Optimal Euclidean Tree Covers

Hsien-Chih Chang
Department of Computer Science, Dartmouth College, Hanover, NH, USA

Jonathan Conroy
Department of Computer Science, Dartmouth College, Hanover, NH, USA

Hung Le
Manning CICS, UMass Amherst, MA, USA

Lazar Milenković
Tel Aviv University, Israel

Shay Solomon
Tel Aviv University, Israel

Cuong Than
Manning CICS, UMass Amherst, MA, USA

Abstract

A $(1+\varepsilon)$-*stretch tree cover* of a metric space is a collection of trees, where every pair of points has a $(1+\varepsilon)$-stretch path in one of the trees. The celebrated *Dumbbell Theorem* [Arya et al. STOC'95] states that any set of n points in d-dimensional Euclidean space admits a $(1+\varepsilon)$-stretch tree cover with $O_d(\varepsilon^{-d} \cdot \log(1/\varepsilon))$ trees, where the O_d notation suppresses terms that depend solely on the dimension d. The running time of their construction is $O_d(n \log n \cdot \frac{\log(1/\varepsilon)}{\varepsilon^d} + n \cdot \varepsilon^{-2d})$. Since the same point may occur in multiple levels of the tree, the *maximum degree* of a point in the tree cover may be as large as $\Omega(\log \Phi)$, where Φ is the aspect ratio of the input point set.

In this work we present a $(1+\varepsilon)$-stretch tree cover with $O_d(\varepsilon^{-d+1} \cdot \log(1/\varepsilon))$ trees, which is optimal (up to the $\log(1/\varepsilon)$ factor). Moreover, the maximum degree of points in any tree is an *absolute constant* for any d. As a direct corollary, we obtain an optimal routing scheme in low-dimensional Euclidean spaces. We also present a $(1+\varepsilon)$-stretch *Steiner* tree cover (that may use Steiner points) with $O_d(\varepsilon^{(-d+1)/2} \cdot \log(1/\varepsilon))$ trees, which too is optimal. The running time of our two constructions is linear in the number of edges in the respective tree covers, ignoring an additive $O_d(n \log n)$ term; this improves over the running time underlying the Dumbbell Theorem.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Tree cover, spanner, Steiner point, routing, bounded-degree, quadtree, net-tree

Digital Object Identifier 10.4230/LIPIcs.SoCG.2024.37

Related Version *Full Version*: https://arxiv.org/abs/2403.17754

Funding Hung Le and Cuong Than are supported by the NSF CAREER Award No. CCF-2237288 and an NSF Grant No. CCF-2121952. Shay Solomon is funded by the European Union (ERC, DynOpt, 101043159). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. Shay Solomon is also supported by the Israel Science Foundation(ISF) grant No.1991/1. Shay Solomon and Lazar Milenković are supported by a grant from the United States-Israel Binational Science Foundation (BSF), Jerusalem, Israel, and the United States National Science Foundation(NSF).

40th International Symposium on Computational Geometry (SoCG 2024).
Editors: Wolfgang Mulzer and Jeff M. Phillips; Article No. 37; pp. 37:1–37:15
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Let M be a given metric space with distance function δ, and X be a finite set of points in M. A *tree cover* for (M, X) is a collection of trees $\mathcal{F}$, each of which consists of (only) points in X as vertices and abstract edges between vertices, such that between every two points x and y in X, one has $\delta_M(x, y) \leq \delta_T(x, y)$ for every tree T in $\mathcal{F}$. A tree cover $\mathcal{F}$ has *stretch* α if for every two points x and y in X, there is a tree T in $\mathcal{F}$ that preserves the distance between x and y up to α factor: $\delta_T(x, y) \leq \alpha \cdot \delta_M(x, y)$. We call such $\mathcal{F}$ an *α-tree cover* of X. In this paper, we will focus on the scenario where M is the d-dimensional Euclidean space for some constant $d = O(1)$. It is not hard to see that, in this case, the edges can be drawn as line segments in $\mathbb{R}^d$ between the corresponding two endpoints, with weights equal to their Euclidean distances. If we relax the condition so that trees in $\mathcal{F}$ may have other points from M (called *Steiner points*) as vertices instead of just points from X, the resulting tree cover is called a *Steiner tree cover*.

Constructions of tree covers, due to their algorithmic significance, are subject to growing research attention [4, 2, 3, 16, 21, 11, 7, 5, 9, 10]; by now generalizations in various metric spaces and graphs are well-explored. The main measure of quality for tree cover is its *size*, that is, the number of trees in a tree cover $\mathcal{F}$. The existence of a small tree cover provides a framework to solve distance-related problems by essentially reducing them to trees. Exemplified applications include distance oracles [9, 10], labeling and routing schemes [25, 18], spanners with small hop diameters [18], and bipartite matching [1].

The celebrated *Dumbbell Theorem* by Arya, Das, Mount, Salowe, and Smid [3] from almost thirty years ago demonstrated that in d-dimensional Euclidean space, any point set X has a tree cover of stretch $1 + \varepsilon$ that uses only $O_d(\varepsilon^{-d} \cdot \log(1/\varepsilon))$ trees.[1] Moreover, the tree cover can be computed within time $O_d\left(n \log n \cdot \frac{\log(1/\varepsilon)}{\varepsilon^d} + n \cdot \varepsilon^{-2d}\right)$, where n is the number of points in X. In the Euclidean plane (when $d = 2$), this gives us a tree cover of size $O(\varepsilon^{-2} \cdot \log(1/\varepsilon))$. The theorem has a long and complex proof, which spans a chapter in the book of Narasimhan and Smid [22]. A few years ago, this theorem was generalized for doubling metrics[2] by Bartal, Fandina, and Neiman [5], who achieved the same bound as [3] via a much simpler construction;[3] the running time of their construction was not analyzed. In the constructions by [3, 5], same point may have multiple copies in different levels of the tree, hence the maximum degree of points[4] may be as large as $\Omega(\log \Phi)$, where Φ is the aspect ratio of the input point set; see Section 1.2 and the full version of the paper for a more detailed discussion.

Since the number of trees provided by the two known constructions [3, 5] matches the packing bound ε^{-d} (up to a logarithmic factor), it is tempting to conjecture that this bound is tight. However, there is a gap between this upper bound and the best lower bound we have, which comes indirectly from $(1 + \varepsilon)$-stretch *spanners*. For any parameter $\alpha \geq 1$, a *Euclidean α-spanner* for any d-dimensional point set is a weighted graph spanning the input point set, whose edge weights are given by the Euclidean distances between the points, that approximates all the original pairwise Euclidean distances within a factor of α. We note that

---

1 The O_d notation suppresses terms that depend solely on the dimension d.

2 The *doubling dimension* of a metric space (M, δ) is the smallest value ddim such that every ball in M can be covered by 2^{ddim} balls of half the radius. A metric δ is called *doubling* if its doubling dimension is constant.

3 In high-dimensional Euclidean spaces the upper bound in [5] improves over that of [3], since the O_d notation in [3] and [5] suppress multiplicative factors of $d^{O(d)}$ and $2^{O(d)}$, respectively.

