Spanners in Planar Domains via Steiner Spanners and non-Steiner Tree Covers[*]

Sujoy Bhore [†] Balázs Keszegh [‡] Andrey Kupavskii [§] Hung Le[¶]
Alexandre Louvet [||] Dömötör Pálvölgyi [**] Csaba D. Tóth[††]

Abstract

We study spanners in planar domains, including polygonal domains, polyhedral terrain, and planar metrics. Previous work showed that for any constant $\varepsilon \in (0,1)$, one could construct a $(2+\varepsilon)$-spanner with $O(n\log(n))$ edges (SICOMP 2019), and there is a lower bound of $\Omega(n^2)$ edges for any $(2-\varepsilon)$-spanner (SoCG 2015). The main open question is whether a linear number of edges suffices and the stretch can be reduced to 2. We resolve this problem by showing that for stretch 2, one needs $\Omega(n\log n)$ edges, and for stretch $2+\varepsilon$ for any fixed $\varepsilon \in (0,1)$, $O(n)$ edges are sufficient. Our lower bound is the first super-linear lower bound for stretch 2.

En route to achieve our result, we introduce the problem of constructing non-Steiner tree covers for metrics, which is a natural variant of the well-known Steiner point removal problem for trees (SODA 2001). Given a tree and a set of terminals in the tree, our goal is to construct a collection of a small number of dominating trees such that for every two points, at least one tree in the collection preserves their distance within a small stretch factor. Here, we identify an unexpected threshold phenomenon around 2 where a sharp transition from n trees to $\Theta(\log n)$ trees and then to $O(1)$ trees happens. Specifically, (i) for stretch $2-\varepsilon$, one needs $\Omega(n)$ trees; (ii) for stretch 2, $\Theta(\log n)$ tree is necessary and sufficient; and (iii) for stretch $2+\varepsilon$, a constant number of trees suffice. Furthermore, our lower bound technique for the non-Steiner tree covers of stretch 2 has further applications in proving lower bounds for two related constructions in tree metrics: reliable spanners and locality-sensitive orderings. Our lower bound for locality-sensitive orderings matches the best upper bound (STOC 2022).

Finally, we study $(1+\varepsilon)$-spanners in planar domains using Steiner points. In planar domains, Steiner points are necessary to obtain a stretch arbitrarily close to 1. Here, we construct a $(1+\varepsilon)$-spanner with an *almost linear dependency on* ε in the number of edges; the precise bound is $O((n/\varepsilon)\cdot\log(\varepsilon^{-1}\alpha(n))\cdot\log\varepsilon^{-1})$ edges, where $\alpha(n)$ is the inverse Ackermann function. Our result generalizes to graphs of bounded genus. For n points in a polyhedral metric, we construct a Steiner $(1+\varepsilon)$-spanner with $O((n/\varepsilon)\cdot\log(\varepsilon^{-1}\alpha(n))\cdot\log\varepsilon^{-1})$ edges.

1 Introduction

Let $\mathcal{M}=(X,\delta_X)$ be a metric space and $P\subseteq X$ be a set of n points in $\mathcal{M}$. A t-spanner of P is an edge-weighted graph $G=(P,E,w)$ such that every edge $(p,q)\in E$ has a weight $w(p,q)=\delta_X(p,q)$ and for every two points $x,y\in P$, $\delta_G(x,y)\le t\cdot\delta_X(x,y)$. Here $\delta_G(x,y)$ denotes the shortest path distance between x and y in G. The parameter t is called the stretch of the spanner G. One of the most well-studied class of spanners are Euclidean spanners, where $\mathcal{M}$ is an Euclidean space. The pioneering work of Chew [32, 33] showed that in the *Euclidean plane* $\mathbb{R}^2$, one can construct a spanner with $O(n)$ edges and $O(1)$ stretch. Over more than three decades, this result has been refined, improved, and extended in various ways. Most notably, for any $\varepsilon\in(0,1)$, one can construct a spanner with $O(n/\varepsilon)$ edges and stretch $1+\varepsilon$ for point sets in the Euclidean plane [35, 53], and the number of edges is tight [60]. In higher dimensions d, one could obtain a similar bound: The number of edges is $O(n/\varepsilon^{d-1})$ [65, 6], which is also tight [60].

[*]The full version of the paper can be accessed at `https://arxiv.org/abs/2404.05045`

[†]Department of Computer Science & Engineering, Indian Institute of Technology Bombay. sujoy@cse.iitb.ac.in
[‡]HUN-REN Alfréd Rényi Institute of Mathematics and ELTE Eötvös Loránd University, Budapest, Hungary. keszegh@renyi.hu
[§]Moscow Institute of Physics and Technology and St. Petersburg State University, Russia, kupavskii@ya.ru
[¶]University of Massachusetts Amherst, hungle@cs.umass.edu
[||] LIPN, Université Sorbonne Paris Nord, alexandre.louvet@mailo.fr
[**]ELTE Eötvös Loránd University and HUN-REN Alfréd Rényi Institute of Mathematics, Budapest, domotor.palvolgyi@ttk.elte.hu
[††]California State University Northridge, Los Angeles, CA, and Tufts University, Medford, MA, csaba.toth@csun.edu

Planar Domains. While spanners for points on the Euclidean plane are well understood, in many practical applications, the domain is planar but not Euclidean. One basic example is the polygonal domain routing in robotics. Here, the metric space $\mathcal{M}$ contains points in a polygon—for example, the floor of a room—and there are (polygonal) obstacles inside the polygon—representing furniture inside the room—called holes. The distance between two points is measured by the shortest path *avoiding the obstacles*; see Figure 1(a). Another important setting is polyhedral terrain. A polyhedral terrain is the graph of a piece-wise linear function $f : D \to \mathbb{R}$ for some convex polygonal region $D \subseteq \mathbb{R}^2$; see Figure 1(b). Polyhedral terrains are central in GIS (geographic information system) to model the surfaces of mountains [49]. Abam, de Berg, and Seraji [2] noted that polyhedral terrain generalizes polygonal domain. It is relatively easy to show that in both settings, achieving a $(2-\varepsilon)$-spanner for any fixed $\varepsilon \in (0,1)$ requires $\Omega(n^2)$ edges; see [1, Theorem 3]. The main problem is to construct a spanner with stretch 2 or $2+\varepsilon$ and a linear number of edges.

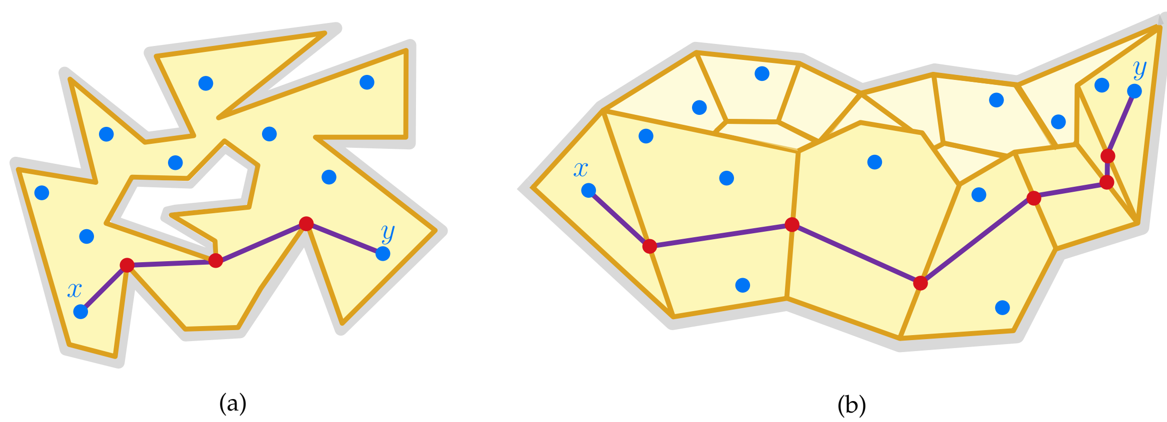

Figure 1: (a) A polygon with holes, blue terminals, and a shortest path between terminals x and y. (b) A polyhedral terrain and a shortest path between two points x and y.

Abam, Adeli, Homapour, and Asadollahpoor [1] constructed a $(5+\varepsilon)$-spanner for any n-point set in a polygonal domain of h holes with $O(n\sqrt{h}\log^2(n))$ edges for any fixed $\varepsilon \in (0,1)$. The number of edges depends on h, which could be as large as n, and furthermore, there is still a gap in the stretch. Their results were significantly generalized and improved by Abam, de Berg, and Seraji [2]. They constructed[1] a $(2+\varepsilon)$-spanner with $O(c(\varepsilon)\cdot n\log n)$ edges, where $c(\varepsilon) = (1+2/\varepsilon)^{O(\log(1/\varepsilon))}$. Note that the dependence on ε is *quasi-polynomial*.

The number of edges of the spanners in both generalized settings [1, 2] remains $\Omega(n\log n)$ for a constant $\varepsilon > 0$, while the number of edges of the spanner in the basic Euclidean setting is $O(n)$. This $\log(n)$ factor gap is due to a fundamental difference in the techniques. The spanner constructions in polygonal domains and polyhedral terrains are based on divide-and-conquer strategy in which $O(n)$ edges will be added in each level of the recursion, resulting in $O(n\log n)$ edges since the recursion depth is $O(\log n)$. On the other hand, in Euclidean spaces, spanner constructions are often non-recursive and directly exploit Euclidean geometry, which is not available in generalized settings.

QUESTION 1.1. *Can we construct a spanner of stretch 2 or $2+\varepsilon$ for any fixed $\varepsilon \in (0,1)$ with $O(n)$ edges? Could the dependence on ε, if necessary, be reduced to be polynomial?*

Both positive and negative answers to Question 1.1 require techniques that are different from those in [1, 2]. First, for stretch 2, we prove an $\Omega(n\log n)$ lower bound on the number of edges. This is the first super-linear lower bound for stretch 2. This lower bound suggests that the number of edges for stretch $2+\varepsilon$ could be super-linear. Our second result shows that this is not the case: We construct a $(2+\varepsilon)$-spanner with $O(n)$ edges for

[1] There was a technical issue in the proof of Abam, de Berg, and Seraji [2], which was recently fixed by de Berg, van Kreveld, and Staals [16].

any constant $\varepsilon \in (0, 1)$, thus completely answering Question 1.1. Our results are summarized in the following theorem.[2]

THEOREM 1.2. *Let $\varepsilon \in (0, 1)$ be a parameter.*

1. *There exists a polyhedral terrain and a set P of n points on the terrain such that any 2-spanner for P must have $\Omega(n \log n)$ edges.*
2. *Given any set P of n points in a polyhedral terrain, we can construct a $(2 + \varepsilon)$-spanner for P with $\tilde{O}(n/\varepsilon^6)$ edges. The number of edges is $O(n)$ for constant ε.*

Our technique for proving Theorem 1.2 builds on a connection to what we call a non-Steiner tree cover for trees, which will be formally defined in Section 1.1. There, we identify a rather surprising threshold phenomenon around stretch 2; see Theorem 1.5. The lower bound (item 1) in Theorem 1.2 will be given in Section 5.2 and the upper bound (item 2) construction will be given in Section 4.

Steiner Spanners. A complementary direction is to study how Steiner points could help constructing spanners in planar domains. Here, Steiner points are points not in the input point set but in the ambient space. In $\mathbb{R}^2$ (and generally in $\mathbb{R}^d$ for any $d \geq 2$), Le and Solomon [60] showed that Steiner points could *quadratically* reduce the dependence of the number of edges of the spanner on $1/\varepsilon$, from $O(n/\varepsilon)$ for non-Steiner spanners to $O(n/\sqrt{\varepsilon})$ for Steiner spanners. Interestingly, the bound $O(n/\sqrt{\varepsilon})$ is tight [18, 60]. Here, we show that for a polyhedral terrain, Steiner points could help in two different ways: They can reduce the stretch from $2 + \varepsilon$ to $1 + \varepsilon$; and also reduce the dependence on ε. We observe that the construction in [2] could be used to construct a Steiner $(1+\varepsilon)$-spanner with $O(n \log(n)/\varepsilon^2)$ edges. Using a completely different technique, we almost remove the $\log(n)$ factor and reduce the dependence on $1/\varepsilon$ to near linear, which we believe is optimal; see more discussion below.

THEOREM 1.3. *Let $\varepsilon \in (0, 1)$ be a parameter. Let P be a set of n points in a polyhedral terrain. We can construct a Steiner $(1 + \varepsilon)$-spanner for P with $O\left((n/\varepsilon) \cdot \log(\varepsilon^{-1}\alpha(n)) \cdot \log \varepsilon^{-1}\right)$ edges, where $\alpha(n)$ is the inverse Ackermann function. The same result holds even when P is on a polyhedral surface of bounded genus.*

The proof of Theorem 1.3 will be given in Section 5. The high-level idea will be given below.

1.1 Key Techniques

1.1.1 Non-Steiner Spanners: Proof of Theorem 1.2

We now describe our technique for constructing $(2 + \varepsilon)$-spanners. As alluded to above, improving upon the previous $O(n \log n)$ bound on the number of edges of a $(2 + \varepsilon)$-spanner requires a novel technique. Our starting point is to understand tree metrics. It is not so hard to see that tree metrics are a special case of polygonal domains. In this case, we are given a tree T and a subset P of n vertices of $V(T)$, and we want to construct *a sparse structure* on P preserving distances between points in P. Here, we introduce a new notion of a sparse structure, called non-Steiner tree cover for tree metrics; the non-Steiner terminology is used to emphasize that the cover does not use Steiner vertices, those that are in $V(T) \setminus P$. Tree covers are a powerful tool that was instrumental for constructing spanners and distance oracles in various metric spaces (such as low-dimensional Euclidean spaces, doubling spaces, planar metrics, and minor-free graphs); see [9, 10, 13, 14, 27, 28, 42, 46, 51].

Non-Steiner Tree Covers for Trees. In this problem, we are given an edge-weighted tree $T = (V_T, E_T, \omega_T)$, and a set of terminals $K \subseteq V_T$. We say that a collection $\mathcal{T} = \{T_1, T_2, \ldots, T_\beta\}$ of β edge-weighted trees is an (α, β)-non-Steiner tree cover if the following hold:

1. **Steiner free.** For every $i \in [\beta]$, $V(T_i) = K$.
2. **Dominating.** For every $i \in [\beta]$, $d_{T_i}(x, y) \geq d_T(x, y)$ for every pair of vertices $x, y \in K$.
3. **Low stretch.** $\min_{i\in[\beta]} d_{T_i}(x, y) \leq \alpha \cdot d_T(x, y)$ for every pair of vertices $x, y \in K$.

[2]The $\tilde{O}(.)$ notation hides logarithmic factors in $1/\varepsilon$.

Parameter α is called the stretch of the cover $\mathcal{T}$, and parameter β is called the size of the cover $\mathcal{T}$.

Our goal is to construct a non-Steiner tree cover with *small stretch and size*. As we will see later, the size-stretch trade-off for the non-Steiner tree cover of tree metrics is central to our construction of spanners in planar domains. While the weight of an edge $e = (x, y)$ in a tree in the non-Steiner tree cover can be theoretically different from $d_T(x, y)$, it is always better to set the weight $w(e) = d_T(x, y)$ since doing so will improve the stretch while preserving all the properties of a non-Steiner tree cover.

The special case of only having exactly one tree in the tree cover is the well-known *Steiner Point Removal* (*SPR*) problem for tree metrics, introduced by Gupta [47]. Gupta achieved stretch at most 8. This result has applications in metric embedding [31, 40] and metric labeling [7]. Gupta also gave a stretch lower bound of $4 - o(1)$. This lower bound was subsequently improved to $8 - o(1)$ [25], matching the upper bound by Gupta [47].

Given that stretch 8 is the best possible for one tree, we ask if we could reduce the stretch by using more than one tree. And more generally:

QUESTION 1.4. *What is the precise trade-off between the number of trees and the stretch?*

Our next result gives a complete answer to Question 1.4.

THEOREM 1.5. *Let $\alpha \geq 1$ be a* stretch parameter, *and $\varepsilon \in (0, 1)$ be any given constant. Let T be an edge-weighted tree and $K \subseteq V(T)$ any set of terminals.*

1. *If $\alpha = 2 + \varepsilon$, then $O(1)$ trees suffice: we can construct a non-Steiner tree cover for K of size $O(1)$ and stretch $2 + \varepsilon$. The number of trees is $O(\varepsilon^{-2} \log(\varepsilon^{-1}))$.*
2. *If $\alpha = 2$, then $O(\log n)$ trees suffice. Furthermore, $\Omega(\log n)$ trees are necessary: there exists a tree and a terminal set such that any tree cover with stretch 2 for the terminals must have $\Omega(\log n)$ trees.*
3. *If $\alpha = 2 - \varepsilon$, then $\Theta(n)$ trees are both necessary and sufficient.*

Here, we also see the threshold phenomenon around stretch 2: (i) for $\alpha = 2 - \varepsilon$, one needs $\Omega(n)$ trees, (ii) for $\alpha = 2$, $\Theta(\log n)$ trees are necessary and sufficient, and (iii) for $\alpha = 2 + \varepsilon$, a constant number of trees suffice. We note that both the lower bound and upper bound construction for the third case in Theorem 1.5 are very simple. The upper bound is obtained by constructing a single star for each terminal to preserve the distances from the terminal to other terminals. The lower bound is realized by the star graph with terminals being the leaves; this example was also considered by Gupta [47]. The upper bound proofs in Theorem 1.5 will be given in Section 2, and the lower bound of $\Omega(\log n)$ trees will be given in Section 5.1.

Next, we discuss the connection between non-Steiner tree cover and spanners in planar domains.

Spanners in Planar Domains. Let $\mathbb{M}_1$ and $\mathbb{M}_2$ be two *families of metric spaces*. We say that $\mathbb{M}_1 \sqsubseteq \mathbb{M}_2$ if for every metric space $\mathcal{M}_1 \in \mathbb{M}_1$ (which could be infinite) and any finite point set $P \in \mathcal{M}_1$, there exists a metric space $\mathcal{M}_2 \in \mathbb{M}_2$ such that P embeds isometrically into $\mathcal{M}_2$. That is, there exists a point set $Q \in \mathcal{M}_2$ and a bijection $f : P \to Q$ such that $\delta_1(x, y) = \delta_2(f(x), f(y))$ for every point pair $x, y \in P$. Here δ_1 and δ_2 are the distance functions of $\mathcal{M}_1$ and $\mathcal{M}_2$, respectively. If $\mathbb{M}_1 \sqsubseteq \mathbb{M}_2$ and $\mathbb{M}_2 \sqsubseteq \mathbb{M}_1$, we write that $\mathbb{M}_1 \cong \mathbb{M}_2$. If $\mathbb{M}_1 \sqsubseteq \mathbb{M}_2$ and $\mathbb{M}_2 \not\sqsubseteq \mathbb{M}_1$, then we write $\mathbb{M}_1 \sqsubset \mathbb{M}_2$.

Let TREE, PLANAR, POLYDOM, TERRAIN, and POLYSURF be the family of shortest-path metrics in edge-weighted trees, edge-weighted planar graphs, polygonal domains, polyhedral terrains, and polyhedral surfaces, respectively. We observe that:

LEMMA 1.6. TREE $\sqsubset$ PLANAR $\cong$ POLYDOM $\cong$ TERRAIN $\sqsubset$ POLYSURF.

Thus, three families of metrics, namely planar metrics, polygonal domains, and polyhedral terrains, are equivalent, and they all strictly contain tree metrics. Abam, de Berg, and Seraji [2] showed that POLYDOM $\sqsubseteq$ TERRAIN by controlling the elevation of polyhedral terrains. We can show that TERRAIN $\sqsubseteq$ PLANAR by looking at the arrangements of the geodesic paths in a polyhedral terrain. Finally, we show that PLANAR $\sqsubseteq$ POLYDOM by using polygonal holes to "fill in" the faces of a planar-embedded graph, thereby proving Lemma 1.6; see Section 4 for details. Note, however, that planar metrics need not embed in Euclidean spaces (without obstacles); see [15, 63, 64].

By Lemma 1.6, we could work with planar metrics instead of polyhedral terrains: given a set of n terminal points in a planar metric, construct a spanner containing the terminals only. (The planar metric might contain more points than the terminal points.) Specifically, we will use a recent (Steiner) tree cover developed by [27]. A Steiner tree cover[3] of a metric $\mathcal{M} = (X, \delta_X)$ is a collection of trees $\mathcal{T}$ such that for every tree $T \in \mathcal{T}$ we have $X \subseteq V(T)$, and $\delta_X(x,y) \leq d_T(x,y)$ for every two points $x, y \in X$. The size of the tree cover $\mathcal{T}$ is the number of trees in $\mathcal{T}$ and the stretch is at most α if $d_T(x,y)) \leq \alpha \cdot \delta_X(x,y)$ for *some* $T \in \mathcal{T}$. Note that by definition, X could be a strict subset of $V(T)$, and the points in $V(T) \setminus X$ are called Steiner points. In all existing tree cover constructions, Steiner points are *copies* of the points in X. A different way to think about this is that a point p in X could appear multiple times in T, and we only keep one copy of p in T as the image of p, and regard other copies as Steiner points.

THEOREM 1.7. (THEOREM 1.2 IN [27]) *Let $G = (V, E, w)$ be an edge-weighted planar graph and $\varepsilon \in (0,1)$ be any given parameter. We can construct a Steiner tree cover $\mathcal{T}$ for the shortest path metric of G such that (a) $\mathcal{T}$ has stretch $1+\varepsilon$ and (b) $\mathcal{T}$ has $\tilde{O}(\varepsilon^{-3})$ trees. Furthermore, Steiner points of every tree in $\mathcal{T}$ are copies of points in X.*

Theorem 1.7 allows us to use non-Steiner tree covers developed in Theorem 1.5 to construct a spanner for points in planar metrics, and hence polyhedral terrain by Lemma 1.6. It is worth noting that our spanner is much more structured than simply having a small number of edges: it is the union of a small number of distance-preserving non-Steiner trees.

Ironically, looking at our series of constructions as a whole, one can see that we first construct a spanner for a point set in a planar domain by adding more Steiner points (Theorem 1.7) and then removing all Steiner points in the final step (Theorem 1.5).

We prove a lower bound (item 1 of Theorem 1.2) for tree metrics in Section 5; Lemma 1.6 implies that the same lower bound holds for polyhedral terrains. Indeed, our lower bound applies to the comb graphs, with terminals being the leaves. Here, we establish a connection between non-Steiner spanners for terminals in the comb graphs and low-hop spanners for points in the line metric. Then, we can slightly adapt the lower bound technique for points in the line metric in [58] to achieve our result.

