

Lower Bounds on Tree Covers

Yu Chen ✉ 

National University of Singapore, Queenstown, Singapore

Zihan Tan ✉ 

University of Minnesota Twin Cities, MN, USA

Hangyu Xu ✉ 

University of Science and Technology of China, Hefei, China

Abstract

Given an n -point metric space (X, d_X) , a tree cover $\mathcal{T}$ is a set of $|\mathcal{T}| = k$ trees on X such that every pair of vertices in X has a low-distortion path in one of the trees in $\mathcal{T}$. Tree covers have been playing a crucial role in graph algorithms for decades, and the research focus is the construction of tree covers with small size k and distortion.

When $k = 1$, the best distortion is known to be $\Theta(n)$. For a constant $k \geq 2$, the best distortion upper bound is $\tilde{O}(n^{\frac{1}{k}})$ and the strongest lower bound is $\Omega(\log_k n)$, leaving a gap to be closed. In this paper, we improve the lower bound to $\Omega(n^{\frac{1}{2^{k-1}}})$.

Our proof is a novel analysis on a structurally simple grid-like graph, which utilizes some combinatorial fixed-point theorems. We believe that they will prove useful for analyzing other tree-like data structures as well.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Data structures design and analysis; Mathematics of computing $\rightarrow$ Combinatorics

Keywords and phrases Tree Covers, Combinatorial Fixed-Point Theorems

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.38

Related Version Full Version: <https://arxiv.org/abs/2508.10376>

Funding Hangyu Xu: Supported in part by NSFC Grant 62272431 and Quantum Science and Technology-National Science and Technology Major Project (Grant No. 2021ZD0302901).

1 Introduction

Given a metric space (X, d_X) on $|X| = n$ points, an (α, k) -tree cover is a collection $\mathcal{T} = \{T_1, \dots, T_k\}$ of k (edge-weighted trees) on X , such that every pair x, y of vertices in X ,
(1) for each $1 \leq i \leq k$, $\text{dist}_{T_i}(x, y) \geq d_X(x, y)$ (that is, each T_i is a *dominating tree*); and
(2) there exists an $1 \leq i \leq k$, such that $\text{dist}_{T_i}(x, y) \leq \alpha \cdot d_X(x, y)$.

The integer $k = |\mathcal{T}|$ is called the *size* of $\mathcal{T}$ and the real number $\alpha \geq 1$ is called the *distortion* or *stretch* of $\mathcal{T}$. The goal is to construct tree covers with small size and distortion.

As a natural and instrumental approach to approximating metrics by trees, tree cover has been studied extensively and has found numerous algorithmic applications, for example, approximating low-dimensional Euclidean spaces [1, 11], routing [3, 17, 2, 8], network design [16], distance oracles [20], Ramsey tree covers [5, 4], Steiner Point Removal [9, 10, 12], and many others.

Despite much attention and effort, several fundamental questions regarding the size-distortion trade-off of tree covers remain unsolved, especially in the regime where k is as small as a constant. When $k = 1$, the best distortion is known to be $\Theta(n)$: the upper bound is achieved by the minimum spanning tree, and the lower bound is proved by the unweighted n -cycle [23]. When k is a small constant, the best upper bound, due to Bartal, Fandina, and Neiman [4], states that the stretch $\tilde{O}(n^{1/k})$ is achievable (actually by a more restricted type



© Yu Chen, Zihan Tan, and Hangyu Xu;

licensed under Creative Commons License CC-BY 4.0

17th Innovations in Theoretical Computer Science Conference (ITCS 2026).

Editor: Shubhangi Saraf; Article No. 38; pp. 38:1–38:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of tree covers, more discussion in Section 1.1). On the other hand, the best lower bound remains $\Omega(\log_k n)$, witnessed by high-girth graphs [7], leaving a big gap. Closing this gap is viewed as one of the major interesting open problems on tree covers:

- (1) “Is there a better lower bound? Or perhaps there exists a tree cover with $O(1)$ trees and distortion $O(\log n)$? What can one say about a tree cover with only 2 trees?” – Bartal, Fandina, and Neiman in [4].
- (2) “A very interesting open question: what is the smallest stretch achieved with 2 trees?” – Le¹.

1.1 Our result

In this paper, we make progress towards closing this gap. Our main result is the following theorem.

► **Theorem 1.** *For every $k \geq 1$, there is an n -point metric, such that any size- k tree cover has stretch $\Omega_k(n^{\frac{1}{2^k-1}})$.*

Specifically, when $k = 2$, our result implies a lower bound of $\Omega(\sqrt{n})$, nearly matching the upper bound of $\tilde{O}(\sqrt{n})$ in [4]. For constant k , we show a polynomial lower bound, refuting the possibility of poly log n stretch. Therefore, Theorem 1 provides an answer to the above questions by Bartal, Fandina, Neiman, and Le.

For proving Theorem 1, we employ a novel approach. At a high-level, we construct a grid-like structure, compute its triangulation, label each point according to its location in the tree cover, and derive the stretch lower bound by Tucker’s Lemma. Our proof reveals a close connection between tree covers and the celebrated fixed-point theorems in combinatorial topology, and we believe that this connection will be useful in analyzing tree covers in different settings and other tree-like data structures.

The stretch upper bound of $\tilde{O}(n^{1/k})$ in [4] is actually achieved by Ramsey tree covers, a restricted type of tree covers that require each vertex to choose a home tree to preserve its distances to all other vertices (formal definition in Section 5). They also showed a stretch lower bound of $\Omega(n^{1/k})$ for Ramsey tree covers. As a side result, we give an alternative proof of this near-optimal lower bound. Our hard instance is a high-dimensional cyclic grid, which is structurally simpler than the recursive construction in [4]. Our proof is similar to the proof of Theorem 1 and also makes use of a variant of Tucker’s Lemma.

2 Technical Overview

In this section, we provide a high-level overview of our techniques, with the focus on the connection between tree covers and combinatorial topology.

Case $k = 1$: examining a tree-based tour on a cycle


We start by reviewing the proof of the $\Omega(n)$ distortion lower bound for the case $k = 1$. We present a proof sketch similar to the one in [23], and an alternative proof can be found in [15].

