, 21], none of them could provide a meaningful bound on the lightness. For instance, the first step in the construction of [10] was to connect the first $n/2$ vertices to the last $n/2$ vertices via a bipartite expander graph. In particular, the total weight of just this step is $\Omega(n^2)$. The method of [21] is to sample $\approx \nu^{-1}$ vertices as star centers, and connect all other vertices to each center. This construction also clearly isn't light, as the total weight of even one such star is $\Omega(n^2)$.

Our construction of an oblivious light ν-reliable spanner for (weighted) P_n is similar to the approach taken by [11]. It starts by sampling a laminar collection of subsets $[n] = V_0 \supseteq V_1 \supseteq V_2 \supseteq \cdots \supseteq V_{\log n}$, where $|V_i|$ contains $\frac{n}{2^i}$ points in expectation. However, the construction of [11] used long range edges: from vertices in V_i to the nearest $\approx 2^{i/2}$ other vertices in V_i, and thus its lightness is polynomial in n.[11]

To ensure bounded lightness, we take a more local approach, and each point $a \in V_i$ adds edges to only the nearest $\ell \approx \nu^{-1}$ points in V_i and V_{i+1} on both its left and right sides. We remark that the connections to the next level are crucial in order to avoid additional logarithmic factors (since unlike [11], we cannot use the exponentially far away vertices, that would have provided high probability for connection of every vertex to the next level). The lightness follows as each edge e of P is expected to be "covered" ℓ^2 times, in each of the $\log n$ levels.

The reliability analysis of our spanner uses the notion of *shadow*, introduced by [10]. For the path P_n, roughly speaking, a vertex u is outside the α-shadow of an attack B, if in all intervals containing u, there is at most an α fraction of failed vertices (in B).

[11] To see why the lightness is polynomial, consider just the level $i = \frac{2}{3}\log n$, then $|V_i| \approx n^{1/3}$, but also the number of connected neighbors is $2^{i/2} = n^{1/3}$, so all $\approx n^{2/3}$ edges between vertices in V_i are added. The average length of these edges is linear in n, so the lightness is $\Omega(n^{2/3})$.

The reliability argument goes as follows: a vertex $a \in [n] \setminus B$ fails and joins B^+ only if there exists a level i in which all its connections to V_{i+1} fail. That is, its ℓ closest vertices in V_{i+1} are in B. But as points are chosen to V_{i+1} independently of B, this is an unlikely event, whose probability can be bounded as a function of the largest α-shadow that does not contain a. To obtain our tight bound, we need a delicate case-analysis for the different regimes of α-shadows.

The stretch analysis is a refinement of [11] stairway approach. A nice feature is that each pair in $[n] \setminus B^+$ will have a shortest path of at most $\log n$ hops in the spanner H.

1.3 Related Work

Light fault-tolerant spanners. Levcopoulos *et. al.* [32] introduced the notion of f-fault-tolerant spanner, where it is guaranteed that for every set F of at most f faulty nodes, $H \setminus F$ is a t-spanner of $X \setminus F$. However, the parameter f has to be specified in advance, and both sparsity and lightness of the spanner must polynomially depend on f. Thus, unlike reliable spanners, it is impossible to construct sparse and light fault-tolerant spanners that can withstand scenarios where, say, half of the nodes fail.

Czumaj and Zhao [16] constructed f-fault-tolerant spanners for points in constant dimensional Euclidean space with optimal $O(f^2)$ lightness (improving over [32] $2^{O(f)}$ lightness). This result was very recently generalized to doubling spaces by Le, Solomon, and Than [31], who obtain $O(f^2)$ lightness (improving over [12] $O(f^2 \log n)$ lightness, and [36] $O(f^2 + f \log n)$ lightness).

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.

