

Optimal Euclidean Tree Covers

Hsien-Chih Chang¹, Jonathan Conroy¹, Hung Le²,
Lazar Milenković³, Shay Solomon³, Cuong Than²

¹Department of Computer Science, Dartmouth College.

²CICS, UMass Amherst.

³Tel Aviv University.

Contributing authors: hsien-chih.chang@dartmouth.edu;
jonathan.conroy.gr@dartmouth.edu; hungle@cs.umass.edu;
milenkovic.lazar@gmail.com; solo.shay@gmail.com; cthan@cs.umass.edu;

Abstract

A $(1 + \epsilon)$ -stretch tree cover of a metric space is a collection of trees, where every pair of points has a $(1 + \epsilon)$ -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 $(1 + \epsilon)$ -stretch tree cover with $O_d(\epsilon^{-d} \cdot \log(1/\epsilon))$ trees, where the O_d notation suppresses terms that depend solely on the dimension d . The running time of their construction is $O_d(n \log n \cdot \frac{\log(1/\epsilon)}{\epsilon^d} + n \cdot \epsilon^{-2d})$. 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 $(1 + \epsilon)$ -stretch tree cover with $O_d(\epsilon^{-d+1} \cdot \log(1/\epsilon))$ trees, which is optimal (up to the $\log(1/\epsilon)$ 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 $(1 + \epsilon)$ -stretch *Steiner* tree cover (that may use Steiner points) with $O_d(\epsilon^{(-d+1)/2} \cdot \log(1/\epsilon))$ 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. *

*A preliminary version of this paper was published in the proceedings of SoCG 2024.

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 , $\delta_M(x, y) \leq \delta_T(x, y)$ 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 [1–10]; 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 [11, 12], spanners with small hop diameters [12], and bipartite matching [13].

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 [14]. A few years ago, this theorem was generalized for doubling metrics² by Bartal, Fandina, and Neiman [8], who achieved the same bound as [3] via a much simpler construction; the running time of their construction was not analyzed.³ In the constructions by [3, 8], same point may have multiple copies in different levels of the tree, hence the maximum degree of points⁴ may be as large as $\Omega(\log \Phi)$, where Φ is the aspect ratio of the input point set; see Sections 1.3 and 4 for a more detailed discussion.

Since the number of trees provided by the two known constructions [3, 8] 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

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

²The *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.

³In high-dimensional Euclidean spaces the upper bound in [8] improves over that of [3], since the O_d notation in [3] and [8] suppress multiplicative factors of $d^{O(d)}$ and $2^{O(d)}$, respectively.

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

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 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 [15, 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 $\mathbb{R}^2$ [15, Theorem 1.4], and in general the $\Omega(n/\varepsilon^{(d-1)/2})$ size lower bound in $\mathbb{R}^d$ [16].

1.1 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. Here we attempt to summarize the best upper and lower bounds known to our knowledge, highlighting the tradeoff between tree cover size and stretch in the previous work. Some of the bounds are not explicitly stated in the cited reference but can be deduced from it. For additional relevant work, refer to [8] and the references therein.

General metrics.

The earliest literature on the notion of tree cover is probably Awerbuch and Peleg [1] and Awerbuch, Kutten, and Peleg [17], focusing on graph metrics. Their main objective is to minimize the number of trees *each vertex belongs to* (in the sparse cover sense) instead of minimize the total number of trees. Thorup and Zwick [11, Corollary 4.4] improved over Awerbuch and Peleg [1] by constructing a Steiner tree cover with stretch $2k - 1$ where every vertex belongs to $O(n^{1/\alpha} \cdot \log^{1-1/k} n)$ trees. Charikar *et al.* [18] studies a similar problem of probabilistically embedding finite metric space into a small number of trees. Many of the earlier work on tree covers are motivated by application in routing [19].

Gupta, Kumar, and Rastogi [4, Theorem 4.3] observed that any tree cover must have size $n^{\Omega(1/\alpha)}$ if the stretch is α ; the lower bound is based on the existence of girth- g graphs with $n^{\Omega(1/g)}$ edges [20][21, Lemma 9]. It is important to emphasize that the tree covers considered in [4] are *spanning* — the trees must be subgraphs of the input graph. Bartal, Fandina, and Neiman established the same lower bound [8, Corollary 13] by reduction from spanners [21]. In a different direction, Dragan, Yan, and Lomonosov [6] studied spanner tree covers with additive stretch on special classes of graphs, such as chordal graphs and co-comparability graphs.

One might relax the condition to allow vertices not presented in the graph (called *Steiner vertices*) to be part of the tree cover. By allowing Steiner vertices, Mendel and Naor [5] showed that any n -point metric space has a Steiner tree cover of size $O(\alpha \cdot n^{1/\alpha})$ and stretch $O(\alpha)$. Bartal, Fandina, and Neiman [8] obtained an inverse tradeoff: any n -point metric space has a Steiner tree cover of size k and stretch $O(n^{1/k} \cdot (\log n)^{1-1/k})$. In particular, this means we can get $O(\log n)$ trees with $O(\log n)$ stretch. While the

lower bound from [4] for spanning tree cover no longer holds when Steiner vertices are allowed, a similar lower bound of $\Omega(n^{4/(3\alpha+2)}/\log n) = n^{\Omega(1/\alpha)}$ for the size of Steiner α -tree covers can be derived from Steiner spanners (also known as *emulators*) [21, Theorem 6], using the same argument in [8].

Doubling metrics.

Chan, Gupta, Maggs, and Zhou [7, Lemma 3.4] constructs Steiner tree covers for doubling metrics [7]. More precisely, if the doubling dimension of the metric space is d , their tree cover uses $O(d \log d)$ Steiner trees and has stretch $O(d^2)$. Bartal, Fandina, and Neiman [8] obtained two separated constructions: one may have tree cover of stretch $O(\alpha)$ and $O(2^{d/\alpha} \cdot d \cdot \alpha)$ Steiner trees for any $\alpha \geq 2$ [8, Theorem 7] using the $O(1)$ -padded hierarchical partition family in [22, Lemma 8], or alternatively a tree cover with $(1 + \varepsilon)$ stretch and $(1/\varepsilon)^{O(d)} \cdot \log(1/\varepsilon)$ trees [8, Theorem 3] using net trees. It is worth emphasizing that the second construction does not use Steiner points. They also established a lower bound on the size of non-Steiner tree cover [8, Corollary 13] by reduction from spanners [21]: there is an n -point metric space with doubling dimension d , such that any α -tree cover requires $\Omega(2^{d/\alpha})$ trees.

Planar and minor-free graphs.

On planar graphs Gupta, Kumar, and Rastogi [4] constructed the first $O(\log n)$ -size (non-Steiner) tree cover with stretch 3. Again this is improved by Bartal, Fandina, and Neiman in two different directions: either one has stretch $O(1)$ and $O(1)$ trees [8, Corollary 9] using the $O(1)$ -padded hierarchical partition family in [23]⁵, or alternatively a tree cover with $(1 + \varepsilon)$ stretch and $O(\varepsilon^{-1} \log^2 n)$ trees [8, Theorem 5], using path separators [24]. Their results naturally extend to minor-free graphs. Recently, the authors get the best of both worlds by constructing a Steiner tree cover with $(1 + \varepsilon)$ stretch using $\tilde{O}(\varepsilon^{-3})$ many trees [9] through the introduction of a new graph partitioning scheme called the *shortcut partition*; the result also extends to minor-free graphs [10].

On planar graphs $\Omega(\sqrt{n})$ trees are required if no stretch is allowed [4]. However in the $(1 + \varepsilon)$ -stretch regime, we are not aware of any existing lower bounds. The strongest lower bound for tree covers on planar graphs we managed to deduce comes from *distance labeling*: Suppose we have a Steiner $(1 + \varepsilon)$ -tree cover using $O(\varepsilon^{-1/(3+\delta)})$ many trees for some $\delta > 0$. Then we can construct an approximate distance labeling scheme by concatenating the $O(\log n \cdot \log(1/\varepsilon))$ -length labeling schemes for all trees [25]. By setting $\varepsilon = 1/n$, we get an exact labeling scheme for unweighted planar graph of length $\tilde{O}(n^{1/(3+\delta)})$, contradicting to an information-theoretical lower bound [26]. This implies that any Steiner $(1 + \varepsilon)$ -tree covers on planar graphs requires at least $\Omega(\varepsilon^{1/3})$ many trees.

Euclidean metrics.

We already discussed tree cover results on Euclidean metrics in the introduction above; here we mentioned a few additional facts.

⁵The constants in [23] imply that the stretch is at least $3^4 = 81$ and the number of trees is at least $3^3 = 27$.

All upper bound constructions on metrics with bounded doubling dimensions immediately apply to Euclidean metrics as well. Surprisingly, relatively few lower bounds have been established in the literature for Euclidean spaces. Early in the introduction we derived an $\Omega(1/\varepsilon^{d-1})$ lower bound for non-Steiner tree cover and an $\Omega(1/\varepsilon^{(d-1)/2})$ lower bound for Steiner tree cover in $\mathbb{R}^d$ by reduction from spanners.

One thing to notice is that in Euclidean spaces, the meaning of Steiner points differs slightly from its graph counterparts: after choosing a Steiner point (which lies in the ambient space $\mathbb{R}^d$), the weight of an edge incident to a Steiner point is determined by its Euclidean distance, unlike in the graph setting one may choose the weight freely (as the Steiner points are artificially inserted and were not part of the graph a priori). One might think that such a distinction cannot possibly make any difference; however, recently Andoni and Zhang [27] proved that $(1+\varepsilon)$ -spanner of subquadratic size exists for arbitrary dimensional Euclidean space by allowing out-of-nowhere Steiner points, while establishing lower bound simultaneously when the Steiner points are required to sit in the Euclidean space. They showed that there are n points in $\mathbb{R}^d$ (for some high dimension d depending on n) where any $(\sqrt{2}-\varepsilon)$ -spanner (with Euclidean Steiner points) requires $\Omega(\varepsilon^4 \cdot n^2 / \log^2 n)$ edges; the lower bound follows from a randomized construction and volume argument. This translates to an almost linear lower bound of $\Omega(\varepsilon^4 \cdot n / \log^2 n)$ on the minimum number of trees required in any Euclidean Steiner tree cover with $(\sqrt{2}-\varepsilon)$ stretch. All Steiner points used in our construction are Euclidean; at the moment, we are unaware of any tree cover construction that obtains a better bound by taking advantage of the non-Euclidean Steiner points.

Ramsey trees.

A stronger notion called the *Ramsey tree cover* has been studied, where every vertex x is associated with a tree T_x in $\mathcal{F}$, such that the distance from x to *every* other vertex is approximated preserved by the same tree T_x . Both the constructions of Mendel and Naor [5] and Bartal, Fandina, and Neiman [8] for general metrics are indeed Ramsey trees. These bounds are essentially tight if the trees are required to be Ramsey; that is, any Ramsey tree cover of stretch α must contain $n^{1/\Omega(\alpha)}$ many tree [8, Corollary 13]. Even when the input metric is planar and doubling, any Ramsey tree cover of stretch α must contain $n^{1/\Omega(\alpha \log \alpha)}$ many tree [8, Theorem 10], and any Ramsey tree cover of size k must has stretch $n^{1/k}$ [8, Theorem 9].