Consider an unweighted cycle $(0, 1, 2, \dots, n-1, 0)$ and a dominating tree T on the same set of vertices. We may assume without loss of generality that for each edge (i, j) in T (with $i < j$), its length/weight is the cycle-distance between vertices i and j , namely

¹ at Sublinear Graph Simplification workshop at Simons Institute, video recording can be accessed at <https://simons.berkeley.edu/talks/hung-le-university-massachusetts-amherst-2024-08-01>

$\text{dist}_T(i, j) = \min \{j - i, n - (j - i)\}$, as otherwise we can decrease the length of this edge without hurting the distortion of T . To reveal the topological difference between a tree and a cycle, we “interpolate” their edges as follows.

1. For each edge (i, j) in T , we subdivide it by $\text{dist}_T(i, j) - 1$ new nodes (called *interpolating nodes*), which are then labeled by consecutive integers connecting i, j in the natural order. For example, edge $(3, 6) \rightarrow (3, 4, 5, 6)$, and edge $(1, n - 3) \rightarrow (1, 0, n - 1, n - 2, n - 3)$. Edges $(3, 4), (4, 5), (5, 6)$ are called *interpolating edges*. See Figure 1 for an illustration.
2. For each i , we define Π_i as the sequence of all nodes (including the old ones and the interpolating ones) we encounter while walking along the path from the old i to the old $i + 1$ in T . Observe that the tree-distance between old nodes $i, i + 1$ is $|\Pi_i|$, the number of interpolating edges in Π_i . Here Π_{n-1} connects the old node $n - 1$ to the old node 0 in the tree.
3. Finally, we define Π as the circular concatenation of sequences $\Pi_0, \Pi_1, \dots, \Pi_{n-1}$, identifying, for each i , the last node of Π_i and the first node of Π_{i+1} (both being the old $i + 1$), so Π is a closed circular tour.



■ **Figure 1** Interpolated tree, sequences Π_i 's, and a bad Π_5 .

It is immediate to verify that all Π_i and Π are continuous, in that any integer j is sided by either $j + 1$ or $j - 1$. The key to the lower bound proof is to *examine the continuous behaviors of Π and Π_i 's, which are extracted from the tree, on the cycle*.

The shortest way to go continuously from i to $i + 1$ on the cycle is to directly use the edge $(i, i + 1)$. We say that Π_i is *good*², if the number of $(i, i + 1)$ appearances in Π_i exceeds that of $(i + 1, i)$ by 1. For example, $\Pi_2 = (2, 3, 4, 5, 4, 3)$ is good, but $(2, 1, 0, n - 1, \dots, 4, 3)$ is not. Intuitively, Π_i is good if it “essentially travels on the short side of cycle”. See Figure 1 for an illustration. The following two properties are easy to prove:

² There is a notion in combinatorial topology called “degree of a mapping”, and our notion “good” is similar to their “degree-0” case. For the purpose of technical overview, we do not go deep in formally presenting the notions.

- (1) if Π_i is good, the number of $(j, j+1)$ appearances equals that of $(j+1, j)$ for each $j \neq i$; and
- (2) if Π_i is bad, it contains all n numbers and therefore has length at least $n-1$.

The crucial observation is that, if some Π_i is bad, then the distortion of T is at least $n-1$, as the cycle-distance between $i, i+1$ is 1 while the tree-distance between $i, i+1$ is $|\Pi_i| \geq n-1$. Therefore, it remains to show that not all Π_i 's are good. Assume for contradiction that they all are. Consider now the closed tour Π . From the first property of good Π_i 's, we know that for each i , the $(i, i+1)$ appearances exceed $(i+1, i)$ by 1. However, as Π is a closed tour that goes continuously around the tree, every time it travels from i to $i+1$ along an interpolating edge $(i, i+1)$, later it has to travel back along the same edge from $i+1$ to i , so the $(i, i+1)$ appearances should equal that of $(i+1, i)$, leading to a contradiction.


To summarize, we interpolate the tree edges, mapping the cycle into a continuous tree-tour Π which consists of sub-trajectories $\Pi_0, \dots, \Pi_{n-1}$, and examine their behaviors from the topological perspective. The key point is that a cycle must contain an edge $(i, i+1)$ whose corresponding trajectory Π_i is bad, and such an edge, which we call a *bad edge*, is the cause of $\Omega(n)$ distortion.

Case $k = 2$: the torus grid

The best distortion upper bound for the $k = 2$ case is $\tilde{O}(\sqrt{n})$ [4]. Therefore, a natural first graph to try proving a matching lower bound on is the $\sqrt{n} \times \sqrt{n}$ grid. However, with a delicate design, two trees can actually cover the grid with distortion $O(1)$. We present a sketch of the construction in Appendix A. Two concurrent works [18, 6] independent to ours give a similar construction with a much detailed presentation.

According to the lesson from the $k = 1$ case, to make it harder for trees to cover, we should plant more cycles on the grid. Specifically, we connect each first-row vertex to their corresponding last-row vertex, and connect each first-column vertex to their corresponding last-column vertex, so the resulting graph is a torus grid. As we show next, this torus grid poses a $\Omega(\sqrt{n})$ distortion lower bound for two trees.

We focus on the following two types of non-contractible cycles: vertical/horizontal cycles, which are cycles that go around the grid in the vertical/horizontal way, respectively. See Figure 2 for an illustration. Both types of cycles are of length at least $\sqrt{n}$.



■ **Figure 2** The torus grid, horizontal and vertical cycles, and T_1 and T_2 -bad edges.

We denote the trees by T_1, T_2 . To prove the $\Omega(\sqrt{n})$ distortion lower bound, it suffices to find a pair of vertices that are $O(1)$ -close in the torus grid but $\Omega(\sqrt{n})$ -far in both trees. The focus is still finding bad edges, similar to the $k = 1$ case. At a high level, the lower bound proof can be summarized as follows.