More on Light spanners. Light spanners were constructed for high dimensional Euclidean and doubling spaces [22, 30]. Subset light spanners were studied for planar and Minor free graphs [27, 28, 29, 15], where the goal is to maintain distances only between a subset of terminals (and the lightness is defined w.r.t. the minimum Steiner tree). Bartal *et. al.* constructed light prioritized and scaling spanner [5], where only a small fraction of the vertex pairs suffer from large distortion. Recently Le and Solomon conducted a systematic study of efficient constructions of light spanners [30] (see also [23, 3]). Finally, light spanners were efficiently constructed in the LOCAL [26], and CONGEST [17] distributed models.

2 Light Reliable Spanner for k-HSTs

In this section we devise a light reliable spanner for the family of k-HSTs (see defenition in the full version [20]). Let T be the tree corresponding to the given k-HST, we refer to its leaves as vertices, and to the interval nodes as nodes. Each node has an arbitrary order on its children. For a node x we denote by $L(x)$ the set of leaves in the subtree rooted at x, and by $L = [n]$ the set of all leaves. For an internal node x in T, let $\deg(x)$ denote the number of children of x. We will assume that $\deg(x) \geq 2$ (as degree 1 nodes are never the least common ancestor, and thus can be contracted). Our goal is to prove the following theorem.

▶ **Theorem 2.** *For any parameters $\nu \in (0, 1/6)$ and $k > 1$, every k-HST T admits an oblivious ν-reliable $(2 + \frac{2}{k-1})$-spanner of size $n \cdot \tilde{O}(\nu^{-1} \cdot \log\log n)^2$ and lightness $\tilde{O}(\nu^{-1} \cdot \log\log n)^2$.*

2.1 Decomposition of T to Heavy Paths

We apply the following decomposition of T into paths, reminiscent of the heavy-path decomposition [35]. Each node $x \in T$ is given a tag, initially $\sigma_x = |L(x)|$, and set $D = \emptyset$. Go over the nodes of T in preorder, and when visiting node x with children $x_1, \dots, x_t$: If there is $1 \le j \le t$ such that $\sigma_{x_j} > (1-\nu/2)\sigma_x$, set $\sigma_{x_j} = \sigma_x$ and add the edge $\{x, x_j\}$ to D. For example, if T contains a path $(y_1, y_2, \dots, y_q)$ where y_1 is the closest vertex to the root, and $|L(y_q)| > (1-\nu/2)|L(y_2)|$ while $|L(y_2)| < (1-\nu/2)|L(y_1)|$ then it will hold that $\sigma_{y_1} \ne \sigma_{y_2} = \sigma_{y_3} = \dots = \sigma_{y_q} = |L(y_2)|$.

We claim that $\sigma_x \ge |L(x)|$ for every node $x \in T$, because we either have equality or x inherit the original tag of one of its ancestors. As $1-\nu/2 > 1/2$, there cannot be two different children of x with more than $|L(x)|/2$ leaves in their subtree, hence there can be at most one child x_j for which an edge is added to D. So indeed D is a decomposition of T into heavy paths (some paths can be singletons). Denote by $\mathcal{Q}$ this collection of paths, and for each $Q \in \mathcal{Q}$, let $f(Q)$ be the lowest vertex (farthest from the root) on Q. We overload this notation, and define $f(x) = f(Q)$, where Q is the heavy path containing x. Let $F = \{f(Q)\}_{Q\in\mathcal{Q}}$ be the set of lowest vertices over all paths.

▷ Claim 3. Each root-to-leaf path W intersects at most $O(\nu^{-1}\log n)$ paths in $\mathcal{Q}$.