1.2 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.⁶ In view of the aforementioned lower bound [15, 16], this is optimal up to the $\log(1/\varepsilon)$ factor. Roughly speaking, we show that the packing bound barrier (incurred in both [3] and [8]) can be replaced by the number of ε -angled cones needed to partition $\mathbb{R}^d$; for more details, refer to Section 1.3.

⁶As 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 [8]. Thus, our results improve over that of [8] 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.

Theorem 1.1. *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 [15, 16], this result too is optimal up to the $\log(1/\varepsilon)$ factor.

Theorem 1.2. *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 $O_d(n \log n + n \cdot \varepsilon^{(-d+1)/2} \cdot \log(1/\varepsilon))$.*

1.2.1 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.

Both the Dumbbell Theorem [3] and the BFN construction [8] 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 [28]), 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 [8] may also incur a point-degree of $\Omega(\log \Phi)$ in any of the trees.⁷

We strengthen Theorem 1.1 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 [15].

Routing.

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

Theorem 1.3. *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 [29], 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 1.1. Having

⁷Even node-degrees may blow up in the construction of [8], but it appears that a simple tweak of their construction can guarantee a node-degree of $\varepsilon^{-O(d)}$.

determined the right tree to route on, our entire routing algorithm is carried out on that tree, while the routing algorithm of [29] is carried out on a spanner; routing in a tree is clearly advantageous over routing in a spanner, also from a practical perspective. Refer to [29] for the definition of the problem and relevant background.

1.3 Technical Highlights

1.3.1 Achieving an optimal bound on the number of trees

The tree cover constructions of [3] and [8] 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 [8] that applies to the wider family of doubling metrics. Here we give a short overview of the simpler construction of [8]; 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 [8] is the standard *hierarchy of 2^w -nets* $\{N_w\}$ [30], which induces a hierarchical *net-tree*.⁸ 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 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 $\ell := \log(1/\varepsilon)$ *congruence classes* $I_0, I_1, \dots, I_{\ell-1}$, where $I_j := \{w \mid w \equiv j \pmod{\ell}\}$. 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 [8], 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 *quadtrees* 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, 8, 13, 15] 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 [8].

⁸The standard notation in the literature on doubling metrics, including [8], 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.

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 [31]: 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 2.4 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, as well as 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 .

1.3.2 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, 8] incur unbounded degree, the Euclidean construction of [3], when restricted to a single level in the hierarchy, achieves an absolute constant degree.⁹

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 [16]. To achieve this goal, we strengthen the aforementioned two-step strategy as follows.

Step 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 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

⁹Although in the original paper of [3] (as well as in [14]) the bound is not an absolute constant, it was shown in [28] that an absolute constant bound can be obtained. Nonetheless, overlaying all levels of the hierarchy leads to a final degree bound of $\Theta(\log \Phi)$.

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; see Section 4.1 for the details. 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.

Step 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, independent of ε . 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 $1 + O(\varepsilon \log(1/\varepsilon))$; see Section 4.2 for the details.

1.4 Organization

In Section 2, we present the construction of tree covers in $\mathbb{R}^2$ with an optimal number of trees in both non-Steiner and Steiner settings, proving Theorem 1.1 and Theorem 1.2 for the plane. In Section 3, we generalize these constructions to $\mathbb{R}^d$ for arbitrary constant d . In Section 4, we reduce the degree of every tree in the (non-Steiner) tree cover an absolute constant. In Section 5, we show some applications of our tree cover to routing, proving Theorem 1.3.

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 the minimum distance between any two points in X is 1.

Lemma 2.1 (Cf. [31, 32]). *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 $D := 2\lceil d/2 \rceil$ and $i = 0, \dots, D$, let $\nu_i := (iL/(D+1), \dots, iL/(D+1))$. Then there exists an index $i \in \{0, \dots, D\}$, 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 2.2. *We call two points (μ, Δ) -far if their distance is in $[\Delta/\mu, \Delta]$.*

Definition 2.3. *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 2.4 (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 $\mu > 1$, 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 $O(d \cdot \log \frac{\gamma d \sqrt{d}}{\varepsilon} \cdot \tau(\varepsilon, \mu))$ with $\mu := 10d\sqrt{d}$.

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 $D := 2\lceil d/2 \rceil$ and let $\mathcal{Q}$ be the quadtree as in Lemma 2.1. For $i \in \{0, \dots, D\}$, let $\mathcal{Q}_i$ be $\mathcal{Q}$ shifted by $-\nu_i = (-iL/(D+1), \dots, -iL/(D+1))$.

Constructing the tree cover. Let $\ell := \log \frac{\gamma d \sqrt{d}}{\varepsilon}$ and let $\mu := 10d\sqrt{d}$. (Assume for simplicity that ℓ is an integer.) Fix some $i \in \{0, \dots, D\}$, $j \in \{0, \dots, \ell - 1\}$, and $k \in \{1, \dots, \tau(\varepsilon, \mu)\}$. We proceed to construct tree $T_{i,j,k}$. Consider the congruence class $I_j := \{z \geq 0 \mid z \equiv j \pmod{\ell}\}$. 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 k th 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 $\Delta_w := 2^w \sqrt{d}$, 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 k th 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 2.5. Let $T_{i,j,k}^w$ be a tree constructed at level w for $i \in \{0, \dots, D\}$, $j \in \{0, \dots, \ell - 1\}$, $k \in \{1, \dots, \tau(\varepsilon, \mu)\}$ and $w \in I_j$.

1. $T_{i,j,k}^w$ is a tree.
2. $T_{i,j,k}^w$ 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 2.4. For the induction step, we have $\gamma\Delta_w + 2\gamma\Delta_{w-\ell} = \gamma\Delta_w + 2\gamma\frac{\Delta_w\varepsilon}{\gamma d \sqrt{d}} \leq 2\gamma\Delta_w$. $\square$

Claim 2.6. The number of trees in the cover is $O(d \log \frac{\gamma d \sqrt{d}}{\varepsilon} \cdot \tau(\varepsilon, \mu))$.

Proof. The trees $T_{i,j,k}$ are 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. $\square$

Claim 2.7. *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 2.1, 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.

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

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.

$$\begin{aligned} \delta_T(p, q) &\leq \delta_T(p, p') + \delta_T(p', q') + \delta_T(q, q') \\ &\leq (1 + \varepsilon) \cdot \|p'q'\| + \delta_T(p, p') + \delta_T(q, q') \\ &\leq (1 + \varepsilon) \cdot (\|p'p\| + \|pq\| + \|qq'\|) + \delta_T(p, p') + \delta_T(q, q') \\ &\leq (1 + \varepsilon) \cdot (\|pq\| + 2\Delta_{w-\ell}) + 2\phi_{w-\ell} \\ &\leq (1 + \varepsilon) \cdot (\|pq\| + 2\Delta_{w-\ell}) + 4\gamma\Delta_{w-\ell} \quad \text{by Claim 2.5} \\ &\leq (1 + \varepsilon) \cdot \left(\|pq\| + 6\Delta_w \cdot \frac{\varepsilon}{d\sqrt{d}} \right) \\ &= (1 + O(\varepsilon)) \cdot \|pq\| \end{aligned}$$

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

Claims 2.5 to 2.7 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 2.4. $\square$

Running Time.

Let $\text{Time}_{\mu, \Delta}(n)$ 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 [33–36]. Constructing a (compressed) quadtree and computing the shifts require $O_d(n \log n)$ time [31]. 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 $\text{Time}_{\mu, \Delta}(k)$ time. We can charge each of the k representatives by $\text{Time}_{\mu, \Delta}(k)/k$. Each of the n points in our point set is charged $\ell = O_d(\log(1/\varepsilon))$ times. In the subsequent sections, we show that function $\text{Time}_{\mu, \Delta}(n)$ satisfies $\text{Time}_{\mu, \Delta}(n) = O_d(n \text{poly}(1/\varepsilon))$, which in particular implies that for any two positive a and b , we have $\text{Time}_{\mu, \Delta}(a) + \text{Time}_{\mu, \Delta}(b) = O_d(\text{Time}_{\mu, \Delta}(a + b))$. Using this inequality, we can bound the total charge across all points by $O_d(\text{Time}_{\mu, \Delta}(n) \cdot \log(1/\varepsilon))$. Hence, the total time complexity is $O_d(n \log n + \text{Time}_{\mu, \Delta}(n) \cdot \log(1/\varepsilon))$.

2.2 Partial Tree Cover Without Steiner Points

This part is devoted to the proof of Theorem 1.1. We present the argument in $\mathbb{R}^2$, and defer the proof for $\mathbb{R}^d$ with $d \geq 3$ to Section 3.1.

Lemma 2.8. *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 $2\Delta \log(4\mu\varepsilon)$.*

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* (shorthand 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 $i \in \{0, \dots, \frac{8\pi\mu}{\varepsilon} - 1\}$. 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 $i \in \{0, \dots, \frac{8\pi\mu}{\varepsilon} - 1\}$. 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 2.9. *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 $i \in \{0, \dots, 8\pi\mu/\varepsilon - 1\}$ such that the angle between the vector θ_i and v is at most $\varepsilon/8\mu$. We write v as a linear




Fig. 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 .

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 $\varepsilon/8\mu$ (and $\Delta/\mu \leq \|xy\| \leq \Delta$), we have

$$|\alpha| \geq \|v\| \cos\left(\frac{\varepsilon}{8\mu}\right) > \frac{\|v\|}{2} \geq \frac{\Delta}{2\mu}, \text{ and}$$

$$|\beta| \leq \|v\| \sin\left(\frac{\varepsilon}{8\mu}\right) \leq \frac{\varepsilon}{8\mu} \|v\| \leq \frac{\varepsilon\Delta}{8\mu}.$$

Let ξ_i be the set of major strip partitions in direction θ_i . As $|\beta| \leq \frac{\varepsilon\Delta}{8\mu}$, and ξ_i consists of shifted strip partitions of width $\frac{\varepsilon\Delta}{2\mu}$, 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. $\square$

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 2.10. *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 $\text{score}_\theta(x)$ 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

528 in direction $\theta^\perp$, without loss of generality $\text{score}_\theta(a) < \text{score}_\theta(b)$ for every $a \in A$ and
 529 $b \in B$. Let $a^* := \arg \max_{a \in A} \text{score}_\theta(a)$. We claim that for any $a \in A$ and $b \in B$,

$$530 \quad \|aa^*\| + \|a^*b\| \leq \|ab\| + \frac{\varepsilon\Delta}{\mu}. \quad (1)$$

531 To show this, consider the line segment ℓ between a and b . Let L be the line in direction
 532 $\theta^\perp$ that passes through a^* . Because $\text{score}_\theta(a) \leq \text{score}_\theta(a^*) \leq \text{score}_\theta(b)$, line L and
 533 segment ℓ intersect at some point a' in the slab S ; see Figure 2. (Note that a' is *not*
 534 the projection of a^* onto ℓ .) The distance $\|a^*a'\|$ can be no greater than the width of
 535 the slab, so $\|a^*a'\| \leq \varepsilon \frac{\Delta}{2\mu}$. By triangle inequality, we have

$$\begin{aligned} 536 \quad \|aa^*\| + \|a^*b\| &\leq (\|aa'\| + \|a'a^*\|) + (\|a^*a'\| + \|a'b\|) \\ 537 \quad &\leq \|aa'\| + \|a'b\| + \varepsilon \frac{\Delta}{\mu} \\ 538 \quad &\leq \|ab\| + \frac{\varepsilon\Delta}{\mu}. \end{aligned}$$

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




Fig. 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 $\text{score}_\theta(a) \leq \text{score}_\theta(a^*) \leq \text{score}_\theta(b)$, and the line L passing through a^* .