Steiner Spanners: Proof of Theorem 1.3. Here, we use Lemma 1.6 again by constructing a Steiner spanner for a set of terminals in a planar metric. For obtaining a linear number of edges at the cost of a polynomial factor in $1/\varepsilon$, Theorem 1.7 suffices: we simply union all the trees in the Steiner tree cover. However, to reduce the dependency on $1/\varepsilon$ to nearly linear (cf. Theorem 1.3), we resort to more advanced tools, including net trees in planar metrics [61], reduction to additive stretch [61, 27], and tree shortcutting [34, 30, 5, 67, 43]. The basic idea is a reduction to constructing a spanner with additive distortion using the net-tree-based spanner technique. We then apply shortest path separators [62, 68] and tree shortcutting to construct an additive spanner with an almost linear dependency on $1/\varepsilon$.

We believe that the linear dependency on $1/\varepsilon$ is optimal. For more restricted types of distance-preserving structures, such as minor-free [57] or aligned planar structures [29][4], the number of edges was shown to be $\Omega(n/\varepsilon)$ where n is the number of points.

1.2 Other Applications Here, we present two more applications of our technique for proving lower bound $\Omega(\log n)$ on the number of trees with stretch 2 in Theorem 1.5. All results claimed here are proven in Section 5.

Reliable Spanners. In this problem, we are given a metric space $\mathcal{M} = (X, \delta_X)$ and a set $P \subseteq X$ of n points, a t-spanner G of P is (deterministic) ν-reliable for a parameter $\nu \in (0,1)$ if for any subset $B \subseteq P$, there exists a set $B^+ \supseteq B$ of size $|B^+| \leq (1+\nu)\,|B|$ such that $G[P \setminus B]$ is a t-spanner of all the points in $P \setminus B^+$. That is, for every $x, y \in P \setminus B^+$, $\delta_{G[P\setminus B]}(x,y) \leq t \cdot \delta_X(x,y)$. Informally, the spanner is reliable if whenever the vertices in a set (B) fail, then it only affects a few other vertices ($B^+ \setminus B$). We say that G is an oblivious ν-reliable spanner if G is drawn from a distribution $\mathcal{D}$ and $\mathbb{E}_{G\sim\mathcal{D}}[|B^+|] \leq (1+\nu)|B|$.

Deterministic reliable spanners were introduced by Bose, Dujmović, Morin, and Smid [21] for point sets in Euclidean spaces. Their results were improved greatly by Buchin, Har-Peled, and Oláh [22], who constructed a deterministic ν-reliable $(1+\varepsilon)$-spanner of $O\left(n(\log n)(\log\log n)^6\right)$ edges for point sets in $\mathbb{R}^d$ for constants ε, d and

[3]Prior work did not include the prefix "Steiner" in the Steiner tree cover terminology, since Steiner points are not their focus. Here we clearly distinguish between Steiner and non-Steiner versions of tree covers.

[4]Refer to page 9 in the arXiv version of [29] for the exact definition of an aligned planar structure.

ν. This almost matches the lower bound $\Omega(n \log n)$ for $d = 1$ by [21]. In a follow-up work [23], the same authors constructed an oblivious ν-reliable $(1+\varepsilon)$-spanner with $O(n(\log \log n))$ edges using *locality sensitive orderings* [26], bypassing the $\Omega(n \log n)$ lower bound for deterministic reliable spanners.

Har-Peled, Mendel, and Oláh [48] studied reliable spanners for metric spaces. One basic problem is to construct reliable t-spanners for tree metrics. Specifically, for constants ν and ε, they constructed oblivious ν-reliable spanners with stretch $3+\varepsilon$ and $O(n \log^2 n \log^2 \Phi \log(\log(\Phi) \log(n)))$ edges where Φ is the aspect ratio of the metric. Note that there exists a tree metric such that any oblivious reliable spanner with stretch $2 - \varepsilon$ for any constant $\varepsilon \in (0,1)$ must have $\Omega(n^2)$ edges [42]. Filtser and Le [42] reduced the stretch to 2 (which is optimal) and the number of edges to $O\left(n \log^5(n)\right)$ by developing a variant of locality-sensitive orderings for tree metrics. The main open problem is how many edges are necessary and sufficient for stretch 2. It is conceivable that, similar to the Euclidean case, $O(n(\log \log n))$ edges suffice for tree metrics, given that the techniques in these settings are quite similar. Here, we show that this is not the case by proving an $\Omega(n \log n)$ lower bound for the number of edges using the technique developed for Theorem 1.5. Our result separates Euclidean metrics from tree metrics.

THEOREM 1.8. *There exists a tree metric T with n points such that any oblivious $\frac{1}{3}$-reliable 2-spanner for $V(T)$ must have $\Omega(n \log n)$ edges.*

Locality Sensitive Ordering. Chan, Har-Peled, and Jones [26] introduced the notion of locality-sensitive ordering (LSO) and showed that for any point set in $\mathbb{R}^d$, one can construct a locality-sensitive ordering, comprised of $O(1)$ orderings for constant dimension d and stretch parameter $\varepsilon \in (0,1)$. Their LSO has many surprising algorithmic applications, including dynamic spanners, dynamic approximate minimum spanning trees, dynamic bichromatic closest pairs, approximate nearest neighbors [26], and reliable spanners [23]. However, for tree metrics, their notion of LSO is too strong: one needs $\Omega(n)$ orderings in an LSO for any fixed $\varepsilon \in (0, \frac{1}{2})$. Filtser and Le [42] introduced a more relaxed version of LSO tailored specifically for tree metrics, called *left-sided LSO*. As defined in [42], a (τ, ρ)*-left-sided LSO* for a tree metric T of n points is a collection Σ of *linear orderings* over *subsets* of $V(T)$ such that (i) every point x in T belongs to at most τ linear orderings, and (ii) for any two points $x, y \in T$, there exists an ordering $\sigma \in \Sigma$ with the following property: for any $x' \preceq_\sigma x$ and $y' \preceq_\sigma y$, $\delta_T(x', y') \leq \rho \delta_T(x, y)$. That is, the distance between any two points *to the left* of x and y in σ is at most $\rho \cdot \delta_T(x, y)$. Parameter τ is the size of the ordering Σ and ρ is the stretch. Filtser and Le [42] showed that tree metrics admit a left-sided LSO with $O(\log n)$ size and stretch $\rho = 1$. A question raised by their work is: Could the size of the ordering be reduced to $O(1)$? We answer this question negatively by showing that $O(\log n)$ size is indeed *the best possible*. We do so by using the technique developed in the proof of Theorem 1.5.

THEOREM 1.9. *There exists a tree metric T with n points such that any (τ, ρ)-left-sided LSO for T with $\rho = 1$ must have $\tau = \Omega(\log n)$, matching the $O(\log n)$ upper bound by Filtser and Le [42].*

1.3 Further Related Work Steiner spanners were studied for point sets in Euclidean plane with obstacles [8], which determines the same metric as a polygonal domain. In this setting, however, the *vertices of the polygonal obstacles belong to the point set*. Arikati et al. [8] constructed a *planar* Steiner t-spanner for L_1 distance with stretch $t = 1 + \varepsilon$ and $O(n/\varepsilon^2)$ edges; and for L_p distance with stretch $t = 2^{(p-1)/p} + \varepsilon$ and $O_\varepsilon(n)$ edges, where the $O_\varepsilon(.)$ notation hides the dependence on ε, which was not explicitly computed in [8]. Specifically for Euclidean distance, the stretch is $\sqrt{2} + \varepsilon$. Our spanner in Theorem 1.3 has stretch $1 + \varepsilon$ and almost linear dependence on $1/\varepsilon$; furthermore, the number of edges of our spanner *does not* depend on the number of vertices of the obstacles.

Kapoor and Li [52] constructed a Steiner $(1 + \varepsilon)$-spanner for points in a polyhedral surface $\mathcal{P}$, which is a higher-genus generalization of polyhedral terrains (cf. Lemma 1.6). Their spanner has $O(\gamma(\mathcal{P})n/\varepsilon)$ edges, where $\gamma(\mathcal{P})$ is the *geodesic dilation factor* of the surface $\mathcal{P}$, which measures how nice $\mathcal{P}$ is. In the worst case, $\gamma(\mathcal{P})$ could be up to $\Theta(n)$, and hence the spanner has a trivial $O(n^2)$ edges. Our spanner in Theorem 1.3 has linearly many edges regardless of $\gamma(\mathcal{P})$.

Another related direction is to study the *complexity* of geodesic spanners in planar and polyhedral domains. An edge in our spanner might be realized by a geodesic path of up to $\Omega(n)$ edges in the input domain. For example, in a polygonal domain, an edge (u, v) in our spanner could only be realized by a geodesic obstacle-avoiding path of many straight-line edges in the domain. The total number of edges in the input domain to "realize" our spanners is called the complexity of the spanners. This question has recently been studied in depth for both non-Steiner spanners [16] and the Steiner version [17]. The main finding is that, for any $k \geq 1$, a

$(2k+\varepsilon)$-spanner for n points in a simple polygon with m vertices has complexity $O(mn^{1/k} + n\log^2 n)$ [16], and that Steiner points do not help reducing the complexity by much [17].

Organization. We establish upper bounds for the sharp threshold phenomenon around stretch $\alpha = 2$ for non-Steiner tree covers (Theorem 1.5) in Section 2. We continue with constructing a Steiner $(1+\varepsilon)$-spanner for a set of terminals in a planar graph (Theorem 3.1) in Section 3; and sketch a generalization to polyhedral surfaces of bounded genus (Theorem 3.14). Equipped with these technical tools, we can construct 2- and $(1+\varepsilon)$-spanners for planar domains (Theorem 1.2 and Theorem 1.3) in Section 4. We conclude in Section 5 with lower bound constructions: for tree covers (item 1 in Theorem 1.5), for spanners in polyhedral domains (item 1 in Theorem 1.2), and for applications in locally-sensitive orderings (Theorem 1.9) and reliable spanners (Theorem 1.8).

2 Non-Steiner Tree Cover for Trees

In this section, we prove Theorem 1.5, which we restate below.

THEOREM 1.5. *Let $\alpha \geq 1$ be a* stretch parameter, *and $\varepsilon \in (0,1)$ be any given constant. Let T be an edge-weighted tree and $K \subseteq V(T)$ any set of terminals.*

1. *If $\alpha = 2+\varepsilon$, then $O(1)$ trees suffice: we can construct a non-Steiner tree cover for K of size $O(1)$ and stretch $2+\varepsilon$. The number of trees is $O(\varepsilon^{-2}\log(\varepsilon^{-1}))$.*
2. *If $\alpha = 2$, then $O(\log n)$ trees suffice. Furthermore, $\Omega(\log n)$ trees are necessary: there exists a tree and a terminal set such that any tree cover with stretch 2 for the terminals must have $\Omega(\log n)$ trees.*
3. *If $\alpha = 2-\varepsilon$, then $\Theta(n)$ trees are both necessary and sufficient.*

By scaling, we may assume that the minimum distance between any two vertices of T is at least 8 and the maximum distance between any two vertices (i.e., the diameter) is at most Δ for an integer Δ. (The distance lower bound of 8 is somewhat arbitrary; any sufficiently large constant works.)

We view the edge-weighted tree $T = (V_T, E_T, w_T)$ as a *continuous tree* by viewing each edge (u,v) as a continuous line segment (i.e., a geometric realization of a 1D cell complex). We still use *vertices* to refer to the vertices of the discrete tree T. The distances between points corresponding to vertices of T are the distances in the tree, and the distance between any two points in the same line segment of an edge is the length of the sub-segment of the edge connecting the two points. We could naturally extend the distance function to measure the distance between *any* two points $p_1, p_2 \in T$ as follows: let (u_i, v_i) be the edge containing p_i for $i = 1,2$; we assign

$$d_T(p_1,p_2) = \min \begin{cases} d_T(p_1,u_1) + d_T(u_1,u_2) + d_T(u_2,p_2), \\ d_T(p_1,u_1) + d_T(u_1,v_2) + d_T(v_2,p_2), \\ d_T(p_1,v_1) + d_T(v_1,v_2) + d_T(v_2,p_2), \\ d_T(p_1,v_1) + d_T(v_1,u_2) + d_T(u_2,p_2). \end{cases} \tag{2.1}$$

In a single tree construction by Gupta [47], one could assume that the terminals are in the leaves of T. However, here we could not make the same assumption when constructing more than one tree in the cover, making our construction more complicated.

To illustrate our new ideas, we first present the construction of a non-Steiner tree cover of stretch $3+\varepsilon$ in Section 2.1; the tree cover has size $O(\varepsilon^{-1}\log\varepsilon^{-1})$. Then, in Section 2.2, we reduce the stretch to $2+\varepsilon$ at the expense of another factor of ε^{-1} in the number of trees using additional insights. In Section 2.3, we construct a tree cover with $O(\log n)$ trees and stretch exactly 2.

2.1 Stretch $3+\varepsilon$ We describe a construction with stretch $3+12\varepsilon$; we could recover stretch $3+\varepsilon$ by scaling ε. Suppose that we root T at a non-terminal point r_0. We start with a brief overview of our construction: Starting from r_0, we recursively partition the (continuous) tree T into (continuous) subtrees, that we call ε-chops (or just chop, for short). We ensuring that each ε-chop contains at most one terminal, and associated with the closest terminal (even if it contains no terminals). The adjacency graph of the ε-chops is also a rooted tree. For every

$i \in \{0, 1, \ldots, k-1\}$, we consider all ε-chops on levels j, for all $j \equiv i \mod k$, and construct a tree on the terminals associated with these ε-chops based on the topology of T. We obtain k trees on the terminals (for $i = 0, 1 \ldots, k$), and then show that they form the required tree cover, with stretch $3 + 12\varepsilon$. We continue with the details.

We assume that $\varepsilon \in (0, 1)$. For every point $x \in T$, let $C(x)$ be the terminal closest to x in the tree T, breaking ties *consistently*, that is, according to some universal linear ordering on the terminals. The closest terminal might not be in the subtree rooted at x. (This is in contrast to Gupta's construction [47] where $C(x)$ was defined to be the closest terminal in the subtree rooted at x.) Let $h(x) = d_T(x, C(x))$, which is the distance from x to its closest terminal. We choose two parameters:

$$k \text{ the smallest integer such that } (1-\varepsilon)^k \leq \varepsilon, \qquad \text{and } p = \lfloor 1/\varepsilon \rfloor + 1. \tag{2.2}$$

Here k will be the number of trees in the tree cover for stretch $3 + 12\varepsilon$ in Section 2.1, and $k \cdot p$ will be the number of trees for stretch $2 + \varepsilon$ in Section 2.2. Note that $k = \Theta(\varepsilon^{-1} \log \varepsilon^{-1})$. Let

$$h_\varepsilon(x) = \begin{cases} \varepsilon \cdot h(x) & \text{if } h(x) \geq (1-\varepsilon)^{k \cdot p} \\ \varepsilon & \text{otherwise.} \end{cases} \tag{2.3}$$

Note that $(1-\varepsilon)^{k \cdot p} \approx \varepsilon^{\Theta(1/\varepsilon)}$ for the choice of k and p in Equation (2.2). Roughly speaking, $h_\varepsilon(x)$ is about $\varepsilon h(x)$ unless when $h(x)$ is smaller than $\varepsilon^{\Theta(1/\varepsilon)}$.

> **ε-Chops.** Let $B_T(r_0, h_\varepsilon(r_0))$ be the ball of radius $h_\varepsilon(r_0)$ centered at r_0, which contains all points in T within distance at most $h_\varepsilon(r_0)$ from r_0. We then "chop" the tree T by removing $B_T(r_0, h_\varepsilon(r_0))$ from T. We call $B_T(r_0, h_\varepsilon(r_0))$ a 0-th level ε-chop rooted at r_0. Note that $B_T(r_0, h_\varepsilon(r_0))$ induces a connected subtree of T. Let $T_1, \ldots, T_a$ be the resulting subtrees of T after removing $B_T(r_0, h_\varepsilon(r_0))$, with roots $r_1 \ldots, r_a$, respectively (see Figure 2(a)). We then recursively chop each tree T_i by removing a ball of radius $h_\varepsilon(r_i)$, which is $B_{T_i}(r_i, h_\varepsilon(r_i))$, from the root r_i for every $i \in [a]$. Each tree $B_{T_i}(r_i, h_\varepsilon(r_i))$ is called a 1st level ε-chop rooted at r_i. We repeat the process for each remaining subtree until every point of T is chopped at some level. The result of the chopping process is recorded by $\mathbb{C} = \{\mathcal{L}_0, \mathcal{L}_1, \ldots\}$ where $\mathcal{L}_i$ contains ε-chops of T at level i.

Ideally, we want every chop at a root r to have radius $\varepsilon h(r)$. However, if we do so, we will never chop a terminal; as the chopping process gets closer to the terminal, the radius of the chop decreases. We remedy this by imposing a chop of radius ε whenever $h(r)$ becomes smaller than $(1-\varepsilon)^{pk}$. If $h_\varepsilon(r) = \varepsilon h(r)$, we call the chop a regular chop; otherwise, $h_\varepsilon(r) = \varepsilon$, and we call the chop a jump chop; see Figure 2(b). The jump chop is a technicality that we introduce to handle the case where internal nodes of the tree could be terminals, as alluded to above.

We can define an ancestor-descendant relationship between two ε-chops: We say that an ε-chop X is an ancestor of an ε-chop Y if the root of X is an ancestor of the root of Y. Furthermore, if $Y \neq X$, we say that X is a proper ancestor of Y.

OBSERVATION 2.1. *Let $\mathcal{L}_i$ and $\mathcal{L}_j$ be two collections of chops at levels $i < j$. Then, for every ε-chop $X \in \mathcal{L}_j$, there exists a unique ε-chop $Y \in \mathcal{L}_i$ such that X is a descendant of Y.*

CLAIM 2.2. *Let X be an ε-chop rooted at x. Then X contains at most one terminal, and X contains a terminal if and only if it is a jump chop. Furthermore, the terminal in X, if any, is the closest terminal to x in T and it is a vertex (of the discrete tree T) in X.*

Proof. If X is a regular chop rooted at x, then the radius of the chop is $\epsilon h(x) < h(x) \leq d(x, t)$ for any terminal t, and thus t cannot belong to X. If X is a jump chop and it contains a terminal t, then, for any other terminal t_0 in X, we have $d(t_0, x) \geq d(t, t_0) - d(x, t_0) \geq 8 - \epsilon$ by the triangle inequality, the definition of the jump chop and the fact that distances between vertices in the tree are at least 8. Thus, clearly, t_0 cannot belong to X and also t is the closest terminal to x. □

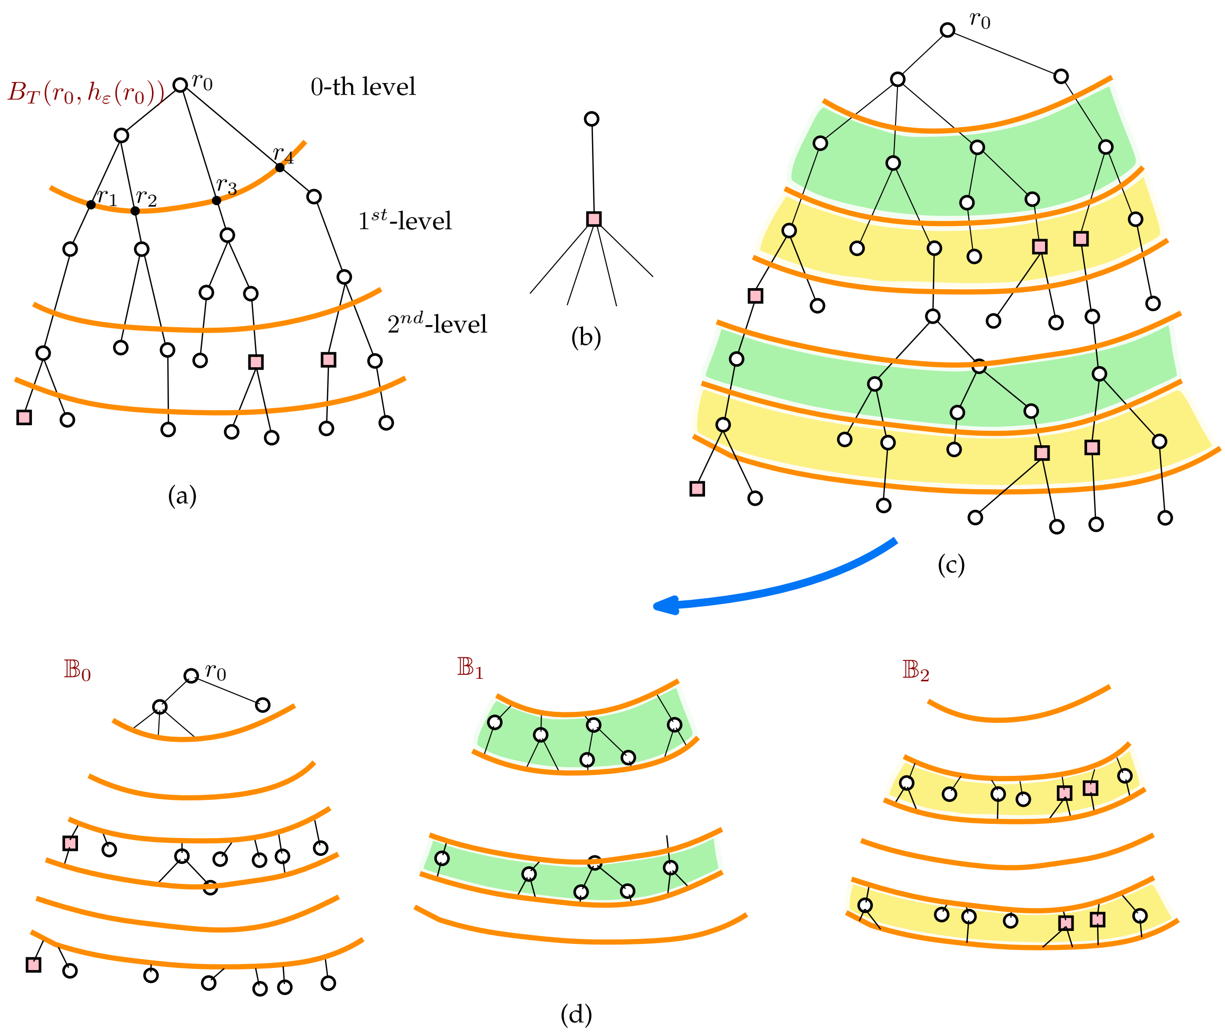

Figure 2: (a) ε-chops; (b) a jump chop; (c) buckets of chops $\mathbb{C}$ for $k = 3$ that is partitioned into 3 buckets $\mathbb{B}_0, \mathbb{B}_1, \mathbb{B}_2$ in (d). Terminals are marked with squares.