Proof. Fix a path $Q \in \mathcal{Q}$. Note that all nodes in Q have the same tag σ_Q. Whenever the path W leaves Q, it will go to some node y with $\sigma_y \le (1-\nu/2)\sigma_Q$. The root has tag n, so after leaving $2\nu^{-1}\ln n$ heavy paths, the tag will be at most

$$n \cdot (1-\nu/2)^{2\nu^{-1}\ln n} < n \cdot e^{-\ln n} = 1 \ ,$$

since the tag of any internal node x is at least $|L(x)|$, we must have reached a leaf. ◁

2.2 Construction

For each node $y \in F$, we independently sample uniformly at random a set Z_y of $\ell = c \cdot \nu^{-1} \cdot \ln\left(\frac{\ln n}{\nu}\right)$ vertices from $L(y)$, where c is a constant to be determined later. If there are less than ℓ vertices in $L(y)$, take $Z_y = L(y)$. For each internal node x in T with children $x_1, \dots, x_t$, and for every $1 \le j \le t$, we add the edges $\{\{y,z\} \ : \ y \in Z_{f(x)}, z \in Z_{f(x_j)}\}$ to the spanner H.

Defining the set B^+. Consider an attack B. We say that an internal node $x \in T$ is *good* if $Z_{f(x)} \setminus B \ne \emptyset$. A leaf u is *safe* if for every ancestor x of u, x is good. In other words, a leaf is safe if every ancestor x sampled a leaf to $Z_{f(x)}$ which is not in B.
Define B^+ as the set of all leaves which are not safe.

2.3 Analysis

Size Analysis. For each internal node x in F and each child x_j of x, we added the bi-clique $Z_x \times Z_{x_j}$, which contains at most ℓ^2 edges. Since the sum of degrees of internal nodes in T is $O(n)$ (recall that all degrees are at least 2), the total number of edges added to H is $O(n \cdot \ell^2) = n \cdot \tilde{O}(\nu^{-1} \cdot \log\log n)^2$.

Weight Analysis. First, we claim that the weight of the MST for the leaves of T equals to

$$\sum_{x\in T} (\deg(x) - 1) \cdot \Gamma_x \ . \qquad (1)$$

This can be verified by running Boruvka's algorithm, say.[12] Every internal node x in F, adds at most $\ell^2 \cdot \deg(x)$ edges of weight at most Γ_x to the spanner. The total weight is thus

$$\sum_{x\in F} \deg(x)\cdot \ell^2 \cdot \Gamma_x = O(w(MST)\cdot \ell^2) = w(MST)\cdot \tilde{O}(\nu^{-1}\cdot \log\log n)^2 \ .$$

Stretch Analysis. The stretch analysis is based on the following lemma.

▶ **Lemma 4.** *Let $u \notin B^+$ be any safe leaf. Then for any ancestor x of u and any $v \in Z_{f(x)} \setminus B$, the spanner H contains a path from u to v of length at most $\left(1+\frac{1}{k-1}\right)\cdot\Gamma_x$ that is disjoint from B.*

Proof. The proof is by induction on $|L(x)|$. The base case is when $x = u$, then $L(u) = \{u\}$ and the statement holds trivially. Let x be an ancestor of u, and take any vertex $v \in Z_{f(x)} \setminus B$. We need to find a path in H of length at most $\left(1+\frac{1}{k-1}\right)\cdot\Gamma_x$ from u to v that is disjoint from B.

Let x_u be the child of x whose subtree contains u. Since u is safe, we know that $Z_{f(x_u)} \setminus B \neq \emptyset$, so take any vertex $u' \in Z_{f(x_u)} \setminus B$. By the induction hypothesis on x_j, there is a path P' in H from u to u' of length at most $\left(1+\frac{1}{k-1}\right)\cdot\Gamma_{x_j}$ disjoint from B (note that indeed $|L(x_j)| < |L(x)|$, as all vertices have degree at least 2). Recall that in the construction step for x, we added all edges from $Z_{f(x)}$ to $Z_{f(x_u)}$, in particular the edge $\{u', v\} \in H$. Note that $v \notin B$, that $u', v \in L(x)$ and therefore $d_T(u', v) \leq \Gamma_x$, and as T is a k-HST we have that $\Gamma_{x_j} \leq \frac{\Gamma_x}{k}$. It follows that the path $P = P' \circ \{u', v\}$ from u to v in H is disjoint from B, and has length at most $\left(1+\frac{1}{k-1}\right)\cdot\Gamma_{x_j} + \Gamma_x \leq \left(\frac{1+\frac{1}{k-1}}{k}\right)\cdot\Gamma_x + \Gamma_x = \left(1+\frac{1}{k-1}\right)\cdot\Gamma_x$ ◀