1. Intuitively, every vertical cycle contains a bad edge for T_1 . Therefore, in order to hit all vertical cycles, the union of all such edges must form a “horizontal cut” (see Figure 2).
2. Similarly, every horizontal cycle contains a bad edge for T_2 . Therefore, in order to hit all horizontal cycles, the union of all such edges must form a “vertical cut”.
3. The horizontal cut has the shape of a horizontal cycle, and so it must intersect the vertical cut. At their intersection we can find the desired pair of vertices. Specifically, as illustrated in Figure 2, a T_1 -bad edge (v, u_1) and a T_2 -bad edge (v, u_2) share a common endpoint v , so $\text{dist}_{T_1}(v, u_1) = \Omega(\sqrt{n})$ and $\text{dist}_{T_2}(v, u_2) = \Omega(\sqrt{n})$. Then
 - (1) if $\text{dist}_{T_2}(v, u_1) = \Omega(\sqrt{n})$, then v, u_1 are $\Omega(\sqrt{n})$ -far in both trees;
 - (2) if $\text{dist}_{T_1}(v, u_2) = \Omega(\sqrt{n})$, then v, u_2 are $\Omega(\sqrt{n})$ -far in both trees;
 - (3) if $\text{dist}_{T_2}(v, u_1), \text{dist}_{T_1}(v, u_2) = o(\sqrt{n})$, then by triangle inequality,

$$\text{dist}_{T_1}(u_1, u_2) \geq \text{dist}_{T_1}(v, u_1) - \text{dist}_{T_1}(v, u_2) = \Omega(\sqrt{n}),$$

and similarly $\text{dist}_{T_2}(u_1, u_2) = \Omega(\sqrt{n})$, and so u_1, u_2 are $\Omega(\sqrt{n})$ -far in both trees.

In fact, instead of the torus grid, our proof of Theorem 1 (presented in Section 4) actually works with the grid structure on the 2-dimensional projective plane: that is, the graph obtained from the $\sqrt{n} \times \sqrt{n}$ grid by connecting (i) for each i , the i -th vertex of the first row to the $(\sqrt{n} - i)$ -th vertex of the last row (rather than the i -th vertex of the last row as in the grid); and (ii) for each i , the i -th vertex of the first column to the $(\sqrt{n} - i)$ -th vertex of the last column (rather than the i -th vertex of the last column as in the grid). The horizontal/vertical cycles and cuts can be defined analogously and analyzed similarly.

General k : the role of Tucker’s Lemma

Based on our approach in the $k = 2$ case, we can set up a similar plan for proving lower bounds for general k : construct a k -dimensional grid with length $n^{1/k}$, make it cyclic by connecting i -dimensional antipodal points for each dimension i , find the “dimension- i cut” consisting of T_i -bad edges, which is essentially a $(k - 1)$ -dimensional plane, and finally analyze the intersection of all k planes.

The issue resides in the very last step (Step 3 in the $k = 2$ case): we are only guaranteed that, around the intersection, for each i , some pair of points cross the dimension- i cut, but that does not guarantee the existence of a pair of points simultaneously crossing all dimension- i cuts. Specifically, we consider the k -dimensional cube that represents the intersection of all dimension-cut planes. We associate with each vertex a vector in $\{+1, -1\}^k$ as follows. For each vertex, for each i , if it lies on the left side (left and right sides can be dictated in an arbitrary way) of the dimension- i cut, then we set the i -th coordinate of its associated vector v as $v_i = +1$, otherwise we set $v_i = -1$. Then, among all vectors in the cube,

1. what we can prove (via the results in [22]): for each i , there exists a pair of vertices whose associated vectors u, v satisfy that $u_i = -v_i$ (such a cube is called a *centrally labeled* cube); and
2. the goal is to show: there is a pair of vertices whose associated vectors u, v satisfy that $u_i = -v_i$ for all i (such a pair of vertices are said to form a *complementary edge*).

In fact, what we can prove does not necessarily imply the goal for $k \geq 3$. For example, the 8 vectors in a 3-dimensional cube can be

$$\{(1, 1, 1) \times 5, (1, -1, -1), (-1, -1, 1), (-1, 1, -1)\}.$$

So the cube is indeed a centrally labeled cube, but there are no complementary edges.

Luckily, the existence of a centrally labeled cube is enough to derive a $\Omega(n^{1/k})$ lower bound for Ramsey tree covers with k trees. We present the details of this proof in Section 5.

But to prove lower bound for general tree covers, we have to find complementary edges, and this is where we need Tucker's lemma, which states that, in a triangulated domain of $\mathbb{R}^d$, in a labeling of its vertices by integers in $\{-d, \dots, -1, 1, \dots, d\}$ that is antipodally symmetric on the domain's boundary, there is a complementary edge. Since the number of possible labels in our setting is 2^k (the size of set $\{+1, -1\}^k$), we have to significantly increase the number of dimensions in our construction from k to $d = 2^{k-1}$, which reduces the length of our grid-like structure from $n^{\frac{1}{k}}$ to $n^{\frac{1}{2^{k-1}}}$. Therefore, we are only able to prove an $\Omega(n^{\frac{1}{2^{k-1}}})$ stretch lower bound for general tree covers with this approach. An additional modification to our grid-like construction is needed in order to achieve the antipodally symmetric property required in Tucker's lemma: for the i -th dimension, instead of connecting pairs of vertices whose connection is parallel to the dimension- i axis, we need to connect pairs of vertices that are antipodally symmetric. In light of this, instead of working with a cube-shape grid, we will work with an octahedron-shape grid, as it simplifies the analysis. We present the details in Section 4.

3 Preliminaries: Simplices, Triangulations, and Tucker's Lemma

A set $\sigma \subseteq \mathbb{R}^d$ is a k -dimensional simplex (or k -simplex) if it is the convex hull $\text{conv}(A)$ of a set A of $k+1$ affinely independent points. Vertices of σ , denoted as $V(\sigma)$, refer to points in A . Simplices of the form $\text{conv}(B)$ where $B \subseteq A$ are called *faces* of σ .

Our proof is based on some combinatorial fixed-point theorems. To formally state them, we need the notion of triangulation and its antipodal symmetry.

► **Definition 2** (Triangulation). *Let X be a subset of $\mathbb{R}^d$ that is homeomorphic to the ℓ_2 -norm unit ball. A finite family Γ of simplices is a triangulation of X , iff $\bigcup_{\sigma \in \Gamma} \sigma = X$, and the following two properties hold:*

- (1) *For each simplex $\sigma \in \Gamma$, all faces of σ are also simplices in Γ .*
 - (2) *For any two simplices $\sigma_1, \sigma_2 \in \Gamma$, the intersection $\sigma_1 \cap \sigma_2$ is a face of both σ_1 and σ_2 .*
- We define the vertices of Γ as $V(\Gamma) = \bigcup_{\sigma \in \Gamma} V(\sigma)$. We define the boundary of Γ to be the triangulation $\partial(\Gamma) \subseteq \Gamma$ containing all simplices in Γ that lie entirely in the boundary $\partial(X)$ of X .*

► **Definition 3** (Antipodal Symmetry). *A triangulation Γ is said to be antipodally symmetric on the boundary if, for all simplices $\sigma \in \partial(\Gamma)$, the opposite simplex $-\sigma$ ($-\sigma = \{-v \mid v \in \sigma\}$) also lies in $\partial(\Gamma)$.*

► **Definition 4** (Labeling of Triangulation). *A labeling of triangulation Γ is an assignment $\ell : V(\Gamma) \rightarrow L$ that assigns an integer of a set $L \subseteq \mathbb{Z}$ to each vertex of Γ .*

With the above definitions, we can now state Tucker's lemma (originally in [24]; see [21] for a survey) and its variant.