Cover construction. We partition the chops in $\mathbb{C}$ into k buckets $\mathbb{B}_0, \ldots \mathbb{B}_{k-1}$ where $\mathbb{B}_i$ contains chops at level i modulo k; see Figure 2(c) and (d). More precisely:

(2.4) $$\mathbb{B}_i = \{\mathcal{L}_j : j \equiv i \mod k\}.$$

For each bucket $\mathbb{B}_i$, we construct a tree T_i as follows:

Constructing T_i. We consider chops in $\mathbb{B}_i$ from lower levels to higher levels. For each ε-chop X in $\mathcal{L}_{i+sk}$ (at level $i + sk$), where s is an integer, and x is the root of X, we do the following. First, if $s = 0$, then we simply add the terminal $C(x)$ corresponding to the root of X. After this step, we have a forest where each component is an isolated vertex. Next, let $s \geq 1$. Let Y be the chop in $\mathcal{L}_{i+(s-1)k}$ that is the ancestor of X; Y exists by Observation 2.1. We then add the terminal $C(x)$ to T_i and connect $C(x)$ to $C(y)$ with an edge. (The terminal $C(y)$ corresponding to the root y of Y was added when the algorithm considered Y in the previous step.) At this point, T_i is a forest—see Figure 3(a) and (b)—where every tree is rooted at the terminal corresponding to trees at level-i chop. Finally, we can designate an (arbitrary) root t of a tree in T_i and make the roots of other trees children of t;

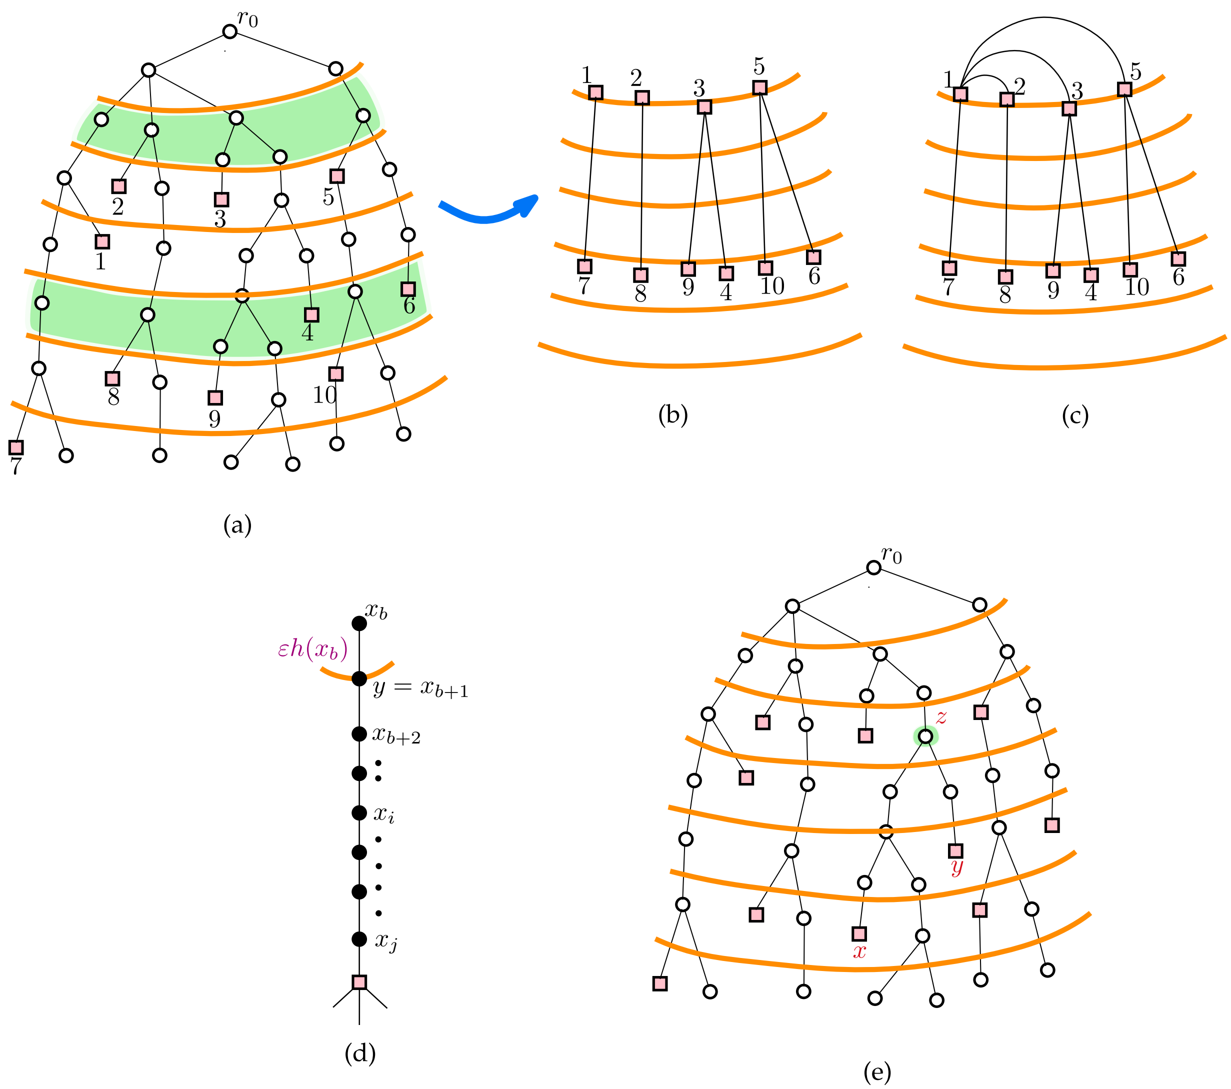

Figure 3: (a) The second bucket $\mathbb{B}_2$; (b) and (c) a tree constructed from $\mathbb{B}_2$; (d) all the ε-chops close enough to a jump chop are segments on the same edge; and (e) $z = \mathcal{LCA}(x, y)$ belongs to a chop at some level $i + sk$.

see Figure 3(c). Now, T_i is a tree that contains only terminals. We then set the weight of every edge (t_1, t_2) in T_i to be $d_T(t_1, t_2)$.

Our tree cover is $\mathcal{T} = \{T_0, T_1, \ldots, T_{k-1}\}$. We note that in this construction, one terminal could appear multiple times in T_i, but the copies of the same terminal will form a subtree of T_i (with edges of weight 0) since we break ties consistently. Thus, we could contract all of them into a single terminal. Here, we keep the copies separate to simplify the stretch analysis.

At this point, it might not be clear why each tree T_i contains all terminals since a terminal t might be chopped by some chop in another bucket $\mathbb{B}_j$ for $j \neq i$ and hence might not be present in any tree in the chops in $\mathbb{B}_i$. However, we will show below that in this case, t will be used to replace the root of some tree in $\mathbb{B}_i$, and hence will be present in T_i.

First, we observe that, as we approach a terminal t in the chopping process, the roots of ε-chops close to t will all be replaced by t.

LEMMA 2.3. *Let t be any terminal in T. Let $X_i, X_{i+1}, \ldots, X_j$ be a sequence of ε-chops at* consecutive levels *rooted at $x_i, x_{i+1}, \ldots, x_j$, respectively, such that: (i) X_j contains t, (ii) X_a is an ancestor of X_{a+1} for any $i \le a \le j-1$, and (ii) $j - i \le pk - 1$. Then $X_i, X_{i+1}, \ldots, X_j$ are (continuous) segments of the same edge (in the discrete T), and $C(x_i) = C(x_{i+1}) = \ldots = C(x_j) = t$.*

Proof. Let us first note that, for any terminal t_0 and $s \in [i, j-1]$ such that X_s is a regular chop, by the triangle inequality we have

$$d_T(x_{s+1}, t_0) \ge d_T(x_s, t_0) - \epsilon h(x_s) \ge (1-\varepsilon) d_T(x_s, t_0).$$

Since X_j contains t, then by Claim 2.2, X_j must be a jump chop and $C(x_j) = t$. Therefore, by the definition of $h_\varepsilon(x_j)$, we have $d_T(x_j, t) = d_T(x_j, C(x_j)) < (1-\varepsilon)^{kp}$. Assume that there is another jump chop among $X_i, \ldots, X_j$, and, moreover, $\ell \in [i, j-1]$ is the largest index of the jump chop. Then $X_{\ell+1}, \ldots, X_{j-1}$ are regular chops. Using the displayed inequality above for $s = j-1, j-2, \ldots, \ell+1$ with $t_0 := t$, we get that

$$d_T(x_{\ell+1}, t) \le (1-\varepsilon)^{\ell+1-j}(1-\varepsilon)^{kp} \le 1 - \epsilon.$$

Here we used that $j - \ell - 1 \le j - i \le pk - 1$. Thus, by the triangle inequality, $d_T(x_\ell, t) \le d_T(x_{\ell+1}, t) + \varepsilon \le 1$. From here, we see that, for any terminal $t_0 \ne t$ we have $d_T(x_\ell, t_0) \ge d_T(t, t_0) - d_T(x_\ell, t) \ge 8 - 1 = 7$ by the triangle inequality. That is, $C(x_\ell) = t$. On the other hand, t is contained in X_j and thus X_ℓ cannot contain t, implying that $d_T(x_\ell, t) > \varepsilon$. Thus, x_ℓ is too far from a terminal, which is a contradiction with the definition of a jump chop.

We conclude that all chops $X_i, \ldots, X_{j-1}$ are regular chops. By the same analysis as above, we have $d_T(x_s, t) \le 1$ and $C(x_s) = t$ for each $s \in [i, j]$. Finally, the vertices x_s, $s \in [i, j-1]$ must lie on the path from t to the root r_0, and, given that $d_T(x_i, t) \le 1$, must lie on the edge (u, t), where u is the parent of t. □

A direct corollary is that every terminal will appear in some tree T_i.

COROLLARY 2.4. *Let t be any terminal in T, and $\mathbb{B}_i$ be a bucket for some $i \in \{0, 1 \ldots, k\}$. Let $s \ge 0$ be the largest integer such that there exists an ε-chop X_{i+sk} at level $i + sk$ containing an ancestor of t, meaning that there exists a point in X_{i+sk} on the path from t to the root of T. Then $C(x_{i+sk}) = t$ where x_{i+sk} is the root of X_{i+sk}.*

Proof. Let $X_{i+sk}, X_{i+1+sk}, \ldots X_{j+sk}$ be the ε-chops at consecutive levels such that X_{j+sk} is the chop containing t, and one chop is the ancestor of the next in the sequence. Let x_{a+sk} be the root of X_{a+sk} for every $i \le a \le j$. By the choice of s, we have $|j - i| \le k - 1$ and hence $|j - i| \le pk - 1$ as $p \ge 1$. By Lemma 2.3, $C(x_{i+sk}) = C(x_{i+1+sk}) = \ldots = C(x_{j+sk}) = t$, as claimed. □

Stretch analysis. Let x, y be any two terminals and $z = \mathcal{LCA}(x, y)$ be the lowest common ancestor of x, y in T. There must be a chop $\mathcal{L}_{i+sk}$ at level $i + sk$ for some $i \in \{0, \ldots, k-1\}$ and $s \ge 0$ such that z belongs to the chop. Let r_{i+sk} be the root of the subtree in $\mathcal{L}_{i+sk}$ containing z. The next claim allows us to focus our attention to analyzing the distance from x to r_{i+sk} (and symmetrically, the distance from y to r_{i+sk}). See Figure 3(e).

CLAIM 2.5. *$d_T(x, y) \ge (1-\varepsilon) d_T(x, r_{i+sk}) + (1-\varepsilon) d_T(y, r_{i+sk}) - 2\varepsilon$.*

Proof. Observe that

$$\begin{aligned} d_T(x, y) &= d_T(x, z) + d_T(z, y) \\ &\ge d_T(x, r_{i+sk}) + d_T(r_{i+sk}, y) - 2 d_T(r_{i+sk}, z) \\ &\ge d_T(x, r_{i+sk}) + d_T(r_{i+sk}, y) - 2 h_\varepsilon(r_{i+sk}). \end{aligned}$$

By definition, $h_\varepsilon(r_{i+sk}) \le \varepsilon + \varepsilon \cdot h(r_{i+sk})$ and furthermore, $h(r_{i+sk}) = d_T(r_{i+sk}, C(r_{i+sk})) \le d_T(r_{i+sk}, x)$ since x is a terminal. The same argument yields $h(r_{i+sk}) \le d_T(r_{i+sk}, y)$. Combining these bounds with the equation above, we obtain

$$\begin{aligned} d_T(x, y) &\ge d_T(x, r_{i+sk}) + d_T(r_{i+sk}, y) - (\varepsilon + \varepsilon d_T(r_{i+sk}, x)) - (\varepsilon + \varepsilon d_T(r_{i+sk}, y)) \\ &\ge (1-\varepsilon) d_T(r_{i+sk}, x) + (1-\varepsilon) d_T(r_{i+sk}, y) - 2\varepsilon, \end{aligned}$$

as claimed. □

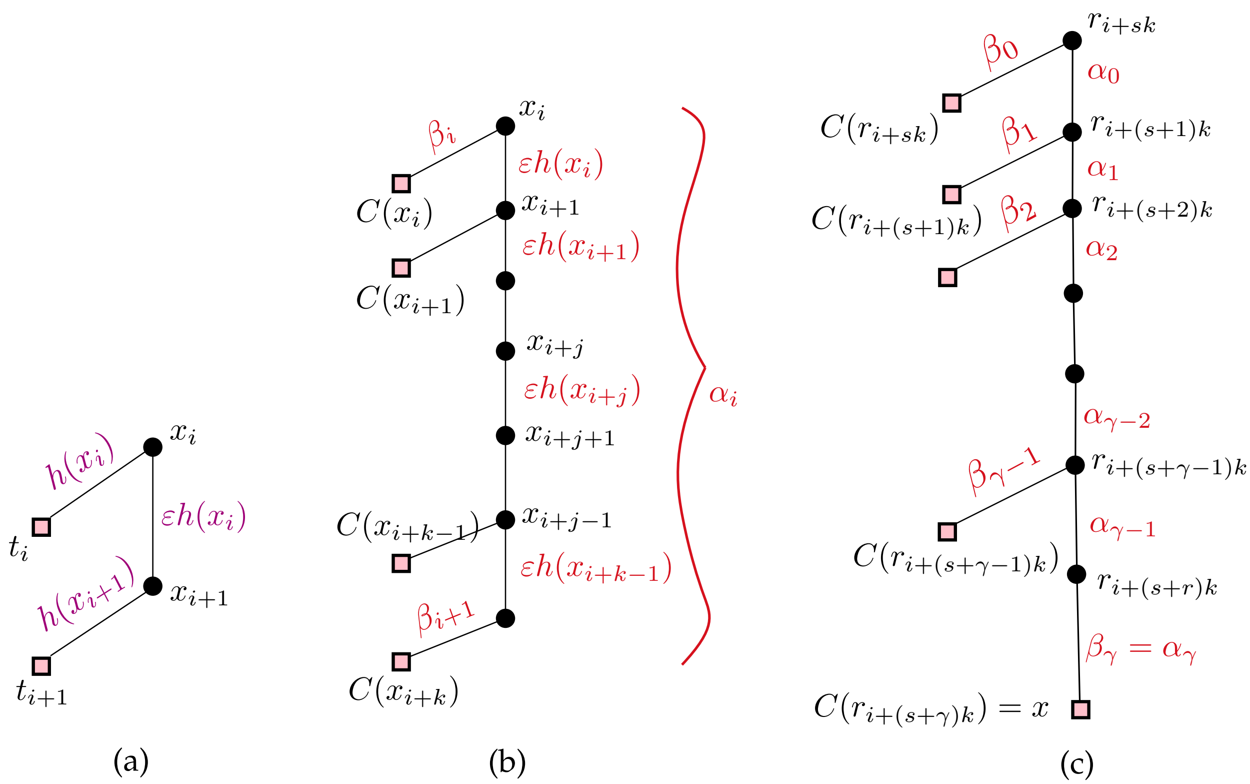

Figure 4: (a) illustration for Claim 2.6; (b) illustration for Lemma 2.7; (c) illustration for Lemma 2.8.

We remark that $d_T(x, y) \geq 8$ and so the 2ε term in Claim 2.5 is less than $\varepsilon d_T(x, y)$.
We have the following property of the regular chops.

CLAIM 2.6. *Let X_i be a regular chop rooted at x at some level i. Let X_{i+1} be a descendant chop of X_i rooted x_{i+1} at level $i+1$. Then $h(x_{i+1}) \geq (1-\varepsilon)h(x_i)$ and $h(x_{i+1}) \leq (1+\varepsilon)h(x_i)$.*

Proof. Let t_i and t_{i+1} be the closest terminal to x_i and x_{i+1}, respectively; see Figure 4(a). Note that $d_T(x_i, x_{i+1}) = \varepsilon h(x_i)$ since X_i is a regular chop. Then by definition of t_i, we have

$$\begin{aligned} h(x_i) &= d_T(x_i, t_i) \leq d_T(x_i, t_{i+1}) \\ &\leq d_T(x_i, x_{i+1}) + d_T(x_{i+1}, t_{i+1}) \\ &= \varepsilon h(x_i) + d_T(x_{i+1}, t_{i+1}) \\ &= \varepsilon h(x_i) + h(x_{i+1}) \ . \end{aligned}$$

For the second inequality, we have

$$h(x_{i+1}) \leq d_T(x_{i+1}, t_i) \leq d_T(x_{i+1}, x_i) + d_T(x_i, t_i) = (1+\varepsilon)h(x_i),$$

as claimed. □

LEMMA 2.7. *Let X_i be a chop at some level-i chop $\mathcal{L}_i$ with root x_i. Let X_{i+k} be a descendant chop of X_i at level-$(i+k)$ chop $\mathcal{L}_{i+k}$ with root x_{i+k}. Let $\beta_i = d_T(x_i, C(x_i))$ and $\alpha_i = d_T(x_i, x_{i+k})$ and $\beta_{i+1} = d_T(x_{i+k}, C(x_{i+k}))$. Then:*

1. *$\beta_i \leq (1+2\varepsilon)\alpha_i$ when $\varepsilon \leq 1/2$.*

2. $\beta_{i+1} \leq \alpha_i + \beta_i$.

Proof. See Figure 4(b). We observe that the second item follows from the triangle inequality:

$$\alpha_i + \beta_i \geq d_T(x_{i+k}, C(x_i)) \geq d_T(x_{i+k}, C(x_{i+k})) = \beta_{i+1}\ .$$

We focus on proving the first item. Let X_{i+j} for $1 \leq j \leq k-1$ be the descendant chops of X_i and ancestor chops of X_{i+k}, where X_{i+j} is at level $i+j$. Let x_{i+j} be the root of X_{i+j}. If any of the chops between X_i and X_{i+k}, excluding X_{i+k}, say X_{i+b} for some $b \in [0, k-1]$, is a jump chop, which contains a terminal t, then $C(x_{i+j}) = t$ for every $i \in [0, b]$ by Lemma 2.3. Furthermore, the path from x_i to x_{i+k} will go through t as X_{i+b} is the ancestor of X_{i+k}. This means $\beta_i \leq \alpha_i$, and item 1 holds.

It remains to consider the case where every chop between X_i and X_{i+k}, excluding X_{i+k}, is regular. This means that each X_{i+j} is formed by chopping the subtree rooted at x_{i+j} with radius $\varepsilon h(x_{i+j})$ for every $j \in [0, k-1]$. Thus, we have $\alpha_i = \varepsilon \sum_{j=0}^{k-1} h(x_{i+j})$. By Claim 2.6, $h(x_{i+j+1}) \geq (1-\varepsilon) h(x_{i+j})$. Consequently, we obtain

$$\begin{aligned}\alpha_i &\geq \varepsilon \left(h(x_i) + (1-\varepsilon)h(x_i) + \ldots + (1-\varepsilon)^{k-1} h(x_i)\right)\\ &\geq \varepsilon\beta_i \left(1 + (1-\varepsilon) + \ldots + (1-\varepsilon)^{k-1}\right)\\ &= \beta_i \left(1 - (1-\varepsilon)^k\right) \geq \beta_i(1-\varepsilon),\end{aligned}$$

using the choice of k in Equation (2.2). This gives $\beta_i \leq \alpha_i/(1-\varepsilon) \leq (1+2\varepsilon)\alpha_i$ when $\varepsilon \leq \frac{1}{2}$. ∎

With Lemma 2.7 at hand, we are ready to bound $d_{T_i}(r_{i+sk}, x)$.

LEMMA 2.8. $d_{T_i}(C(r_{i+sk}), x) \leq (3+4\varepsilon) d_T(r_{i+sk}, x)$. *Similarly,* $d_{T_i}(C(r_{i+sk}), y) \leq (3+4\varepsilon) d_T(r_{i+sk}, y)$.

Proof. Let $\{r_{i+(s+1)k}, r_{i+(s+2)k}, \ldots, r_{i+(s+\gamma)k}\}$ be the set of all the roots of the chops in $\mathbb{B}_i$ that are on the path from r_{i+sk} to x. See Figure 4(c). Note that if $\gamma = 0$, then by Corollary 2.4, $C(r_{i+sk}) = x$ and hence $d_{T_i}(C(r_{i+sk}), x) = 0$, so the lemma trivially holds. We now assume that $\gamma \geq 1$.

Corollary 2.4 yields $C(r_{i+(s+\gamma)k}) = x$. By construction, we have

$$\begin{aligned}d_{T_i}(x, C(r_{i+sk})) &= d_{T_i}(C(r_{i+(s+1)k}), C(r_{i+sk})) + \ldots + d_{T_i}(C(r_{i+(s+\gamma)k}), C(r_{i+(s+\gamma-1)k}))\\ &= d_T(C(r_{i+(s+1)k}), C(r_{i+sk})) + \ldots + d_T(C(r_{i+(s+\gamma)k}), C(r_{i+(s+\gamma-1)k})).\end{aligned} \tag{2.5}$$

Let $\beta_a = d_T(r_{i+(s+a)k}, C(r_{i+(s+a)k}))$ for every $0 \leq a \leq \gamma$ and $\alpha_a = d_T(r_{i+(s+a)k}, r_{i+(s+a+1)k})$ for every $0 \leq a \leq \gamma - 1$. Let $\alpha_\gamma = d_T(r_{i+(s+\gamma)k}, x)$; and note that $\alpha_\gamma = \beta_\gamma$ by Corollary 2.4. Now Lemma 2.7 and the fact that $\alpha_\gamma = \beta_\gamma$ imply

$$\beta_a \leq (1+2\varepsilon)\alpha_a \quad \forall\, 0 \leq a \leq \gamma. \tag{2.6}$$

Combing Equation (2.5) and the triangle inequality, we obtain

$$\begin{aligned}d_{T_i}(x, C(r_{i+sk})) &\leq (\beta_0 + \alpha_0 + \beta_1) + (\beta_1 + \alpha_1 + \beta_2) + \ldots + (\beta_{\gamma-1} + \alpha_{\gamma-1} + \beta_\gamma)\\ &= \sum_{a=0}^{\gamma-1} \alpha_a + 2\sum_{a=1}^{\gamma-1} \beta_a + \beta_\gamma + \beta_0\\ &\leq \sum_{a=0}^{\gamma} \alpha_a + 2\sum_{a=0}^{\gamma} \beta_a\\ &\leq (3+4\varepsilon)\sum_{a=0}^{\gamma} \alpha_a \qquad \text{(by Equation (2.6))}\\ &= (3+4\varepsilon) d_T(x, r_{i+sk}),\end{aligned} \tag{2.7}$$

as desired. ∎

We can now bound the stretch of $\mathcal{T}$.