Fix a pair of leaves $u, v \notin B^+$, and let $x = \mathrm{lca}(u, v)$. Since both are safe, $Z_{f(x)} \setminus B \neq \emptyset$, and pick any $z \in Z_{f(x)} \setminus B$. By Lemma 4 there are paths in H from u to z and from v to z, both disjoint from B, of combined length at most

$$2\cdot\left(1+\frac{1}{k-1}\right)\cdot\Gamma_x = \left(2+\frac{2}{k-1}\right)\cdot d_T(u,v) \ .$$

Reliability Analysis. For every $x \in T$, denote by $B^{(x)}$ the set of all vertices in $u \in L(x) \setminus B$, such that there is an ancestor z of u in the subtree rooted at x for which $Z_{f(z)} \subseteq B$. In other words, those are the leaves (outside B) who are not safe due to a bad ancestor in the subtree rooted at x.

We say that a node $x \in T$ is *brutally attacked* if $|B \cap L(x)| \geq (1-\nu)\cdot|L(x)|$, that is at least a $1-\nu$ fraction of the decedent leaves of x are in the attack B. Denote by $B_1^{(x)} \subseteq B^{(x)}$ the set of vertices $u \in L(x) \setminus B$ that have a brutally attacked ancestor y in the subtree rooted at x. Denote by $B_2^{(x)} = B^{(x)} \setminus B_1^{(x)}$ the rest of the vertices in $B^{(x)}$.

We next argue that the number of vertices added to B^+ (in the worst case) due to brutally attacked nodes is bounded by $O(\nu)\cdot|B|$. Let A_{ba} be the set of T nodes which are brutally attacked, and they are maximal w.r.t. the order induced by T. That is, $x \in A_{\mathrm{ba}}$ if and only

[12] In Boruvka's algorithm, we start with all vertices as singleton components. In each iteration, every component adds to the MST the edge of smallest weight leaving it (breaking ties consistently). For a k-HST, we use a small variation – only components which are the deepest leaves in the HST participate in the current iteration. We claim that the connected components after the j-th iteration correspond to nodes of height j above the leaves. Thus, in the j-th iteration, any node x of height j will add $\deg(x)-1$ edges with weight Γ_x each, that connect the components corresponding to its children.

if x is brutally attacked, while for every ancestor y of x, y is not brutally attacked. Clearly, for every $x \in A_{\text{ba}}$ it holds that $|B_1^{(x)}| \le |L(x) \setminus B| \le \nu \cdot |L(x)| \le \frac{\nu}{1-\nu} \cdot |L(x) \cap B|$. In total, for the root r of T it holds that

$$|B_1^{(r)}| = \sum_{x \in A_{\text{ba}}} |B_1^{(x)}| \le \sum_{x \in A_{\text{ba}}} \frac{\nu}{1-\nu} \cdot |L(x) \cap B| \le \frac{\nu}{1-\nu} \cdot |B| \le 2\nu \cdot |B| \ .$$

Next we bound the damage done (in expectation) due to non brutally attacked nodes. Denote $\beta = \frac{1}{\ln \ln n}$. We will prove for any node $x \in T$ which is not a heavy child, by induction on $|L(x)|$ that

$$\mathbb{E}[|B_2^{(x)}|] \le \max\left\{0, \nu \cdot \beta \cdot \ln \ln(|L(x)|) \cdot |B \cap L(x)|\right\} \ . \tag{2}$$