542 We can now prove Lemma 2.8.

543 *Proof of Lemma 2.8.* Let ξ be the set of major strip partitions defined above. Let P
 544 be an arbitrary major strip partition in ξ , and let $P^\perp$ be the associated minor strip
 545 partition in $\xi^\perp$. For each pair of strips S_1 and S_2 in $P^\perp$, we define *tree* T_{P,S_1,S_2} as
 546 follows: For every strip S in P , apply Claim 2.10 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 tree¹⁰ 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 2.9, 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 2.10 implies that tree $T = T_{P,S_1,S_2}$ satisfies $\|ab\| \leq \delta_T(a, b) \leq (1+\varepsilon) \cdot \|ab\|$.

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 $2 \cdot (\Delta + \log(2\mu/\varepsilon) \cdot \Delta) = 2\Delta \log(4\mu/\varepsilon)$. $\square$

We can plug this partial tree cover into the reduction of Lemma 2.4 (with $\tau(\varepsilon, \mu) = O_\mu(\frac{1}{\varepsilon})$ and $\gamma = O(\log \frac{1}{\varepsilon})$) to prove Theorem 1.1 in the plane, obtaining a non-Steiner tree cover of stretch $(1 + \varepsilon)$ and size $O(\varepsilon^{-1} \log \frac{1}{\varepsilon})$ for points in the plane.

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) -partial tree cover is:

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

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

2.3 Partial tree cover with Steiner points

This part is devoted to the proof of Theorem 1.2 for $\mathbb{R}^2$; the argument for dimension $d \geq 3$ is deferred to Section 3.2.

¹⁰we use a balanced binary tree instead of a star to help guarantee bounded degree; see Section 4

Lemma 2.11. *Let X be a set of points in $\mathbb{R}^2$ with diameter Δ . For every constant μ and every $\varepsilon \leq 1/72$, there is a Steiner (μ, Δ) -partial tree cover with stretch $(1 + \varepsilon)$ for X with $O(1/\sqrt{\varepsilon})$ 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 2.12. *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. $\square$

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, $[a_0, a_1], [a_2, a_3], \dots, [a_{k-1}, a_k]$, 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. This ensures that the diameter of T is at most 3Δ .

Claim 2.13. *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 $i \in \{0, \dots, k\}$ such that $\delta_{T_S^i}(p, q) \leq (1 + \varepsilon) \cdot \|pq\|$.*

Proof. By Observation 2.12, 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 line segment ℓ of S into k intervals $[a_0, a_1], [a_2, a_3], \dots, [a_{k-1}, a_k]$ 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 pq . Hence, $\|a_i r'\| \leq \|a_i r\| \leq \frac{\sqrt{\varepsilon}\Delta}{2\mu}$ and similarly $\|r' r\| \leq \frac{\sqrt{\varepsilon}\Delta}{2\mu}$. By construction, we have:

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

Observe that $\|pr'\| \geq \|pr\| - \|r'r\| \geq \frac{1}{2} \cdot \frac{\Delta}{3\sqrt{2}\mu} - \frac{\sqrt{\varepsilon}\Delta}{2\mu} \geq \frac{\Delta}{12\sqrt{2}\mu}$, for $\varepsilon \leq 1/72$. Thus, $\|r'a_i\| \leq \frac{\sqrt{\varepsilon}\Delta}{2\mu} \leq 6\sqrt{2} \cdot \sqrt{\varepsilon} \cdot \|pr'\|$. Similarly, $\|r'a_i\| \leq 6\sqrt{2} \cdot \sqrt{\varepsilon} \cdot \|r'q\|$. Combining with Equation 2, we get:

$$\begin{aligned} \delta_{T_S^i}(p, q) &\leq \sqrt{\|pr'\|^2 + 72\varepsilon\|pr'\|^2} + \sqrt{\|r'q\|^2 + 72\varepsilon\|r'q\|^2} \\ &\leq \|pr'\|\sqrt{1 + 72\varepsilon} + \|r'q\|\sqrt{1 + 72\varepsilon} \\ &\leq \sqrt{1 + 72\varepsilon}\|pq\| \\ &\leq (1 + 72\varepsilon)\|pq\| \end{aligned}$$

Stretch $1 + \varepsilon$ can be achieved by appropriate scaling. $\square$

We now prove Lemma 2.11. 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 2.13 that $\mathcal{T}$ is a Steiner (μ, Δ) -partial tree cover for X with stretch $(1 + \varepsilon)$.

Plugging this partial tree cover into the reduction of Lemma 2.4 (with $\tau(\varepsilon, \mu) = O_\mu(\frac{1}{\sqrt{\varepsilon}})$ and $\gamma = 3$) proves Theorem 1.2 for points in the plane.

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 $\text{Time}_{\mu, \Delta}(|X|) = O(|X|/\sqrt{\varepsilon})$. Therefore, the time complexity of constructing the tree cover is $O(n \log n + n\varepsilon^{-1/2} \log(1/\varepsilon))$.

3 Tree Cover in Higher Dimensions

3.1 Non-Steiner tree covers

We now prove an analog of Lemma 2.8 in $\mathbb{R}^d$, for any constant $d = O(1)$.

Lemma 3.1. *Let X be a set of points in $\mathbb{R}^d$ with diameter Δ . For every constant $\mu > 0$ there is a (μ, Δ) -partial tree cover for X with stretch $(1 + \varepsilon)$ and size $O_d(\varepsilon^{1-d})$, where each tree has diameter at most $O_d(\Delta \log(4\mu\varepsilon))$.*

The definition of strip partition and the sets ξ and $\xi^\perp$, are different in $\mathbb{R}^d$ than in $\mathbb{R}^2$. Let θ be a vector. An $\mathbb{R}^d$ -strip with direction θ and width w is a convex region $S \subset \mathbb{R}^d$ such that there is a line ℓ in S such that every point in the strip is within distance at most $w/2$ of ℓ . The line ℓ is called the *spine* of the strip. An $\mathbb{R}^d$ -strip partition is a partition of $\mathbb{R}^d$ into $\mathbb{R}^d$ -strips. For the construction of the major strip partitions ξ , we use the following well-known lemma, slightly adapted from a version in the textbook by Narasimhan and Smid [14].

Lemma 3.2 (Cf. Lemma 5.2.3 of [14]). *Let ε be a number in $(0, 1)$. There is a set $\mathcal{V}$ of vectors in $\mathbb{R}^d$ such that (1) $\mathcal{V}$ contains $O_d(\varepsilon^{1-d})$ vectors, and (2) for any vector v in $\mathbb{R}^d$, there is some vector $v' \in \mathcal{V}$ such that the angle between v and v' is at most ε .*

We also use a variant of the shifted quadtree construction of Chan [31] (which follows immediately from our Lemma 2.1).

Lemma 3.3 (Cf. [31]). *For any constant $\Delta > 0$, there is a set $\mathcal{P}$ of partitions of $\mathbb{R}^d$ into hypercubes of side length $(4\lceil d/2 \rceil + 2)\Delta$ such that (1) there are $O(d)$ partitions in $\mathcal{P}$, and (2) for every pair of points $x, y \in \mathbb{R}^d$ with $\|xy\| \leq \Delta$, there is some partition $P \in \mathcal{P}$ where x and y are in the same hypercube in P .*

Let θ be a vector in $\mathbb{R}^d$. We define X_θ to be the hyperplane orthogonal to θ . We can view X_θ as a copy of $\mathbb{R}^{d-1}$. Let P_θ be an arbitrary partition of X_θ into $\mathbb{R}^{d-1}$ -hypercubes with side length $\varepsilon \frac{\Delta}{2\mu d}$. This partition induces an $\mathbb{R}^d$ -strip partition with direction θ and width $\varepsilon \frac{\Delta}{2\mu}$: for every hypercube R in the partition P_θ , the corresponding strip is defined by $\{r + \alpha \cdot \theta : r \in R, \alpha \in \mathbb{R}\}$. We denote this strip partition as $S(P_\theta)$. The fact that $S(P_\theta)$ has width $\varepsilon \frac{\Delta}{2\mu}$ follows from the fact that every point in a $\mathbb{R}^{d-1}$ -hypercube of side length $\varepsilon \frac{\Delta}{2\mu d}$ is within distance $\varepsilon \frac{\Delta}{4\mu}$ of the center point of the hypercube.

We now define the set ξ of major strip partitions. Let $\mathcal{V}$ be the set of vectors provided by Lemma 3.2, setting the parameter $\varepsilon' = \frac{\varepsilon}{10\mu d^2}$. For every $\theta \in \mathcal{V}$, let $\mathcal{P}_\theta$ denote the set of partitions of X_θ into $\mathbb{R}^{d-1}$ -hypercubes of side length $\varepsilon \frac{\Delta}{2\mu d}$, as guaranteed by Lemma 3.3. The set ξ_θ contains the $(\theta, \varepsilon \frac{\Delta}{2\mu})$ -strip partitions $S(P_\theta)$ associated with every $P_\theta \in \mathcal{P}_\theta$. Define $\xi = \bigcup \xi_\theta$. The following observation is immediate from Lemma 3.3:

Observation 3.4. *Let θ be a vector in $\mathbb{R}^d$. If x and y are two points whose projections onto X_θ are within distance $\varepsilon \frac{\Delta}{10\mu d^2}$, then there is some strip partition in ξ_θ with a strip containing both x and y .*

We now define $\xi^\perp$, the set of minor strip partitions. For every $\theta \in \mathcal{V}$, let $\theta^\perp$ be some arbitrary vector that is orthogonal to θ . Let $P_{\theta^\perp}$ be an arbitrary partition of $X_{\theta^\perp}$ into $\mathbb{R}^{d-1}$ -hypercubes with side length $\frac{\Delta}{2\mu d}$. Define $\xi^\perp$ to be the set containing the $(\theta^\perp, \frac{\Delta}{2\mu})$ -strip partition $S(P_{\theta^\perp})$ for every $\theta \in \mathcal{V}$. With these modified definition of ξ and $\xi^\perp$, the claims from the $\mathbb{R}^2$ case generalize naturally. We restate the necessary claims below.