4 The *degree* of a point is the number of edges incident to it.

an α-spanner can be obtained directly by taking the union of all trees in any α-tree cover for the input point set. The $\Omega(n \cdot \varepsilon^{-d+1})$ size lower bound for $(1+\varepsilon)$-spanners [19, Theorem 1.1] directly implies that any $(1+\varepsilon)$-tree cover must contain $\Omega(\varepsilon^{-d+1})$ trees. This is an ε^{-1}-factor away from the packing bound. In particular, in the Euclidean plane, there is a gap between the upper bound of $O(\varepsilon^{-2})$ and the lower bound of $\Omega(\varepsilon^{-1})$. One can extend the notions of spanner by introducing Steiner points as well, which are additional points that are not part of the input. A weaker $\Omega(\varepsilon^{(-d+1)/2})$ lower bound can be obtained for *Steiner tree cover*, from the $\Omega(n/\sqrt{\varepsilon})$ size lower bound for Steiner $(1+\varepsilon)$-spanner in $\mathbb{R}^2$ [19, Theorem 1.4], and the $\Omega(n/\varepsilon^{(d-1)/2})$ size lower bound in general $\mathbb{R}^d$ [6].

Short survey on tree covers. There are many papers published on tree covers in recent years, with subtle variations in their definitions due to differences in main objectives and applications. In the full version of the paper, we attempt to summarize the best upper and lower bounds known to our knowledge, including some results that were not in any earlier literature.

1.1 Main Results

We improve the longstanding bound on the number of trees for Euclidean tree cover by a factor of $1/\varepsilon$, for any constant-dimensional Euclidean space.[5] In view of the aforementioned lower bound [19, 6], this is optimal up to the $\log(1/\varepsilon)$ factor. Roughly speaking, we show that the packing bound barrier (incurred in both [3] and [5]) can be replaced by the number of ε-angled cones needed to partition $\mathbb{R}^d$; for more details, refer to Section 1.2.

▶ **Theorem 1.** *For every set of points in $\mathbb{R}^d$ and any $0 < \varepsilon < 1/20$, there exists a tree cover with stretch $1+\varepsilon$ and $O_d(\varepsilon^{-d+1} \cdot \log(1/\varepsilon))$ trees. The running time of the construction is $O_d(n \log n + n \cdot \varepsilon^{-d+1} \cdot \log(1/\varepsilon))$.*

We note our construction is faster than that of the Dumbbell Theorem of [3] by more than a multiplicative factor of ε^{-d}.

In addition, we demonstrate that the bound on the number of trees can be quadratically improved using *Steiner points*; in $\mathbb{R}^2$ we can construct a Steiner tree cover with stretch $1+\varepsilon$ using only $O(1/\sqrt{\varepsilon})$ many trees. The result generalizes for higher dimensions. In view of the aforementioned lower bound [19, 6], this result too is optimal up to the $\log(1/\varepsilon)$ factor.

▶ **Theorem 2.** *For every set of points in $\mathbb{R}^d$ and any $0 < \varepsilon < 1/20$, there exists a Steiner tree cover with stretch $1+\varepsilon$ and $O_d(\varepsilon^{(-d+1)/2} \cdot \log(1/\varepsilon))$ trees. The running time of the construction is $O_d(n \log n + n \cdot \varepsilon^{(-d+1)/2} \cdot \log(1/\varepsilon))$.*

1.1.1 Bounded degree tree cover

Although the number of trees in the tree cover is the most basic quality measure, together with the stretch, another important measure is the *degree*. One can optimize the maximum degree of a point in any of the trees, or to optimize the maximum degree of a point over all trees – both these measures are of theoretical and practical importance.

[5] As with [3], the O_d notation in our bound suppresses a multiplicative factor of $d^{O(d)}$, which should be compared to the multiplicative factor of $O(1)^d$ suppressed in the bound of [5]. Thus, our results improve over that of [5] only under the assumption that ε is sufficiently small with respect to the dimension d; this assumption should be acceptable since the focus of this work, as with the great majority of the work on Euclidean spanners, is low-dimensional Euclidean spaces.

Both the Dumbbell Theorem [3] and the BFN construction [5] use *copies* of the same point in multiple trees, and even in different levels of the same tree. Consequently, each point may have up to $\log \Phi$ copies, which can be viewed as distinct *nodes* of the tree, where Φ is the aspect ratio of the input point set. The Dumbbell trees have bounded *node*-degree (which is improved to degree 3 in [24]), but the maximum *point*-degree in any tree could still be $\Theta(\log \Phi)$ after reidentifying all the copies of the points. The construction of [5] may also incur a point-degree of $\Omega(\log \Phi)$ in any of the trees.[6]

In the full version of the paper, we strengthen Theorem 1 by achieving a constant degree for each point in any of the trees; in fact, our bound is an absolute constant in any dimension. As a result, the maximum degree of a point over all trees is $O_d(\varepsilon^{-d+1} \cdot \log(1/\varepsilon))$; this is optimal up to the $\log(1/\varepsilon)$ factor, matching the average degree (or size) lower bound of spanners mentioned above [19].

Routing. We highlight one application of our bounded degree tree cover to efficient routing.

▶ **Theorem 3.** *For any set of points in $\mathbb{R}^d$ and any $0 < \varepsilon < 1/20$, there is a compact routing scheme with stretch $1+\varepsilon$ that uses routing tables and headers with $O_d(\varepsilon^{-d+1} \log^2(1/\varepsilon) \cdot \log n)$ bits of space.*

Our routing scheme uses smaller routing tables compared to the routing scheme of Gottlieb and Roditty [15], which uses routing tables of $O(\varepsilon^{-d} \log n)$ bits. At a high level, we provide an efficient reduction from the problem of routing in low-dimensional Euclidean spaces to that in trees; more specifically, we present a new labeling scheme for determining the right tree to route on in the tree cover of Theorem 1. Having determined the right tree to route on, our entire routing algorithm is carried out on that tree, while the routing algorithm of [15] is carried out on a spanner; routing in a tree is clearly advantageous over routing in a spanner, also from a practical perspective. The details are omitted due to space constraints; refer to [15] for the definition of the problem and relevant background.

1.2 Technical Highlights

1.2.1 Achieving an optimal bound on the number of trees

The tree cover constructions of [3] and [5] achieve the same bound of $O(\varepsilon^{-d} \cdot \log(1/\varepsilon))$ on the number of trees, which is basically the packing bound $O(\varepsilon^{-d})$. The Euclidean construction of [3] is significantly more complex than the construction of [5] that applies to the wider family of doubling metrics. Here we give a short overview of the simpler construction of [5]; then we describe our Euclidean construction, aiming to focus on the geometric insights that we employed to breach the packing bound barrier.

The starting point of [5] is the standard *hierarchy of 2^w-nets* $\{N_w\}$ [17], which induces a hierarchical *net-tree*.[7] Each net N_w is greedily partitioned into a collection of $\Theta(\frac{2^w}{\varepsilon})$-*sub-nets* $N_{w,t}$, which too are hierarchical. For a fixed level w, the number of sub-nets $\{N_{w,t}\}$ is bounded by the packing bound $O(\varepsilon^{-d})$, and each of them is handled by a different tree via a straightforward clustering procedure. Naïvely this introduces a $\log \Phi$ factor to the number

[6] Even node-degrees may blow up in the construction of [5], but it appears that a simple tweak of their construction can guarantee a node-degree of $\varepsilon^{-O(d)}$.