LEMMA 2.9. *The stretch of $\mathcal{T}$ is at most $3 + 12\varepsilon$ when $\varepsilon \in (0, \frac{1}{8})$.*

Proof. We continue analyzing the stretch between two terminals x and y. By Claim 2.5 and Lemma 2.8, we have

$$\begin{aligned} d_T(x,y) &\geq \frac{(1-\varepsilon)}{3+4\varepsilon}\left(d_{T_i}(x, C(r_{i+sk})) + d_{T_i}(y, C(r_{i+sk}))\right) - 2\varepsilon \\ &\geq \frac{(1-\varepsilon)}{3+4\varepsilon} d_{T_i}(x,y) - 2\varepsilon \\ &\geq \frac{(1-\varepsilon)}{3+4\varepsilon} d_{T_i}(x,y) - \varepsilon d_T(x,y) \quad \text{(as } d_T(x,y) \geq 8\text{)}. \end{aligned}$$

This gives

$$d_{T_i}(x,y) \leq \frac{(1+\varepsilon)(3+4\varepsilon)}{1-\varepsilon} d_T(x,y) \leq (3+12\varepsilon) d_T(x,y)$$

when $\varepsilon \leq 1/8$. ∎

2.2 Stretch $2+\varepsilon$ The construction in the previous section has stretch at most $3+\varepsilon$ due to the sum $2\sum_{a=0}^{t}\beta_a$, called β-sum, in Equation (2.7); each β_a is approximately α_a. The key idea is to reduce the contribution of the β-sum by doing an even more fine-grained bucketing of each $\mathbb{B}_i$ modulo p for $p \approx 1/\varepsilon$ in Equation (2.2). (So far, we have not really used p.) Importantly, the fined-grained bucketing allows us to bound $2\sum_{a=0}^{t}\beta_a \approx \sum_{a=0}^{t}\alpha_a + \frac{1}{p}(\sum_{a=0}^{t}\beta_a)$, which is approximately $(1+\varepsilon)\sum_{a=0}^{t}\alpha_a$ for $p \approx 1/\varepsilon$. Plugging this into Equation (2.7), we obtain an improved version of Lemma 2.8, where the stretch is $2 + O(\varepsilon)$. This ultimately leads to stretch $2 + O(\varepsilon)$ in our non-Steiner tree cover.

Cover Construction. Recall that $p = \lfloor 1/\varepsilon \rfloor + 1$. We construct k buckets of chops $\mathbb{B}_0, \ldots, \mathbb{B}_{k-1}$ as described in Equation (2.4). We then further partition each bucket $\mathbb{B}_i$ for every $i \in [0, k-1]$ into p buckets $\mathbb{B}_{i,0}, \mathbb{B}_{i,1}, \ldots, \mathbb{B}_{i,p-1}$ as follows. For each $j \in [0, p-1]$, let

$$\mathbb{B}_{i,j} = \{\mathcal{L}_{i+(j+sp)\cdot k} : s \geq 0\}. \tag{2.8}$$

We then obtain a set of $k \cdot p$ buckets of the form $\mathbb{B}_{i,j}$ where $i \in [0, k-1]$ and $j = [0, p-1]$. For each bucket $\mathbb{B}_{i,j}$, we construct a tree $T_{i,j}$ in the same way we construct T_i in the previous section. Our final tree cover is $\mathcal{T} = \{T_{i,j} : i \in [0, k-1], j \in [0, p-1]\}$.

> **Constructing $T_{i,j}$.** Consider the ε-chops in $\mathbb{B}_{i,j}$ from lower levels to higher levels. Let X be an ε-chop at level $i + (j + sp)k$ for some integers $s \geq 0, j \geq 0$; assume for now that $s \geq 1$. Let x be the root of X. Let Y the ε-chop at level $i + (j + (s-1)p)k$ that is the child of X in $\mathbb{B}_{i,j}$. We then add terminal $C(x)$ to $T_{i,j}$, and connect $C(x)$ to $C(y)$ with an edge. If $s = 0$, then we simply add a terminal $C(x)$ corresponding to the root of every tree X in the chop at level $i + j$. Finally, we connect the roots corresponding to $s = 0$ to form a tree. We set of the weight of every edge (t_1, t_2) in $T_{i.j}$ to be $d_T(t_1, t_2)$.

One can think of the construction of $T_{i,j}$ as skipping connections by exactly p consecutive levels in the tree T_i constructed in the previous section; see Figure 5. (Again, one terminal could have multiple copies in a tree of the cover, which could then be resolved by contraction.) By Lemma 2.3, every terminal will appear in $T_{i,j}$. We now focus on the stretch analysis.

Stretch Analysis. We consider any two terminals x, y and $z = \mathcal{LCA}(x,y)$. There must be some chop $\mathcal{L}_{i+(j+sp)k}$ such that z belongs to the chop for some $s \geq 0$. To simplify the presentation, w.l.o.g., we assume that $i = 0$, $j = 0$, and $s = 0$. Thus, the chop containing z is $\mathcal{L}_0$. Let r_0 be the root of the subtree containing z. Claim 2.5 carries over for this setting, with r_0 in place of r_{i+sk}.

CLAIM 2.10. *$d_T(x,y) \geq (1-\varepsilon)d_T(x,r_0) + (1-\varepsilon)d_T(y,r_0) - 2\varepsilon$.*

We now prove an analogue of Lemma 2.8.

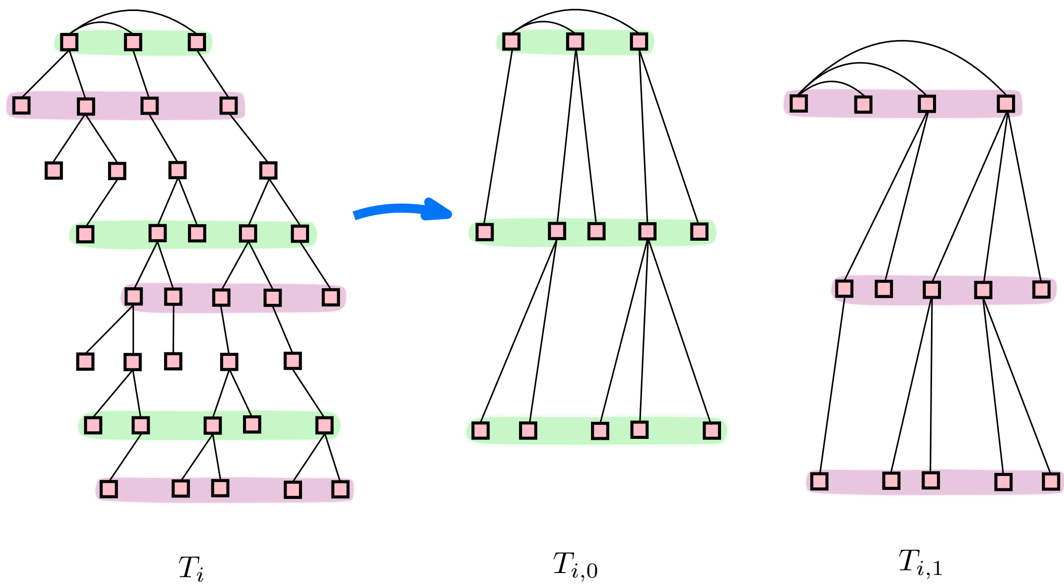

Figure 5: Construct two trees $T_{i,0}$ and $T_{i,1}$ from T_i by skipping 3 levels; here $p = 3$.

LEMMA 2.11. *$d_{T_{0,0}}(C(r_0), x) \le (2 + 2\varepsilon)d_T(r_0, x)$. Similarly, $d_{T_{0,0}}(C(r_0), y) \le (2 + 2\varepsilon)d_T(r_0, y)$.*

Proof. We focus on $d_{T_{0,0}}(C(r_0), x)$. Let $r_{\gamma pk}$, for some integer $\gamma \ge 0$, be the root of a tree in a chop in $\mathbb{B}_{0,0}$ closest to x such that $r_{\gamma pk}$ is on the path from r_0 to x; see Figure 6(b). Let $r_k, r_{2k}, \ldots, r_{\gamma pk-1}$ be all the roots of consecutive chops in $\mathbb{B}_0$ (not in $\mathbb{B}_{0,0}$) that are on the path from r_0 to $r_{\gamma pk}$; see Figure 6(a). We note that $C(r_{\gamma pk}) = x$ and $C(r_0), C(r_{pk}), C(r_{2pk}), \ldots, C(r_{\gamma pk})$ is a path from $C(r_0)$ to x in $T_{0,0}$. If $\gamma = 0$, then $d_{T_{0,0}}(x, C(r_0)) = 0$ and the lemma trivially holds. It remains to consider the case $\gamma \ge 1$.

For every $a \in [0, \gamma \cdot p]$, let $\beta_a = d_T(r_{ak}, C(r_{ak}))$ and $\alpha_a = d_T(r_{ak}, r_{(a+1)k})$, with $r_{(\gamma p+1)k}$ defined to be x so that $d_T(r_{\gamma pk}, r_{(\gamma p+1)k})$ is well-defined. Note that $\alpha_{\gamma p} = \beta_{\gamma p}$. By Lemma 2.7, we have

$$\begin{aligned} \beta_a &\le (1 + 2\varepsilon)\alpha_a \qquad \forall a \in [0, \gamma p], \\ \beta_{a+1} &\le \alpha_a + \beta_a \qquad \forall a \in [0, \gamma p - 1]. \end{aligned} \tag{2.9}$$

Next, we observe that

$$d_{T_{0,0}}(C(r_0), x) \le \sum_{a=0}^{\gamma p-1} \alpha_a + \beta_0 + 2\sum_{\ell=1}^{\gamma-1} \beta_{\ell p} \qquad \text{and} \qquad \sum_{a=0}^{\gamma p} \alpha_a = d_T(r_0, x). \tag{2.10}$$

For every $\ell \in [1, \gamma]$, we define $\bar{\ell} \in [1, p-1]$ such that $\beta_{\bar{\ell}+(\ell-1)p} = \min\{\beta_{1+(\ell-1)p}, \beta_{2+(\ell-1)p}, \ldots, \beta_{\ell p-1}\}$. See Figure 6(c). We make the following two claims:

CLAIM 2.12. $\beta_{\ell p} \le \beta_{\bar{\ell}+(\ell-1)p} + \sum_{a=\bar{\ell}+(\ell-1)p}^{\ell p-1} \alpha_a$

By applying Equation (2.9) repeatedly, we have

$$\begin{aligned} \beta_{\ell p} &\le \beta_{\ell p-1} + \alpha_{\ell p-1} \\ &\le \beta_{\ell p-2} + \alpha_{\ell p-2} + \alpha_{\ell p-1} \\ &\le \beta_{\bar{\ell}+(\ell-1)p} + \sum_{a=\bar{\ell}+(\ell-1)p}^{\ell p-1} \alpha_a, \end{aligned}$$

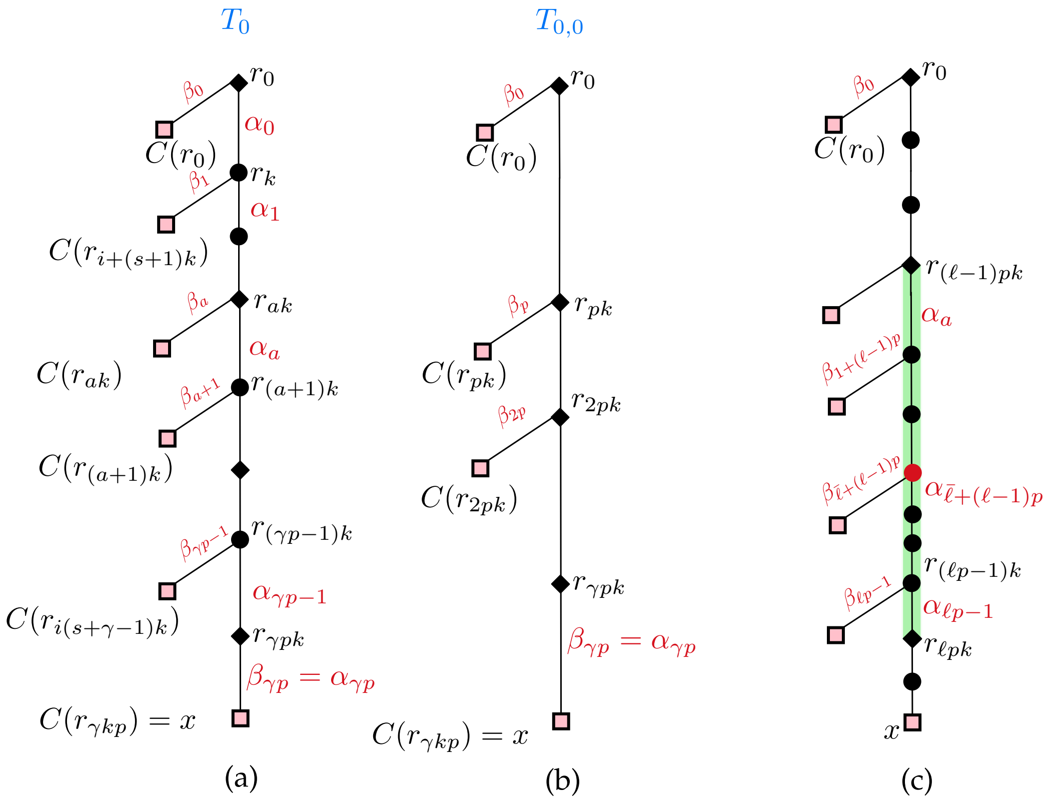

Figure 6: (a) In T_0, each root r_{ak} will be replaced by the corresponding closest terminal; (b) the roots of the ε-chops in the construction $T_{0,0}$, each root will also be replaced with a corresponding closest interval; (c) the root corresponding to $\bar{\ell}$ is highlighted red.

implying Claim 2.12.

CLAIM 2.13. $\beta_{\bar{\ell}+(\ell-1)p} \leq \frac{1+\varepsilon}{p-1}\sum_{a=1+(\ell-1)p}^{\ell p-1} \alpha_a$.

By the definition of $\bar{\ell}$, we have

$$\beta_{\bar{\ell}+(\ell-1)p} \leq \frac{1}{p-1}\sum_{a=1+(\ell-1)p}^{\ell p-1} \beta_a \leq \frac{1+\varepsilon}{p-1}\sum_{a=1+(\ell-1)p}^{\ell p-1} \alpha_a$$

by Equation (2.9), implying Claim 2.13.

We continue with bounding $d_{T_{0,0}}(x, C(r_0))$. By Claim 2.12, we have

$$\begin{aligned}
\sum_{\ell=1}^{\gamma} \beta_{\ell p} &\leq \sum_{\ell=1}^{\gamma}\left(\beta_{\bar{\ell}+(\ell-1)p} + \sum_{a=\bar{\ell}+(\ell-1)p}^{\ell p-1} \alpha_a\right) \\
&\leq \sum_{\ell=1}^{\gamma}\left(\beta_{\bar{\ell}+(\ell-1)p} + \sum_{a=1+(\ell-1)p}^{\ell p-1} \alpha_a\right) \qquad \text{(since } \bar{\ell} \geq 1\text{)} \\
&\leq \sum_{\ell=1}^{\gamma}\sum_{a=1+(\ell-1)p}^{\ell p-1}\left(1+\frac{1+\varepsilon}{p-1}\right)\alpha_a \qquad \text{(by Claim 2.13)}
\end{aligned} \tag{2.11}$$

By Equation (2.10), we have

$$\begin{aligned} d_{T_{0,0}}(C(r_0), x) &\leq \sum_{a=0}^{\gamma p-1} \alpha_a + \sum_{\ell=0}^{\gamma} \beta_{\ell p} + \sum_{\ell=1}^{\gamma} \beta_{\ell p} \\ &\leq \sum_{a=0}^{\gamma p-1} \alpha_a + \sum_{\ell=0}^{\gamma} (1+2\varepsilon)\alpha_{\ell p} + \sum_{\ell=1}^{\gamma} \beta_{\ell p} \qquad \text{(by Equation (2.9))} \\ &\leq \sum_{a=0}^{\gamma p-1} \alpha_a + \sum_{\ell=0}^{\gamma} (1+2\varepsilon)\alpha_{\ell p} + \sum_{\ell=1}^{\gamma} \sum_{a=1+(\ell-1)p}^{\ell p-1} \left(1 + \frac{1+2\varepsilon}{p-1}\right) \alpha_a \qquad \text{(by Equation (2.11))} \\ &\leq \sum_{a=0}^{\gamma p-1} \alpha_a + (1+2\varepsilon) \left(\sum_{\ell=0}^{\gamma} \alpha_{\ell p} + \sum_{\ell=1}^{\gamma} \sum_{a=1+(\ell-1)p}^{\ell p-1} \alpha_a \right) \qquad \text{(since } p = \lfloor 1/\varepsilon \rfloor + 1 \text{ and } \varepsilon \leq 1) \\ &\leq (2+2\varepsilon) \sum_{a=0}^{\gamma p} \alpha_a \\ &\leq (2+2\varepsilon) d_T(r_0, x) \qquad \text{(by Equation (2.10))} \end{aligned}$$

as desired. □

Now we can bound the stretch of $\mathcal{T}$.

LEMMA 2.14. *The stretch of $\mathcal{T}$ is at most $2 + 10\varepsilon$ when $\varepsilon \in (0, \frac{1}{8})$.*

Proof. By Claim 2.10 and Lemma 2.11, we have

$$\begin{aligned} d_T(x, y) &\geq \frac{(1-\varepsilon)}{2+2\varepsilon} \left(d_{T_{0,0}}(x, C(r_0)) + d_{T_{0,0}}(y, C(r_0)) \right) - 2\varepsilon \\ &\geq \frac{(1-\varepsilon)}{2+2\varepsilon} d_{T_{0,0}}(x, y) - 2\varepsilon \\ &\geq \frac{(1-\varepsilon)}{2+2\varepsilon} d_{T_{0,0}}(x, y) - \varepsilon d_T(x, y) \quad \text{(as } d_T(x, y) \geq 8). \end{aligned}$$

This gives

$$d_{T_{0,0}}(x, y) \leq \frac{(1+\varepsilon)(2+3\varepsilon)}{1-\varepsilon} d_T(x, y) \leq (2+10\varepsilon) d_T(x, y)$$

when $\varepsilon \leq 1/8$. □

2.3 Stretch 2 In this subsection, we construct a tree cover of stretch 2 that has $O(\log n)$ trees. We say that a vertex v is a centroid of a tree T if every connected component of $T \setminus \{v\}$ has at most $n/2$ vertices. Initially, $\mathcal{T} = \emptyset$. For each vertex u, we denote by $C(u)$ the closest terminal to u. (If u is a terminal, then $C(u) = u$.) Our construction is recursive.

1. **Step 1.** If T contains a single terminal (the base case), then we simply return a singleton tree. Otherwise, we find a centroid v of T. Then we make a star X_v with center $C(v)$, and for every terminal $u \in K \setminus \{C(v)\}$, we add an edge $(C(v), u)$ to X_v and set the weight $w_{X_v}(C(v), u) = d_T(u, C(v))$. Then we add X_v to $\mathcal{T}$.

2. **Step 2.** Let $\bar{T}_1, \bar{T}_2, \ldots, \bar{T}_\kappa$ be all the connected components of $T \setminus \{v\}$. We recursively construct a non-Steiner tree cover $\bar{\mathcal{T}}_j$ for each component $\bar{T}_j$, where $j \in [\kappa]$. Let $s = \max_{j \in [\kappa]} |\bar{\mathcal{T}}_j|$. By making duplicate copies if necessary, we assume that every cover $\bar{\mathcal{T}}_j$ contains exactly s trees, denoted by $\{Y_1^j, Y_2^j, \ldots, Y_s^j\}$. Then, we create s trees $\{Z_1, Z_2 \ldots, Z_s\}$: for each $a \in [s]$, the a-th tree Z_a is formed by taking all the a-th trees $Y_a^1, Y_a^2, \ldots, Y_a^\kappa$, one from each tree cover; we then connect Y_a^j to Y_a^1 for every $j \in \{2, \ldots, \kappa\}$ by adding an edge (t_j, t_1) from an (arbitrary) terminal $t_j \in Y_a^j$ to an arbitrary terminal $t_1 \in Y_a^1$. By adding all the edges (t_j, t_1), we effectively connect every Y_a^j to Y_a^1, and finally get a tree Z_a. The weight of the edge (t_j, t_1) is $w_{Z_a}(t_j, t_1) = d_T(t_j, t_1)$. We then add all the trees $Z_1, \ldots, Z_s$ to $\mathcal{T}$.

It follows directly from the construction that every tree in $\mathcal{T}$ is non-Steiner. Furthermore, for every tree $X \in \mathcal{T}$, every edge $(x, y) \in X$ has a weight $w_X(x, y) = d_T(x, y)$. Thus, by the triangle inequality, X is dominating. It remains to bound the number of trees, as well as the stretch of $\mathcal{T}$.

Bounding the Number of Trees in $\mathcal{T}$. Let $s(n)$ be the number of trees in $\mathcal{T}$ when applying the above algorithm to a tree T with n vertices. Then we have

$$s(n) \leq s(n/2) + 1,$$

where the $+1$ term is due to the tree X_v in Step 1, and $s(n/2)$ is an upper bound for the size of all the covers $\bar{\mathcal{T}}_1, \bar{\mathcal{T}}_2, \ldots, \bar{\mathcal{T}}_\kappa$, as each connected component of $T \setminus \{v\}$ has at most $n/2$ vertices. Solving the above recurrence gives $s(n) = O(\log n)$.

Analyzing the Stretch. Let x and y be any two terminals in T. If x and y are in the same component of $T \setminus \{v\}$, then we get stretch 2 by induction. If x and y are in different components of $T \setminus \{v\}$, then we have

$$\begin{aligned} d_{X_v}(x, y) &= w_{X_v}(x, C(v)) + w_{X_v}(y, C(v)) \\ &= d_T(x, C(v)) + d_T(y, C(v)) \\ &\leq d_T(x, v) + d_T(v, C(v)) + d_T(y, v) + d_T(v, C(v)) \quad \text{(by triangle inequality)} \\ &= d_T(x, y) + 2d_T(v, C(v)) \\ &\leq d_T(x, y) + d_T(v, x) + d_T(v, y) \quad \text{(by definition, } d_T(v, C(v)) \leq \min\{d_T(v, x), d_T(v, y)\}) \\ &= 2d_T(x, y), \end{aligned}$$

implying that the stretch is at most 2 for the pair x, y.

3 Steiner Spanners for Terminals in Planar Graphs

Recall that a metric space (X, d_X) is *planar* if there exists an edge-weighted planar graph $G = (V, E, w)$ such that $X \subseteq V$ and d_X is the shorted-path metric of G restricted to X.