► **Theorem 5** (Tucker’s Lemma). *Let Γ be a triangulation of the domain $X \subset \mathbb{R}^d$ that is antipodally symmetric on the boundary. Let $\ell : V(\Gamma) \rightarrow \{-d, \dots, -2, -1, 1, 2, \dots, d\}$ be a labeling such that $\ell(v) = -\ell(-v)$ for all $v \in \partial(X)$. Then Γ contains a complementary edge, that is, an edge (v, v') with $\ell(v) = -\ell(v')$.*

We can further use 2^d labels, instead of $2d$ labels in Tucker’s lemma, to obtain the following extension of Tucker’s lemma (Theorem 1.9 in [14]):

► **Theorem 6** ([14]). *Let Γ be a triangulation of $X \subset \mathbb{R}^d$ that is antipodally symmetric on the boundary of the domain X . Let $\ell : V(\Gamma) \rightarrow \{+1, -1\}^k$ be a labeling such that $\ell(v) = -\ell(-v)$ for all $v \in \partial(X)$. Then Γ contains a neutral simplex, that is, a simplex σ such that for all $1 \leq i \leq n$, there exist $v, v' \in V(\sigma)$ such that $\ell_i(v) = -\ell_i(v')$, where $\ell_i(v)$ is the i^{th} coordinate of $\ell(v)$.*

Theorem 6 can be proved by either applying Borsuk-Ulam Theorem [19] (a topology approach) or by extending the framework of Ky Fan [13] to more possible labelings [14] (a combinatorial approach). We refer the interested readers to [14] for a survey of relevant approaches.

4 Proof of Theorem 1

In this section, we provide the proof of Theorem 1. We will construct a metric (X, d_X) , and show that for any k dominating trees $T_1, \dots, T_k$ on X , there is a pair of vertices u, v in X , such that

$$\min_{1 \leq i \leq k} \{\text{dist}_{T_i}(u, v)\} = \Omega\left(\frac{n^{\frac{1}{2^{k-1}}}}{2^{k-1}}\right) \cdot d_X(u, v).$$

4.1 Construction of the Hard Instance

We first define an edge-weighted graph $\tilde{G}$ as follows.


- (1) Vertex set: $V(\tilde{G}) = \{x \in \mathbb{Z}^d : \|x\|_1 \leq r\}$;
- (2) Edges set: $E(\tilde{G}) = E_0 \cup E_{\text{sp}}$, where edges in E_{sp} are called *special edges*, and
 - (i) $E_0 = \{(u, v) : \|u - v\|_1 = 1\}$,
 - (ii) $E_{\text{sp}} = \{(u, v) : u = -v, \|u\|_1 = r\}$;
- (3) Edge weights: $w_e = \begin{cases} 1, & \text{if } e \in E_0; \\ 0, & \text{if } e \in E_{\text{sp}}. \end{cases}$

Our hard instance G is the graph obtained by contracting all the edges in E_{sp} within $\tilde{G}$. Specifically, we denote the two vertices contracted by the one with positive value at the first non-zero coordinate. By choosing $d = 2^{k-1}$ and $r = \Theta(n^{1/d})$, we can ensure that $|V(G)| = O(n)$. The metric (X, d_X) is simply the shortest path distance metric on all vertices of G .

4.2 Reduction to Tucker’s Lemma

► **Lemma 7.** *There is a triangulation Γ' of $\{x \in \mathbb{R}^d : \|x\|_1 \leq r\}$, such that $V(\Gamma') = V(\tilde{G})$, and for every edge (u, v) in Γ' , $\|u - v\|_\infty \leq 1$.*

Proof. Let $\{e_1, \dots, e_d\}$ be the standard basis vectors in $\mathbb{R}^d$. Given a simplex σ , we denote by $\mathcal{C}(\sigma)$ the set of all faces of σ , i.e., $\mathcal{C}(\sigma) = \bigcup_{A \subseteq V(\sigma)} \text{conv}(A)$. For each vertex $v = \sum_{i=1}^d \alpha_i e_i \in \{x \in \mathbb{R}^d : \|x\|_1 \leq 1\}$ where $\alpha_i \neq 0$ for all i , we define $\sigma(v) = \text{conv}(\{\text{sign}(\alpha_i) \cdot e_i : i = 1, \dots, d\})$ (see Figure 3). Specially, we define $\sigma(v) = \emptyset$ if $\alpha_i = 0$ for some i . We further define $\Gamma = \bigcup_{\|v\| \leq 1} \mathcal{C}(\sigma(v))$, as illustrated in Figure 3.



■ **Figure 3** The ℓ_1 norm ball grid and its decomposition, and the triangulations of ℓ_1 -norm unit balls in 2 and 3 dimensions.

For every vertex $v \in \{x \in \mathbb{R}^d : \|x\|_1 \leq 1\}$, if $\alpha_i \neq 0$ for all i , by $v \in \sigma(v)$, we have $v \in \bigcup_{\sigma \in \Gamma} \sigma$. Otherwise, $v \in \text{conv}(\{\text{sign}(\alpha_i) \cdot e_i : i = 1, \dots, d\}) \in \Gamma$, where $\text{sign}(x) = \begin{cases} 1, & \text{if } x \geq 0; \\ -1, & \text{if } x < 0. \end{cases}$. Thus $v \in \bigcup_{\sigma \in \Gamma} \sigma$ for every $v \in \{x \in \mathbb{R}^d : \|x\|_1 \leq 1\}$. It follows that $\bigcup_{\sigma \in \Gamma} \sigma = \{x \in \mathbb{R}^d : \|x\|_1 \leq 1\}$. By the definition of Γ , for each simplex $\sigma \in \Gamma$, all the faces of σ are also simplices in Γ . Furthermore, for every two simplices $\sigma_1, \sigma_2 \in \Gamma$, $\sigma_1 \cap \sigma_2 = \text{conv}(V(\sigma_1) \cap V(\sigma_2)) \in \mathcal{C}(\sigma_1) \cap \mathcal{C}(\sigma_2)$. This concludes that Γ is a triangulation of $\{x \in \mathbb{R}^d : \|x\|_1 \leq 1\}$.