Claim 2.9. *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. The proof of Claim 2.9 is similar to the $\mathbb{R}^2$ case. We break $v = y - x$ into a component parallel to θ and a component that lies in the hyperplane orthogonal to θ ; the former has length $\alpha > \frac{\Delta}{2\mu}$ and the latter has length $\beta \leq \varepsilon \frac{\Delta}{10\mu d^2}$. Observation 3.4 (together with the upper-bound on β) guarantees that there is some major strip partition in direction θ in which x and y are in the same strip. The lower bound on α implies that x and y are in different strips of the associated minor strip partition in direction $\theta^\perp$. $\square$

Claim 2.10. *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. In the proof of Claim 2.10, the only difference is that the line L in the 2D case is replaced by a hyperplane L orthogonal to θ . To show that $\|a^*a\| \leq \varepsilon \frac{\Delta}{2\mu}$, we argue as follows. Hyperplane L intersects the spine of the strip at some point s ; as the width of the strip is $\varepsilon \frac{\Delta}{2\mu}$, every point in L that is in the strip (which includes a^* and a) is within distance $\varepsilon \frac{\Delta}{4\mu}$ of s . Triangle inequality proves the claim, and the rest of the proof carries over. $\square$

Proof of Lemma 3.1. The proof of Lemma 3.1 carries over almost exactly from Lemma 2.8. The size of ξ is $O_d(\varepsilon^{1-d})$. For every major strip partition $P \in \xi$, there are $\binom{(2\mu d+1)^{d-1}}{2} = O_d(1)$ pairs of strips in the corresponding minor strip partition of $\xi^\perp$, and thus the tree cover $\mathcal{T}$ contains $O_d(\varepsilon^{1-d})$ trees. The stretch bound carries over without modification.

The diameter of each of the trees is $O_d(\Delta \cdot \log 1/\varepsilon)$, following a similar argument as in the 2-dimensional case. Every major strip partition is induced by a set of $\mathbb{R}^{d-1}$ -hypercubes with side length $\varepsilon \frac{\Delta}{2\mu^d}$. Since the diameter of the point set is Δ , the number of hypercubes required is at most $\left(\frac{2\mu^d}{\varepsilon}\right)^{d-1}$.

The height of the binary tree is at most $(d-1) \log \frac{2\mu^d}{\varepsilon}$ and each edge is of length at most Δ . The diameter is at most $2(\Delta + \Delta(d-1) \log \frac{2\mu^d}{\varepsilon}) = O_d(\Delta \cdot \log 1/\varepsilon)$. $\square$

Together with the reduction to a fixed scale (Lemma 2.4, with $\tau(\varepsilon, \mu) = O_{\mu,d}(\varepsilon^{1-d})$ and $\gamma = O_d(\log 1/\varepsilon)$), we obtain Theorem 1.1.

Running Time: The running time analysis is similar to the 2D case. The inner product between each point to each direction vector can be precomputed in $O_d(\varepsilon^{1-d})|X|$ time. For each major strip S , there are at most $\binom{(2\mu^d+1)^{d-1}}{2}$ pair of minor strips that intersect $S \cap X$. Hence, the total running time to construct a (μ, Δ) -partial tree cover is:

$$\text{Time}_{\mu,\Delta}(|X|) = \varepsilon^{1-d}|X| + \binom{(2\mu^d+1)^{d-1}}{2} \cdot \sum_{S \text{ is a major strip}} |S \cap X| = O_d(\varepsilon^{1-d}|X|)$$

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

3.2 Tree covers with Steiner points

The construction for d -dimensional Euclidean space is a direct generalization of two-dimensional case. Consider a hypercube of side length Δ . We divide the hypercube by each coordinate into slabs of height $\frac{\Delta}{3\sqrt{d}\mu}$; all other sides have length Δ . One can think of each slab as a d -dimensional rectangle, joined by two $(d-1)$ -hypercubes that are at a distance $\frac{\Delta}{3\sqrt{d}\mu}$ away from each other. Analog of Obs. 2.12 follows.

Observation 3.5. *For any two points $p, q \in X$ such that p and q are (μ, Δ) -far, there exists a slab such that p and q are from different sides of it.*

Each tree is constructed similarly to the two-dimensional case. For each slab, we find a $\frac{\sqrt{\varepsilon}\Delta}{2\mu}$ -net for the $(d-1)$ -hypercube at the middle of each slab. For each net point u , we create a tree connecting u to every vertex in X . The proof follows similarly. The total number of trees is

$$O\left(d \cdot 3\sqrt{d}\mu \cdot \left(\frac{2\mu}{\sqrt{\varepsilon}}\right)^{d-1}\right) = O\left(d^{3/2}2^d\mu^d\varepsilon^{(1-d)/2}\right).$$

The diameter can be bounded by $O_d(\Delta)$, as in the 2-dimensional argument. Together with the reduction to a fixed scale (Lemma 2.4, with $\tau(\varepsilon, \mu) = O_{\mu,d}(\varepsilon^{(1-d)/2})$ and $\gamma = 3$), this proves Theorem 1.2.

Running Time: Creating the set of slabs requires time equal to the number of slabs, which is $d \cdot 3\sqrt{d}\mu = O(d^{3/2})$. For each slab, we find a $\frac{\sqrt{\varepsilon}\Delta}{2\mu}$ -net of a $(d-1)$ -hypercube,

which has $(2\mu/\sqrt{\varepsilon})^{d-1}$ points. For each Steiner point, we connect it to every point in X in $O(|X|)$ time. Hence, the total time complexity is:

$$\text{Time}_{\mu,\Delta}(|X|) = O(d^{3/2} \cdot (2\mu/\sqrt{\varepsilon})^{d-1} \cdot |X|) = O_d(\varepsilon^{(1-d)/2} |X|).$$

Then, the time complexity of constructing the tree cover is $O_d(n \log n + n\varepsilon^{(1-d)/2} \log 1/\varepsilon)$.

4 Constant degree constructions

In this section, we prove the following theorem.

Theorem 4.1. *For every set of points in $\mathbb{R}^d$ and any $0 < \varepsilon < 1/16$, there exists a tree cover with stretch $1 + \varepsilon$ and $O_d(\varepsilon^{-(d-1)} \log 1/\varepsilon)$ trees such that every metric point has degree at most 11 in each of the trees.*

Our tree cover construction is a collection of trees, each of which possibly uses a copy of the same point many times. Each tree is constructed iteratively, going from smaller scales to the larger ones. In Section 4.1 we use the degree reduction technique due to [7] for each tree in the cover. This allows us to bound the degree in terms of the number of trees in the cover and the degree at a single scale in the construction. In Section 4.2, we show that the degree at a single scale is constant.

4.1 Bounding the degree of metric points

Consider a single tree in the cover. Our tree cover construction from Section 2.2 does not use Steiner points, but it still might consider the same point from the metric X across multiple levels of construction. Even if at each level of the construction every node has a bounded degree (which we show how to achieve in Section 4.2) the degree of each metric point might still be unbounded. To remedy this, we apply the degree reduction technique of [7].

Start from the tree cover construction from Section 2 and fix a tree $T = (V, E)$ from the cover. Let $\ell = \log(d\sqrt{d}/\varepsilon)$ be the same as in Section 2.1. Without loss of generality, assume the tree was constructed in the congruence class $I_j := \{z \mid z \equiv j \pmod{\ell}\}$. Assume that at the every level of the construction, the edges of the tree are oriented from the parents to the children so that the outdegree of each node is α and indegree is 1. We show how to bound the degree of every point of the metric with respect to T .

Let $i^*(v)$ be the highest quadtree level at which point v is considered as a representative. For every edge (u, v) in T , we orient it from u to v if $i^*(u) < i^*(v)$. If $i^*(u) = i^*(v)$, break the ties according to the tree structure, from children towards parent. We use $\hat{E}$ to denote the set of arcs obtained in this way. Note that $|\hat{E}| = |E|$, since we do not change any edges. Next, we describe the modification of $\hat{E}$, where we replace some edges of $\hat{E}$ and obtain the set of $\tilde{E}$. Let u be a vertex at level i and let $\hat{E}_i$ be the set of edges used in the tree constructed at level i . Let $M_i(u)$ be the set of endpoints of edges in $\hat{E}_i$ oriented into u . Let $\mathcal{I}_u := \{i \mid M_i(u) \neq \emptyset\}$. Suppose that the indices in $\mathcal{I}_u$ are ordered increasingly. Next we modify arcs going into u as follows. Keep $M_{i_1}(u)$ directed into u . For $j > 1$ we pick an arbitrary vertex $w \in M_{i_j - \ell}(u)$ and for each point $v \in M_{i_j}(u)$ replace arc (v, u) by an arc (v, w) .

Claim 4.2. *If at every level T has an outdegree α , then every metric point has outdegree α . Moreover, every node with an outgoing edge at level i ceases to be considered at levels higher than i .*

Proof. Consider an arc (u, v) , i.e., an edge (u, v) directed from u towards v . This means that $i^*(u) \leq i^*(v)$. Let i be the level at which (u, v) was added to T . Recall that the edges are added while handling a single quadtree cell at level i and only one point from the cell is chosen a representative for the subsequent handling of level $i + \ell$. If $i^*(u) < i^*(v)$, this means that v is a representative and u does not exist on any subsequent level starting from $i + \ell$. If $i^*(u) = i^*(v)$, then neither u nor v are representatives, since the representative exists at a level higher than $i^*(u)$. Hence, u does not exist on any subsequent level starting from $i + \ell$. In conclusion, u can have outgoing edges only at a single level of construction. $\square$

Claim 4.3. *If at every level T has an outdegree α and indegree β , then every metric point has degree with respect to T at most $\alpha + \beta + \alpha\beta$.*

Proof. Consider a metric point w . There are at most α edges directed out of w by Claim 4.2. Out of the edges that were directed into w in $\hat{E}$, there are only edges from $M_{i_1}(w)$ that remained directed into w . There are at most β such edges. Finally, some new edges might have been attached to w due to the modification into $\tilde{E}$. Consider an arc (w, u) directed out of w ; there is a unique level i_j where w is in $M_{i_j-\ell}(u)$. Only edges of the form (v, u) at level i_j can be redirected to (v, w) by the modification process; each such v must be an in-neighbor of u at level i_j . A counting argument shows that for each arc (w, u) going out of w , there are at most β new arcs attached to w , each attribute to an in-neighbor of u in $\hat{E}$ before the modification; and there are α possible choices of u , all being the out-neighbors of w . Putting everything together, the bound on the degree is $\alpha + \beta + \alpha\beta$. $\square$

We next show that the modification of the edges of T does not create cycles.

Claim 4.4. *The modified tree does not contain cycles.*

Proof. Suppose towards contradiction that the modified T contains a cycle and let (v, w) be the first edge in the modification process whose insertion caused a cycle. Recall that the arc (v, w) gets inserted in place of arc (w, u) , where $v \in M_{i_j}(u)$ and $w \in M_{i_j-\ell}(u)$, for some levels i_j and $i_j - \ell$. Since (v, w) introduces a cycle, this means that T contains an alternative path between v and w . By Claim 4.2, node w does not exist at level higher than $i_j - \ell$. Hence, the path appears at some level lower than $i_j - \ell$. In the original tree T , there were edges (v, u) and (w, u) . Together with the path between w and v , this creates a cycle in the original T , a contradiction. $\square$

Next, we show that the stretch does not increase by more than a $(1 + O(\varepsilon))$ factor.

Claim 4.5. *Let $\tilde{d}_T$ be the metric induced by $\tilde{T}$. Then, for every $u, v \in V(T)$, it holds $\|uv\| \leq (1 + O(\varepsilon))d_T(u, v)$.*

Proof. It suffices to show that for an arc (v, u) that is removed from $\hat{E}$ it holds $\tilde{d}_T(u, v) \leq (1 + O(\varepsilon))d_T(u, v)$. Let $s = \lfloor j/\ell \rfloor$. By construction, since (v, u) is removed, there exists points $v_0, v_1, \dots, v_s$, such that $v = v_0$, $(v_s, u) \in \tilde{E}$ and for $0 \leq k < s$:

(v_k, v_{k+1}) $\in \tilde{E}$, and $v_k \in M_{j-k\ell}(u)$. Recall that we use $\Delta_w = 2^w \sqrt{d}$ to denote the diameter of a quadtree at level w .