[7] The standard notation in the literature on doubling metrics, including [5], uses index i instead of w to refer to levels or distance scales; however, this paper focuses on Euclidean constructions, and we view it instructive to use a different notation.

of trees, each corresponding to a level (Φ is the aspect ratio of the point set). The key observation to remove the dependency on the aspect ratio is that two far apart levels are more or less *independent*, and one can pretty much use the same collection of trees for both. More precisely, the levels are partitioned into $\ell \coloneqq \log(1/\varepsilon)$ *congruence classes* $I_0, I_1, \ldots, I_{\ell-1}$, where $I_j \coloneqq \{w \mid w \equiv j \pmod{\ell}\}$. Since distances across different levels of the same class I_j differ by at least a factor of $1/\varepsilon$, it follows that all sub-nets $\{N_{w,t}\}_{w\in I_j}$ can be handled by a single tree via a greedy hierarchical clustering. Now the total number of trees is the number of sub-nets in one level, which is $O(\varepsilon^{-d})$, times the number of congruence classes $\log(1/\varepsilon)$.

Taking a bird's eye view of the construction of [5], the following two-step strategy is used to handle pairwise distances within each congruence class I_j:

1. *Reduce the problem from the entire congruence class I_j to a single level $w \in I_j$.* This is done by a simple greedy procedure.
2. *Handle each level $w \in I_j$ separately.* This is done by a simple greedy clustering to the sub-nets $\{N_{w,t}\}$.

In Euclidean spaces, we shall use *quadtree* which is the natural analog of the hierarchical net-tree. We too employ the trick of partitioning all levels in the hierarchy to congruence classes [7, 5, 19, 1] and handle each one separately, and follow the above two-step strategy. However, the way we handle each of these two steps deviates significantly from [5].

Step 1: Reduce the problem to a single level. At any level w, we handle every quadtree cell of width 2^w separately. Every cell is partitioned into subcells from level $w-\ell$ of width $\varepsilon \cdot 2^w$, and each non-empty cell contains a single representative assigned by the construction at level $w-\ell$. At level w, we construct a *partial* $(1+\varepsilon)$*-tree cover*, which roughly speaking only preserves distances between pairs of representatives that are at distance roughly 2^w from each other; this is made more precise in the description of Step 2 below. Let $\tau(\varepsilon)$ be the number of trees required for such a partial tree cover. To obtain a tree cover for all points in the current level-w cell, we simply merge the aforementioned partial tree cover constructed for the level-$(w-\ell)$ representatives with the tree cover obtained previously for the points in the subcells. Finally, we choose one of those level-$(w-\ell)$ representatives as the level-w representative of the current cell, and proceed to level $w+\ell$ of the construction.

To achieve the required stretch bound, it is sufficient to guarantee that for every pair of points (p,q), some quadtree cell of side-length proportional to $\|pq\|$ would contain both p and q. Alas, this is impossible to achieve with a single quadtree. To overcome this obstacle, we use a result by Chan [8]: there exists a collection of $\Theta(d)$ carefully chosen *shifts* of the input point set, such that in at least one shift there is a quadtree cell of side-length at most $\Theta(d) \cdot \|pq\|$ that contains both p and q. The number of trees in the cover grows by a factor of $O_d(1)$. Consequently, if each cell can be handled using $\tau(\varepsilon)$ trees, then ranging over all the $\log(1/\varepsilon)$ congruence classes and all the shifts, the resulting tree cover consists of $\tau(\varepsilon) \cdot \log(1/\varepsilon) \cdot O_d(1)$ trees; see Lemma 7 for a more precise statement. The full details of the reduction are in Section 2.1.

Step 2: Handling a single level. Handling a single level is arguably the more interesting step, since this is where we depart from the general packing bound argument that applies to doubling metrics, and instead employ a more fine-grained geometric argument. We next give a high-level description of the tree cover construction for a single level w. For brevity, in this discussion, we focus on the 2-dimensional construction that does not use Steiner points. The full details of this construction, as well as of the generalization for higher dimension and the Steiner tree cover construction, are given in Sections 2.2 and 2.3.

We consider a single 2-dimensional quadtree cell of side-length $\Delta := 2^w$ at level w, which is subdivided into subcells of side-length $\varepsilon \cdot 2^w$. Every level-$(w-\ell)$ cell has a representative and our goal is to construct a partial tree cover for any pair of representatives that are at a distance between $\Delta/10$ and Δ. (The final constants are slightly different; here we choose 10 for simplicity.) To this end, we select a collection of $\Theta(1/\varepsilon)$ *directions*. For each direction ν, we partition the plane into *strips* of width $\varepsilon\Delta$, each strip parallel to ν. We then shift each such partition orthogonally by $\varepsilon\Delta/2$; we end up with a collection of $2 \cdot \Theta(1/\varepsilon)$ partitions, two for each direction. We call these partitions the *major strip partitions*. Observe that for every pair of representative points p and q, there is at least one major strip partition in some direction, such that both p and q are contained in the same strip. Crucially, we show that for every strip S in a partition P, there is a collection of $O(1)$ trees that preserves distances between all points p and q in strip S that are at distance between $\Delta/100$ and Δ. The key observation is that, since the strips in the same partition P are *disjoint* by design, the $O(1)$-many trees for each strip of P can be combined into $O(1)$ forests. Thus the total number of forests is $O(1/\varepsilon)$.

To construct a collection of trees preserving distances within a single strip S, we first subdivide the strip S. If S is in direction ν, we partition S into sub-strips *orthogonal to* ν, each of width $\Delta/20$. We call this a *minor strip partition*. Observe that if points p and q are at distance $\geq \Delta/10$, they are in different sub-strips of the minor strip partition. For every pair of sub-strips S_1 and S_2 in the minor strip partition, we construct a single tree that preserves distances between points in S_1 and S_2 to within a factor of $1+\varepsilon$. There are $O(1)$ sub-strips in the minor strip partition, so overall only $O(1)$ trees are needed for any strip S.

1.2.2 Bounding the degree

The tree cover construction described above achieves the optimal bound on the number of trees, but the degree of points could be arbitrarily large. While the previous tree cover constructions [3, 5] incur unbounded degree, the Euclidean construction of [3], when restricted to a single level in the hierarchy, achieves an absolute constant degree.[8]

In our construction, when restricted to a single level, the degree of points can be easily bounded by $O(1/\varepsilon^2)$. However, in contrast to [3], our goal is to achieve this bound for the entire tree, *across all levels* of the hierarchy. In particular, if we achieve this goal, the total degree of each point over all trees will be $O(\varepsilon^{-1} \cdot \log(1/\varepsilon)))$ ($O(\varepsilon^{-d+1} \cdot \log(1/\varepsilon))$ in general), which is optimal (up to logarithmic factor) due to the aforementioned lower bound [6].
To achieve this goal, we strengthen the aforementioned two-step strategy as follows. The details appear in the full version of the paper.

1. In the reduction from the entire congruence class I_j to a single level $w \in I_j$, the challenge is not to overload the same representative point over and over again across different levels of I_j. To this end, we refine a *degree reduction technique*, originally introduced by Chan et al. [7] to achieve a bounded degree for $(1+\varepsilon)$-stretch net-tree spanners in arbitrary doubling metrics. The technique of [7] is applied on a *bounded-arboricity* net-tree spanner, first by orienting its edges to achieve bounded out-degree for all points. Then, apply a greedy *edge-replacement* process, where the edges are scanned in nondecreasing order of their level (or weight), and any incoming edge (u, v) leading to a high-in-degree point v is replaced by an edge leading to an incoming neighbor w of v in a sufficiently lower level, with $\|wv\| \leq \varepsilon\|uv\|$. It is shown

[8] Although in the original paper of [3] (as well as in [22]) the bound is not an absolute constant, it was shown in [24] that an absolute constant bound can be obtained. Nonetheless, overlaying all levels of the hierarchy leads to a final degree bound of $\Theta(\log \Phi)$.

that this process terminates with a bounded-degree spanner, where the degree bound is quadratic in the out-degree bound (arboricity) of the original spanner, and the stretch bound increases only by an additive factor of $O(\varepsilon)$.