$\{x \in \mathbb{R}^d : \|x\|_1 \leq r\}$ can be naturally partitioned into ℓ_1 -norm unit balls, each of which can be triangulated in the same manner as Γ , as illustrated in Figure 3. It gives a triangulation of $\{x \in \mathbb{R}^d : \|x\|_1 \leq r\}$, denoted by Γ' . Note that $V(\Gamma') = V(\tilde{G})$, and for every edge (u, v) in Γ' , $\|u - v\|_\infty \leq 1$. ◀

Given a dominating tree T of G . As in Section 2, we may assume without loss of generality that for each edge (u, v) in T , its length is the distance between u and v in G , namely $\text{dist}_T(u, v) = \text{dist}_G(u, v)$. For each edge (u, v) in T , we subdivide it by $\text{dist}_T(u, v) - 1$ *interpolating nodes*, which are then labeled by consecutive vertices in the shortest path between u, v in G (arbitrarily chosen if there are more than one shortest path). Edges connecting the neighboring interpolating nodes are called *interpolating edges*.

We define $\Pi(u, v)$ as the sequence of all the above nodes from u to v . Note that an interpolating edge may not be an edge in $\tilde{G}$. For each such edge (x, y) , observe that either $\|x\|_1 = r$ or $\|y\|_1 = r$. We may assume without loss of generality that $\|x\|_1 = r$. One can observe that $(-x, y)$ is included in the edge set of $\tilde{G}$. Thus we subdivide (x, y) by $(x, -x)$ and $(-x, y)$. By the above subdivision process, we obtain a new sequence consisting of edges in $\tilde{G}$, and we denote it by $\tilde{\Pi}(u, v)$.

For vertices $v, v' \in G$, let $(v_0(=v), v_1, \dots, v_\ell(=v'))$ be the v - v' path in T . We define $P_{\tilde{G}}^T(v, v')$ as the sequential concatenation of sequences $\tilde{\Pi}(v_0, v_1), \tilde{\Pi}(v_1, v_2), \dots, \tilde{\Pi}(v_{\ell-1}, v_\ell)$, identifying, for each i , the last node of $\tilde{\Pi}(v_{i-1}, v_i)$ and the first node of $\tilde{\Pi}(v_i, v_{i+1})$.

Now for a vertex pair $v, v' \in \tilde{G}$, if v, v' are also in G , the definition of $P_{\tilde{G}}^T(v, v')$ is the same as before. Otherwise, we may assume v' is in $\tilde{G} \setminus G$, namely, $\|v'\| = r$ with negative value at the first non-zero coordinate. If $P_{\tilde{G}}^T(v, -v')$ is ended by $(v', -v')$, then we truncate it and obtain a sequence ended by v' . Otherwise, we concatenate it by $(-v', v')$. The new sequence is defined as $P_{\tilde{G}}^T(v, v')$.

Denote the length of a path P as $\ell(P)$. Observe that the tree-distance between vertices v, v' is $\ell(P_G^T(v, v'))$ for each vertex pair v, v' in G .

► **Definition 8 (Tree Index).** Given a tree T and a vertex $v \in V(\tilde{G})$. Denote by c the vertex corresponding to the vector $\vec{0}$. We define an index $i_T(v)$ of v for T as:

$$i_T(v) = \begin{cases} 1, & P_G^T(c, v) \text{ contains an odd number of special edges;} \\ -1, & P_G^T(c, v) \text{ contains an even number of special edges.} \end{cases}$$


For the trees $T_1, \dots, T_k$ of the tree cover, we define tree index of v as $i(v) = (i_{T_1}(v), \dots, i_{T_k}(v))$.

By definition, $i(v) = -i(-v)$ for all $\|v\|_1 = r$. Note that tree indices on $V(\tilde{G})$ is a labeling $V(\Gamma') \rightarrow \{+1, -1\}^k$ of the triangulation Γ' . We now show that neighboring vertices with opposite indices have large tree distances.

► **Lemma 9.** For a tree T and vertices u, v with $\|u - v\|_\infty \leq 1$, if $i_T(u) = -i_T(v)$, then $\text{dist}_T(u, v) \geq 2r - d$.

The proof of Lemma 9 requires the following claim.

► **Claim 10.** Given a vertex pair $u, v \in V(\tilde{G})$ and a path P between u and v in $\tilde{G}$. If $|E(P) \cap E_{\text{sp}}|$ is an even number, then $\ell(P) \geq \|u - v\|_1$.



■ **Figure 4** Construction of $\tilde{P}$ by alternately flipping P . The two paths have the same length.

Proof. Without loss of generality, we assume

$$P = (v_{00}(=u), v_{01}, \dots, v_{0s_0}, v_{10}, \dots, v_{1s_1}, v_{20}, \dots, v_{(t-1)s_{t-1}}, v_{t0}, \dots, v_{ts_t}(=v)),$$

where for each $j = 0, \dots, t-1$, $v_{js_j} = -v_{(j+1)0}$ and $\|v_{js_j}\|_1 = r$. Then we construct

$$\tilde{P} = (v_{00}, v_{01}, \dots, v_{0s_0}, -v_{10}, \dots, -v_{1s_1}, v_{20}, \dots, (-1)^{t-1}v_{(t-1)s_{t-1}}, (-1)^t v_{t0}, \dots, (-1)^t v_{ts_t}).$$

It can be observed that $\tilde{P}$ is a path in $\tilde{G} \setminus E_{\text{sp}}$, namely containing no edges in E_{sp} . Since $t = |P \cap E_{\text{sp}}|$, $2 \mid t$, and so $(-1)^t v_{ts_t} = v_{ts_t} = v$. Therefore, $\ell(P) = \ell(\tilde{P}) \geq \|u - (-1)^t v_{ts_t}\|_1 = \|u - v\|_1$. ◀

Proof of Lemma 9. Since $i_T(u) = -i_T(v)$, without loss of generality we can assume that $2 \nmid |P_G^T(c, u) \cap E_{\text{sp}}|$ and $2 \mid |P_G^T(c, v) \cap E_{\text{sp}}|$. We consider the tree T rooted at c . If u and v are in different subtrees, $P_G^T(u, v)$ is the concatenation of $P_G^T(u, c)$ and $P_G^T(c, v)$. Otherwise, we may assume u is an ancestor of v , then $P_G^T(c, u)$ is a subsequence of $P_G^T(c, v)$. In both cases,

$$|P_G^T(u, v) \cap E_{\text{sp}}| \equiv |P_G^T(c, u) \cap E_{\text{sp}}| + |P_G^T(c, v) \cap E_{\text{sp}}| \pmod{2},$$

and so $2 \nmid |P_G^T(u, v) \cap E_{\text{sp}}|$.