Observation 4.6. *For every $0 \leq k < s$, $\|uv_{k+1}\| \leq \varepsilon^\ell \|uv_k\|$.*

Applying Observation 4.6 inductively, we can prove that $\|uv_k\| \leq \varepsilon^{\ell k} \|uv_0\|$ for every $0 \leq k \leq s$. We can also bound $\|v_k v_{k+1}\| \leq \|v_k u\| + \|uv_{k+1}\| \leq (1 + \varepsilon) \|v_k u\| \leq (1 + \varepsilon) \varepsilon^{\ell k} \|uv_0\|$. By triangle inequality $\|uv\|$ can be upper bounded by the length of the path $\langle v_0, v_1, \dots, v_s, u \rangle$.

$$\begin{aligned} \tilde{d}_T(u, v) &\leq \|v_s u\| + \sum_{0 \leq k < s} \|v_k v_{k+1}\| \\ &\leq \varepsilon^{-\ell s} \|uv_0\| + \sum_{0 \leq k < s} (1 + \varepsilon) \varepsilon^{\ell k} \|uv_0\| \\ &\leq (1 + O(\varepsilon)) \|uv_0\| \quad \square \end{aligned}$$

In the next subsection, we show that $\alpha = 1$ and $\beta = 5$. Plugging in Claim 4.3, we obtain the bound of $\alpha + \beta + \alpha\beta = 11$ on the degree of every node in the metric.

4.2 Bounding the degree of tree nodes

Recall our construction from Section 2.2. Consider a single strip S and the sets A and B as in the proof of Claim 2.10. Let θ be the direction of strip S . The tree handling the distances between A and B is a star rooted at a point $a^* := \arg \max_{a \in A} \text{score}_\theta(a)$. We describe three different constructions. The first construction achieves a constant degree but it requires scaling after which the number of trees grow by a factor of roughly $\log(1/\varepsilon)^d$. The second construction achieves degree of roughly 2^d and does not require any scaling. The third construction achieves degree 5 and requires scaling so that the number of trees grows by a factor of roughly $d^{d-1} = O_d(1)$.

Constant degree, simple attempt.

Let $A' := \langle a^*, a_1, a_2, \dots \rangle$ be the set of points in A , sorted in decreasing order with respect to score_θ . We make a balanced binary tree T_A rooted at a_1 such that for every node a_i and its parent a_j , $\text{score}_\theta(a_i) \leq \text{score}_\theta(a_j)$. To do so, we mark a_1 as *visited* and make it the root of T_A . Next, we scan the points $a_2, a_3, \dots$ in order, make a_i child of the node in T_A that was visited earliest and still has 0 or 1 children, then mark a_i visited. We similarly construct T_B satisfying that for every node b_i and its parent b_j , $\text{score}_\theta(b_i) \geq \text{score}_\theta(b_j)$. Finally, let T be the tree rooted at a^* having subtrees T_A and T_B as its children.

Recall that, during the reduction to a single scale (Lemma 2.4), we only construct the partial tree covers of Section 2.2 on $1/\varepsilon^{O_d(1)}$ representative points contained in a quadtree cell. Since T_A is balanced, it has height $O(\log |A|) = O(\log(1/\varepsilon))$. Similarly, the height of T_B is $O(\log(1/\varepsilon))$. In other words, between any node in A and any node in B there exists a path in T consisting of at most $O(\log(1/\varepsilon))$ edges.

We next prove the bound on the stretch between two points $a \in A$ and $b \in B$ that are (μ, Δ) -far. The proof follows the lines of Claim 2.10. Consider line ab and let




Fig. 3. The binary trees T_A and T_B , constructed greedily from point sets A and B

853 $a = c_1, c_2, \dots, c_k = b$ be the points on the path from a to b in T . For $1 \leq i \leq k$, let c'_i be
 854 the intersection of ab with a line¹¹ orthogonal to θ that passes through c_i ; as the width
 855 of the strip is no more than $\varepsilon\Delta/\mu \leq \varepsilon\|ab\|$, we have $\|c'_i c_i\| \leq \varepsilon\|ab\|$. By construction,
 856 every path in T between a point in A and a point in B goes up the subtree T_A , passes
 857 through the root a^* and goes down the tree T_B . In other words, we have that for
 858 every $i \in \{1, 2, \dots, k-1\}$ it holds that $\text{score}_\theta(c_i) \leq \text{score}_\theta(c_{i+1})$ and $\text{score}_\theta(c'_i) \leq$
 859 $\text{score}_\theta(c'_{i+1})$. In addition, $k = O(\log 1/\varepsilon)$. We have $\|ab\| = \sum_{1 \leq i \leq k-1} \|c'_i c'_{i+1}\|$. Thus,
 860 the length of the path in T is

$$\begin{aligned}
 861 \quad d_T(a, b) &= \sum_{1 \leq i \leq k-1} \|c_i c_{i+1}\| \\
 862 \quad &\leq \sum_{1 \leq i \leq k-1} (\|c'_i c'_{i+1}\| + 2\|c'_{i+1} c_{i+1}\|) + \|c'_1 c_1\| \quad \text{by triangle inequality} \\
 863 \quad &\leq \sum_{1 \leq i \leq k-1} (\|c'_i c'_{i+1}\| + 2\varepsilon\|ab\|) + \varepsilon\|ab\| \\
 864 \quad &= \|ab\| + O(\varepsilon \log(1/\varepsilon))\|ab\|.
 \end{aligned}$$

865 The above stretch argument guarantees that T preserves path between any point
 866 in X and any point in Y up to a factor of $1 + O(\varepsilon \log(1/\varepsilon))$. Applying the same degree
 867 reduction step for every strip in the strip partition and for every strip partition in the
 868 family ζ , we obtain a tree cover with $O(1/\varepsilon)$ trees and stretch $O(1 + \varepsilon \log(1/\varepsilon))$. To
 869 complete the argument, we need to scale the parameters. Let $\varepsilon' := O(\varepsilon \log(1/\varepsilon))$, so
 870 that the tree cover has stretch $1 + \varepsilon'$. The number of trees expressed in terms of ε' is
 871 $O(\frac{\log^d(1/\varepsilon')}{\varepsilon'})$, for $\varepsilon < 1/16$.

¹¹In higher dimensions, we consider the hyperplane orthogonal to θ .

Degree $O_d(1)$.

Let $A' = \langle a^*, a_1, a_2, \dots \rangle$ be the set of points in A , sorted in decreasing order with respect to score_θ . Recall that the direction of the major strip is θ and the direction of the minor strip is $\theta^\perp$. We build a binary tree T_A rooted at a_1 as follows. Let the interval corresponding to a_1 be $[0, \varepsilon \frac{\Delta}{2\mu})$. Recall that the width of the strip is $\varepsilon \frac{\Delta}{2\mu}$. Let $\mathcal{I}$ be the set of active intervals, consisting of two elements: $[0, \varepsilon \frac{\Delta}{4\mu})$, corresponding to the future left child of a_1 (if any) and $[\varepsilon \frac{\Delta}{4\mu}, \varepsilon \frac{\Delta}{2\mu})$ corresponding to the future right child of a_1 (if any). The elements of $\mathcal{I}$ form a partition of $[0, \varepsilon \frac{\Delta}{2\mu})$ at all times. Scan the points $a_2, a_3, \dots$ in order and perform the following. Let a_i be the currently scanned point and let d_i be its distance from the left border of the strip. Go over all the intervals in $\mathcal{I}$ and see which one contains d_i . (Such an interval exists because $\mathcal{I}$ forms a partition of $[0, \varepsilon \frac{\Delta}{2\mu})$.) Let $[l_i, r_i)$ be such an interval. Add a_i at the corresponding place in the tree. Let $m_i := (l_i + r_i)/2$. Create two new intervals: $[l_i, m_i)$ corresponding to the left child of d_i and $[m_i, r_i)$, corresponding to the right child of d_i . Note that after this, $\mathcal{I}$ still forms a partition of $[0, \varepsilon \frac{\Delta}{2\mu})$. This concludes the description of T_A . The tree T_B is constructed analogously. Finally, the tree T is obtained by attaching the roots of T_A and T_B as the left and right child of a^* .

We next analyze the stretch. Consider two points $a \in A$ and $b \in B$ that are (μ, Δ) -far. Let $c_1 = a^*, c_2 = a_1, c_3, \dots, c_p$ be the path from a^* (which is the root of T) to a in T and let $d_1 = a^*, d_2 = b_1, d_3, \dots, d_q = b$ be the path from a^* to b in T . For two points x and y , let $x = x_\theta \cdot \theta + x_\perp \cdot \theta^\perp$ and similarly $y = y_\theta \cdot \theta + y_\perp \cdot \theta^\perp$. Let $\|xy\|_\theta = |x_\theta - y_\theta|$ and $\|xy\|_\perp = |x_\perp - y_\perp|$. Using this notation, we observe that $\|ab\|_\theta = \sum_{i=1}^{p-1} \|c_i c_{i+1}\|_\theta + \sum_{i=1}^{q-1} \|d_i d_{i+1}\|_\theta$. The second observation is that $\sum_{i=1}^{p-1} \|c_i c_{i+1}\|_\perp = O(\varepsilon \Delta / \mu)$. This is because $\|c_i c_{i+1}\|_\perp$ form a geometrically decreasing sequence. Similarly, $\sum_{i=1}^{q-1} \|d_i d_{i+1}\|_\perp = O(\varepsilon \Delta / \mu)$. Using these two observations, we can upper bound the distance between a and b in T as follows.

$$\begin{aligned}
d_T(a, b) &= \sum_{i=1}^{p-1} \|c_i c_{i+1}\| + \sum_{i=1}^{q-1} \|d_i d_{i+1}\| \\
&\leq \sum_{i=1}^{p-1} \|c_i c_{i+1}\|_\theta + \sum_{i=1}^{p-1} \|c_i c_{i+1}\|_\perp + \sum_{i=1}^{q-1} \|d_i d_{i+1}\|_\theta + \sum_{i=1}^{q-1} \|d_i d_{i+1}\|_\perp \\
&\leq \|ab\|_\theta + \sum_{i=1}^{p-1} \|c_i c_{i+1}\|_\perp + \sum_{i=1}^{q-1} \|d_i d_{i+1}\|_\perp \\
&\leq \|ab\| + O(\varepsilon \Delta / \mu) \\
&\leq (1 + \varepsilon) \|ab\|
\end{aligned}$$

The argument for higher dimensions carries over almost exactly. The intervals used in the argument become $\mathbb{R}^{d-1}$ -hypercubes. Consider a tree node a and an interval $I_a \subset \mathbb{R}^{d-1}$ corresponding to it. We partition the interval I_a into 2^{d-1} subintervals of twice the smaller side length. Those subintervals correspond to the children of a . To argue the stretch, we split the distance between points a and b in T into two

components: one along the vector θ and the remaining orthogonal part that lies in $\mathbb{R}^{d-1}$. The component along θ is at most $\|ab\|$ and the component in $\mathbb{R}^d$ is at most $O(\varepsilon)\|ab\|$, due to the geometrically decreasing interval sizes.