THEOREM 3.1. *Let $\varepsilon \in (0, 1)$ be a parameter. Let T be a set of n points* (terminals*) in a planar metric. We can construct a Steiner $(1 + \varepsilon)$-spanner for T with $O((n/\varepsilon) \cdot \max\{\alpha(n), \log \varepsilon^{-1}\} \cdot \log \varepsilon^{-1})$ edges, where $\alpha(n)$ is the inverse Ackermann function.*

In this section, we prove Theorem 3.1. We start in Section 3.1 with reviewing a classical spanner construction based on a net trees, in planar metrics [61]. The problem of constructing $(1 + \varepsilon)$-spanners is reduced to additive spanners [61, 27] on each level of a net tree in Section 3.1, and further to additive spanners in each subgraph in a cover decomposition in Section 3.2. Finally in Section 3.3, we construct the required additive spanners for bounded-diameter planar graphs using shortest path separators [43, 54, 68] and tree shortcutting [5, 20, 30].

3.1 Net Trees Based Spanners Let $G = (V, E, w)$ be an edge-weighted planar graph, and let $T \subset V$ be a set of n *terminals*. In this section, we work with the planar metric (T, d_G), where d_G is the shortest path distance between the terminals in G. The aspect ratio of the metric is the ratio of the maximum to the minimum distance between distinct vertices: $\rho = \max_{x,y \in T} d_G(x, y) / \min_{x,y \in T, x \neq y} d_G(x, y)$. Without loss of generality, we may assume that minimum distance between distinct terminals is 1, i.e., $\min_{x,y \in T, x \neq y} d_G(x, y) = 1$. In particular, the maximum distance between terminals is ${}_G(V) = \rho$.

An δ-net in a metric space (X, d) is a subset $N \subset X$ such that for every $x \in X$ there exists $y \in N$ such that $d(x, y) \leq \delta$ (i.e., the closed balls of radius δ centered in N cover X) and $\min_{x,y \in N, x \neq y} d(x, y) \geq \delta$ (i.e., the open balls of radius $\delta/2$ centered in N are pairwise disjoint).

For a given $\varepsilon \in (0, \frac{1}{4})$, we construct a hierarchy of nets on the terminals $T = N_0 \supseteq N_1 \supseteq \ldots \supseteq N_{\lceil \log_2 \rho \rceil}$, where N_i is a 2^i-net. This hierarchy induces a net-tree $\mathcal{T}$, where level i of the tree is the net N_i, and the parent of a vertex $v \in N_i$ is $v \in N_{i+1}$ (if $v \in N_{i-1}$) or another vertex $u \in N_{i+1}$ such that $d_G(u, v) \leq 2^{i+1}$. Every vertex $v \in T$ has a unique ancestor in the net N_i, that we denote by $v^{(i)} \in N_i$. Using geometric series, we obtain

$$d_G(v, v^{(i)}) \leq \sum_{j=1}^{i} d_G(v^{(j-1)}, v^{(j)}) \leq \sum_{j=1}^{i} 2^j < 2^{i+1}. \tag{3.12}$$

We construct a spanner $H_{\rm net}$ for the planar metric (T, d_G) as follows: At level $i = 0$, we have $N_0 = T$ and we connect every terminal $u \in T$ to all other terminal $v \in T$ such that $d_G(u, v) \le \frac{18}{\varepsilon}$. For every level $i \in \{1, 2, \ldots, \lceil \log_2 \rho \rceil\}$, let $\Delta_i = 2^i/\varepsilon$; and connect every terminal $u \in N_i$ to all other terminal $v \in N_i$ such that

$$8\,\Delta_i \le d_G(u, v) \le 19\,\Delta_i.$$

Using a standard proof by induction, we show that $H_{\rm net}$ is a $(1+\varepsilon)$-spanner on T.

LEMMA 3.2. *For $\varepsilon \in (0, \frac{1}{4})$, the graph $H_{\rm net}$ is a $(1+\varepsilon)$-spanner for the metric (T, d_G).*

Proof. Consider the $\binom{n}{2}$ point pairs $\{x, y\} \subset T$ sorted in nondecreasing order by distance $d_G(x, y)$. Let $\{x_j, y_j\}$ denote the j-th pair. We prove, by induction on j, that $d_{H_{\rm net}}(x_j, y_j) \le (1+\varepsilon) d_G(x_j, y_j)$.

In the base case, $\{x_1, y_1\}$ is a closest pair in G, and we have $d_G(x_1, y_1) = 1$ by assumption. The edge xy was added to $H_{\rm net}$ at level 0, and so $d_{H_{\rm net}}(x_1, y_1) = d_G(x_1, y_1)$. In fact, the same argument holds for all pairs $\{x_j, y_j\} \subset V$ with $d_G(x_j, y_j) \le \frac{18}{\varepsilon}$.

For the induction step, consider a pair $\{x_j, y_j\} \subset V$ with $d_G(x_j, y_j) > \frac{18}{\varepsilon}$, and assume that $d_H(x, y) \le (1+\varepsilon) d_G(x, y)$ for all pairs $\{x, y\} \subset V$ such that $d_G(x, y) < d_G(x_j, y_j)$. Then x_j and y_j are not adjacent at level 0. Since $d_G(x_j, y_j) \le_G (V) \le \rho$, there exists $i \in \{1, 2, \ldots, \lceil \log_2 \Delta \rceil\}$ such that

$$9\,\Delta_i < d_G(x_j, y_j) \le 18\,\Delta_i.$$

Considering the ancestors of x_j and y_j at level i, Equation (3.12) and the triangle inequality yield

$$d_G(x_j^{(i)}, y_j^{(i)}) \le d_G(x_j^{(i)}, x_j) + d_G(x_j, y_j) + d_G(y_j, y_j^{(i)}) \le 18\,\Delta_i + 2 \cdot 2^{i+1} \le (4 + 4\varepsilon)\,\Delta_i \le 19\,\Delta_i,$$
$$d_G(x_j^{(i)}, y_j^{(i)}) \ge d_G(x_j, y_j) - d_G(x_j^{(i)}, x_j) - d_G(y_j, y_j^{(i)}) \ge 9\,\Delta_i - 2 \cdot 2^{i+1} \ge (2 - 4\varepsilon)\,\Delta_i \ge 8\,\Delta_i,$$

for $\varepsilon \in (0, \frac{1}{4})$. Consequently, the edge $x_j^{(i)} y_j^{(i)}$ has been added to $H_{\rm net}$ at level i.

By the induction hypothesis, $H_{\rm net}$ contains paths $\pi(x_j, x_j^{(i)})$ and $\pi(x_j, x_j^{(i)})$ of length at most $(1+\varepsilon)2^{i+1}$. Concatenate the path $\pi(x_j, x_j^{(i)})$, the edge $x_j^{(i)} y_j^{(i)}$, and the path $\pi(y_j, y_j^{(i)})$. We obtain a path in $H_{\rm net}$ between x_j and y_j of length

$$\begin{aligned}
w(\pi(x_j, x_j^{(i)})) + d_G(x_j^{(i)}, y_j^{(i)}) + w(\pi(y_j, y_j^{(i)})) &\le d_G(x_j^{(i)}, y_j^{(i)}) + 2(1+\varepsilon)2^{i+1} \\
&\le d_G(x_j^{(i)}, x_j) + d_G(x_j, y_j) + d_G(y_j, y_j^{(i)}) + 4(1+\varepsilon)2^i \\
&\le d_G(x_j, y_j) + 2 \cdot 2^{i+1} + 4(1+\varepsilon)2^i \\
&= d_G(x_j, y_j) + (8 + 4\varepsilon)2^i \\
&\le d_G(x_j, y_j) + (8 + 4\varepsilon)\varepsilon\Delta_i \\
&\le d_G(x_j, y_j) + (8 + 4\varepsilon)\varepsilon \cdot \frac{1}{9}\, d_G(x_j, y_j) \\
&= \left(1 + \frac{8 + 4\varepsilon}{9} \cdot \varepsilon\right) \cdot d_G(x_j, y_j) \\
&\le (1+\varepsilon) \cdot d_G(x_j, y_j),
\end{aligned}$$

as required. This completes the induction step, hence the entire proof. ◻

Reduction to Additive Spanners in Net-Trees We reduce the problem to additive spanners. Specifically, we show that Lemma 3.3 below implies Theorem 3.1. The proof of Lemma 3.3 is presented in Sections 3.2 and 3.3

LEMMA 3.3. *For every $i \in \mathbb{N}$, there exists a spanner H_i on N_i such that*

1. *for all $x, y \in N_i$, if $d_G(x, y) = \Theta(\Delta_i)$, then $d_{H_i}(x, y) \le d_G(x, y) + \varepsilon\Delta_i$, and*
2. $|E(H_i)| \le O\Big(|N_i|\, \varepsilon^{-1} \cdot \log(\varepsilon^{-1}\alpha(n))\Big)$,

where $\alpha(.)$ denotes the inverse Ackermann function.

Spanner construction. Let H_i be the additive spanners provided by Lemma 3.3 for each level N_i of a net tree $\mathcal{T}$; and let $H = \bigcup_{i\in\mathbb{N}} H_i$.

LEMMA 3.4. *For $\varepsilon \in (0, \frac{1}{4})$, the graph H is a $(1+2\varepsilon)$-spanner for the metric (T, d_G).*

Proof. Let $x, y \in T$ be a pair of terminals in the edge-weighted graph $G = (V, E, w)$. By Lemma 3.2, the net-based spanner H_{net} contains an path $\pi_{xy} = (x = v_0, v_1, v_2, \ldots, v_k = y)$ of weight at most $(1+\varepsilon)d+_G(x, y)$.

Each edge e of H_{net} was added at some level of the net-tree $\mathcal{T}$. Recall that at level 0, we added edges of length in the range $[1, 18/\varepsilon]$. We can also partition the edges at level 0 into $O(\log \varepsilon^{-1})$ subsets such that in each subset the ratio between the edge lengths is bounded by a constant: For every edge e of H_{net}, there is an index $i \in \{-\lceil \log_2(18/\varepsilon)\rceil, \ldots, \lceil \log_2 \rho\rceil\}$ such that $8\Delta_i \leq w(e) \leq 19\Delta_i$. In particular, for every $j \in \{1, \ldots, k\}$, there exists $i(j) \in \{-\lceil \log_2(18/\varepsilon)\rceil, \ldots, \lceil \log_2 \rho\rceil\}$ such that $8\Delta_{i(j)} \leq d_G(v_{j-1}, v_j) \leq 19\Delta_{i(j)}$.

By Lemma 3.3, for every $j \in \{1, \ldots, k\}$, the additive spanner $H_{i(j)}$ contains a path $\pi(v_{j-1}, v_j)$ from v_{j-1} to v_j of weight $w(\pi(v_{j-1}, v_j)) \leq d_G(v_{j-1}, v_j) + \varepsilon\Delta_{i(j)} \leq (1 + \varepsilon/8) \cdot d_G(v_{j-1}, v_j)$. The concatenation of the paths $\pi(v_0, v_1), \ldots, \pi(v_{k-1}, v_k)$ is a path in H of weight at most

$$\begin{aligned}
d_H(x,y) &\leq \sum_{j=1}^{k} w(\pi(v_{j-1}, v_j)) \\
&\leq \sum_{j=1}^{k} \left(1 + \frac{\varepsilon}{8}\right) \cdot d_G(v_{j-1}, v_j) \\
&= \left(1 + \frac{\varepsilon}{8}\right) w(\pi_{xy}) \\
&\leq \left(1 + \frac{\varepsilon}{8}\right)(1+\varepsilon) d_G(x,y) \\
&= \left(1 + \frac{9}{8}\varepsilon + \frac{\varepsilon^2}{8}\right) d_G(x,y) \\
&< (1+2\varepsilon) d_G(x,y),
\end{aligned}$$

for any $\varepsilon \in (0, \frac{1}{4})$, as claimed. □

3.2 Reduction to Planar Metrics of Bounded Diameter Consider the planar metric (N_i, d_G) on a single level of the net-tree. In this section, we reduce the problem of construction of H_i (claimed in Lemma Lemma 3.3) to subspaces of N_i of bounded diameter.

DEFINITION 3.5. *A (β, s, Δ)-sparse cover for a graph G is a collection $\mathcal{C} = \{C_1, \ldots, C_t\}$ of subgraphs of G (called* clusters*) such that*

1. *$(C_j) \leq \Delta$;*
2. *for every $v \in V(G)$, there is a cluster $C_j \in \mathcal{C}$ such that $B_G(v, \Delta/\beta) \subseteq C_j$ (that is, C_j contains all vertices at distance at most Δ/β from v); and*
3. *every $v \in N_i$ is contained in at most s clusters (that is, $|\{j : v \in C_j \in \mathcal{C}\}| \leq s$).*

Sparse covers were introduced by Awerbuch and Peleg [11]. For planar graphs, Busch et al. [24] showed that one can construct (β, s, Δ)-sparse cover for any Δ with constant τ and s.

THEOREM 3.6. ([24]) *There exist absolute constants $\beta, s \in O(1)$ such that for every $\Delta > 0$ and every planar graphs, a (β, s, Δ)-sparse cover can be constructed in polynomial time.*

Recall that for each level i of the net tree, the minimum distance between any two points of the net N_i is at least $2^i = \varepsilon \cdot \Delta_i$. Using Theorem 3.6 with parameter $\Delta = 20\beta\Delta_i$, we obtain a (β, s, Δ_i)-sparse cover $\mathcal{C}_i = \{C_1, \ldots, C_t\}$ of the planar graph G. We only care about the clusters $C_j \in \mathcal{C}$ where $|N_i \cap C_j| \geq 2$.

3.3 Recursive Shortest-Path Separators Consider a single cluster, $C_j \in \mathcal{C}_i$, which is a planar graph with diameter $O(\Delta_i) = O(2^i/\varepsilon)$, and recall the distance between any two net points in N_i is at least $\varepsilon\Delta_i$. In this section, we construct a $(1+\varepsilon)$-spanner for $N_i \cap C_j$ when $|N_i \cap C_j| \geq 2$.

Shortest Path Separators. We recursively partition C_j along shortest path separators until each subgraph contains at most one net point in $N_i \cap C_j$. A *balanced separator* (for short, *separator*) of a graph G is a set of vertices $S \subset V(G)$ such that every connected component of $G - S$ has at most $\frac{2}{3} \cdot |V(G)|$ vertices. According to a celebrated result by Lipton and Tarjan [62], every n-vertex planar graph admits a balanced separator of size $O(\sqrt{n})$ A recursive partition of planar (or minor-free) graphs along balanced separators is well-known powerful technique; see [12, 27, 28, 41, 50, 55, 68] for examples. Goodrich [45] noticed that one can choose balanced separators in planar graph as the vertices of a *fundamental cycle*, which his composed of two shortest paths from a common endpoint such that the other two endpoints of the paths are adjacent in G (if G is a triangulation) or at least incident to a common face (in general). Such a balanced separator is called a *shortest path separator*. Since can use a shortest path tree to recursively partition a planar graph [54, 68], and maintain the additional property that the each subgraph in the recursion is bounded by $O(1)$ shortest paths. We use the terminology presented in [43]:

DEFINITION 3.7. *Given an edge-weighted graph $G = (V, E, w)$, a vertex $r \in V$, and a parameter $\eta > 0$, an η-rooted shortest path decomposition (for short, η-RSPD) with root r, denoted by Φ, is a binary tree with the following properties:*
Each node $\alpha \in \Phi$ is associated with a subgraph G_α of G, called a piece, *such that:*

(P1) The subtree of Φ rooted at α has height $O(\log|V(G_\alpha)|$.

(P2) For every piece G_α, there is a set of boundary vertices *$Q_\alpha \subset V(G_\alpha)$ such that every path between a vertex $u \in V(G_\alpha)$ and $v \in V(G) \setminus V(G_\alpha)$ in G contains a vertex in Q_α. The vertices in $V(G_\alpha) \setminus Q_\alpha$ are called* internal vertices *of G_α.*

(P3) For every piece G_α, all boundary vertices in Q_α are contained in at most η shortest paths in G with a common endpoint r.

(P4) If α is the root of Φ, then $G - \alpha = G$; if α is a leaf of Φ, then G_α has at most η internal vertices. Otherwise, α is an internal node of Φ with exactly two children β_1 and β_2. It holds that $G_\alpha = G_{\beta_1} \cup G_{\beta_2}$ and $V(G_{\beta_1}) \cap V(G_{\beta_1}) \subseteq Q_{\beta_1} \cap Q_{\beta_2}$.

Let $G = (V, E, w)$ be an edge-weighted planar graph. We may assume that G is a triangulation (i.e., an edge-maximal planar graph) by adding edges of sufficiently large weight (that do not change the shortest path distance). Thorup [68, Section 2.5] showed that for a triangulated wedge-weighed planar graph with n vertices, one can compute a an η-RSPD with $\eta = O(1)$ in $O(n \log n)$ time (an RSPD is called *frame separator decomposition* in Thorup's paper).

Note that if α and β are siblings in Φ, then G_α and G_β may share boundary vertices. That is, $V(G) = \bigcup\{V(G_\alpha) : \alpha \in \Phi \text{ is a leaf node}\}$ is not a partition. In order to avoid duplication, for each net point $v \in N_i \cap G$, we specify a lowest node $\varphi(v) \in \Phi$ such that $v \in V(G_{\varphi(v)})$. For every node $\alpha \in \Phi$, let N_α denote the set of net points $v \in N_i$ such that $\varphi(v)$ is a descendant of α (possibly $\varphi(v) = \alpha$). With this notation, $N_i = \bigcup\{N_\alpha : \alpha \in \Phi \text{ is a leaf node}\}$ is a partition; and we have $|N_i| = \sum_{\text{leaf } \alpha \in \Phi} |N_\alpha|$.

Filtser and Le [41] used the η-RSDP for an (exact) emulator for tree metrics with treewidth $O(\log\log n)$ and hop-diameter $O(\log\log n)$; not for Steiner spanners. However, they proved the following lemma.

LEMMA 3.8. (LEMMA 4 IN [43]) *Let α and β be two nodes in Φ. Let P_{uv} be any path between two vertices u and v in G such that $u \in Q_\alpha$ and $v \in Q_\beta$. Let $(\alpha = \lambda_1, \lambda_2, \ldots, \lambda_k = \beta)$ be a set of nodes in the unique path $\Phi[\alpha, \beta]$ such that $\lambda_{i+1} \in \Phi[\lambda_i, \beta]$ for any $1 \leq i \leq k-1$. Then, there exists a sequence of vertices $(u = x_1, x_2, \ldots, x_k = v)$ such that $x_i \in P_{uv} \cap Q_{\lambda_i}$ and $x_{i+1} \in P_{uv} \cap P[x_i, v]$ for any $1 \leq i \leq k-1$.*

REMARK 3.9. *Assume that $\Phi[\alpha, \beta] = (\alpha = \lambda_1, \lambda_2, \ldots, \lambda_k = \beta)$ is the (unique) path in Φ between α and β. Then Lemma 3.8 states that the shortest path P_{uv} contains some boundary vertices from $Q_{\lambda_1, \lambda_2}, \ldots, Q_{\lambda_k}$ is this order (a vertex can belong to several boundary sets). However, it does* not *say that the shortest path is contained in the union of these pieces, $\bigcup_{i=1}^{k} G_{\lambda_i}$. It is possible that P_{uv} passes though some additional pieces—but this will not affect our construction.*

Shortcut Edges in RSPD. Let Φ be an η-RSPD for G with a constant η. It is a binary tree of height $O(\log n)$, and so its diameter is $O(\log n)$. We augment Φ with *shortcut* edges in order to reduce its diameter.

Chung and Garay [34] initiated the studied the minimum number of shortcut edges for paths and trees. Chazelle [30] showed that a n-vertex tree can be augmented with m new edges to reduce the diameter to $O(\alpha(n, m))$, where $\alpha(n, m)$ is the two-parameter inverse Ackermann function. This bound is the best possible [69], and was later rediscovered several times [5, 20]; see also [19, 43, 58, 67] for algorithmic aspects and generalizations. Alternatively, the diameter of a tree with n vertices can be reduced to $2k$ by adding $O(n\alpha_k(n))$ edges, where $\alpha_k(n)$ is the inverse of a certain Ackermann-style function at the $\lfloor k/2 \rfloor$th level of the primitive recursive hierarchy: Specifically, $\alpha_0(n) = \lceil n/2 \rceil$, $\alpha_1(n) = \lceil \sqrt{n} \rceil$, $\alpha_2(n) = \lceil \log n \rceil$, $\alpha_3(n) = \lceil \log \log n \rceil$, $\alpha_4(n) = \log^* n$, etc.. Moreover, $\alpha_{2\alpha(n)+2}(n) \leq 4$, where $\alpha(n)$ is the one-parameter inverse Ackermann function, which is an extremely slowly growing function [5, 58, 66].

We apply the above results to the η-RSPD Φ carefully. For our purposes, we distinguish between two types of shortcut edges: type 1 shortcut edges between internal nodes of Φ; and type 2 shortcut edges between a leaf and an internal node. For our purposes (discussed below), a shortcut edge of type 1 costs roughly $O(\varepsilon^{-2})$; and one of type 2 costs roughly $O(\varepsilon^{-1})$. For this reason, we prefer shortcut edges of type 2. We add shortcut edges to Φ as follows.

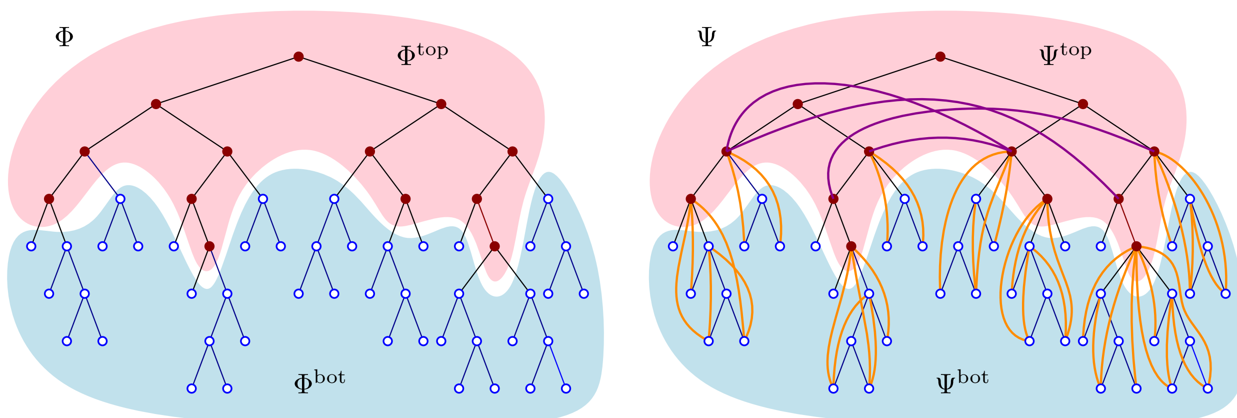

Figure 7: Left: a binary tree Φ, with Φ^{top} and Φ^{bot}. Right: The augmented graph Ψ, with Ψ^{top} and Ψ^{bot}.