The base case where $|L(x)| \le \nu^{-1}$ holds trivially as $B_2^{(x)} = \emptyset$. Indeed, consider a descendent leaf $v \notin B$ of x. For every ancestor internal node y of v, which is a descendent of x, it holds that $f(y) = y$ (y does not have heavy children as $|L(y)| - 1 = (1 - \frac{1}{|L(y)|}) \cdot |L(y)| < (1 - \frac{\nu}{2}) \cdot |L(y)|$). In particular $v \in Z_{f(y)} \setminus B$. It follows that $v \notin B_2^{(x)}$, and thus $B_2^{(x)} = \emptyset$. In general, let $x \in T$ be an inner node, which is not a heavy child. Denote $m = |L(x)| > \nu^{-1}$. x is the first vertex in a heavy path $Q = (x = y_1, y_2, ..., y_s) \in \mathcal{Q}$. Let $x_1, \dots, x_t$ be the children of all the nodes in Q. Observe that none of $x_1, \dots, x_t$ is a heavy child, and that $L(x_1), \dots, L(x_t)$ is a partition of $L(x)$. The main observation is that all the vertices in Q use the same sample $Z_{f(x)}$, so a leaf u is in $B_2^{(x)}$ if at least one the following holds:

1. $u \in B_2^{(x_j)}$ for some $1 \le j \le t$, or
2. $Z_{f(x)} \subseteq B$.

We conclude that
$$\mathbb{E}[|B_2^{(x)}|] \le \sum_{j=1}^{t} \mathbb{E}[|B_2^{(x_j)}|] + |L(x)| \cdot \Pr[Z_{f(x)} \subseteq B] \ . \tag{3}$$

In what follows we bound each of the two summands. For the first, we use the induction hypothesis on x_j (clearly $|L(x_j)| < m = |L(x)|$), to get that

$$\mathbb{E}\left[\left|B_2^{(x_j)}\right|\right] \le \max\left\{0, \nu \cdot \beta \cdot \ln \ln(|L(x_j)|) \cdot |B \cap L(x_j)|\right\} \ .$$

By definition of a heavy path, for every $1 \le j \le t$, $|L(x_j)| \le (1-\nu/2) \cdot \sigma_Q = (1-\nu/2) \cdot m$. It holds that $(1-\frac{\nu}{2}) \cdot m \ge (1-\frac{\nu}{2}) \cdot \nu^{-1} \ge \nu^{-1} - \frac{1}{2} \ge 5.5$, and in particular, $\ln \ln\left((1-\frac{\nu}{2}) \cdot m\right) > 0$. It follows that

$$\begin{aligned} \sum_{j=1}^{t} \mathbb{E}[|B_2^{(x_j)}|] &\le \textstyle\sum_{j=1}^{t} \nu \cdot \beta \cdot \ln \ln\left(\left(1-\frac{\nu}{2}\right) \cdot m\right) \cdot |B \cap L(x_j)| \\ &= \nu \cdot \beta \cdot \ln \ln\left(\left(1-\frac{\nu}{2}\right) \cdot m\right) \cdot |B \cap L(x)| \ . \end{aligned} \tag{4}$$

For the second summand, we now analyze the probability of the event $Z_{f(x)} \subseteq B$. If $|B \cap L(x)| \ge (1-\nu) \cdot |L(x)|$, then x is brutally attacked and thus $B_2^{(x)} = \emptyset$ and (2) holds. We thus can assume $|B \cap L(x)| < (1-\nu) \cdot |L(x)|$. By the heavy path decomposition, it holds that $|L(f(x))| > (1-\frac{\nu}{2}) \cdot m$. In the case that $|L(f(x))| \le \ell$ we take $Z_{f(x)} = L(f(x))$, and as $|L(f(x))| > (1-\frac{\nu}{2}) \cdot m > (1-\nu)m > |B \cap L(x)|$, there must be a vertex in $Z_{f(x)} \setminus B$. In particular, $\Pr\left[Z_{f(x)} \subseteq B\right] = 0$. Otherwise, we have that $|L(f(x))| > \ell$. As $Z_{f(x)}$ is chosen from $L(f(x))$ independently of B, by lemma proved in the full version [20], the probability that all of the ℓ vertices in $Z_{f(x)}$ are chosen from $B \cap L(f(x))$ is at most