Finally, we bound the diameter of each of the trees. Using analysis similar to the one used for the stretch, we conclude that the diameter of a tree corresponding to a single strip is at most $\|ab\|(1 + O(\varepsilon)) \leq 2\|ab\|$. The trees of different major strips in a major strip partition are connected via a binary tree. As in Section 3.1, the height of the binary tree is at most $\log \frac{4\mu d}{\varepsilon}$. Hence, the overall degree is $2^{d-1} + 2$. The diameter of the tree is at most $2(2\Delta + \Delta \log \frac{4\mu d}{\varepsilon}) \leq 2\Delta \log \frac{16\mu d}{\varepsilon}$.

Constant degree.

We next explain a tweak which leads to degree 5. Instead of constructing a 2^{d-1} -ary tree for each strip we can work with a binary tree. Tree T_A is built as follows. Let $[0, \varepsilon \frac{\Delta}{\mu d}]^{d-1}$ be the interval corresponding to a_1 . We assign *level* to each node in the tree, ranging from 1 to $d - 1$. The level of a_1 is 1. The future children of a_1 are at level 2. In general, the children of a node at level $i < d - 1$ are at level $i + 1$ and the children of a node at level $d - 1$ are at level 1. The set of active intervals $\mathcal{I}$ consists of $[0, \varepsilon \frac{\Delta}{2\mu d}] \times [0, \varepsilon \frac{\Delta}{\mu d}]^{d-2}$, corresponding to the left child of a_1 and $[\varepsilon \frac{\Delta}{2\mu d}, \varepsilon \frac{\Delta}{\mu d}] \times [0, \varepsilon \frac{\Delta}{\mu d}]^{d-2}$. Once again, we maintain the property that $\mathcal{I}$ is a partition of $[0, \varepsilon \frac{\Delta}{\mu d}]^{d-1}$. Scan the points $a_2, a_3, \dots$ in that order and let a_i be the currently scanned point and d_i the $(d - 1)$ -dimensional vector of distances from each of the sides of the strip. Find the interval $I = [l_1, r_1] \times [l_2, r_2] \times \dots \times [l_{d-1}, r_{d-1}]$ in $\mathcal{I}$ where d_i belongs to and place a_i at the corresponding place in the tree. Let $j \in \{1, 2, \dots, d\}$ be the level of a_i . Let $m_j := (l_j + r_j)/2$. Split I into $I_l = [l_1, r_1] \times [l_2, r_2] \times \dots \times [l_j, m_j] \times \dots \times [l_{d-1}, r_{d-1}]$ corresponding to the left child of a_i and $I_r = [l_1, r_1] \times [l_2, r_2] \times \dots \times [m_j, r_j] \times \dots \times [l_{d-1}, r_{d-1}]$ corresponding to the right child of a_i . Replace I with I_l and I_r in $\mathcal{I}$. This concludes the description of the binary tree.

The stretch argument remains almost the same, except that $\sum_{i=1}^{p-1} \|c_i c_{i+1}\|_{\perp} = O(d\varepsilon\Delta/\mu)$, which is d times larger than before. The reason is that every d hops down the tree, we incur an additive stretch of $O(\varepsilon\Delta/\mu)$ after which the additive stretch reduces by a factor of two. Using the same argument as before, we conclude that $d_T(a, b) \leq (1 + O(\varepsilon d))\|ab\|$. By scaling the stretch, we get that the number of trees increases by a factor of d^{d-1} .

5 Application to Routing

In this section, we show an application of our tree cover to compact routing scheme; in particular, we prove Theorem 1.3.

We start by giving some background on the problem. A *compact routing scheme* is a distributed algorithm for sending messages or packets of information between points in the network. Specifically, a packet has an origin and it is required to arrive at a destination. Every node in the network contains a *routing table*, which stores local routing-related information, and a unique *label*, sometimes also called *address*. In the beginning, the network is preprocessed and every node is assigned a routing table and a label. Given a destination node v , routing algorithm is initiated at source u and is

given the label of v . Based on the local routing table of u and the label of v , it has to decide on the next node w to which the packet should be transmitted. More formally, the algorithm outputs the *port number* leading to its neighbor w . Each packet has a message *header* attached to it, which contains the label of the destination node v , but may also contain other helpful information. Upon receiving the packet the algorithm at node w has at its disposal the local routing table of w and the information stored in the header. This process continues until the packet arrives at its destination, which is node v . The *stretch* of the routing scheme is the ratio between the distance packet traveled in the network and the distance in the original metric space.

We consider routing in metric spaces, where each among n points in the metric corresponds to a network node. In the preprocessing stage, we choose a set of links that induces an *overlay network* over which the routing must be performed. The goal is to have an overlay network of small size, whilst also optimizing the tradeoff between the maximum storage per node (that is, the size of routing tables, labels, and headers) and the stretch. In addition, one may try to further optimize the time it takes for every node to determine (or output) the next port number along the path, henceforth *decision time*, and other quality measures, such as the maximum degree in the overlay network.

There are two different models, based on the way labels are chosen: *labeled*, where the designer is allowed to choose (typically $\text{polylog}(n)$) labels, and *name-independent*, where an adversary chooses labels. Similarly, depending on who is choosing the port numbers, there is a *designer-port* model, where the designer can choose the port number, and the *fixed-port* model, where the port numbers are chosen by an adversary. Our routing scheme works in the labeled, fixed-port model. For an additional background on compact routing schemes, we refer the reader to [37–40].

5.1 Routing in trees

We first explain the interval routing scheme due to [41]. Let T be a given routed tree. We first preprocess the tree by performing a DFS on it and marking for every node u the timestamp at which it got visited, l_u . For every node u , let h_u be the maximum l_w among the children w of u . The label of node u consists of l_u and requires $\lceil \log n \rceil$ bits of storage. The routing table of node u consists of the port number leading to its parent in T (unless u is a root), and for each child w_i of u , the port number leading to w_i together with $\langle l_{w_i}, h_{w_i} \rangle$. This requires $\deg_T(V) \cdot O(\log n)$ bits. Specifically, it requires $O(\log n)$ bits for trees of constant degree, which is the case for our construction. To route from some node u to a destination v , the routing algorithm has routing table of u and the label of v at its disposal. For every child w_i of u , if l_v falls in the interval $\langle l_{w_i}, h_{w_i} \rangle$, the algorithm outputs the corresponding port to w_i and otherwise the algorithm outputs port to the parent of u . Note that in bounded degree trees the aforementioned routing algorithm needs to inspect only a constant number of entries in order to decide on the next port.

5.2 Routing in Euclidean spaces

To route in a Euclidean space, first construct a non-Steiner tree cover $\mathcal{T}$ with bounded degree, using Theorem 4.1. The routing table of each point consists of its routing table in each of the trees in the cover, which takes $O_d(\varepsilon^{-(d-1)} \log^2 1/\varepsilon \cdot \log n)$ bits, since each tree is of a constant degree. The label for each point consists of its label in each of the trees in $\mathcal{T}$, which overall takes $O_d(\varepsilon^{-(d-1)} \log 1/\varepsilon \cdot \log n)$ bits, together with an additional label of $O_d(\varepsilon^{-(d-1)} \log^2 1/\varepsilon \log n)$ bits described in the next section (“identifying a distance-preserving tree”). Overall this label takes $O_d(\varepsilon^{-(d-1)} \log^2 1/\varepsilon \cdot \log n)$ bits. To route from a point x to some other point y , the algorithm first identifies a tree in $\mathcal{T}$ that preserves the xy distance up to a $(1 + \varepsilon)$ factor: this step is described in the next section. After that, the routing algorithm proceeds on the single tree as described before.

5.3 Identifying a distance-preserving tree

Given two points x and y in $\mathbb{R}^d$, we now describe how to identify a tree in $\mathcal{T}$ that preserves the distance between x and y up to $1 + \varepsilon$ stretch. The total size of this label will be $O_d(\varepsilon^{-1} \log^2 1/\varepsilon \cdot \log n)$.

Review of tree cover construction.

We first recall the construction of Theorem 1.1. We have a collection of compressed quadrees Q_i (for every $i \in [O_d(1)]$) and congruence classes $j \in [\ell]$ (where $\ell = O_d(\log 1/\varepsilon)$). For ease of notation, let $Q_{i,j}$ denote the tree obtained by starting with Q_i and then contracting away all nodes except those at level w for $w \equiv j \pmod{\ell}$. We refer to $Q_{i,j}$ as a *contracted quadtree*. Notice that if C is a cell in the contracted quadtree $Q_{i,j}$ with diameter Δ , then the children of C in $Q_{i,j}$ have diameter $O_d(\varepsilon\Delta)$. For every shift i and congruence class j , we construct a set of trees as follows: for every cell C in $Q_{i,j}$, we arbitrarily choose a set of $1/\varepsilon^{O_d(1)}$ representative points, one from each *child cell* of C in $Q_{i,j}$; we construct a partial tree cover on the representative points; and we merge these partial tree covers together into a final set of trees. Our proof of correctness guarantees that, for any pair of points x and y , there is some contracted quadtree $Q_{i,j}$ and some cell C in $Q_{i,j}$ with diameter Δ , such that the two representative points x' and y' are (μ, Δ) -far. There is some tree in the partial tree cover of C that preserves $\|x'y'\|$ up to a factor $1 + \varepsilon$, and this tree corresponds to the tree in the final tree cover of Theorem 1.1 that preserves $\|xy\|$ up to a factor $1 + O(\varepsilon)$.

In Section 4.2, we constructed a tree cover in which each partial tree cover had bounded-degree. The construction is identical to that of Theorem 1.1, except that we use a slightly modified construction for the partial tree cover on the representative points (modified from Section 2.2).

In Section 4.1, we used the result of Section 4.2 to get a bounded-degree tree cover (proving Theorem 4.1). The trees constructed in this section are in one-to-one correspondence with the trees constructed in Section 4.2: if a tree T in the cover of Section 4.2 preserves the distance between two points x and y up to a factor $1 + \varepsilon$, the corresponding transformed tree T' from Section 4.2 will preserve the distance up to a factor $1 + O(\varepsilon)$.

For simplicity, we describe how to identify a distance-preserving tree in the tree cover of Section 4.2: for any x and y , we will find a tree T such that $\delta_T(x, y) \leq (1 + \varepsilon)\|xy\|$. As described above, these trees are in one-to-one correspondence with the bounded-degree trees of Theorem 4.1 (which is the tree cover we actually use for routing).

Our labeling scheme will consist of a short label for each tree $Q_{i,j}$. For each $Q_{i,j}$, this label will let us identify a cell whose partial tree cover preserves the $\|xy\|$ distance (if such a cell exists), as well as the index of the corresponding distance-preserving tree. To construct this label for $Q_{i,j}$, we will need the following simple observation:

Observation 5.1. *Let x and y be points in X , and let $Q_{i,j}$ be a contracted quadtree. Suppose there is a cell $\hat{C}$ of $Q_{i,j}$ such that $\hat{C}$ contains both x and y , and the representatives $\hat{x}$ and $\hat{y}$ are $(\mu, \text{diam}(\hat{C}))$ -far. Then, in the smallest-diameter cell C that contains both x and y , the representatives x' and y' are $(\mu, \text{diam}(C))$ -far. In other words, if we view x and y as leaves of $Q_{i,j}$, the lowest common ancestor of x and y guarantees that the representatives are $(\mu, \text{diam}(C))$ -far.*

5.3.1 Identifying a valid partial tree cover

In this subsection, we describe a labeling scheme that lets us identify a distance-preserving tree in a *partial* tree cover.