We decompose the tree Φ as follows. Let λ be a parameter to be optimized later (we shall choose $\lambda = \varepsilon^{-1}\alpha(n)/\log(\varepsilon^{-1}\alpha(n))$). Let Φ^{top} denote the subtree of Φ induced by all nodes α with $|V(G_\alpha)| \geq \lambda$, and Φ^{bot} the forest induced by all other nodes of Φ and the leaves of Φ^{top}. Then Φ^{top} has $O(n/\lambda)$ nodes; and the height of (each tree in) Φ^{bot} is $O(\log \lambda)$. Now we augment Φ to a graph Ψ with shortcut edges as follows (see Fig. 7):

1. Reduce the diameter of Φ^{top} to $O(1)$ [5, 20, 30];

2. augment Φ^{bot} by connecting every leaf of the forest Φ^{bot} to all of its ancestors in Φ^{bot}.

Denote by Ψ^{top} and Ψ^{bot}, resp., the subgraph of Ψ induced by the vertices of Φ^{top} and Φ^{bot}.

LEMMA 3.10. *The graph Ψ^{top} has $O(n\alpha(n)/\lambda)$ edges; and the forest Ψ^{bot} has $O(n \log \lambda)$ edges incident to its leaves.*

Proof. Since Ψ^{top} is a binary tree on $O(n/\lambda)$ nodes, it has $O(n/\lambda)$ edges. The first augmentation phase adds $O((n/\lambda)\alpha(n/\lambda)) \leq O(n\alpha(n)/\lambda)$ shortcut edges, and so Ψ^{top} has $O(n\alpha(n)/\lambda)$ edges.

The forest Φ^{bot} has $O(n)$ leaves, each of which is incident to only one edge in Φ. Since the height of Φ^{bot} is $O(\log \lambda)$, the second augmentation step adds $O(n \log \lambda)$ new edges to every leaf. Overall, Ψ^{bot} has $O(n \log \lambda)$ edges incident to leaves. $\square$

Portals along Shortest Path Separators. Recall that every node α corresponds to a piece G_α, and the boundary vertices Q_α of G_α all lie in $\eta = O(1)$ shortest paths of G.

We place Steiner points, that we call *portals*, along the shortest path in Q_α as follows; see Fig. 8. Let α be a node of Φ, and suppose that the vertices in Q_α lie in the shortest paths $P_1, P_2, \ldots, P_\eta$ in G. The length of each path is at most $(G) \leq \Delta_i$. For each j, $1 \leq j \leq \eta$, we place portals at the two endpoint of P_j, and recursively place portals at internal nodes until any two consecutive portals at at distance at most $\frac{\varepsilon}{10}\Delta_1$ apart or are adjacent along P_j. Let S_α denote the set of portals (i.e., Steiner points) over all η paths. It follows that we place $O(\varepsilon^{-1})$ portals along each shortest path, and so $|S_\alpha| \leq O(\eta \cdot \varepsilon^{-1}) = O(\varepsilon^{-1})$.

We can now define the Steiner spanner $H_{i,j}$ for $N_i \cap C_j$. Let the vertex set of H_i be $N_i \cup \bigcup_{\alpha\in\Phi} S_\alpha$, that is, the net points in $N_i \cap C_j$ and all portals defined above. We add the following edges to $H_{i,j}$, each with the same weight as in G:

1. For every edge $\alpha\beta$ in Ψ^{top}, add a complete bipartite graph between the portals S_α and S_β;
2. for every edge $\alpha\beta$ of Ψ, where α is a leaf of Ψ^{bot}, add a complete bipartite graph between N_α and S_β;
3. for every leaf α, add a complete graph among its internal net vertices $N_i \cap (G_\alpha \setminus Q_\alpha)$ and a complete bipartite graph between N_α and S_α.

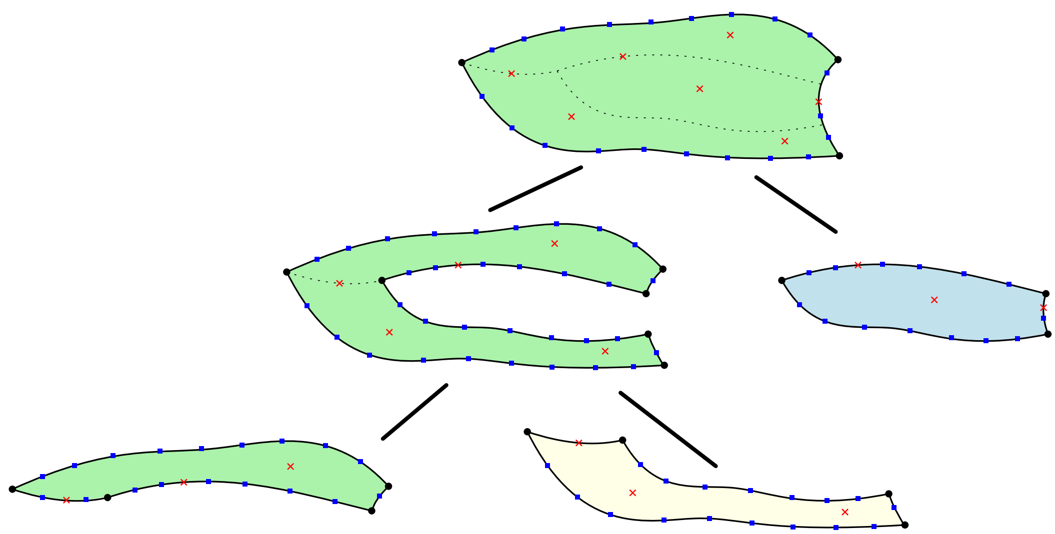

Figure 8: A recursive partition of a plane graph along shortest paths into pieces, until there are at most $\eta = O(1)$ terminals in the interior of each piece. Terminals (red crosses) and portals (blue squares).

We are now ready to prove Lemma 3.3.

LEMMA 3.11. *We have* $|E(H_{i,j})| \leq O\Big(|N_i \cap C_j| \cdot \varepsilon^{-1} \log(\varepsilon^{-1}\alpha(n))\Big)$,

Proof. By Lemma 3.10, the graph Ψ^{top} has $O(|N_i \cap C_j|\,\alpha(n)/\lambda)$ edges, and for each such edge we add a complete bipartite graph with $O(\varepsilon^{-2})$ edges to H_i. This contributes $O(\varepsilon^{-2}\,|N_i \cap C_j|\,\alpha(n)/\lambda)$ edges to H_i. By Lemma 3.10, the graph Ψ has $O(|N_i \cap C_j| \log \lambda)$ edges incident to leaves of Ψ^{bot}, and for each such edge we add a star with $O(\varepsilon^{-1})$ edges to $H_{i,j}$. This contributes $O(\varepsilon^{-1}|N_i \cap C_j| \log \lambda)$ such edges to $H_{i,j}$.

We choose $\lambda := \varepsilon^{-1}\alpha(n)/\log(\varepsilon^{-1}\alpha(n))$ to balance the above two contributions. They both amount to $O\Big(\varepsilon^{-1}|N_i \cap C_j| \log(\varepsilon^{-1}\alpha(n))\Big)$.

Finally, each leaf node $\alpha \in \Phi$ has at most $\eta = O(1)$ internal vertices. The complete graphs among internal net vertices contribute at most $\eta \cdot |N_i \cap C_j| = O(|N_i \cap C_j|)$ edges contribute a total of $O(\eta^2 \cdot |N_i \cap C_j) = O(|N_i \cap C_j)$ edges. □

LEMMA 3.12. *For all* $x, y \in N_i \cap C_j$, *if* $d_G(x, y) = \Theta(\Delta_i)$, *then* $d_{H_{i,j}}(x, y) \leq d_G(x, y) + \varepsilon\Delta_i$.

Proof. Let $x, y \in N_i$ such that $d_G(x, y) = \Theta(\Delta_i)$. Then $x \neq y$, and the η-RSPC has two leaves $\varphi(x), \varphi(y) \in \Phi$ such that $x \in N_{\varphi(x)}$ and $y \in N_{\varphi(y)}$. We distinguish between three cases:

Case 1: $\varphi(x) = \varphi(y)$**.** If both x and y are internal vertices of $G_{\varphi(\alpha)}$, then the complete graph on $N_\alpha \cap (G_{\varphi(\alpha)} \setminus Q_{\varphi(\alpha)})$ contains the edge xy. Otherwise we may assume w.l.o.g. that x is a boundary vertex, that is, $x \in Q_{\varphi(\alpha)}$. Then x lies on a shortest path in $Q_{\varphi(\alpha)}$, and there is a portal $s \in S_{\varphi(\alpha)}$ such that $d_G(x,s) \le \frac{\varepsilon}{10}\Delta_i$. The complete bipartite graph between $N_{\varphi(\alpha)}$ and $S_{\varphi(\alpha)}$ contains the edge sy. Consequently,

$$\begin{aligned} d_{H_{i,j}}(x,y) &\le d_{H_i}(x,s) + d_{H_i}(s,y) = d_G(x,s) + d_G(s,y) \\ &\le \Big(d_G(x,y) + d_G(s,y)\Big) + d_G(s,y) = d_G(x,y) + 2d_G(s,y) \\ &\le d_G(x,y) + 2\cdot\frac{\varepsilon}{10}\Delta_i < d_G(x,y) + \varepsilon\Delta_i, \end{aligned}$$

as required.

Case 2: $\varphi(x) \ne \varphi(y)$ **and** $\mathcal{LCA}\{\varphi(x),\varphi(y)\}$ **is in** Φ^{bot}**.** In this case, by construction, Ψ contains the path $\Psi[\varphi(x), \mathcal{LCA}\{\varphi(x),\varphi(y)\}, \varphi(y)]$. By Lemma 3.8, the shortest path P_{xy} in G contains a sequence of vertices $(x = x_0, x_1, x_2 = y)$ in this order such that $x_1 \in Q_{\mathcal{LCA}\{\varphi(x),\varphi(y)\}}$. The construction of portals ensures that there exists a portal $s_1 \in S_{\mathcal{LCA}\{\varphi(x),\varphi(y)\}}$ with $d_G(x_1, s_)\le \varepsilon\Delta_i$. Consequently,

$$\begin{aligned} d_{H_{i,j}}(x_y) &\le d_{H_i}(x,s_1) + d_{H_i}(s_1,y) \\ &= d_G(x,s_1) + d_G(s_1,y) \\ &\le \Big(d_G(x,s_1) + d_G(s_1,x_1)\Big) + \Big(d_G(x_1,s_1) + d_G(s_1,y)\Big) \\ &\le d_G(x,x_1) + d_G(x_1,y) + 2\cdot\frac{\varepsilon}{10}\cdot\Delta_i \\ &< d_G(x,y) + \varepsilon\cdot\Delta_i, \end{aligned}$$

as required.

Case 3: $\varphi(x) \ne \varphi(y)$ **and** $\mathcal{LCA}\{\varphi(x),\varphi(y)\}$ **is not in** Φ^{bot}**.** Nodes $\varphi(x)$ and $\varphi(y)$ have ancestors α and β, resp., that are leaves in Φ^{top}. Furthermore, $\alpha \ne \beta$ or else we would be in Case 2. Due to the shortcut edges, the distance between α and β is at most 4 in Ψ^{top}. Consequently, Ψ contains a path $\Psi[\varphi(x),\varphi(y)] = (\varphi(x) = \lambda_0, \lambda_1, \ldots, \lambda_k = \varphi(y))$ of length at most $k \le 6$. By Lemma 3.8, the shortest path P_{xy} in G contains a sequence of vertices $(x = x_0, x_1, \ldots, x_k = y)$ in this order such that $x_j \in Q_{\lambda_j}$. The construction of portals ensures that there exist portals $s_j \in S_{\lambda_j}$ with $d_G(x_j, s_j) \le \varepsilon\Delta_i$ for all $j \in \{1, \ldots, k-1\}$. The construction of $H_{i,j}$ guarantees that the edges xs_1, $s_1s_2, \ldots, s_{k-2}s_{k-1}$, $s_{k-1}y$ are present in H_i. Consequently,

$$\begin{aligned} d_{H_{i,j}}(x_y) &\le d_{H_i}(x,s_1) + d_{H_i}(s_1,s_2) + \ldots + d_{H_i}(s_{k-2},s_{k-1}) + d_{H_i}(s_{k-1},y) \\ &= d_G(x,s_1) + d_G(s_1,s_2) + \ldots + d_G(s_{k-2},s_{k-1}) + d_G(s_{k-1},y) \\ &\le \Big(d_G(x,s_1) + d_G(s_1,x_1)\Big) + \Big(d_G(x_1,s_1) + d_G(s_1,s_2) + d_G(s_2,x_2)\Big) + \ldots + \Big(d_G(x_{k-1},s_{k-1}) + d_G(s_{k-1},y)\Big) \\ &\le \sum_{j=1}^{k} d_G(x_{i-1},x_i) + 2(k-1)\cdot\frac{\varepsilon}{10}\cdot\Delta_i \\ &\le d_G(x,y) + \varepsilon\cdot\Delta_i, \end{aligned}$$

as required. □

The combination of Lemma 3.11 and Lemma 3.12 implies Lemma 3.3, that we restate for convenience:

LEMMA 3.3. *For every $i \in \mathbb{N}$, there exists a spanner H_i on N_i such that*

1. *for all $x, y \in N_i$, if $d_G(x,y) = \Theta(\Delta_i)$, then $d_{H_i}(x,y) \le d_G(x,y) + \varepsilon\Delta_i$, and*
2. $|E(H_i)| \le O\Big(|N_i|\,\varepsilon^{-1}\cdot\log(\varepsilon^{-1}\alpha(n))\Big)$,

where $\alpha(.)$ denotes the inverse Ackermann function.

Proof. At every level $i \in \mathbb{N}$, we have constructed a (β, s, Δ_i)-sparse cover $\mathcal{C}_i = (C_1, \ldots, C_{t(i)})$ of G; and for every $j \in \{1, \ldots, t\}$, we have constructed and an additive spanner $H_{i,j}$ for $N_i \cap C_j$. Let $H_i = \bigcup_{j=1}^{t} H_{i,j}$. We claim that

1. **(stretch condition)** for all $x, y \in N_i$, if $\Delta_i \leq d_G(x,y) \leq \frac{\Delta_i}{\beta}$, then $d_{H_i}(x,y) \leq d_G(x,y) + \varepsilon\Delta_i$, and
2. **(size condition)** $|E(H_i)| \leq O\Big(|N_i| \cdot \varepsilon^{-1}\log(\varepsilon^{-1}\alpha(n))\Big)$,

where $\alpha(.)$ denotes the inverse Ackermann function.

Stretch analysis. Let $x, y \in N_i$ such that $\Delta_i \leq d_G(x,y) \leq \frac{\Delta_i}{\beta}$. By the definition of (β, s, Δ_i)-sparse covers, there exists a cluster $C_j \in \mathcal{C}_i$ such that $B_G(x, \Delta/\beta) \subseteq C_j$. Then $x, y \in C_j$. By Lemma 3.12, we have $d_{H_i}(x,y) \leq d_{H_i,j}(x,y) \leq d_G(x,y) + \varepsilon\Delta_i$, as required.

Size analysis. By the definition of (β, s, Δ_i)-sparse covers, each net point $v \in N_i$ is contained in at most $s = O(1)$ clusters in $\mathcal{C}_i$. Consequently, $\sum_{j=1}^{t} |N_i \cap C_j| \leq s \cdot |N_i| = O(|N_i|)$. Summation of the bound in Lemma 3.11 now yields

$$|E(H_i)| \leq \sum_{j=1}^{t} |E(H_{i,j})| \leq \sum_{j=1}^{t} O\Big(|N_i \cap C_j| \cdot \varepsilon^{-1}\log(\varepsilon^{-1}\alpha(n))\Big) \leq O\Big(|N_i| \cdot \varepsilon^{-1}\log(\varepsilon^{-1}\alpha(n))\Big),$$

as required. □

3.4 Proof of Theorem 3.1 We can now put the pieces of the puzzle together and prove Theorem 3.1. Recall the definition of the graph H. For each level N_i of a net tree $\mathcal{T}$, Lemma 3.3 yields an additive spanner H_i; and we put $H = \bigcup_{i \in \mathbb{N}} H_i$. We already know (Lemma 3.4) that H is a $(1+2\varepsilon)$-spanner for the metric (T, d_G). It remains to bound the number of edges in H.

LEMMA 3.13. *The graph H has $O\Big(n \cdot \log(\varepsilon^{-1}\alpha(n)) \cdot \varepsilon^{-1}\log\varepsilon^{-1}\Big)$ edges*

Proof. We are given a set T of n terminals in an edge-weighted planar graph $G = (V, E, w)$. We defined $\lceil \log_2(18/\varepsilon) \rceil + \lceil \log_2 \varrho \rceil + 1$ levels, constructed an (β, s, Δ_i)-sparse cover $\mathcal{C}_i = (C_1, \ldots, C_{t(i)})$ on each level, and the created an additive spanner $H_{i,j}$ on the net points $N_i \cap C_j$ if $|N_i \cap C_j| \geq 2$ (Lemma 3.3). We may assume that $H_{i,j}$ is the empty graph when $|N_i \cap C_j| \leq 1$. Then by Lemma 3.11, $H_{i,j}$ has

$$\max\{0, |N_i \cap C_j| - 1|\} \cdot O\left(\varepsilon^{-1}\log(\varepsilon^{-1}\alpha(n))\right)$$

edges for all $j \in \{1, 2, \ldots, t(i)\}$. To complete the proof of Lemma 3.13, it is enough to show that

$$\sum_{i=-\lceil \log_2(18/\varepsilon) \rceil}^{\lceil \log_2 \rho \rceil} \sum_{j=1}^{t(i)} \max\{0, |N_i \cap C_j| - 1|\} = O\left(\frac{n}{\varepsilon}\right). \tag{3.13}$$

The proof of Equation (3.13) is based on the following key observation: We have $(N_i \cap C_j) \leq 20\beta\Delta_i = 20\beta 2^i/\varepsilon$, and N_i is a 2^i-net (i.e., the minimum distance between net points in N_i is 2^i). This implies that at most one point of $N_i \cap C_j$ is present in N_{i+M}, where $M := \lceil \log_2 20\beta/\varepsilon \rceil$. That is,

$$\max\{0, |N_i \cap C_j| - 1|\} \leq \Big|(N_i \setminus N_{i+M}) \cap C_j\Big| + 1 \leq 2 \cdot \Big|(N_i \setminus N_{i+M}) \cap C_j\Big|. \tag{3.14}$$

Based on this observation, we partition the levels into M groups: Specifically, for every $g \in \{0, 1, \ldots, M-1\}$, group g comprises all levels i such that $i \equiv g \pmod M$. Recall that at each level i, every net point in N_i is contained in at most $s = O(1)$ clusters of the (β, s, Δ_i)-sparse cover $\mathcal{C}_i$. Combined with Equation (3.13), this gives

$$\sum_{j=1}^{t(j)} \max\{0, |N_i \cap C_j| - 1|\} \leq 2s \cdot \Big|N_i \setminus N_{i+M}\Big|.$$

Summation over all levels in any group $g \in \{0, 1, \ldots, M-1\}$ yields

$$\sum_{i \equiv g \pmod M} \sum_{j=1}^{t(i)} \max\{0, |N_i \cap C_j| - 1|\} \leq 2s \cdot \sum_{i \equiv g \bmod M} \Big|N_i \setminus N_{i+M}\Big| \leq 2s \cdot |N_0| = O(n).$$

Finally, summation over all $M = O(\varepsilon^{-1})$ groups gives

$$\sum_{i=-\lceil \log_2(18/\varepsilon)\rceil}^{\lceil \log_2 \rho\rceil} \sum_{j=1}^{t(i)} \max\{0, |N_i \cap C_j| - 1|\} = \sum_{g=0}^{M-1} \left(\sum_{i \equiv g \pmod M} \sum_{j=1}^{t(i)} \max\{0, |N_i \cap C_j| - 1|\} \right)$$
$$\leq \sum_{g=0}^{M-1} O(n) = O\left(\frac{n}{\varepsilon}\right),$$

as required. □

The combination of Lemmas 3.4 and 3.13 readily implies Theorem 3.1.

3.5 Generalization to Graphs of Bounded Genus In this section, we generalize Theorem 3.1 to polyhedral metrics.

THEOREM 3.14. *Let $\varepsilon \in (0,1)$ be a parameter. Let T be a set of n points (*terminals*) in a polyhedral metric. We can construct a Steiner $(1+\varepsilon)$-spanner for T with $O((n/\varepsilon)\cdot\log(\varepsilon^{-1}\alpha(n))\cdot\log\varepsilon^{-1})$ edges, where $\alpha(n)$ is the inverse Ackermann function.*

Every step of the proof of Theorem 3.1 generalizes to graphs of bounded genus. We briefly sketch the key differences. In Section 3.1, the construction of net-tree based spanners and Lemma 3.2–Lemma 3.4 hold for any metric space (X, d_X). In Section 3.2, we used (β, s, Δ)-sparse covers to reduce the problem to planar graphs of bounded diameter. Klein et al. [56] constructed (β, s, Δ)-sparse covers for minor-free classes of graphs, where β and s depend on $\mathcal{H}$ (i.e., β and s are constants for fixed $\mathcal{H}$). Busch et al. [24] later constructed (β, s, Δ)-sparse covers for planar graphs with better constants; and Abraham et al. [4] for $K_{r,r}$-free graphs for any $r \in \mathbb{N}$.

In Section 3.3, we used a η-RSPD (η-rooted shortest path decomposition) for planar graphs of bounded diameter, based on Thorup [68, Section 2.5]. For graphs of constant genus g, we first reduce this step to planar graphs at the expense of increasing the diameter by a factor of $O(g)$. [5] A cut graph is a graph C embedded on a surface S such that $S \setminus C$ is homeomorphic to a closed disk [37]. It is NP-hard to find the *shortest* cut subgraph in a given graph embedded in an oriented surface [39, 36]. However, a cut graph has a very simple structure: It consists of a tree T with g cross edges, and can be computed in $O(n)$ time [38, 39, 44]. If we start with a shortest path tree T, we obtain a cut graph C that consists of $O(g)$ pairwise noncrossing shortest paths. When we cut the surface S (and the graph G) along C, the edges and vertices on C are duplicated. If we start with a graph G of genus g, we obtain a *planar* graph G'.

LEMMA 3.15.

1. *Every shortest path in G' is the union of $O(g)$ shortest paths in G, and*
2. *every shortest path in G is the union of $O(g)$ shortest paths in G'.*

Proof. The cut graph C of G is a union of $O(g)$ noncrossing shortest paths $\{\gamma_1, \ldots, \gamma_k\}$. By the optimal substructure property, every subpath of a shortest path is a shortest path.