We would like to apply this technique on every tree in the tree cover *separately*; if instead we were to apply it on the union of the trees, we would create cycles; resolving them blows up the number of trees in the cover. We demonstrate that by working on each tree separately, not only does the greedy edge-replacement process reduce the degree in each tree to an absolute constant, but it also keeps the tree cycle-free as well as provides the required stretch bound. In fact, it turns out to be advantageous to operate on each tree separately rather than on their union, since this way the out-degree bound in a single tree reduces to 1, which directly improves the total degree bound *over all trees* to be linearly depending on $1/\varepsilon$ rather than quadratically. This is the key to achieving an optimal degree bound both within each tree as well as over all trees.

2. When handling a single level individually, the degree of points can be easily bounded by $O(1/\varepsilon^2)$ as mentioned. However, we would like to achieve an absolute constant bound at each level. Recall that, for every pair of sub-strips S_1 and S_2 in the minor strip partition of some strip S, we construct a single tree that preserves distances between points in S_1 and S_2 to within a factor of $1+\varepsilon$; this tree is in fact a *star*. Perhaps surprisingly, every such star can be transformed into a *binary tree* via a simple greedy procedure, with the stretch bound increased by just a factor of $1 + O(\varepsilon \log(1/\varepsilon))$.

1.3 Organization

In Section 2, we present the construction of tree covers with an optimal number of trees in both non-Steiner and Steiner settings, proving Theorem 1 and Theorem 2. In the full version of the paper, we reduce the degree of every tree in the (non-Steiner) tree cover an absolute constant; and we show some applications of our tree cover to routing, proving Theorem 3.

2 Optimal Tree Covers for Euclidean Spaces

2.1 Reduction to Partial Tree Cover

Let X be a set of points in $\mathbb{R}^d$. For any two points p and q in X, we use $\|pq\|$ to denote their Euclidean distance. Without loss of generality we assume that the minimum distance between any two points in X is 1.

▶ **Lemma 4** (Cf. [8, 14]). *Let $L > 0$ be an arbitrary real parameter. Consider any two points $p, q \in [0, L)^d$, and let $\mathcal{T}$ be the infinite quadtree of $[0, 2L)^d$. For $D \coloneqq 2\lceil d/2 \rceil$ and $i = 0, \ldots, D$, let $\nu_i \coloneqq (iL/(D+1), \ldots, iL/(D+1))$. Then there exists an index $i \in \{0, \ldots, D\}$, such that $p + \nu_i$ and $q + \nu_i$ are contained in a cell of $\mathcal{T}$ with side-length at most $(4\lceil d/2 \rceil + 2) \cdot \|pq\|$.*

▶ **Definition 5.** *We call two points (μ, Δ)-*far *if their distance is in $[\Delta/\mu, \Delta]$.*

▶ **Definition 6.** *A (μ, Δ)-*partial tree cover *for $X \subset \mathbb{R}^d$ with stretch $(1+\varepsilon)$ is a tree cover with the following property: for every two (μ, Δ)-far points p and q, there is a tree T in the cover such that $\delta_T(p, q) \leq (1+\varepsilon) \cdot \|pq\|$.*

▶ **Lemma 7** (Reduction to partial tree cover). *Let X be a set of points in $\mathbb{R}^d$, and let ε be a number in $(0, 1/20)$. Suppose that for every $\mu > 0$, every set of points in $\mathbb{R}^d$ with diameter Δ admits a (μ, Δ)-partial tree cover with stretch $(1+\varepsilon)$, size $\tau(\varepsilon, \mu)$ and diameter of each tree at most $\gamma\Delta$ for some $\gamma \geq 1$. Then X admits a tree cover with stretch $(1+\varepsilon)$ and size $O(d \cdot \log \frac{\gamma \cdot d\sqrt{d}}{\varepsilon} \cdot \tau(\varepsilon, \mu))$ with $\mu \coloneqq 10d\sqrt{d}$.*

Proof. Assume without loss of generality that the smallest coordinate of a point in X is 0 and let L be the largest coordinate in X. Let $D \coloneqq 2\lceil d/2 \rceil$ and let $\mathcal{Q}$ be the quadtree as in Lemma 4. For $i \in \{0, \ldots, D\}$, let $\mathcal{Q}_i$ be $\mathcal{Q}$ shifted by $-\nu_i = (-iL/(D+1)), \ldots, -iL/(D+1))$.

Constructing the tree cover. Let $\ell \coloneqq \log \frac{\gamma d \sqrt{d}}{\varepsilon}$ and let $\mu \coloneqq 10d\sqrt{d}$. (Assume for simplicity that ℓ is an integer.) Fix some $i \in \{0, \ldots, D\}$, $j \in \{0, \ldots, \ell - 1\}$, and $k \in \{1, \ldots, \tau(\varepsilon, \mu)\}$. We proceed to construct *tree* $T_{i,j,k}$. Consider the *congruence class* $I_j \coloneqq \{z \geq 0 \mid z \equiv j \pmod{\ell})\}$. The following construction is done for every $z \in I_j$ in increasing order. Consider the level-w quadtree $\mathcal{Q}_i$, with cells of width 2^w. If $w < \ell$, for each level-w cell C, construct the kth among $\tau(\varepsilon, \mu)$ trees from the $(\mu, 2^w)$-partial tree cover on the points in C, and root it at an arbitrary point in C. For $w \geq \ell$, consider the subdivision of level-w cell into subcells of level $w - \ell$. Let X' be a subset of X consisting of all the roots of the previously built subtrees in subcells of levels $w - \ell$. Let $\Delta_w \coloneqq 2^w\sqrt{d}$, and observe that Δ_w is an upper-bound on the diameter of X'. Construct a (μ, Δ_w)-partial tree cover for X' with $\tau(\varepsilon, \mu)$ trees, and let T be the kth tree of the $\tau(\varepsilon, \mu)$ trees constructed. Take the previously built subtrees rooted at X', and construct a new tree by identifying their roots with the vertices of T. Root this new tree arbitrarily. The tree $T_{i,j,k}$ is the final tree obtained after iterating over every $z \in I_j$.

We prove the following two claims inductively.

▷ Claim 8. Let $T^w_{i,j,k}$ be a tree constructed at level w for $i \in \{0, \ldots, D\}$, $j \in \{0, \ldots, \ell - 1\}$, $k \in \{1, \ldots, \tau(\varepsilon, \mu)\}$ and $w \in I_j$.

1. $T^w_{i,j,k}$ is a tree.
2. $T^w_{i,j,k}$ has diameter ϕ_w at most $2\gamma\Delta_w$.

Proof. We prove the claim by induction over the level $w \in I_j$.