Lemma 5.2. *Let $X \subset \mathbb{R}^d$ be a point set with diameter Δ . For any constant $\mu = O_d(1)$, there is a labeling scheme with $O_d(1)$ -bit labels, such that given the labels of any two points $x, y \in X$, we can either certify that x and y are $(\mu/4, \Delta)$ -far or that they are not (μ, Δ) -far.*

Proof. Let $x \in \mathbb{R}^d$ be a point with coordinates $x[1], \dots, x[d]$. The label of x consists of d parts: for each coordinate $i \in \{1, \dots, d\}$, the label stores difference between $x[i]$ and $\min_{x' \in X} x'[i]$ rounded to a multiple of $\Delta/(8\mu d)$; that is, we store $\lfloor \frac{x[i] - \min_{x' \in X} x'[i]}{\Delta/(8\mu d)} \rfloor$. Because the maximum such difference is Δ (because the diameter of X is Δ), the label takes $O_d(1)$ bits in total.

Given the labels of any two points x and y , we can compute, for each coordinate, an estimate of their difference within accuracy $\pm \Delta/(4\mu d)$. Thus, we can estimate ℓ_2 distance between x and y within an accuracy of $\pm \Delta/(4\mu)$. If this estimated distance is at least $\Delta/2\mu$, the $\|xy\| \geq \Delta/(4\mu)$, and so x and y are $(\mu/4, \Delta)$ -far. Otherwise, if the estimated distance is smaller than $\Delta/(2\mu)$, we have $\|xy\| < \Delta/\mu$, and so x and y are not (μ, Δ) -far. $\square$

Notice that if a partial tree cover consists of $O(1/\varepsilon^{d-1})$ trees, one tree in the cover can be identified with $O_d(\log 1/\varepsilon)$ bits. We will allow these “IDs” of the trees to be fixed in advance.

Lemma 5.3. *Let $\mathcal{T}'$ be a (μ, Δ) -partial tree cover for a point set $X \subset \mathbb{R}^d$, constructed as in Section 4.2, with $\mu = O_d(1)$. Let $\text{ID} : \mathcal{T}' \rightarrow \{0, 1\}^k$ be a function that maps trees to unique identifiers. There is a labeling scheme for X with $O_d(1/\varepsilon^{d-1}(\log 1/\varepsilon + k))$ -bit labels, such that given the labels of any two points $x, y \in X$, we can either return $\text{ID}(T)$ for some tree $T \in \mathcal{T}'$ that preserves the distance $\|xy\|$ up to a $1 + O(\varepsilon)$ factor, or we can certify that x and y are not (μ, Δ) -far.*

Proof. Recall from Sections 2.2 and 4.2 that the construction of the partial tree cover $\mathcal{T}'$ proceeds by constructing $O_d(1/\varepsilon^{d-1})$ major strip partitions and $O_d(1)$ minor strip partitions. The major strip partitions have width $\varepsilon \frac{\Delta}{2\mu} = \Theta_d(\varepsilon\Delta)$; thus, in each major strip partition, there are $1/\varepsilon^{O(d)}$ strips that contain points in X . The minor strip partitions have width $\frac{\Delta}{2\mu} = O_d(\Delta)$; thus, in each minor strip partition, there are $O_d(1)$ strips that contain points in X . Every triple consisting of a major strip partition P and two minor strips in the associated minor strip partition $P^\perp$ corresponds to some tree in the partial tree cover.

Label. For every point $x \in X$, the label consists of four parts:

- For each of the $O_d(1/\varepsilon^{d-1})$ major strip partitions, store a $O_d(\log 1/\varepsilon)$ -bit label identifying which strip in the major strip partition contains x .
- Similarly, for each of the $O_d(1)$ minor strip partitions, store a $O_d(1)$ -bit label identifying which strip in the minor strip partition contains x .
- Store the $O_d(1)$ -bit label of Lemma 5.2.
- For each of the $O_d(1/\varepsilon^{d-1})$ triples consisting of a major strip partition P and two minor strips in the associated partition $P^\perp$, store the k -bit identifier (given by $\text{ID}(\cdot)$) of the corresponding tree.

Size. The total size of all four parts is $O_d(1/\varepsilon^{d-1}(\log 1/\varepsilon + k))$.

Label correctness. Suppose we have the labels of two points $x, y \in X$. First, we use the label of Lemma 5.2 to determine either (1) x and y are $(\mu/4, \Delta)$ -far, or (2) x and y are not (μ, Δ) -far. In the latter case, we are done; we have a certificate that x and y are not (μ, Δ) -far.

In the former case, we use the labels to check if there is some major strip partition P such that the points x and y are in the same strip of P , and x and y are in different strips of the associated minor strip partition $P^\perp$. If there is no such strip, then Claim 2.9 implies that x and y are not (μ, Δ) -far, and we are done. Suppose there is such a strip. By the construction in Section 4.2, this triple of major strip and minor strips corresponds to a tree T in the partial tree cover $\mathcal{T}'$. Claim 2.10 implies that $\delta_T(a, b) \leq \|ab\| + \frac{\varepsilon\Delta}{\mu}$. As a and b are $(\mu/4, \Delta)$ -far, we have $\delta_T(a, b) \leq (1 + 4 \cdot \varepsilon)\|ab\|$. Thus, we have identified a tree in $\mathcal{T}'$ that preserves the distance between a and b up to a $1 + O(\varepsilon)$ factor. $\square$

5.3.2 LCA Labeling Tools

Before constructing our distance label, we need a preliminary result on LCA labeling. For any two vertices x and y in a tree T , let $\text{lca}(x, y)$ denote the *lowest common ancestor* of x and y . For any vertex in the tree, we say its weight is the number of descendants. We say a vertex is *heavy* if its weight is greater than half the weight of its parent, otherwise it is *light*. For any vertex x , let $\text{Apices}[T, x] = \{a_1, \dots, a_{O(\log n)}\}$ denote the parents of light ancestors of x . We remark on two important facts: (1) there are $O(\log n)$ vertices cells in $\text{Apices}[T, x]$, and (2) the LCA of x and y is in $\text{Apices}[T, x] \cup \text{Apices}[T, y]$. These facts are used in existing LCA labeling schemes. We will modify the labeling scheme of Alstrup, Halvorsen, and Larsen:

Lemma 5.4 (Corollary 4.17 of [42]). *Let T be a tree, and let $L : V(T) \rightarrow \{0, 1\}^k$ be a function that indicates some predefined k -bit names for the vertices of T . There is*

a labeling scheme on the vertices of T that uses $O(k \log n)$ bits, such that given labels of any two vertices x and y , we can compute $L(\text{lca}(x, y))$.

We will use a variant of their labeling scheme.

Lemma 5.5. *Let T be a tree. For every vertex x , let $L_x : V(T) \rightarrow \{0, 1\}^k$ be a function that indicates some predefined k -bit names for the vertices of T . There is a labeling scheme on the vertices of T that uses $O(k \log n)$ bits, such that given labels of two leaves x and y , we can compute:*

- $L_x(\text{lca}(x, y))$, if $\text{lca}(x, y) \in \text{Apices}[T, x]$
- $L_y(\text{lca}(x, y))$, if $\text{lca}(x, y) \in \text{Apices}[T, y]$

If $\text{lca}(x, y) \in \text{Apices}[T, x] \cap \text{Apices}[T, y]$, then we can compute both labels $(L_x(\text{lca}(x, y)), L_y(\text{lca}(x, y)))$.

Sketch. We first review the labeling scheme of [42]. For every vertex x , the label of x consists of two parts. The first part (cf. [42, Corollary 4.17]) is just a lookup table: for every vertex $a \in \text{Apices}[T, x]$, we record the k -bit name $L(a)$. The second part encodes information about the root-to- x path in the tree: in particular (cf. [42, Lemma 4.13]), given labels for x and y , we can use the second part of the label to detect whether $\text{lca}(x, y)$ is in $\text{Apices}[T, x]$ — and to look up $L(\text{lca}(x, y))$ in the lookup table, if $\text{lca}(x, y)$ is in fact in $\text{Apices}[T, x]$.

To obtain Lemma 5.5, we simply change the lookup table in the label of x to store $L_x(\cdot)$ instead of $L(\cdot)$. The proof of [42] guarantees that we can return $L_x(\text{lca}(x, y))$ whenever $\text{lca}(x, y) \in \text{Apices}[T, x]$, and symmetrically $L_y(\text{lca}(x, y))$ whenever $\text{lca}(x, y) \in \text{Apices}[T, y]$. $\square$

We are now ready to describe the label to identify a distance-preserving tree.

5.3.3 Labeling scheme

Let $\mathcal{T}$ be the tree cover of Theorem 4.1, of size $O_d(\varepsilon^{-(d-1)} \log(1/\varepsilon))$.

Label.

Let $Q_{i,j}$ be a contracted quadtree used in the construction of $\mathcal{T}$. For each cell C in $Q_{i,j}$, assign an arbitrary ordering to its $1/\varepsilon^{O_d(1)}$ children (so that we can specify a child of C with $O_d(\log 1/\varepsilon)$ bits.) Let x be a vertex in X , and treat x as a leaf of $Q_{i,j}$. For every cell $C \in \text{Apices}(Q_{i,j}, x)$, we define $L_x(C)$ to be a label consisting of three parts:

- (L1) Store $O_d(\log 1/\varepsilon)$ bits to identify which child of C is an ancestor of x .
- (L2) Store $O_d(\log 1/\varepsilon)$ bits to identify which child of C is heavy (if there is a heavy child).
- (L3) Let x' be the representative point for x . Let $\mathcal{T}'$ denote the partial tree cover at cell C , and for each tree $T' \in \mathcal{T}'$, define $\text{ID}(T')$ to be the $O_d(\log 1/\varepsilon)$ -bit identifier of the tree $T \in \mathcal{T}$ of the final tree cover that contains T' . Store the label of x' from Lemma 5.3 (using ID): with this label, for any two points we can either find a tree in $\mathcal{T}$ that preserves the distances of the representative points up to a factor $1 + O(\varepsilon)$, or we certify that the representative points are not (μ, Δ) -far.

Lemma 5.5 gives us an LCA label for x . For each contracted quadtree $Q_{i,j}$, store this label.

Size.

The label $L_x(C)$ consists of $O(\varepsilon^{-(d-1)} \log(1/\varepsilon))$ bits. Thus, Lemma 5.5 gives us labels of size $O(\varepsilon^{-(d-1)} \log(1/\varepsilon) \cdot \log n)$. There are $O_d(\log 1/\varepsilon)$ quadrees $Q_{i,j}$, so the label size is $O_d(\varepsilon^{-1} \log^2(1/\varepsilon) \cdot \log n)$ in total.

Decoding.