(1) Let P' be a shortest path in G'. Then P' has a connected intersection with each path γ_i, which is a shortest path in G. Every component of $P' \setminus C$ is a shortest path in G. Overall, P' has $O(g)$ subpaths along the the paths γ_i, consequently $P' \setminus C$ has $O(g)$ components. All $O(g)$ subpaths of P' are shortest paths in G.

(2) Let P be a shortest path in G. Then P' has a connected intersection with each path γ_i, which is a shortest path in both G and G'. Furthermore, every component of $P \setminus C$ is a shortest path in both G and G'. Overall, P decomposes into $O(g)$ subpaths that are shortest paths in both G and G'. □

[5] We note that for graphs of bounded genus, Abraham and Gavoille [3] constructed a weaker decomposition than η-RSPD: Specifically, property (P3) of η-RSPD stipulates that for every node $\alpha \in \Phi$, the boundary vertices Q_α are contained in at most η shortest paths of G_α. In [3], the boundary vertices are covered in at most η *successive* shortest paths, that is, Q_α is contained in paths $\gamma_1, \ldots, \gamma_k$ where $k \leq \eta$, and γ_i is a shortest path in the graph $G_\alpha \setminus \bigcup_{j<i} \gamma_j$. Unfortunately, the removal of one or more paths may increase the diameter—and is unsuitable for our purposes.

By Lemma 3.15(1), we have $(G') \leq O(g \cdot (G))$. We can compute recursive shortest paths separators in G' as in Section 3.3; but the very first separator is C, which consists of $O(g)$ shortest paths. Given a shortest path P_{uv} in G, it is a union of $O(g)$ shortest paths in G' by Lemma 3.15(2), each of which can be traced in the recursive decomposition of G' by Lemma 3.8. Our construction of portals and Steiner spanners now works for G' as described in Section 3.3 but the number of portals (and Steiner points) increases by a factor of $O(g)$. With these modifications, the proof of Lemma 3.3 in Section 3.3 and Theorem 3.1 in Section 3.4 go through with a constant $O(g)$ factor increase for any constant genus g.

4 Spanners in Planar Domains

In this section, we prove (item 2 in) Theorem 1.2 and Theorem 1.3 using tools we develop in Section 2 and Section 3. First, we show Lemma 1.6, which implies that planar metrics and polyhedral domain are equivalent.

LEMMA 1.6. TREE $\sqsubset$ PLANAR $\cong$ POLYDOM $\cong$ TERRAIN $\sqsubset$ POLYSURF.

Proof. Observe that TREE $\ncong$ PLANAR since (the metric induced by) the unweighted cycle graph C_n with n vertices cannot be embedded isometrically into any tree metrics. Thus, TREE $\sqsubset$ PLANAR.

Abam, de Berg, and Seraji [2] showed that POLYDOM $\sqsubseteq$ TERRAIN by controlling the elevation of polyhedral terrains. Next, we show that TERRAIN $\sqsubseteq$ PLANAR.

Let P be a set of points on the polyhedral terrain. Our goal is to show that there exists a planar metric induced by an edge-weighted planar graph G and a subset of points $Q \subseteq V(G)$ such that there exists an isometry from P to Q.

We start by taking the arrangements of geodesic paths between all points in P; that is, we draw the shortest path between any two points in P on the terrain. We say that two lines that meet at a point a are intersecting if they are not equal on one of the two sides of a as shown in Figure 9.

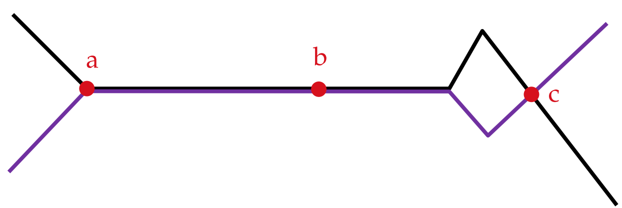

Figure 9: The black and purple lines intersect in a and c but not in b

We can modify paths so that any two of them only intersect at at most two points. Suppose two shortest paths p and q intersect at three points a, b and c. By the suboptimality principle, the section of p and q between a and b are the same length. We modify q to be equal to p between a and b. This does not change the length of q. After doing the same between b and c, p and q only intersect in a and c.

We choose an ordering $\sigma : P^2 \mapsto \left[\binom{|P|}{2}\right]$ of the paths and modify the path as described above such that, if two distinct paths p and q with $\sigma(p) < \sigma(q)$ intersect in more than two points, we replace the sections of q with the sections of p. This gives an arrangement of geodesic paths such that any two paths intersect at at most two points.

We add Steiner points at all intersection points of the geodesic paths and denote by G the graph obtained. Then G is planar by construction since no pair of edges are intersecting. The point set Q contains the vertices corresponding to P. Since all the shortest distances between points in P are preserved in G, the natural mapping from every point in P to its copy in Q is an isometry, showing that TERRAIN $\sqsubseteq$ PLANAR.

Next, we show that PLANAR $\sqsubseteq$ POLYDOM. Let G be an edge-weighted planar graph realizing a planar metric $\mathcal{M}_1$ and P be a subset of vertices of G. We draw G on the plane so that every edge is a polygonal curve, with its length being the weight of the edge. Let R be a sufficiently big rectangle that encloses the drawing of G. Removing all points (in the drawing) of G from R, we get a set of polygonal regions, where each region corresponds to a face of G, except one region which corresponds to the intersection of the infinite face of G and R. We now regard each polygonal region as a hole, a.k.a., a polygonal obstacle, thereby obtaining a polygonal domain $\mathcal{M}_2$. Let Q be the points in $\mathbb{R}^2$ that correspond to the vertices of P in the drawing. Clearly, the geodesic distance between any two points of Q in $\mathcal{M}_2$ is the shortest path distance in G, and hence (the metric induced by) P can be embedded isometrically int $\mathcal{M}_2$.

Lastly, we show that TERRAIN $\sqsubset$ POLYSURF. Consider a point set P on a polyhedral surface, which is a piece-wise linear function $f : D \to \mathbb{R}$ for some convex polygonal region $D \subset \mathbb{R}^2$. Since P is finite, all shortest paths between points in P lie in a compact subset of the terrains. Consequently, we may assume w.l.o.g. that D is compact and (after scaling) lies in a unit square, $D \subset [0,1]^2$; and $f > 0$. Extend f to a larger square domain $[-1,2]^2 \subset \mathbb{R}^2$, and let consider the solid $S = \{(x,y,z) \in \mathbb{R}^3 : (x,y) \in D \text{ and } 0 \leq z \leq f(x,y)\}$. Now the boundary of S is a polyhedral surface that contains the terrain (as well as all points in P), and the shortest paths among P are the same in both metrics.

To see that TERRAIN $\not\sqsubseteq$ POLYSURF, let $\mathbb{M}_1$ be the shortest path metric of $K_{3,3}$ (with unit edge weights). It is in POLYSURF, as $K_{3,3}$ can be realized in $\mathbb{R}^3$ with noncrossing polygonal arcs of unit length, which are shortest paths in some polyhedral surface of sufficiently high genus. However, in any realization of $K_{3,3}$ on a polyhedral terrain $\mathbb{M}_2$, two shortest paths will cross. Assume that the a_1b_1- and a_2b_2-paths cross. Then $\delta_2(a_1,a_2) + \delta(b_1,b_2) \leq \delta_2(a_1,b_1) + \delta_2(a_2,b_2) = 2$, in the metric δ_2 of the terrain, and so $\delta_2(a_1,a_2) < 2$ and $\delta(b_1,b_2) < 2$, while $\delta_1(a_1,a_2) = \delta_1(b_1,b_2) = 2$ in the metric of $K_{3,3}$. □

We are now ready to construct the spanners claimed in the second item in Theorem 1.2 and in Theorem 1.3.

THEOREM 1.2. *Let $\varepsilon \in (0,1)$ be a parameter.*

1. *There exists a polyhedral terrain and a set P of n points on the terrain such that any 2-spanner for P must have $\Omega(n \log n)$ edges.*
2. *Given any set P of n points in a polyhedral terrain, we can construct a $(2+\varepsilon)$-spanner for P with $\tilde{O}(n/\varepsilon^6)$ edges. The number of edges is $O(n)$ for constant ε.*

Proof. [Proof of item 2] By Lemma 1.6, it suffices to construct a $(2+4\varepsilon)$-spanner for a set of point P in a planar metric realized by an edge-weighted planar graph G. One can obtain a $(2+\varepsilon)$-spanner by simply scaling ε. We first construct a tree cover $\mathcal{T}$ for G with $O(\varepsilon^{-3}\log(\varepsilon^{-1}))$ trees using Theorem 1.7. The union of trees of $\mathcal{T}$ is a $(1+\varepsilon)$-spanner of G. In each tree in $\mathcal{T}$, we remove the Steiner points using Theorem 1.5. This gives a family of $O(\varepsilon^{-3}\log(\varepsilon^{-1}))$ trees on P, each with $O(\varepsilon^{-2}\log(\varepsilon^{-1}))$ vertices. The Steiner points removal operation increases the stretch by a factor of $2+\varepsilon$ or less. The union of trees is therefore a non-Steiner spanner for the shortest path metric over P, with at most $(1+\varepsilon)(2+\varepsilon) \leq (2+4\varepsilon)$ stretch and $O(n\varepsilon^{-5}\log^2(\varepsilon^{-1}))$ edges. □

THEOREM 1.3. *Let $\varepsilon \in (0,1)$ be a parameter. Let P be a set of n points in a polyhedral terrain. We can construct a Steiner $(1+\varepsilon)$-spanner for P with $O\left((n/\varepsilon) \cdot \log(\varepsilon^{-1}\alpha(n)) \cdot \log \varepsilon^{-1}\right)$ edges, where $\alpha(n)$ is the inverse Ackermann function. The same result holds even when P is on a polyhedral surface of bounded genus.*

Proof. By Lemma 1.6, it suffices to construct a Steiner $(1+\varepsilon)$-spanner for a set P of n points in a planar metric realized by an edge-weighted planar graph G. Here, we apply Theorem 3.1 on G with P as our set of terminals. (Alternatively, we can apply Theorem 3.14 on a polyhedral surface of bounded genus.) The number of edges of the spanner is $O((n/\varepsilon) \cdot \log(\varepsilon^{-1}\alpha(n)) \cdot \log \varepsilon^{-1})$. □

5 Lower Bounds

5.1 Stretch 2 Tree Cover We now prove that $\Omega(\log n)$ trees are sometimes necessary in any non-Steiner tree cover with stretch 2, as claimed in item 2 in Theorem 1.5. Instead of proving a bound on the number of trees directly, we show a stronger lower bound (Lemma 5.1 below): any 2-spanner that does not contain any Steiner points must have $\Omega(n \log n)$ edges. As the union of k trees in a Steiner-tree cover with stretch 2 gives a 2-spanner with $O(nk)$ edges, it follows that $k = \Omega(\log n)$.

LEMMA 5.1. *There exists a weighted tree T and a subset of vertices $S \subseteq V(T)$ with n points in T such that any non-Steiner 2-spanner $G = (S, E, w)$ for S must have $|E| = \Omega(n \log n)$.*

Let P_n be the unweighted path graph with n vertices $\{1,2,\ldots,n\} = [n]$. We say that an edge-weighted graph $H = ([n], E_H, w_H)$ is a 2-hop t-spanner for P_n if for every two points $x, y \in P_n$, $d_{P_n}(x,y) \leq d_H(x,y)$ and there exists a path Q_{xy} containing at most two edges such that $w_H(Q_{x,y}) \leq t \cdot d_{P_n}(x,y)$. In [58], the authors showed that any 2-hop $\frac{3}{2}$-spanner for P_n must have $\Omega(n \log n)$ edges. Here, we observe that their proof actually implies that any 2-hop t-spanner for any constant $t \geq 1$ must have $\Omega(n \log n)$ edges; we reproduce their poof with the required changes for completeness. Our proof of Lemma 5.1 requires that $t = 2$.

LEMMA 5.2. (ADAPTED FROM [58]) *Let $H = ([n], E_H, w_H)$ be any 2-hop t-spanner for P_n with $t \geq 1$ and $n \geq 2$, then $|E_H| = \Omega(n \log n/t)$.*

Proof. We will show a slightly stronger statement by induction. Think of P_n as n integer points on the line, and we allow H to contain edges with integer endpoints outside P_n. (But H only needs to preserve distances between points of P_n, not the distances to the integer points outside P_n.) Our goal is to lower bound the minimum number of edges, denoted by $T(n)$, in a 2-hop t-spanner on $[n]$ with *both endpoints in* P_n.

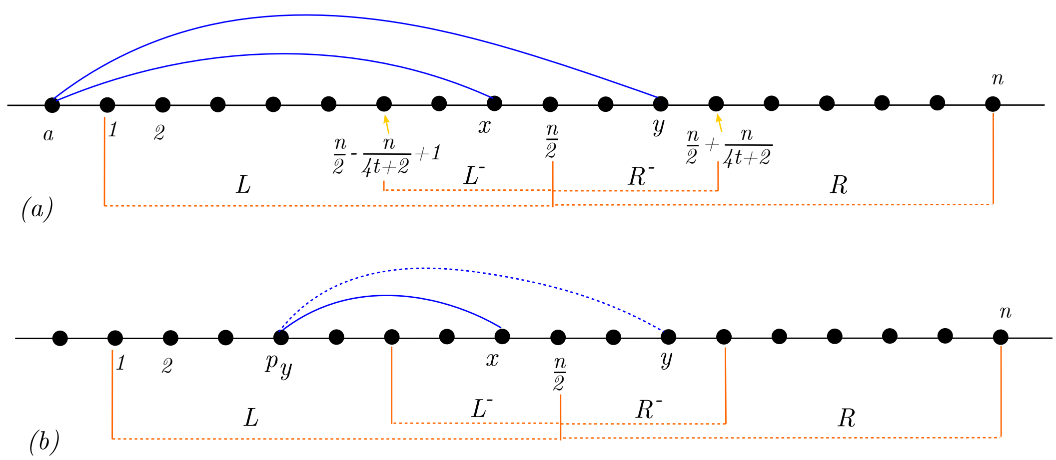

Figure 10: (a) if Q_{xy} contains an integer point $a \leq 0$, then $w(Q_{xy}) > 2|xy|$; (b) A cross edge (p_y, y) if x is not incident to a cross edge.

Let $L = \{1, \ldots, \lfloor n/2 \rfloor\}$ be the left half of P_n and $R = P_n \setminus L$ be the right half; see Figure 10. Note that E_H is a 2-hop t-spanner for both L and R. By induction, E_H has at least $T(\lfloor n/2 \rfloor)$ edges with both endpoints in L and E_H has at least $T(|R|) \geq T(\lfloor n/2 \rfloor)$ edges with both endpoints in R. Note that these two sets are disjoint. Let E_C be the set of edges with exactly one endpoint in L and another in R, called cross edges. Then we have

$$T(n) \geq 2T(\lfloor n/2 \rfloor) + |E_C|. \tag{5.15}$$

We now bound the number of cross edges by $\Omega(n/t)$. To avoid notational clutter, we will drop the floors and ceilings, and assume that n is even. Let $L^- = \{n/2 - n/(4t+2) + 1, n/2 - n/(4t+2) + 2, \ldots, n/2\}$ be the last $n/(4t+2)$ points of L and similarly, $R^- = \{n/2+1, n/2+2, \ldots, n/2 + n/(4t+2)\}$ be the first $n/(4t+2)$ of R.

We claim that for any two points $x \in L^-$ and $y \in R^-$, any t-spanner path Q_{xy} in H of the pair (x,y) must contain only integer points in P_n. Since otherwise, w.l.o.g, assume that Q_{xy} contains a point $a \leq 0$, see Figure 10(a), then

$$\frac{w(Q_{xy})}{|xy|} > \frac{n/2 - n/(4t+2)}{|xy|} = \frac{n/2 - n/(4t+2)}{2n/(4t+2)} = t, \tag{5.16}$$

contradicting that Q_{xy} is a t-spanner path.

Now we continue to bound $|E_C|$. If every vertex in L^- is incident to a cross edge, which is an edge with an endpoint in R, then $|E_C| \geq |L^-| = n/(4t+2)$. Otherwise, there is a point $x \in L^-$ that has no crossing edge. For any $y \in R^-$, there must be a 2-hop path between x and y of the form $\{x, p_y, y\}$ for some p_y in $L \setminus L^-$ (because the 2-hop path does not leave P_n as claimed above). See Figure 10(b). Thus, there is a cross edge (p_y, y) for every $y \in R^-$, meaning that the number of cross edges is $|E_C| \geq |R^-| = n/(4t+2)$. In both cases, we have $|E_C| \geq n/(4t+2)$ and Equation (5.15) becomes

$$T(n) \geq 2T(\lfloor n/2 \rfloor) + \Omega(n/t)\ . \tag{5.17}$$

The recurrence in Equation (5.17) solves to $T(n) = \Omega(\frac{n}{t} \log n)$. □

Proof. [Proof of Lemma 5.1] Our lower bound is obtained by a reduction to a lower bound for 2-hop 2-spanner for the path graph. We construct a tree T as follows. Let P_n be the path with n vertices, where every edge has weight 1. Let T be obtained by attaching to each vertex $i \in P_n$ a distinct vertex s_i via an edge of weight $w_T(s_i, i) = M$ for an integer $M \geq n$. See Figure 11(a) and (b). We call the resulting graph the comb graph, denoted by $\textsc{Comb}_n$.

Let $S = \{s_1, \ldots, s_n\}$. Note that $d_T(s_i, s_j) = 2M + |j - i|$ for any $j \neq j$, and in particular, $d_T(s_i, s_j) \geq 2M$. Let $G = (S, E, w)$ be any non-Steiner 2-spanner for S.

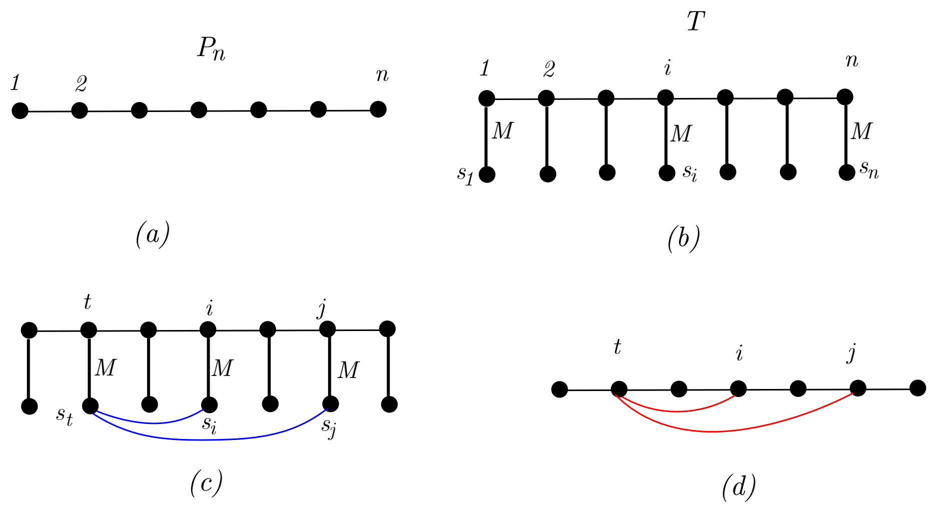

Figure 11: (a) path P_n and (b) tree T. Here $M \geq n$. (c) A path Q_{ij} in G between s_i and s_j containing blue edges and (d) the corresponding 2-hop path Z_{ij} in H between i and j containing red edges.

We construct another graph $H = ([n], E_H, w_H)$ as follows: for every edge (s_i, s_j) in G, we add an edge (i, j) to H of weight $w_H = |j - i|$. We claim that H is a 2-hop 2-spanner of P_n; this will imply that $|E| \geq |E_H| = \Omega(n \log n)$ by Lemma 5.2.

Let i, j be any two vertices in H with $j > i$, and two corresponding vertices s_i and s_j in G, respectively. If there is an edge $(s_i, s_j) \in G$, then $(i, j) \in E_H$ and hence (i, j) has a path (which is an edge) of stretch 1 in H.

We now assume that there is no direct edge from i to j. Let Q_{ij} be the shortest path from s_i to s_j in G. Note that $d_T(s_i, s_j) = 2M + (j - i)$. Since G is a 2-spanner of S, we have

$$w(Q_{ij}) \leq 4M + 2(j - i). \tag{5.18}$$

First, we claim that Q_{ij} contains exactly one point s_t with $t \notin \{i, j\}$. Suppose otherwise, $Q_{i,j}$ contains at least two points s_{t_1}, s_{t_2}. Assume w.l.o.g. that t_1 is closer to i than t_2 on Q_{ij}.Then we have

$$\begin{aligned} w(Q_{ij}) &\geq d_T(s_i, s_{t_1}) + d_T(s_{t_1}, s_{t_2}) + d_T(t_2, s_j) \\ &\geq 2M + 2M + 2M = 6M > 4M + 2(j - i), \end{aligned}$$

since $M \geq n$, contradicting Equation (5.18). Thus, $Q_{ij} = (s_i, s_t, s_j)$ for some $t \neq i, j$; see Figure 11(c). Then $w(Q_{ij}) = 4M + |i - t| + |j - t|$. By Equation (5.18), we have:

$$|i - t| + |j - t| \leq 2(j - i). \tag{5.19}$$

Let $Z_{ij} = \{i, t, j\}$; see Figure 11(d). Clearly, Z_{ij} is a 2-hop path in H by the construction of H. Furthermore, $w_H(Z_{ij}) = |i - t| + |j - t|$ and hence by Equation (5.19), Z_{ij} has stretch 2. Thus, H is a 2-hop 2-spanner of P_n as claimed. □

5.2 Spanners in Polyhedral Terrains with Stretch 2 We now prove the lower bound for stretch 2 as claimed in item 1 in Theorem 1.2.

THEOREM 1.2. *Let $\varepsilon \in (0,1)$ be a parameter.*

1. *There exists a polyhedral terrain and a set P of n points on the terrain such that any 2-spanner for P must have $\Omega(n \log n)$ edges.*
2. *Given any set P of n points in a polyhedral terrain, we can construct a $(2+\varepsilon)$-spanner for P with $\tilde{O}(n/\varepsilon^6)$ edges. The number of edges is $O(n)$ for constant ε.*

Proof. [Proof of item 1] By Lemma 1.6, any lower bound for non-Steiner spanners in tree metrics will imply the same lower bound for point sets in a polyhedral terrain. Thus, by Lemma 5.1, there exists a set of n points P such that any 2-spanner for P must have $\Omega(n \log n)$ edges. □

5.3 Locality-Sensitive Ordering In this section, we prove an $\Omega(\log n)$ lower bound on the size of the left-sided LSO stated in Theorem 1.9. We restate the theorem for convenience.