$$\Pr\left[Z_{f(x)} \subseteq B\right] = \frac{\binom{|B\cap L(f(x))|}{\ell}}{\binom{|L(f(x))|}{\ell}} \leq O(\sqrt{\ell}) \cdot \left(\frac{|B \cap L(f(x))|}{|L(f(x))|}\right)^{\ell}$$
$$\leq O(\sqrt{\ell}) \cdot \left(\frac{1-\nu}{1-\frac{\nu}{2}}\right)^{\ell-1} \cdot \frac{|B \cap L(f(x))|}{m}$$
$$\overset{(*)}{\leq} \frac{\nu^2 \cdot \beta}{4 \cdot \ln n} \cdot \frac{|B \cap L(f(x))|}{m} \leq \frac{\nu^2 \cdot \beta}{4 \cdot \ln m} \cdot \frac{|B \cap L(x)|}{m} \ , \tag{5}$$

where the inequality $^{(*)}$ uses that $\frac{1-\nu}{1-\frac{\nu}{2}} \leq 1 - \frac{\nu}{2} \leq e^{-\nu/2}$, and taking a large enough constant c in the definition of ℓ. By plugging (4) and (5) into (3) we conclude that,

$$\mathbb{E}\left[\left|B_2^{(x)}\right|\right] \leq \sum_{j=1}^{t} \mathbb{E}[|B_2^{(x_j)}|] + m \cdot \Pr[Z_{f(x)} \subseteq B]$$
$$\leq \nu \cdot \beta \cdot \ln\ln\left(\left(1 - \frac{\nu}{2}\right) \cdot m\right) \cdot |B \cap L(x)| + \frac{\nu^2 \cdot \beta}{4 \cdot \ln m} \cdot |B \cap L(x)|$$
$$\overset{(**)}{\leq} \nu \cdot \beta \cdot \ln\ln m \cdot |B \cap L(x)| \ ,$$

which concludes the proof of (2), and thus the induction step. It remains to validate $^{(**)}$:

$$\ln\ln m - \ln\ln\left((1 - \frac{\nu}{2}) \cdot m\right) = \ln \frac{\ln m}{\ln\left((1-\frac{\nu}{2}) \cdot m\right)} \geq \ln \frac{\ln m}{\ln m - \ln(1 + \frac{\nu}{2})}$$
$$\geq \ln\left(1 + \frac{\ln(1+\frac{\nu}{2})}{\ln m}\right) \geq \frac{\ln(1+\frac{\nu}{2})}{2\ln m} \geq \frac{\nu}{4\ln m} \ ,$$

using $\ln(1+x) \geq \frac{x}{2}$ for $0 < x < 1$. Finally, by applying (2) on the root r of T, we get that

$$\mathbb{E}[|B^+ \setminus B|] = \mathbb{E}[|B_1^{(r)}| + |B_2^{(r)}|] \leq (2\nu + \nu \cdot \beta \cdot \ln\ln n) \cdot |B| = 3\nu \cdot |B| \ .$$

Theorem 2 follows by rescaling ν by a factor of 3.

References

1 Mohammad Ali Abam, Mark de Berg, Mohammad Farshi, and Joachim Gudmundsson. Region-fault tolerant geometric spanners. *Discret. Comput. Geom.*, 41(4):556–582, 2009. Preliminary version published in SODA 2007. `doi:10.1007/s00454-009-9137-7`.

2 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`.

3 Stephen Alstrup, Søren Dahlgaard, Arnold Filtser, Morten Stöckel, and Christian Wulff-Nilsen. Constructing light spanners deterministically in near-linear time. In Michael A. Bender, Ola Svensson, and Grzegorz Herman, editors, *27th Annual European Symposium on Algorithms, ESA 2019, September 9-11, 2019, Munich/Garching, Germany*, volume 144 of *LIPIcs*, pages 4:1–4:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. `doi:10.4230/LIPIcs.ESA.2019.4`.

4 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`.

5 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`.