1. The base case holds because the graph $T_{i,j,k}$ is initialized as a tree. For the induction step, consider some level $w \in I_j$ that is at least ℓ. At this stage we construct a tree T with vertex set consisting of representatives of the level $w - \ell$, and attach the trees rooted at each of the representatives we constructed previously to T. This is clearly a tree and the induction step holds.
2. The base case holds because the diameter of each tree is at most $\gamma\Delta_w$, as guaranteed in the statement of Lemma 7. For the induction step, we have $\gamma\Delta_w + 2\gamma\Delta_{w-\ell} = \gamma\Delta_w + 2\gamma\frac{\Delta_w\varepsilon}{\gamma d\sqrt{d}} \leq 2\gamma\Delta_w$. ◁

▷ Claim 9. The number of trees in the cover is $O(d \log \frac{\gamma d\sqrt{d}}{\varepsilon} \cdot \tau(\varepsilon, \mu))$.

Proof. The tree cover consists of trees $T_{i,j,k}$ ranging over $(D+1) \cdot \ell \cdot \tau(\varepsilon, \mu) = (2\lceil d/2 \rceil + 1) \cdot \log \frac{\gamma d\sqrt{d}}{\varepsilon} \cdot \tau(\varepsilon, \mu)$ indices. ◁

▷ Claim 10. For every two points $p, q \in X$, there is a tree T in the cover such that $\delta_T(p, q) \leq (1 + \varepsilon) \cdot \|pq\|$, where $\delta_T(p, q)$ is the distance between p and q in T.

Proof. By Lemma 4, there exists a cell C in one of the $D + 1$ quadtrees which contains both p and q and has side-length $2^w \leq (4\lceil d/2 \rceil + 2) \cdot \|pq\| \leq 5d \cdot \|pq\|$. Let $\mathcal{Q}_i$ be such a quadtree, where $0 \leq i \leq D$, and let $0 \leq j \leq \ell - 1$ be such that $j \equiv w \pmod{\ell}$. Observe that p and q are (μ, Δ_w)-far. If $w < \ell$, we constructed a (μ, Δ_w)-partial tree cover of C, so the claim holds. Otherwise suppose $w \geq \ell$. Recall that in the construction of the tree cover, we considered a subdivision of a level-w cell (of side-length 2^w) into smaller subcells of level $w - \ell$. For each subcell we choose a representative and constructed a tree cover on top of them. Let p' (resp. q') denote the representative of p (resp. q) in the subcell at level $w - \ell$. We claim that p'

and q' are (μ, Δ_w)-far, where $\Delta_w = 2^w\sqrt{d}$ denotes the diameter of the cell at level w. The bound $\|p'q'\| \leq \Delta_w$ follows from the fact that p' and q' are both in cell C. The distance can be lower-bounded as follows.

$$\begin{aligned}\|p'q'\| &\geq \|pq\| - 2\Delta_{w-\ell} \geq \frac{2^w}{5d} - 2\cdot\frac{\varepsilon\cdot 2^w}{\gamma d\sqrt{d}}\sqrt{d} \\ &= 2^w\left(\frac{1}{5d} - \frac{2\varepsilon}{\gamma d}\right) \geq \frac{2^w}{10d} = \frac{\Delta_w}{\mu} \qquad \text{as } \varepsilon < \frac{1}{20},\ \gamma \geq 1, \text{ and } \mu = 10d\sqrt{d}\end{aligned}$$

In other words, the representatives p' and q' are (μ, Δ_w)-far, meaning that one of the $\tau(\varepsilon, \mu)$ trees T in the partial tree cover for cell C will preserve the stretch between p' and q' up to a factor of $(1+\varepsilon)$. The distance between p and q in this tree can be upper bounded as follows.

$$\begin{aligned}\delta_T(p,q) &\leq \delta_T(p,p') + \delta_T(p',q') + \delta_T(q,q') \\ &\leq (1+\varepsilon)\cdot\|p'q'\| + \delta_T(p,p') + \delta_T(q,q') \\ &\leq (1+\varepsilon)\cdot(\|p'p\| + \|pq\| + \|qq'\|) + \delta_T(p,p') + \delta_T(q,q') \\ &\leq (1+\varepsilon)\cdot(\|pq\| + 2\Delta_{w-\ell}) + 2\phi_{w-\ell} \\ &\leq (1+\varepsilon)\cdot(\|pq\| + 2\Delta_{w-\ell}) + 4\gamma\Delta_{w-\ell} \qquad \text{by Claim 8} \\ &\leq (1+\varepsilon)\cdot\left(\|pq\| + 6\Delta_w\cdot\frac{\varepsilon}{d\sqrt{d}}\right) \\ &= (1+O(\varepsilon))\cdot\|pq\|\end{aligned}$$

Stretch $1+\varepsilon$ can be obtained by appropriate scaling. ◁

Claims 8–10 imply that the resulting construction is a tree cover with stretch $(1+\varepsilon)$ and $O(d\log\frac{\gamma d\sqrt{d}}{\varepsilon}\cdot\tau(\varepsilon,\mu))$ trees, as required. This concludes the proof of Lemma 7. ◀

Running Time. Let $\mathrm{Time}_{\mu,\Delta}(n)$ be the time needed to construct a (μ, Δ)-partial tree cover for a given set of points of size n. In this paper, we assume that all algorithms are analyzed using the real RAM model [13, 20, 23, 12]. Constructing a (compressed) quadtree and computing the shifts require $O_d(n\log n)$ time [8]. For each non-trivial node in the quadtree (a trivial node is a node that have only one child), we select a representative point, and then compute a (μ, Δ)-partial tree cover of the representative points corresponding to descendants of the node at $\ell = O(\log(1/\varepsilon))$ levels lower. Computing this (μ, Δ)-partial tree cover on k representatives takes $\mathrm{Time}_{\mu,\Delta}(k)$ time. We can charge each of the k representatives by $\mathrm{Time}_{\mu,\Delta}(k)/k$. Each of the n points in our point set is charged $\ell = O_d(\log(1/\varepsilon))$ times. Assuming that $\mathrm{Time}_{\mu,\Delta}(a) + \mathrm{Time}_{\mu,\Delta}(b) \leq \mathrm{Time}_{\mu,\Delta}(a+b)$, we can bound the total charge across all points by $O_d(\mathrm{Time}_{\mu,\Delta}(n)\cdot\log(1/\varepsilon))$. Hence, the total time complexity is $O_d(n\log n + \mathrm{Time}_{\mu,\Delta}(n)\cdot\log(1/\varepsilon))$.

2.2 Partial Tree Cover Without Steiner Points

This part is devoted to the proof of Theorem 1. We present the argument in $\mathbb{R}^2$, and defer the proof for $\mathbb{R}^d$ with $d \geq 3$ to the full version of the paper.