Suppose we have labels for x and y . For each $Q_{i,j}$, we use Lemma 5.5 to find information about $C := \text{lca}(x, y)$. There are two cases:

- **Case 1: C is in $\text{Apices}[Q_{i,j}, x] \cap \text{Apices}[Q_{i,j}, y]$.** In this case, we have access to both $L_x(C)$ and $L_y(C)$. We use the (L1) parts of labels $L_x(C)$ and $L_y(C)$ to identify the two children C_x and C_y of C that contain x and y , respectively.
- **Case 2: (Without loss of generality) C is only in $\text{Apices}[Q_{i,j}, x]$.** Let C_y denote the child of C that is an ancestor of y . Because C is not the parent of a light ancestor of y , we know that the child C_y is heavy. Use the (L2) part of $L_x(C)$ to identify the child C_y . As before, use the (L1) part of label $L_x(C)$ to identify the child C_x that is an ancestor of x .

Having identified C_x and C_y , we can now use the (L3) part of label $L_x(C)$ to determine whether there is a tree that preserves the distance between the representatives of C_x and C_y up to a $1 + O(\varepsilon)$ factor. If there is such a tree, return it; otherwise, check the next $Q_{i,j}$.

By the proof of correctness of Theorem 1.1, there is some contracted quadree $Q_{i,j}$ with a cell in which the representatives of x and y are (μ, Δ) -far; further, our Observation 5.1 guarantees that it suffices to check only the LCA of x and y in each contracted quadree. Thus, this process (iterating over all contracted quadrees, and checking the LCA of each) will eventually find a quadree cell in which the representatives are (μ, Δ) -far, and thus (by Lemma 5.3), the (L3) part of the label will return a tree that preserves the distance of the representative points. By the proof of Claim 2.7, this tree preserves the distance between x and y up to a $1 + O(\varepsilon)$ factor.

Acknowledgement.

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).

References

- [1] Awerbuch, B., Peleg, D.: Routing with polynomial communication-space trade-off. SIAM J. Discret. Math. **5**(2), 151–162 (1992) <https://doi.org/10.1137/0405013>

- 1199 [2] Alon, N., Karp, R.M., Peleg, D., West, D.: A graph-theoretic game and its appli-
1200 cation to the k -server problem. *SIAM Journal on Computing* **24**(1), 78–100 (1995)
1201 <https://doi.org/10.1137/s0097539792224474>
- 1202 [3] Arya, S., Das, G., Mount, D.M., Salowe, J.S., Smid, M.: Euclidean spanners: short,
1203 thin, and lanky. In: *Proceedings of the Twenty-seventh Annual ACM Symposium*
1204 *on Theory of Computing*, pp. 489–498 (1995)
- 1205 [4] Gupta, A., Kumar, A., Rastogi, R.: Traveling with a Pez dispenser (or, routi-
1206 ng issues in MPLS). In: *Proceedings 42nd IEEE Symposium on Foundations of*
1207 *Computer Science, FOCS' 01* (2001). <https://doi.org/10.1109/sfcs.2001.959889> .
1208 <https://doi.org/10.1109/sfcs.2001.959889>
- 1209 [5] Mendel, M., Naor, A.: Ramsey partitions and proximity data structures. In: *2006*
1210 *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS' 06)*
1211 (2006). <https://doi.org/10.1109/focs.2006.65>
- 1212 [6] Dragan, F.F., Yan, C., Lomonosov, I.: Collective tree spanners of graphs. *SIAM*
1213 *Journal on Discrete Mathematics* **20**(1), 240–260 (2006)
- 1214 [7] Chan, T.-H.H., Gupta, A., Maggs, B.M., Zhou, S.: On hierarchical routing in
1215 doubling metrics. *ACM Transactions on Algorithms (TALG)* **12**(4), 1–22 (2016)
- 1216 [8] Bartal, Y., Fandina, O.N., Neiman, O.: Covering metric spaces by few trees.
1217 *Journal of Computer and System Sciences* **130**, 26–42 (2022)
- 1218 [9] Chang, H., Conroy, J., Le, H., Milenkovic, L., Solomon, S., Than, C.: Covering
1219 planar metrics (and beyond): $O(1)$ trees suffice. In: *FOCS*, pp. 2231–2261. IEEE,
1220 USA (2023)
- 1221 [10] Chang, H., Conroy, J., Le, H., Milenkovic, L., Solomon, S., Than, C.: Short-
1222 cut partitions in minor-free graphs: Steiner point removal, distance oracles, tree
1223 covers, and more. In: *SODA*, pp. 5300–5331. SIAM, USA (2024)
- 1224 [11] Thorup, M., Zwick, U.: Approximate distance oracles. *Journal of the ACM*
1225 (JACM) **52**(1), 1–24 (2005)
- 1226 [12] Kahalon, O., Le, H., Milenković, L., Solomon, S.: Can’t see the forest for the trees:
1227 navigating metric spaces by bounded hop-diameter spanners. In: *Proceedings of*
1228 *the 2022 ACM Symposium on Principles of Distributed Computing*, pp. 151–162
1229 (2022)
- 1230 [13] Agarwal, P.K., Chang, H.-C., Raghvendra, S., Xiao, A.: Deterministic, near-linear
1231 ε -approximation algorithm for geometric bipartite matching. In: *Proceedings of*
1232 *the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1052–
1233 1065 (2022)

- 1234 [14] Narasimhan, G., Smid, M.: Geometric Spanner Networks. Cambridge University
1235 Press, USA (2007)
- 1236 [15] Le, H., Solomon, S.: Truly optimal Euclidean spanners. *SIAM Journal on*
1237 *Computing* (0), 19–135 (2022)
- 1238 [16] Bhore, S., Tóth, C.D.: Euclidean steiner spanners: Light and sparse. *SIAM*
1239 *Journal on Discrete Mathematics* **36**(3), 2411–2444 (2022)
- 1240 [17] Awerbuch, B., Kutten, S., Peleg, D.: On buffer-economical store-and-forward
1241 deadlock prevention. *IEEE transactions on communications* **42**(11), 2934–2937
1242 (1994)
- 1243 [18] Charikar, M., Chekuri, C., Goel, A., Guha, S., Plotkin, S.: Approximating a finite
1244 metric by a small number of tree metrics. In: *Proceedings 39th Annual Symposium*
1245 *on Foundations of Computer Science* (Cat. No. 98CB36280), pp. 379–388 (1998).
1246 IEEE
- 1247 [19] Thorup, M., Zwick, U.: Compact routing schemes. In: *Proceedings of the Thir-*
1248 *teenth Annual ACM Symposium on Parallel Algorithms and Architectures*, pp.
1249 1–10 (2001)
- 1250 [20] Margulis, G.A.: Explicit group-theoretical constructions of combinatorial schemes
1251 and their application to the design of expanders and concentrators. *Problemy*
1252 *Peredachi informatsii* **24**(1), 51–60 (1988)
- 1253 [21] Althöfer, I., Das, G., Dobkin, D., Joseph, D., Soares, J.: On sparse spanners of
1254 weighted graphs. *Discrete & Computational Geometry* **9**(1), 81–100 (1993)
- 1255 [22] Abraham, I., Bartal, Y., Neiman, O.: Advances in metric embedding theory.
1256 *Advances in Mathematics* **228**(6), 3026–3126 (2011)
- 1257 [23] Krauthgamer, R., Lee, J.R., Mendel, M., Naor, A.: Measured descent: a new
1258 embedding method for finite metrics. *Geometric and Functional Analysis* **15**(4),
1259 839–858 (2005) <https://doi.org/10.1007/s00039-005-0527-6>
- 1260 [24] Abraham, I., Gavoille, C.: Object location using path separators. In: *Proceedings*
1261 *of the Twenty-fifth Annual ACM Symposium on Principles of Distributed Com-*
1262 *puting. PODC '06*, pp. 188–197 (2006). <https://doi.org/10.1145/1146381.1146411>
1263 . Full version: <https://www.cse.huji.ac.il/~ittaia/papers/AG-TR.pdf>
- 1264 [25] Freedman, O., Gawrychowski, P., Nicholson, P.K., Weimann, O.: Optimal dis-
1265 tance labeling schemes for trees. In: *Proceedings of the ACM Symposium on*
1266 *Principles of Distributed Computing* (2017). [https://doi.org/10.1145/3087801.](https://doi.org/10.1145/3087801.3087804)
1267 [3087804](https://doi.org/10.1145/3087801.3087804) . <https://doi.org/10.1145/3087801.3087804>
- 1268 [26] Gavoille, C., Peleg, D., Pérennes, S., Raz, R.: Distance labeling in graphs. *Journal*

- of Algorithms **53**(1), 85–112 (2004) <https://doi.org/10.1016/j.jalgor.2004.05.002>
- [27] Andoni, A., Zhang, H.: Sub-quadratic $(1+\epsilon)$ -approximate Euclidean Spanners, with Applications (2023)
- [28] Smid, M.: Notes on Binary Dumbbell Trees. Unpublished notes (2012)
- [29] Gottlieb, L., Roditty, L.: Improved algorithms for fully dynamic geometric spanners and geometric routing. In: Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms. SODA’08, pp. 591–600 (2008). <https://doi.org/10.5555/1347082.1347148>
- [30] Gupta, A., Krauthgamer, R., Lee, J.R.: Bounded geometries, fractals, and low-distortion embeddings. In: FOCS, pp. 534–543. IEEE Computer Society, USA (2003)
- [31] Chan, T.M.: Approximate nearest neighbor queries revisited. Discret. Comput. Geom. **20**(3), 359–373 (1998)
- [32] Gao, Z., Har-Peled, S.: Almost optimal locality sensitive orderings in Euclidean space. CoRR **abs/2310.12792** (2023)
- [33] Fortune, S., Van Wyk, C.J.: Efficient exact arithmetic for computational geometry. In: Proceedings of the Ninth Annual Symposium on Computational Geometry. SCG ’93, pp. 163–172. Association for Computing Machinery, New York, NY, USA (1993). <https://doi.org/10.1145/160985.161015> . <https://doi.org/10.1145/160985.161015>
- [34] Li, C., Pion, S., Yap, C.-K.: Recent progress in exact geometric computation. The Journal of Logic and Algebraic Programming **64**(1), 85–111 (2005)
- [35] Salesin, D., Stolfi, J., Guibas, L.: Epsilon geometry: building robust algorithms from imprecise computations. In: Proceedings of the Fifth Annual Symposium on Computational Geometry, pp. 208–217 (1989)
- [36] Erickson, J., Van Der Hoog, I., Miltzow, T.: Smoothing the gap between np and er. SIAM Journal on Computing (0), 20–102 (2022)
- [37] Peleg, D.: Proximity-preserving labeling schemes. J. Graph Theory **33**(3), 167–176 (2000) <https://doi.org/10.5555/1379811.1379818>
- [38] Thorup, M., Zwick, U.: Compact routing schemes. In: SPAA, pp. 1–10. ACM, USA (2001)
- [39] Fraigniaud, P., Gavoille, C.: Routing in trees. In: ICALP. Lecture Notes in Computer Science, vol. 2076, pp. 757–772. Springer, USA (2001)
- [40] Abraham, I., Gavoille, C., Goldberg, A.V., Malkhi, D.: Routing in networks with

- low doubling dimension. In: Proc. of 26th ICDCS, p. 75 (2006)
- [41] Santoro, N., Khatib, R.: Labelling and implicit routing in networks. *Comput. J.* **28**(1), 5–8 (1985)
- [42] Alstrup, S., Halvorsen, E.B., Larsen, K.G.: Near-optimal labeling schemes for nearest common ancestors. In: Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 972–982 (2014). SIAM