We can assume without loss of generality that $P_G^T(u, v)$ consecutively pass through the vertex w and $-w$ with $\|w\|_1 = r$, such that the first part split from $P_G^T(u, v)$ (denoted by $P(u, w)$) contains no vertex pair i_1, i_2 with $i_1 = -i_2$ and $\|i_1\|_1 = r$. The other part split from $P_G^T(u, v)$ is denoted as $P(-w, v)$. It follows that $2 \mid |P(u, w) \cap E_{\text{sp}}|$ and $2 \mid |P(-w, v) \cap E_{\text{sp}}|$. By Claim 10, we have $\ell(P(u, w)) \geq \|u - w\|_1$ and $\ell(P(-w, v)) \geq \|-w - v\|_1$. Therefore,

$$\begin{aligned} \text{dist}_T(u, v) &= \ell(P_G^T(u, v)) = \ell(P(u, w)) + \ell(P(-w, v)) \\ &\geq \|u - w\|_1 + \|-w - v\|_1 \geq \|w - (-w)\|_1 - \|u - v\|_1 \\ &\geq 2r - d \cdot \|u - v\|_\infty \geq 2r - d. \end{aligned} \quad \blacktriangleleft$$

Observe that Theorem 5 requires only d pairs of opposite labels, rather than $\{-d, \dots, -2, -1, 1, 2, \dots, d\}$. As we mentioned before, $i(v) = -i(-v)$ for all vertices v with $\|v\|_1 = r$. Thus, by Theorem 5, the triangulation Γ' contains a complementary edge. That is, there exists a pair of neighboring vertices $u, v \in \Gamma'$ such that $i(u) = -i(v)$ and $\|u - v\|_\infty \leq 1$.

From Lemma 9, for each tree T_j in the tree cover $\{T_1, \dots, T_k\}$, $\text{dist}_{T_j}(u, v) \geq 2r - d$. Note that $\text{dist}_G(u, v) \leq \|u - v\|_1 \leq d \cdot \|u - v\|_\infty \leq d$. Therefore,

$$\min_{1 \leq j \leq k} \{\text{dist}_{T_j}(u, v)\} = \Omega\left(\frac{r}{d}\right) \cdot \text{dist}_G(u, v) = \Omega\left(\frac{n^{\frac{1}{2k-1}}}{2^{k-1}}\right) \cdot \text{dist}_G(u, v).$$

5 Lower Bound for Ramsey Tree Covers

Given a metric space (X, d_X) on $|X| = n$ points, a *Ramsey tree cover* is a collection $\mathcal{T} = \{T_1, \dots, T_k\}$ of (edge-weighted) trees on X , such that

- (1) for each $1 \leq i \leq k$ and every pair x, y of vertices in X , $\text{dist}_{T_i}(x, y) \geq d_X(x, y)$ (that is, each T_i is a *dominating tree*); and
 - (2) for every vertex x in X , there exists a tree $T_i \in \mathcal{T}$ (called the *home tree* of x), such that, for every vertex $y \in X$, $\text{dist}_{T_i}(x, y) \leq \alpha \cdot d_X(x, y)$.
- $k = |\mathcal{T}|$ is called the *size* of $\mathcal{T}$ and $\alpha \geq 1$ is called the *distortion* or *stretch* of $\mathcal{T}$.

In this section, we provide an alternative proof of the near-optimal bound of $\Omega_k(n^{1/k})$ for Ramsey tree covers of size k . We use a metric (X, d_X) similar to the one introduced in Section 4, and show that for every k dominating trees $T_1, \dots, T_k$ on X , there is a vertex u in X , such that

$$\min_{1 \leq i \leq k} \left\{ \max_{v \in X \setminus \{u\}} \frac{\text{dist}_{T_i}(u, v)}{d_X(u, v)} \right\} = \Omega\left(\frac{n^{\frac{1}{k}}}{k}\right).$$

5.1 Construction of the Hard Instance

We first define an edge-weighted graph $\tilde{G}$ as follows.

- (1) Vertex set: $V(\tilde{G}) = \{x \in \mathbb{Z}^d : \|x\|_\infty \leq r\}$;

(2) Edges set: $E(\tilde{G}) = E_0 \cup E_{\text{sp}}$, where edges in E_{sp} are called *special edges*, and

(i) $E_0 = \{(u, v) : \|u - v\|_1 = 1\}$,

(ii) $E_{\text{sp}} = \{(u, v) : u = -v, \|u\|_\infty = r\}$;


(3) Edge weights: $w_e = \begin{cases} 1, & \text{if } e \in E_0; \\ 0, & \text{if } e \in E_{\text{sp}}. \end{cases}$

Our hard instance G is the graph obtained by contracting all the edges in E_{sp} within $\tilde{G}$. Specifically, we denote the two vertices contracted by the one with positive value at the first non-zero coordinate. By choosing $d = k$ and $r = \Theta(n^{1/d})$, we can ensure that $|V(G)| = O(n)$. The metric (X, d_X) is simply the shortest path distance metric on all vertices of G .

5.2 Reduction to Theorem 6

► **Lemma 11.** *There is a triangulation Γ' of $\{x \in \mathbb{R}^d : \|x\|_\infty \leq r\}$, such that $V(\Gamma') = V(\tilde{G})$, and for every edge (u, v) in Γ' , $\|u - v\|_\infty \leq 1$.*

Proof. For a simplex σ , we denote by $\mathcal{C}(\sigma)$ the set of all its faces, i.e., $\mathcal{C}(\sigma) = \bigcup_{A \subset V(\sigma)} \text{conv}(A)$. For a simplex σ and a vertex v affinely independent with $V(\sigma)$, denote $\text{Ext}(\sigma, v) = \text{conv}(V(\sigma) \cup \{v\})$ as the simplex obtained by extending σ with v .