▶ **Lemma 11.** *Let X be a set of points in $\mathbb{R}^2$ with diameter Δ. For every constant $\mu > 0$ there is a (μ, Δ)-partial tree cover for X with stretch $(1+\varepsilon)$ and size $O(1/\varepsilon)$, where each tree has diameter at most $2\Delta\log(4\mu\varepsilon)$.*

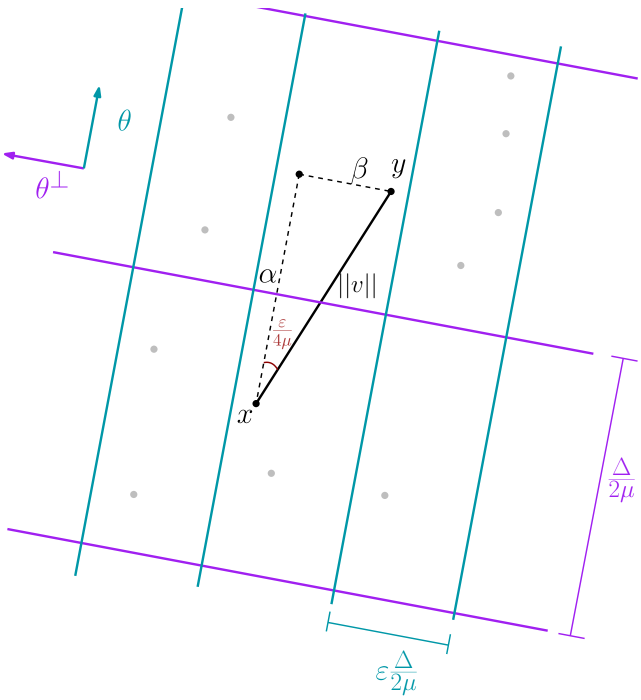

Figure 1 A major strip partition (in blue) in direction θ, and a minor strip partition (in purple) in direction $\theta^\perp$. Points x and y, and the vector v broken into components parallel to and orthogonal to θ.

The construction relies on partitioning the plane into strips. Let θ be a unit vector. We define a *strip in direction* θ to be a region of the plane bounded by two lines, each parallel to θ. The *width* of the strip is the distance between its two bounding lines. We define the *strip partition with direction* θ *and width* w (shorthanded as (θ, w)*-strip partition*) to be the unique partition of $\mathbb{R}^2$ into strips of direction θ and width w, where there is one strip that has a bounding line intersecting the point $(0, 0)$. Let $\theta^\perp := (-\theta_y, \theta_x)$ be a unit vector perpendicular to θ. A (θ, w)-strip partition with *shift* s is obtained by shifting the boundary lines of the (θ, w) strip partition by $s \cdot \theta^\perp$.

Consider the following family of strip partitions: Let $\theta_i := (\cos(i \cdot \frac{\varepsilon}{4\mu}), \sin(i \cdot \frac{\varepsilon}{4\mu}))$ be the unit vector with angle $i \cdot \frac{\varepsilon}{4\mu}$, for $i \in \{0, \ldots, \frac{8\pi\mu}{\varepsilon} - 1\}$. Let set ξ_i contains (1) the $(\theta_i, \varepsilon\frac{\Delta}{2\mu})$-strip partition with shift 0, and (2) the $(\theta_i, \varepsilon\frac{\Delta}{2\mu})$-strip partition with shift $\varepsilon\frac{\Delta}{4\mu}$. Let $\xi := \bigcup_i \xi_i$. We call the strip partitions of ξ the *major strip partitions*. Clearly, ξ contains $16\pi\mu/\varepsilon = O(1/\varepsilon)$ major strip partitions. We define $\theta_i^\perp$ to be a vector orthogonal to θ_i; and we define $\xi^\perp$ to be the set of all $(\theta_i^\perp, \frac{\Delta}{2\mu})$-strip partitions with shift 0, for every $i \in \{0, \ldots, \frac{8\pi\mu}{\varepsilon} - 1\}$. We call the shift partitions of $\xi^\perp$ the *minor strip partitions*. Every set ξ_i of major strip partitions is associated with a minor strip partition; notice that every major strip partition has an ε-factor smaller width to its orthogonal minor strip partition. See Figure 1.

▷ Claim 12. For any two points $x, y \in X$ such that x and y are (μ, Δ)-far, there exists some major strip partition $P \in \xi$ such that (1) the points x and y are in the same strip of P; and (2) in the associated minor strip partition $P^\perp \in \xi^\perp$, the points x and y are in different strips.

Proof. Let v denote the vector $y - x$. There exists some $i \in \{0, \ldots, 8\pi\mu/\varepsilon - 1\}$ such that the angle between the vector θ_i and v is at most $\varepsilon/4\mu$. We write v as a linear combination of

θ_i and a vector $\theta_i^\perp$ orthogonal to θ: $v = \alpha \cdot \theta_i + \beta \cdot \theta_i^\perp$. As the angle between v and θ_i is at most $\varepsilon/4\mu$ (and $\Delta/\mu \leq \|xy\| \leq \Delta$), we have

$$|\alpha| \geq \|v\| \cos\left(\frac{\varepsilon}{4\mu}\right) > \frac{\|v\|}{2} \geq \frac{\Delta}{2\mu}, \text{ and}$$

$$|\beta| \leq \|v\| \sin\left(\frac{\varepsilon}{4\mu}\right) \leq \frac{\varepsilon}{4\mu}\|v\| \leq \frac{\Delta}{4\mu}.$$

Let ξ_i be the set of major strip partitions in direction θ_i. As $|\beta| \leq \frac{\Delta}{4\mu}$, and ξ_i consists of shifted strip partitions of width $\frac{\Delta}{2\mu}$, there is some major strip partition $P \in \xi_i$ in which x and y are in the same strip. Further, every strip in the associated minor strip partition $P^\perp$ has width $\frac{\Delta}{2\mu}$, so the fact that $|\alpha| > \frac{\Delta}{2\mu}$ implies that x and y are in different strips of $P^\perp$. This proves the claim. ◁

For every major strip partition in ξ, we now construct a tree which preserves approximately distances between points that lie in the same major strip but different minor strips. The following is the key claim.

▷ Claim 13. Let S be a strip from a major strip partition in ξ, with direction θ. Let S_1 and S_2 be two strips from a minor strip partition in ξ, both with direction $\theta^\perp$. Then there is a tree T on $X \cap S$ such that for every $a \in X \cap S_1 \cap S$ and $b \in X \cap S_2 \cap S$, $\|ab\| \leq \delta_T(a,b) \leq \|ab\| + \frac{\varepsilon\Delta}{\mu}$. In particular, if x and y are (μ,Δ)-far, then $\|ab\| \leq \delta_T(a,b) \leq (1+\varepsilon) \cdot \|ab\|$.

Proof. For any point $x \in \mathbb{R}^2$, we define $\mathrm{score}_\theta(x)$ to be the inner product $\langle x, \theta \rangle$. Let $A := X \cap S_1 \cap S$ and $B := X \cap S_2 \cap S$. As A and B belong to different minor strips in direction $\theta^\perp$, without loss of generality $\mathrm{score}_\theta(a) < \mathrm{score}_\theta(b)$ for every $a \in A$ and $b \in B$. Let $a^* := \arg\max_{a \in A} \mathrm{score}_\theta(a)$. We claim that for any $a \in A$ and $b \in B$,

$$\|aa^*\| + \|a^*b\| \leq \|ab\| + \frac{\varepsilon\Delta}{\mu}. \tag{1}$$

To show this, consider the line segment ℓ between a and b. Let L be the line in direction $\theta^\perp$ that passes through a^*. Because $\mathrm{score}_\theta(a) \leq \mathrm{score}_\theta(a^*) \leq \mathrm{score}_\theta(b)$, line L and segment ℓ intersect at some point a' in the slab S; see Figure 2. (Note that a' is *not* the projection of a^* onto ℓ.) The distance $\|a^*a'\|$ can be no greater than the width of the slab, so $\|a^*a'\| \leq \varepsilon\frac{\Delta}{2\mu}$. By triangle inequality, we have

$$\begin{aligned}
\|aa^*\| + \|a^*b\| &\leq (\|aa'\| + \|a'a^*\|) + (\|a^*a'\| + \|a'b\|) \\
&\leq \|aa'\| + \|a'b\| + \varepsilon\frac{\Delta}{\mu} \\
&\leq \|ab\| + \frac{\varepsilon\Delta}{\mu}.
\end{aligned}$$

Let T be the star centered at a^*, with an edge to every other point $x \in A \cup B$; the weight of the edge between a^* and x is $\|a^*x\|$. For any $a \in A$ and $b \in B$, we clearly have $\|ab\| \leq \delta_T(a,b)$, and Equation (1) guarantees that $\delta_T(a,b) \leq \|ab\| + \frac{\varepsilon\Delta}{\mu}$. ◁

We can now prove Lemma 11.