6 Greg Bodwin. An alternate proof of near-optimal light spanners. *CoRR*, abs/2305.18647, 2023. doi:10.48550/arXiv.2305.18647.

7 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.

8 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.

9 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.

10 Kevin Buchin, Sariel Har-Peled, and Dániel Oláh. A spanner for the day after. *Discret. Comput. Geom.*, 64(4):1167–1191, 2020. doi:10.1007/s00454-020-00228-6.

11 Kevin Buchin, Sariel Har-Peled, and Dániel Oláh. Sometimes reliable spanners of almost linear size. *J. Comput. Geom.*, 13(1):178–196, 2022. doi:10.20382/jocg.v13i1a6.

12 T.-H. Hubert Chan, Mingfei Li, Li Ning, and Shay Solomon. New doubling spanners: Better and simpler. *SIAM J. Comput.*, 44(1):37–53, 2015. doi:10.1137/130930984.

13 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.

14 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.

15 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. arXiv:2009.05039.

16 Artur Czumaj and Hairong Zhao. Fault-tolerant geometric spanners. *Discret. Comput. Geom.*, 32(2):207–230, 2004. doi:10.1007/s00454-004-1121-7.

17 M. Elkin, A. Filtser, and O. Neiman. Distributed construction of light networks. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*, PODC'20, pages 483–492, 2020. doi:10.1145/3382734.3405701.

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

19 Arnold Filtser. Labeled nearest neighbor search and metric spanners via locality sensitive orderings. In Erin W. Chambers and Joachim Gudmundsson, editors, *39th International Symposium on Computational Geometry, SoCG 2023, June 12-15, 2023, Dallas, Texas, USA*, volume 258 of *LIPIcs*, pages 33:1–33:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.SoCG.2023.33.

20 Arnold Filtser, Yuval Gitlitz, and Ofer Neiman. Light, reliable spanners. *CoRR*, abs/2307.16612, 2023. doi:10.48550/arXiv.2307.16612.

21 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.

22 Arnold Filtser and Ofer Neiman. Light spanners for high dimensional norms via stochastic decompositions. *Algorithmica*, 84(10):2987–3007, 2022. doi:10.1007/s00453-022-00994-0.

23 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.

24 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.

25 Sariel Har-Peled, Manor Mendel, and Dániel Oláh. Reliable spanners for metric spaces. *ACM Trans. Algorithms*, 19(1):7:1–7:27, 2023. `doi:10.1145/3563356`.

26 Iyad A. Kanj, Ljubomir Perkovic, and Ge Xia. Computing lightweight spanners locally. In Gadi Taubenfeld, editor, *Distributed Computing, 22nd International Symposium, DISC 2008, Arcachon, France, September 22-24, 2008. Proceedings*, volume 5218 of *Lecture Notes in Computer Science*, pages 365–378. Springer, 2008. `doi:10.1007/978-3-540-87779-0_25`.

27 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`.

28 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`.

29 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`.

30 Hung Le and Shay Solomon. A unified framework for light spanners. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 295–308. ACM, 2023. `doi:10.1145/3564246.3585185`.

31 Hung Le, Shay Solomon, and Cuong Than. Optimal fault-tolerant spanners in euclidean and doubling metrics: Breaking the $\Omega(\log n)$ lightness barrier. *CoRR*, abs/2306.11226, 2023. To appear in FOCS 2023. `doi:10.48550/arXiv.2306.11226`.

32 Christos Levcopoulos, Giri Narasimhan, and Michiel H. M. Smid. Efficient algorithms for constructing fault-tolerant geometric spanners. In Jeffrey Scott Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 186–195. ACM, 1998. `doi:10.1145/276698.276734`.

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

34 D. Peleg and A. A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989. `doi:10.1002/jgt.3190130114`.

35 Daniel Dominic Sleator and Robert Endre Tarjan. A data structure for dynamic trees. *J. Comput. Syst. Sci.*, 26(3):362–391, 1983. `doi:10.1016/0022-0000(83)90006-5`.

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