■ **Figure 5** The ℓ_∞ norm ball grid and its decomposition, and the triangulations of ℓ_∞ -norm unit balls in 2 and 3 dimensions.

We define Γ via the following inductive construction:

For $d = 2$, we define $\Gamma_0 = \{(u, v) \in \mathbb{Z}^2 : \|u\|_\infty = \|v\|_\infty = 1, \|u - v\|_1 = 1\}$. Denote by c the vertex corresponding to vector $\vec{0}$. Observe that the set $\{x \in \mathbb{R}^2 : \|x\|_\infty \leq 1\}$ can be triangulated by $\Gamma = \bigcup_{\sigma \in \Gamma_0} \mathcal{C}(\text{Ext}(\sigma, c))$. See Figure 5.

For general $d > 2$, we define

$$F := \left\{ \{x = (x_1, \dots, x_d) \in \mathbb{R}^d : \|x\|_\infty = 1, x_j = \beta\} : j = 1, \dots, d, \beta = \pm 1 \right\},$$

namely, the $d-1$ dimensional faces of the ℓ_∞ -norm unit ball. By induction, for each face f in F , there exists a triangulation of f . We denoted it as $\Gamma(f)$. We further define $\Gamma_{d-1} = \bigcup_{f \in F} \Gamma(f)$. Observe that $\{x \in \mathbb{R}^d : \|x\|_\infty \leq 1\}$ can be triangulated by $\Gamma = \bigcup_{\sigma \in \Gamma_{d-1}} \mathcal{C}(\text{Ext}(\sigma, c))$. (See Figure 5)

By induction, one can verified that Γ is a triangulation of $\{x \in \mathbb{R}^d : \|x\|_\infty \leq 1\}$.

The set $\{x \in \mathbb{R}^d : \|x\|_\infty \leq r\}$ can be naturally partitioned into ℓ_∞ -norm unit balls, each of which can be triangulated in the same way as in Γ , as illustrated in Figure 5. It gives a triangulation of $\{x \in \mathbb{R}^d : \|x\|_\infty \leq r\}$ denoted by Γ' . Note that $V(\Gamma') = V(\tilde{G})$, and, for every edge (u, v) in Γ' , $\|u - v\|_\infty \leq 1$. $\blacktriangleleft$

Given a dominating tree T of G . For a pair v, v' of vertices in $\tilde{G}$, we define $P_{\tilde{G}}^T(v, v')$ using the same definition as in Section 4. For each vertex $v \in \tilde{G}$, we define the tree index of v analogously. By definition, $i(v) = -i(-v)$ for all $\|v\|_\infty = r$. Note that tree indices on $V(\tilde{G})$ is a labeling $V(\Gamma') \rightarrow \{+1, -1\}^k$ of the triangulation Γ' .

By the following lemma similar to Lemma 9, we show that neighboring vertices with opposite indices have large tree distances. The proofs are almost the same as the corresponding proofs (with small changes) in Lemma 9. We present the proofs here for the sake of completeness.

► **Lemma 12.** *For a tree T and vertices u, v with $\|u - v\|_\infty \leq 1$, if $i_T(u) = -i_T(v)$, then $\text{dist}_T(u, v) \geq 2r - d$.*

Similarly, the proof of Lemma 12 requires the following lemma.

▷ **Claim 13.** Given a vertex pair $u, v \in V(\tilde{G})$ and a path P between u and v in $\tilde{G}$. If $|E(P) \cap E_{\text{sp}}|$ is an even number, then $\ell(P) \geq \|u - v\|_1$.

Proof. Without loss of generality, we assume

$$P = (v_{00}(=u), v_{01}, \dots, v_{0s_0}, v_{10}, \dots, v_{1s_1}, v_{20}, \dots, v_{(t-1)s_{t-1}}, v_{t0}, \dots, v_{ts_t}(=v)),$$

where for each $j = 0, \dots, t-1$, $v_{js_j} = -v_{(j+1)0}$ and $\|v_{js_j}\|_\infty = r$. Then we construct

$$\tilde{P} = (v_{00}, v_{01}, \dots, v_{0s_0}, -v_{10}, \dots, -v_{1s_1}, v_{20}, \dots, (-1)^{t-1}v_{(t-1)s_{t-1}}, (-1)^t v_{t0}, \dots, (-1)^t v_{ts_t}).$$

It can be observed that $\tilde{P}$ is a path in $\tilde{G} \setminus E_{\text{sp}}$, namely containing no edges in E_{sp} . Since $t = |P \cap E_{\text{sp}}|$, $2 \mid t$. Then $(-1)^t v_{ts_t} = v_{ts_t} = v$. Therefore, $\ell(P) = \ell(\tilde{P}) \geq \|u - (-1)^t v_{ts_t}\|_1 = \|u - v\|_1$. $\triangleleft$

Proof of Lemma 12. Since $i_T(u) = -i_T(v)$, by the same analysis in the proof of Lemma 9, $2 \nmid |P_{\tilde{G}}^T(u, v) \cap E_{\text{sp}}|$. We can assume without loss of generality that $P_{\tilde{G}}^T(u, v)$ consecutively pass through the vertex w and $-w$ with $\|w\|_\infty = r$, such that the first part split from $P_{\tilde{G}}^T(u, v)$ (denoted by $P(u, w)$) contains no vertex pair i_1, i_2 with $i_1 = -i_2$ and $\|i_1\|_\infty = r$. The other part split from $P_{\tilde{G}}^T(u, v)$ is denoted as $P(-w, v)$. It follows that $2 \mid |P(u, w) \cap E_{\text{sp}}|$ and $2 \mid |P(-w, v) \cap E_{\text{sp}}|$. By Claim 13, we have $\ell(P(u, w)) \geq \|u - w\|_1$ and $\ell(P(-w, v)) \geq \| -w - v\|_1$. Therefore,

$$\begin{aligned} \text{dist}_T(u, v) &= \ell(P_{\tilde{G}}^T(u, v)) = \ell(P(u, w)) + \ell(P(-w, v)) \\ &\geq \|u - w\|_1 + \| -w - v\|_1 \geq \|w - (-w)\|_1 - \|u - v\|_1 \\ &\geq 2r - d \cdot \|u - v\|_\infty \geq 2r - d. \end{aligned} \quad \blacktriangleleft$$

As we mentioned before, $i(v) = -i(-v)$ for all vertices v with $\|v\|_\infty = r$. Thus, by Theorem 6, the triangulation Γ' contains a neutral simplex. That is, there exists a simplex σ such that for all $1 \leq j \leq d$ ($d = k$), there are $v, v' \in V(\sigma)$ satisfying $i_{T_j}(v) = -i_{T_j}(v')$.