Proof of Lemma 11. Let ξ be the set of major strip partitions defined above. Let P be an arbitrary major strip partition in ξ, and let $P^\perp$ be the associated minor strip partition in $\xi^\perp$.

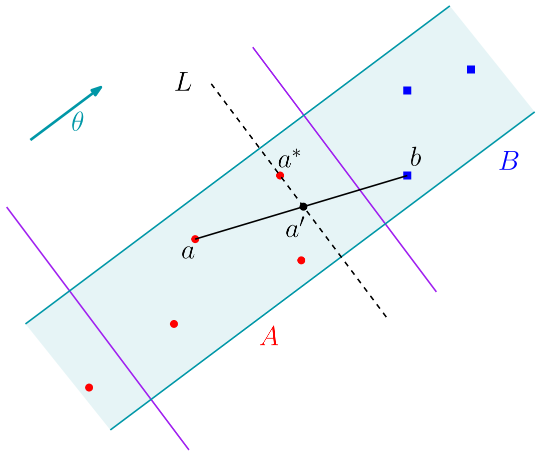

Figure 2 Point sets A and B, both in the same major strip (blue) but in different minor strips (purple). The points a, a^*, and b, with $\text{score}_\theta(a) \le \text{score}_\theta(a^*) \le \text{score}_\theta(b)$, and the line L passing through a^*.

For each pair of strips S_1 and S_2 in $P^\perp$, we define *tree* T_{P,S_1,S_2} as follows: For every strip S in P, apply Claim 13 to construct a tree T_S on (a subset of) $X \cap S$ that preserves distances between $X \cap S_1$ and $X \cap S_2$; and let T_{P,S_1,S_2} be the tree obtained by joining together the trees T_S from all strips S in P. To join the trees, we build a balanced binary tree from the roots of T_S for all strips S in P. The *tree cover* $\mathcal{T}$ consists of the set of all trees T_{P,S_1,S_2}, for every major strip partition $P \in \xi$ and every pair of strips S_1, S_2 in the associated minor strip partition $P^\perp$.

To bound the size of $\mathcal{T}$, observe that (1) there are at most $\frac{8\pi\mu}{\varepsilon} \cdot 2 = O(1/\varepsilon)$ major strip partitions containing points in X, and (2) for every strip S in a major strip partition, at most $2\mu + 1 = O(1)$ strips in the associated minor strip partition contain points in $X \cap S$ (recall that point set X has diameter Δ). Thus $\mathcal{T}$ contains $\frac{16\pi\mu}{\varepsilon} \cdot \binom{2\mu+1}{2} = O(1/\varepsilon)$ trees.

To bound the stretch, let a and b be arbitrary points in X. By Claim 12, there exists some major strip partition $P \in \xi$ such that (1) a and b are in the same strip in P; and (2) a and b are in different strips S_1 and S_2 of the associated minor strip partition $P^\perp$. Thus Claim 13 implies that tree T_{P,S_1,S_2} satisfies $\|ab\| \le \delta_T(a,b) \le (1+\varepsilon) \cdot \|ab\|$.

To bound the diameter, let P be a major strip partition and let S be a major strip in P. Observe that T_S is a star and the distance from the root of T_S to any other point in T_S is at most Δ. The roots of trees corresponding to strips in P are connected by a binary tree by construction. Each edge of this binary tree is of length at most Δ. The number of strips in P is upper bounded by $2\mu/\varepsilon$. Hence, the height of the binary tree is at most $\log(2\mu/\varepsilon)$. This means that the diameter of the resulting tree is at most $2 \cdot (\Delta + \log(2\mu/\varepsilon) \cdot \Delta) = 2\Delta \log(4\mu/\varepsilon)$. ◀

Running Time. The inner product between each point with each vector θ_i can be precomputed using $O(|X| \cdot \frac{4\mu}{\varepsilon})$ operations. For a major strip S, finding the maximum point in the intersection between S and each of its minor strip only need time proportional to the number of points in $S \cap X$. Those points are chosen as roots of the stars corresponding to S. For each root, constructing the corresponding star requires $O(|S \cap X|)$ time. There are $\binom{2\mu+1}{2}$ roots for each major strip. Hence, the total time complexity of constructing the (μ, Δ)-partitial tree cover is:

$$|X| \cdot \frac{4\mu}{\varepsilon} + \binom{2\mu+1}{2} \sum_{\text{major strip } S} |S \cap X| = O(|X| \cdot \varepsilon^{-1})$$

Therefore, the time complexity of constructing the tree cover is $O_d(n \log n + n\varepsilon^{-1} \log(1/\varepsilon))$.

2.3 Partial tree cover with Steiner points

This part is devoted to the proof of Theorem 2 for $\mathbb{R}^2$; the argument for dimension $d \geq 3$ is deferred to the full version of the paper.

▶ **Lemma 14.** *Let X be a set of points in $\mathbb{R}^2$ with diameter Δ. For every constant $\mu > 0$, there is a Steiner (μ, Δ)-partial tree cover with stretch $(1+\varepsilon)$ for X with $1/\sqrt{\varepsilon}$ trees, where each tree has diameter at most 3Δ.*

Consider a square of side-length Δ containing X, and let μ be an arbitrary constant. Divide the square into vertical slabs of width $\frac{\Delta}{3\sqrt{2}\mu}$ and height Δ, and into horizontal slabs of width Δ and height $\frac{\Delta}{3\sqrt{2}\mu}$.

▶ **Observation 15.** *For any two points $p, q \in X$ such that p and q are (μ, Δ)-far, there exists either a horizontal or a vertical slab such that p and q are from different sides of the slab.*

Proof. Suppose towards contradiction there are two adjacent horizontal slabs containing both p and q and also two adjacent vertical slabs containing both p and q. The distance between p and q is at most $\|pq\| \leq 2 \cdot \frac{\Delta}{3\sqrt{2}\mu} \cdot \sqrt{2} < \frac{\Delta}{\mu}$, contradicting the assumption that p and q are (μ, Δ)-far. ◀

For each horizontal (resp. vertical) slab S, we consider the horizontal (resp. vertical) line segment ℓ that cuts the slab into two equal-area parts. The length of ℓ is Δ. Let $k := \lfloor 2\mu/\sqrt{\varepsilon} \rfloor$ be an integer. We partition ℓ into k intervals, called $[a_0, a_1], [a_2, a_3], \ldots [a_{k-1}, a_k]$, each of length $\sqrt{\varepsilon}\Delta/2\mu$. For each point a_i, we construct *tree* T_S^i by adding edges between a_i and every point in X. Finally, connect the points a_i using a straight line and let T be the resulting tree. The diameter of this tree is at most 3Δ.