THEOREM 1.9. *There exists a tree metric T with n points such that any (τ, ρ)-left-sided LSO for T with $\rho = 1$ must have $\tau = \Omega(\log n)$, matching the $O(\log n)$ upper bound by Filtser and Le [42].*

Proof. Let COMB_n be the comb graph constructed in the proof of Lemma 5.1 with terminal set S to be the leaves of COMB_n. Let Σ be a $(\tau, 1)$-left-sided LSO for COMB_n. For any linear ordering $\sigma \in \Sigma$, let σ_S be the linear ordering obtained by removing all vertices not in S from σ. Thus, σ_S only contains (a subset of) vertices in S. Let $\Sigma_S = \{\sigma_S | \sigma \in \Sigma\}$ be the resulting set of linear orderings.

We now construct a non-Steiner 2-spanner for S, denoted by G, as follows. For each ordering $\sigma_S \in \Sigma_S$, let $v^*_{\sigma_S}$ be the leftmost endpoint of σ_S. That is, $v^*_{\sigma_S}$ is the smallest vertex in the ordering σ_S. Then for every $x \in \sigma_S$, we add an edge $(v^*_{\sigma_S}, x)$ to G. The weight of every edge in G is the distance between its endpoints in COMB_n. This completes the construction of G.

We now argue that G is a spanner with stretch 2 for S. Let s_1 and s_2 be any two vertices in S. By the definition of left-sided LSO, there exists a linear ordering $\sigma_S \in \Sigma$ containing both s_1 and s_2 such that for any $x \preceq_{\sigma_S} s_1$ and $y \preceq_{\sigma_S} s_2$, we have $d_{\text{COMB}_n}(x, y) \le d_{\text{COMB}_n}(s_1, s_2)$. Applying this property to $x = v^*_{\sigma_S}$ and $y = s_2$, yields $d_{\text{COMB}_n}(v^*_{\sigma_S}, s_2) \le d_{\text{COMB}_n}(s_1, s_2)$. By a symmetric argument, we alos have $d_{\text{COMB}_n}(v^*_{\sigma_S}, s_1) \le d_{\text{COMB}_n}(s_1, s_2)$. Thus, $d_{\text{COMB}_n}(v^*_{\sigma_S}, s_1) + d_{\text{COMB}_n}(v^*_{\sigma_S}, s_2) \le 2 \cdot d_{\text{COMB}_n}(s_1, s_2)$, which gives $d_G(s_1, s_2) \le 2 \cdot d_{\text{COMB}_n}(s_1, s_2)$ since G contains both $(v^*_{\sigma_S}, s_1)$ and $(v^*_{\sigma_S}, s_2)$.

Observe that for every ordering σ_S, we add at most $|\sigma_S| - 1$ edges to G, where $|\sigma_S|$ is the number of points in σ_S. Thus, on average, we add at most 1 edge per vertex to G. This means that the total number of edges we add to G is at most τn since every vertex appears in at most τ linear orderings in Σ. By Lemma 5.1, we obtain $\tau n = \Omega(n \log n)$, implying that $\tau = \Omega(\log n)$, as claimed. □

5.4 Reliable Spanners To establish a lower bound for reliable spanners in trees, we use a *density-sensitive* version of Lemma 5.2 developed in [59]. Let P be a point set in an interval $[0, L]$ for some $L \ge 1$. We say that P satisfies the unit interval condition if every unit sub-interval of $[0, L]$ contains at most one point in P. Le, Milenković, and Solomonn [59] proved the following.

LEMMA 5.3. (ADAPTED FROM LEMMA 12 [59]) *Let P be a set of $n \ge 2$ points in the interval $[0, L]$ satisfying the unit interval condition. Let H be any Steiner 2-hop t-spanner for P with $t \ge 1$, then $|E_H| = \Omega(n^2 \log n / L)$.*

Le, Milenković, and Solomonn [59] stated their lower bound in Lemma 5.3 for stretch $1+\varepsilon$ where $\varepsilon \in (0, 1/4]$. However, following the same modification we made in the proof of Lemma 5.2, the lower bound holds for *any constant stretch*. Note that in Lemma 5.3, we allow H to contain points in $[0, L]$ that are not in P. The lower bound in Lemma 5.2 also applies to Steiner spanners, but we do not need this property for the proof of Lemma 5.1. However, for proving lower bounds on reliable spanners, we do need the lower bound to hold even for Steiner spanners. We now prove Theorem 1.8 which we restate below.

THEOREM 1.8. *There exists a tree metric T with n points such that any oblivious $\frac{1}{3}$-reliable 2-spanner for $V(T)$ must have $\Omega(n \log n)$ edges.*

Proof. Let $\textsc{Comb}_n$ be the comb graph (with $2n$ vertices) and $\mathcal{D}$ be a distribution of oblivious ν-reliable 2-spanners for $\textsc{Comb}_n$ with $\nu = 1/3$. Our goal is to show that there exists a graph in the support of $\mathcal{D}$ with $\Omega(n \log n)$ edges.

Let the attack set B contain all the internal nodes of $\textsc{Comb}_n$. By definition, $\mathbb{E}_{G \sim \mathcal{D}}[|B^+|] \le (1+\nu)|B|$. Thus, there exists a graph $G \in \mathcal{D}$ such that $|B^+| \le (1+\nu)|B| = 4n/3$. Next we will show that $|E(G)| = \Omega(n \log n)$.

Observe that there are at least $2n - 4n/3 \ge 2n/3$ (leaf) vertices that are not in B^+. Let this set of leaves be A. By definition, $G[V(\textsc{Comb}_n) \setminus B]$ is a 2-spanner for vertices in A. Note that $V(\textsc{Comb}_n) \setminus B$ only contain leaves of $\textsc{Comb}_n$.

Similar to the proof of Lemma 5.1, we construct another graph $H = ([n], E_H, w_H)$ as follows: for every edge (s_i, s_j) in G between two leaves s_i, s_j of $\textsc{Comb}_n$, we add an edge (i,j) to H of weight $w_H = |j - i|$. Let $A_H = \{j | s_j \in A\}$ be the set of points corresponding to vertices in A.

Since $G[V(\textsc{Comb}_n) \setminus B]$ is a 2-spanner for vertices in A, H is a Steiner 2-spanner for A_H. By translation, we can assume that the vertices of H are in the interval $[n, 2n]$, and thus, any Steiner point in H is inside $[0, 3n]$. Since A_H satisfies the unit interval condition with $L = 3n$, by Lemma 5.3, we have

$$|E(H)| = \Omega\left(\frac{|A_H|^2 \log(|A_H|)}{3n}\right). \tag{5.20}$$

Combined with $|A_H| = |A| \ge 2n/3$, this implies $|E(G)| = \Omega(n \log n)$, as desired. ∎

Acknowledgements. This research project started at the *Focused Week: Geometric Spanners*, held at the the Erdős Center in Budapest, Hungary, October 23-29, 2023, supported in part by the ERC grant 882971 "GeoScape". Balázs Keszegh was supported by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences, by the National Research, Development and Innovation Office – NKFIH under the grant K 132696 and FK 132060, by the ÚNKP-23-5 New National Excellence Program of the Ministry for Innovation and Technology from the source of the National Research, Development and Innovation Fund, and by the ERC Advanced Grant "ERMiD" and by the Ministry of Innovation and Technology of Hungary from the National Research, Development and Innovation Fund, financed under the ELTE TKP 2021-NKTA-62 funding scheme. Hung Le was supported by the NSF CAREER award CCF-2237288 and the NSF grant CCF-2121952 and a Google Research Scholar award. Dömötör Pálvölgyi was supported by the János Bolyai Research Scholarship of the Hungarian Academy of Sciences, by the ÚNKP-23-5 New National Excellence Program of the Ministry for Innovation and Technology, and by the ERC Advanced Grant "ERMiD". Research by Tóth was supported in part by the NSF award DMS-2154347.

References

[1] Mohammad Ali Abam, Marjan Adeli, Hamid Homapour, and Pooya Zafar Asadollahpoor. Geometric spanners for points inside a polygonal domain. In *Proc. 31st Symposium on Computational Geometry (SoCG)*, volume 34 of *LIPIcs*, pages 186–197. Schloss Dagstuhl, 2015. doi:10.4230/LIPIcs.SOCG.2015.186.

[2] Mohammad Ali Abam, Mark de Berg, and Mohammad Javad Rezaei Seraji. Geodesic spanners for points on a polyhedral terrain. *SIAM Journal on Computing*, 48(6):1796–1810, 2019. doi:10.1137/18m119358x.

[3] Ittai Abraham and Cyril Gavoille. Object location using path separators. In *Proc. 25th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 188–197, 2006. doi:10.1145/1146381.1146411.

[4] Ittai Abraham, Cyril Gavoille, Dahlia Malkhi, and Udi Wieder. Strong-diameter decompositions of minor free graphs. *Theory Comput. Syst.*, 47(4):837–855, 2010. doi:10.1007/S00224-010-9283-6.

[5] Noga Alon and Baruch Schieber. Optimal preprocessing for answering on-line product queries. Technical report, The Moise and Frida Eskenasy Institute of Computer Science, Tel Aviv University, 1987.

[6] Ingo Althöfer, Gautam Das, David P. Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete Computational Geometry*, 9(1):81–100, 1993. doi:10.1007/BF02189308.

[7] Aaron Archer, Jittat Fakcharoenphol, Chris Harrelson, Robert Krauthgamer, Kunal Talwar, and Éva Tardos. Approximate classification via earthmover metrics. In *Proc. 15th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1079–1087, 2004. URL: http://dl.acm.org/citation.cfm?id=982792.982952.

[8] Srinivasa Arikati, Danny Z. Chen, L. Paul Chew, Gautam Das, Michiel Smid, and Christos D. Zaroliagis. Planar spanners and approximate shortest path queries among obstacles in the plane. In *Proc. 4th European Sympos. Algorithms (ESA)*, volume 1136 of *LNCS*, pages 514–528. Springer, 1996. doi:10.1007/3-540-61680-2_79.

[9] Sunil Arya, Gautam Das, David M Mount, Jeffrey S Salowe, and Michiel Smid. Euclidean spanners: Short, thin, and lanky. In *Proc. 27th ACM Symposium on Theory of Computing (STOC)*, pages 489–498, 1995. doi:10.1145/2488608.248869.

[10] Baruch Awerbuch, Shay Kutten, and David Peleg. On buffer-economical store-and-forward deadlock prevention. *IEEE Transactions on Communications*, 42(11):2934–2937, 1994. doi:10.1109/26.328973.

[11] Baruch Awerbuch and David Peleg. Network synchronization with polylogarithmic overhead. In *Proc. 31st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 514–522, 1990. doi:10.1109/FSCS.1990.89572.

[12] Baruch Awerbuch and David Peleg. Sparse partitions (extended abstract). In *Proc. 31st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 503–513, 1990. doi:10.1109/FSCS.1990.89571.

[13] Baruch Awerbuch and David Peleg. Routing with polynomial communication-space trade-off. *SIAM J. Discret. Math.*, 5(2):151–162, may 1992. doi:10.1137/0405013.

[14] Yair Bartal, Ora Nova Fandina, and Ofer Neiman. Covering metric spaces by few trees. *Journal of Computer and System Sciences*, 130:26–42, 2022. doi:10.1016/j.jcss.2022.06.001.

[15] MohammadHossein Bateni, Erik D. Demaine, MohammadTaghi Hajiaghayi, and Mohammad Moharrami. Plane embeddings of planar graph metrics. *Discret. Comput. Geom.*, 38(3):615–637, 2007. doi:10.1007/S00454-007-1353-4.

[16] Sarita de Berg, Marc van Kreveld, and Frank Staals. The complexity of geodesic spanners. In *Proc. 39th Symposium on Computational Geometry (SoCG)*, volume 258 of *LIPIcs*, pages 16:1–16:16. Schloss Dagstuhl, 2023. doi:10.4230/LIPIcs.SoCG.2023.16.

[17] Sarita de Berg, Tim Ophelders, Irene Parada, Frank Staals, and Jules Wulms. The complexity of geodesic spanners using Steiner points, 2024. arXiv:2402.12110.

[18] Sujoy Bhore and Csaba D. Tóth. Euclidean Steiner spanners: Light and sparse. *SIAM Journal on Discrete Mathematics*, 36(3):2411–2444, 2022. doi:10.1137/22m1502707.

[19] Davide Bilò. Almost optimal algorithms for diameter-optimally augmenting trees. *Theor. Comput. Sci.*, 931:31–48, 2022. doi:10.1016/J.TCS.2022.07.028.

[20] Hans L. Bodlaender, Gerard Tel, and Nicola Santoro. Trade-offs in non-reversing diameter. *Nord. J. Comput.*, 1(1):111–134, 1994.

[21] Prosenjit Bose, Vida Dujmović, Pat Morin, and Michiel Smid. Robust geometric spanners. *SIAM Journal on Computing*, 42(4):1720–1736, 2013. doi:10.1137/120874473.

[22] 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.

[23] 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.

[24] Costas Busch, Ryan LaFortune, and Srikanta Tirthapura. Sparse covers for planar graphs and graphs that exclude a fixed minor. *Algorithmica*, 69(3):658–684, 2014. doi:10.1007/S00453-013-9757-4.

[25] T.-H. Hubert Chan, Donglin Xia, Goran Konjevod, and Andrea Richa. A tight lower bound for the Steiner point removal problem on trees. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, volume 4110 of *LNCS*, pages 70–81. Springer, 2006. doi:10.1007/11830924_9.

[26] Timothy M. Chan, Sariel Har-Peled, and Mitchell Jones. On locality-sensitive orderings and their applications. *SIAM Journal on Computing*, 49(3):583–600, 2020. doi:10.1137/19m1246493.

[27] Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Covering planar metrics (and beyond): O(1) trees suffice. In *Proc. 64th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 2231–2261, 2023. doi:10.1109/FOCS57990.2023.00139.

[28] Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Shortcut partitions in minor-free graphs: Steiner point removal, distance oracles, tree covers, and more. In *Proc. 35th ACM-SIAM Sympos. on Discrete Algorithms (SODA)*, pages 5300–5331, 2024. doi:10.1137/1.9781611977912.191.

[29] Hsien-Chih Chang, Robert Krauthgamer, and Zihan Tan. Almost-linear ϵ-emulators for planar graphs. In *Proc. 54th ACM Symposium on Theory of Computing (STOC)*, pages 1311–1324, 2022. doi:10.1145/3519935.3519998.

[30] Bernard Chazelle. Computing on a free tree via complexity-preserving mappings. *Algorithmica*, 2:337–361, 1987. doi:10.1007/BF01840366.

[31] Chandra Chekuri, Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Embedding k-outerplanar graphs into ℓ_1. *SIAM Journal on Discrete Mathematics*, 20(1):119–136, 2006. doi:10.1137/s0895480102417379.

[32] L. Paul Chew. There is a planar graph almost as good as the complete graph. In *Proc. 2nd ACM Symposium on Computational Geometry (SoCG)*, pages 169–177, 1986. doi:10.1145/10515.10534.

[33] L. Paul Chew. There are planar graphs almost as good as the complete graph. *Journal of Computer and System Sciences*, 39(2):205–219, 1989. doi:10.1016/0022-0000(89)90044-5.

[34] Fan R. K. Chung and Michael R. Garey. Diameter bounds for altered graphs. *J. Graph Theory*, 8(4):511–534, 1984. doi:10.1002/JGT.3190080408.

[35] Kenneth Clarkson. Approximation algorithms for shortest path motion planning. In *Proc. 19th ACM Symposium on*

Theory of Computing (STOC), pages 56–65, 1987. doi:10.1145/28395.28402.
[36] Vincent Cohen-Addad, Éric Colin de Verdière, Dániel Marx, and Arnaud de Mesmay. Almost tight lower bounds for hard cutting problems in embedded graphs. *J. ACM*, 68(4):30:1–30:26, 2021. doi:10.1145/3450704.
[37] Éric Colin de Verdière. Computational topology of graphs on surfaces. In Jacob E. Goodman, Joeseph O'Rourke, and Csaba D. Tóth, editors, *Handbook of Discrete and Computational Geometry*. CRC Press, Boca Raton, FL, 3rd edition, 2017. doi:10.1201/9781315119601.
[38] David Eppstein. Dynamic generators of topologically embedded graphs. In *Proc. 14th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 599–608, 2003. URL: http://dl.acm.org/citation.cfm?id=644108.644208.
[39] Jeff Erickson and Sariel Har-Peled. Optimally cutting a surface into a disk. *Discret. Comput. Geom.*, 31(1):37–59, 2004. doi:10.1007/S00454-003-2948-Z.
[40] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *Journal of Computer and System Sciences*, 69(3):485–497, 2004. doi:10.1016/j.jcss.2004.04.011.
[41] Arnold Filtser. Scattering and sparse partitions, and their applications. In *Proc. 47th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 168 of *LIPIcs*, pages 47:1–47:20. Schloss Dagstuhl, 2020. doi:10.4230/LIPICS.ICALP.2020.47.
[42] Arnold Filtser and Hung Le. Locality-sensitive orderings and applications to reliable spanners. In *Proc. 54th ACM Symposium on Theory of Computing (STOC)*, pages 1066–1079, 2022. doi:10.1145/3519935.3520042.
[43] Arnold Filtser and Hung Le. Low treewidth embeddings of planar and minor-free metrics. In *Proc. 63rd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1081–1092, 2022. doi:10.1109/FOCS54457.2022.00105.
[44] John R. Gilbert, Joan P. Hutchinson, and Robert Endre Tarjan. A separator theorem for graphs of bounded genus. *J. Algorithms*, 5(3):391–407, 1984. doi:10.1016/0196-6774(84)90019-1.
[45] Michael T. Goodrich. Planar separators and parallel polygon triangulation. *J. Comput. Syst. Sci.*, 51(3):374–389, 1995. doi:10.1006/JCSS.1995.1076.
[46] A. Gupta, A. Kumar, and R. Rastogi. Traveling with a Pez dispenser (or, routing issues in MPLS). In *Proc. 42nd IEEE Symposium on Foundations of Computer Science (FOCS)*, 2001. doi:10.1109/sfcs.2001.959889.
[47] Anupam Gupta. Steiner points in tree metrics don't (really) help. In *Proc. 12th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 220–227, 2001. URL: http://dl.acm.org/citation.cfm?id=365411.365448.
[48] Sariel Har-Peled, Manor Mendel, and Dániel Oláh. Reliable spanners for metric spaces. *ACM Transactions on Algorithms*, 19(1):1–27, 2023. doi:10.1145/3563356.
[49] David Ian Heywood, Sarah Catharine Cornelius, and S. J. Carver. *An Introduction to Geographical Information Systems*. Pearson Prentice Hall, 2011.
[50] Lujun Jia, Guolong Lin, Guevara Noubir, Rajmohan Rajaraman, and Ravi Sundaram. Universal approximations for TSP, Steiner tree, and set cover. In *Proc. 37th ACM Symposium on Theory of Computing (STOC)*, pages 386–395, 2005. doi:10.1145/1060590.1060649.
[51] Omri Kahalon, Hung Le, Lazar Milenković, and Shay Solomon. Can't see the forest for the trees: Navigating metric spaces by bounded hop-diameter spanners. In *Proc. ACM Symposium on Principles of Distributed Computing (PODC)*, pages 151–162, 2022. doi:10.1145/3519270.3538414.
[52] Sanjiv Kapoor and Xiang-Yang Li. Geodesic spanners on polyhedral surfaces. In *Proc. 20th International Sympos. on Algorithms and Computation (ISAAC)*, volume 5878 of *LNCS*, pages 213–223. Springer, 2009. doi:10.1007/978-3-642-10631-6_23.
[53] J. Mark Keil. Approximating the complete Euclidean graph. In *Proc. 1st Scandinavian Workshop on Algorithm Theory (SWAT)*, volume 318 of *LNCS*, pages 208–213. Springer, 1988. doi:10.1007/3-540-19487-8_23.
[54] Philip N. Klein. Preprocessing an undirected planar network to enable fast approximate distance queries. In *Proc. 13th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 820–827, 2002. URL: https://dl.acm.org/doi/10.5555/545381.545488.
[55] Philip N. Klein, Shay Mozes, and Christian Sommer. Structured recursive separator decompositions for planar graphs in linear time. In *Proc. 45th ACM Sympos. on Theory of Computing (STOC)*, pages 505–514, 2013. doi:10.1145/2488608.2488672.
[56] Philip N. Klein, Serge A. Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proc. 25th ACM Symposium on Theory of Computing (STOC)*, pages 682–690, 1993. doi:10.1145/167088.167261.
[57] Robert Krauthgamer, Huy L. Nguyen, and Tamar Zondiner. Preserving terminal distances using minors. *SIAM Journal on Discrete Mathematics*, 28(1):127–141, 2014. doi:10.1137/120888843.
[58] Hung Le, Lazar Milenković, and Shay Solomon. Sparse Euclidean spanners with tiny diameter: A tight lower bound. In *Proc. 38th Symposium on Computational Geometry (SoCG)*, pages 54:1–54:15, 2022. doi:10.4230/LIPIcs.SoCG.2022.54.
[59] Hung Le, Lazar Milenković, and Shay Solomon. Sparse Euclidean spanners with optimal diameter: A general and

robust lower bound via a concave inverse-Ackermann function. In *Proc. 39th Symposium on Computational Geometry (SoCG)*, pages 47:1–47:17, 2023. doi:10.4230/LIPICS.SOCG.2023.47.

[60] Hung Le and Shay Solomon. Truly optimal Euclidean spanners. *SIAM Journal on Computing*, pages FOCS19–135–199, 2022. doi:10.1137/20m1317906.

[61] Hung Le and Christian Wulff-Nilsen. Optimal approximate distance oracle for planar graphs. In *Proc. 62nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 363–374, 2022. doi:10.1109/focs52979.2021.00044.

[62] Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979. doi:10.1137/0136016.

[63] Ilan Newman and Yuri Rabinovich. A lower bound on the distortion of embedding planar metrics into Euclidean space. *Discret. Comput. Geom.*, 29(1):77–81, 2003. doi:10.1007/S00454-002-2813-5.

[64] Satish Rao. Small distortion and volume preserving embeddings for planar and Euclidean metrics. In *Proc. 15th ACM Symposium on Computational Geometry (SoCG)*, pages 300–306, 1999. doi:10.1145/304893.304983.

[65] Jim Ruppert and Raimund Seidel. Approximating the d-dimensional complete Euclidean graph. In *Proc. 3rd Canadian Conference on Computational Geometry (CCCG)*, pages 207–210, 1991.

[66] Shay Solomon. Sparse Euclidean spanners with tiny diameter. *ACM Trans. Algorithms*, 9(3):28:1–28:33, 2013. doi:10.1145/2483699.2483708.

[67] Mikkel Thorup. Parallel shortcutting of rooted trees. *J. Algorithms*, 23(1):139–159, 1997. doi:10.1006/JAGM.1996.0829.

[68] Mikkel Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *J. ACM*, 51(6):993–1024, 2004. doi:10.1145/1039488.1039493.

[69] Andrew Chi-Chih Yao. Space-time tradeoff for answering range queries (extended abstract). In *Proc. 14th ACM Symposium on Theory of Computing (STOC)*, pages 128–136, 1982. doi:10.1145/800070.802185.