▷ Claim 16. For any two points $p, q \in X$ such that p and q are (μ, Δ)-far, there exists a slab s and an integer $i \in \{0, \ldots, k\}$ such that $\delta_{T_S^i}(p, q) \leq (1+\varepsilon) \cdot \|pq\|$.

Proof. By Observation 15, there exists a slab S such that p and q are in different sides of it. Without loss of generality assume that S is horizontal. By construction, we partition the middle interval ℓ of S into k intervals $[a_0, a_1], [a_2, a_3], \ldots [a_{k-1}, a_k]$ each of length $\sqrt{\varepsilon}\Delta/2\mu$. Let r be the intersection between pq and ℓ, and let a_i be the closest point to r. Let r' be the projection of a_i to r'. Hence, $||a_i r'|| \leq ||a_i r|| \leq \sqrt{\varepsilon}\Delta/2\mu$. Using the triangle inequality, we have:

$$\delta_{T_S^i}(p, q) \leq ||pa_i|| + ||a_i q|| = \sqrt{||pr'||^2 + ||r'a_i||^2} + \sqrt{||r'q||^2 + ||r'a_i||^2}. \tag{2}$$

Observe that $||pr'|| \geq \Delta/2\mu$. Thus, $||r'a_i|| \leq \sqrt{\varepsilon}\Delta/2\mu \leq ||pr'||\sqrt{\varepsilon}$. Similarly, $||r'a_i|| \leq ||r'q||\sqrt{\varepsilon}$. Combining with Equation 2, we get:

$$\begin{aligned}\delta_{T_S^i}(p, q) &\leq ||pa_i|| + ||a_i q|| = \sqrt{||pr'||^2 + \varepsilon||pr'||^2} + \sqrt{||r'q||^2 + \varepsilon||r'q||^2} \\ &\leq \sqrt{1+\varepsilon} \cdot (||pr'|| + ||r'q||) \leq (1+\varepsilon) \cdot ||pq||.\end{aligned}$$ ◁

We now prove Lemma 14. Let $\mathcal{T}$ be the set containing trees T_s^i for every horizontal or vertical slabs s and every index $i \in [0, k]$. There are $O(\mu) = O(1)$ horizontal and vertical slabs, so $\mathcal{T}$ contains $O(k) = O(1/\sqrt{\varepsilon})$ trees. It follows immediately from Claim 16 that $\mathcal{T}$ is a Steiner (μ, Δ)-partial tree cover for X with stretch $(1+\varepsilon)$.

Running Time. For a set X, creating the set of slabs can be done in $O(1)$ time. For each slab, finding a net of the middle line takes $O(1/\sqrt{\varepsilon})$ time. For each Steiner point, it requires $O(|X|)$ time to create a tree connecting that point to everyone in X. Totally, the time complexity is $O(|X|/\sqrt{\varepsilon})$. Therefore, the time complexity of constructing the tree cover is $O_d(n\log n + n\varepsilon^{-1/2}\log(1/\varepsilon))$.

References

1 Pankaj K Agarwal, Hsien-Chih Chang, Sharath Raghvendra, and Allen Xiao. Deterministic, near-linear ε-approximation algorithm for geometric bipartite matching. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1052–1065, 2022.

2 Noga Alon, Richard M. Karp, David Peleg, and Douglas West. A graph-theoretic game and its application to the k-server problem. *SIAM Journal on Computing*, 24(1):78–100, 1995. doi:10.1137/s0097539792224474.

3 Sunil Arya, Gautam Das, David M Mount, Jeffrey S Salowe, and Michiel Smid. Euclidean spanners: short, thin, and lanky. In *Proceedings of the twenty-seventh annual ACM symposium on Theory of computing*, pages 489–498, 1995.

4 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.

5 Yair Bartal, Ora Nova Fandina, and Ofer Neiman. Covering metric spaces by few trees. *Journal of Computer and System Sciences*, 130:26–42, 2022.

6 Sujoy Bhore and Csaba D Tóth. Euclidean steiner spanners: Light and sparse. *SIAM Journal on Discrete Mathematics*, 36(3):2411–2444, 2022.

7 T.-H. Hubert Chan, Anupam Gupta, Bruce M. Maggs, and Shuheng Zhou. On hierarchical routing in doubling metrics. *ACM Transactions on Algorithms (TALG)*, 12(4):1–22, 2016.

8 Timothy M. Chan. Approximate nearest neighbor queries revisited. *Discret. Comput. Geom.*, 20(3):359–373, 1998.

9 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, and Shay Solomon. Covering planar metrics (and beyond): O(1) trees suffice. *arXiv preprint arXiv:2306.06215*, 2023. Accepted to FOCS 2023.

10 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. *arXiv preprint arXiv:2308.00555*, 2023. Accepted to SODA 2024.

11 Feodor F Dragan, Chenyu Yan, and Irina Lomonosov. Collective tree spanners of graphs. *SIAM Journal on Discrete Mathematics*, 20(1):240–260, 2006.

12 Jeff Erickson, Ivor Van Der Hoog, and Tillmann Miltzow. Smoothing the gap between np and er. *SIAM Journal on Computing*, pages FOCS20–102 – FOCS20–138, 2022.

13 Steven Fortune and Christopher J. Van Wyk. Efficient exact arithmetic for computational geometry. In *Proceedings of the Ninth Annual Symposium on Computational Geometry*, SCG '93, pages 163–172, New York, NY, USA, 1993. Association for Computing Machinery. doi:10.1145/160985.161015.

14 Zhimeng Gao and Sariel Har-Peled. Almost optimal locality sensitive orderings in Euclidean space. *CoRR*, abs/2310.12792, 2023.

15 L. Gottlieb and L. Roditty. Improved algorithms for fully dynamic geometric spanners and geometric routing. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'08, pages 591–600, 2008. doi:10.5555/1347082.1347148.

16 A. Gupta, A. Kumar, and R. Rastogi. Traveling with a Pez dispenser (or, routing issues in MPLS). In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science,FOCS' 01*, 2001. doi:10.1109/sfcs.2001.959889.

17 Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *FOCS*, pages 534–543. IEEE Computer Society, 2003.

18 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 *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, pages 151–162, 2022.

19 Hung Le and Shay Solomon. Truly optimal Euclidean spanners. *SIAM Journal on Computing*, pages FOCS19–135 – FOCS19–199, 2022.

20 Chen Li, Sylvain Pion, and Chee-Keng Yap. Recent progress in exact geometric computation. *The Journal of Logic and Algebraic Programming*, 64(1):85–111, 2005.

21 Manor Mendel and Assaf Naor. Ramsey partitions and proximity data structures. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS' 06)*, 2006. doi:10.1109/focs.2006.65.

22 Giri Narasimhan and Michiel Smid. *Geometric spanner networks*. Cambridge University Press, 2007.

23 David Salesin, Jorge Stolfi, and Leonidas Guibas. Epsilon geometry: building robust algorithms from imprecise computations. In *Proceedings of the fifth annual symposium on Computational geometry*, pages 208–217, 1989.

24 Michiel Smid. Notes on binary dumbbell trees. Unpublished notes, 2012. URL: https://people.scs.carleton.ca/~michiel/dumbbelltrees.pdf.

25 Mikkel Thorup and Uri Zwick. Approximate distance oracles. *Journal of the ACM (JACM)*, 52(1):1–24, 2005.