We arbitrarily choose a vertex u in $V(\sigma)$. For each tree T_j in the tree cover $\{T_1, \dots, T_k\}$, there exist $v, v' \in V(\sigma)$ such that $i_{T_j}(v) = -i_{T_j}(v')$. By Lemma 12, this implies that $\text{dist}_{T_j}(v, v') \geq 2r - d$. Since $\text{dist}_{T_j}(u, v) + \text{dist}_{T_j}(u, v') \geq \text{dist}_{T_j}(v, v')$, without loss of generality, we can assume that $\text{dist}_{T_j}(u, v) \geq \frac{r}{2}$. Note that $\text{dist}_G(u, v) \leq \|u - v\|_1 \leq d\|u - v\|_\infty \leq d$. Therefore,

$$\min_{1 \leq i \leq k} \left\{ \max_{v \in V(G) \setminus \{u\}} \frac{\text{dist}_{T_i}(u, v)}{\text{dist}_G(u, v)} \right\} = \Omega\left(\frac{r}{d}\right) = \Omega\left(\frac{n^{\frac{1}{k}}}{k}\right).$$

References


- 1 Sunil Arya, Gautam Das, David M Mount, Jeffrey S Salowe, and Michiel Smid. Euclidean spanners: short, thin, and lanky. In *Proceedings of the twenty-seventh annual ACM symposium on Theory of computing*, pages 489–498, 1995. doi:10.1145/225058.225191.
- 2 Baruch Awerbuch, Shay Kutten, and David Peleg. On buffer-economical store-and-forward deadlock prevention. *IEEE transactions on communications*, 42(11):2934–2937, 2002.
- 3 Baruch Awerbuch and David Peleg. Routing with polynomial communication-space trade-off. *SIAM Journal on Discrete Mathematics*, 5(2):151–162, 1992. doi:10.1137/0405013.
- 4 Yair Bartal, Ora Nova Fandina, and Ofer Neiman. Covering metric spaces by few trees. *Journal of Computer and System Sciences*, 130:26–42, 2022. doi:10.1016/J.JCSS.2022.06.001.
- 5 Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. On metric ramsey-type phenomena. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 463–472, 2003. doi:10.1145/780542.780610.
- 6 Artur Bikeev, Andrey Kupavskii, and Maxim Turevskii. Tree covers of size 2 for the euclidean plane, 2025. doi:10.48550/arXiv.2508.16875.
- 7 Béla Bollobás. *Extremal graph theory*. Courier Corporation, 2004.
- 8 T-H Hubert Chan, Anupam Gupta, Bruce M Maggs, and Shuheng Zhou. On hierarchical routing in doubling metrics. *ACM Transactions on Algorithms (TALG)*, 12(4):1–22, 2016. doi:10.1145/2915183.
- 9 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Covering planar metrics (and beyond): $O(1)$ trees suffice. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2231–2261. IEEE, 2023. doi:10.1109/FOCS57990.2023.00139.
- 10 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenkovic, Shay Solomon, and Cuong Than. Resolving the steiner point removal problem in planar graphs via shortcut partitions. *arXiv preprint arXiv:2306.06235*, 2023. doi:10.48550/arXiv.2306.06235.
- 11 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenković, Shay Solomon, and Cuong Than. Optimal euclidean tree covers. In *40th International Symposium on Computational Geometry (SoCG 2024)*, pages 37:1–37:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.SocG.2024.37.
- 12 Hsien-Chih Chang, Jonathan Conroy, Hung Le, Lazar Milenković, Shay Solomon, and Cuong Than. Shortcut partitions in minor-free graphs: Steiner point removal, distance oracles, tree covers, and more. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5300–5331. SIAM, 2024. doi:10.1137/1.9781611977912.191.
- 13 Ky Fan. A generalization of tucker’s combinatorial lemma with topological applications. *Annals of Mathematics*, 56(3):431–437, 1952.
- 14 Elyot Grant and Will Ma. A geometric approach to combinatorial fixed-point theorems. In *The Seventh European Conference on Combinatorics, Graph Theory and Applications: EuroComb 2013*, pages 463–468. Springer, 2013.
- 15 Anupam Gupta. Steiner points in tree metrics don’t (really) help. In *Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 220–227, 2001. URL: <http://dl.acm.org/citation.cfm?id=365411.365448>.
- 16 Anupam Gupta, Mohammad T Hajiaghayi, and Harald Räcke. Oblivious network design. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 970–979, 2006.
- 17 Anupam Gupta, Amit Kumar, and Rajeev Rastogi. Traveling with a pez dispenser. In *42nd Annual Symposium on Foundations of Computer Science: Proceedings: October 14-17, 2001, Las Vegas, Nevada, USA*, page 148. IEEE Computer Society, 2001.
- 18 Hung Le, Lazar Milenković, Shay Solomon, and Tianyi Zhang. Covering the euclidean plane by a pair of trees, 2025. doi:10.48550/arXiv.2508.11507.
- 19 Jiří Matoušek, Anders Björner, and Günter M Ziegler. *Using the Borsuk-Ulam theorem: lectures on topological methods in combinatorics and geometry*. Springer, 2003.

- 20 Manor Mendel and Assaf Naor. Ramsey partitions and proximity data structures. *Journal of the European Mathematical Society*, 9(2):253–275, 2007.
- 21 Frédéric Meunier. *Pleins étiquetages et configurations équilibrées: aspects topologiques de l’optimisation combinatoire*. PhD thesis, Université Joseph-Fourier-Grenoble I, 2006.
- 22 Oleg R Musin. Sperner type lemma for quadrangulations. *arXiv preprint*, 2014. [arXiv:1406.5082](#).
- 23 Yuri Rabinovich and Ran Raz. Lower bounds on the distortion of embedding finite metric spaces in graphs. *Discrete & Computational Geometry*, 19(1):79–94, 1998. [doi:10.1007/PL00009336](#).
- 24 A. W. Tucker. Some topological properties of disk and sphere. In *Proceedings of the First Canadian Math. Congress (Montreal, 1945)*, pages 285–309, 1945.

A

 A $(O(1), 2)$ -Tree Cover for Grid

Consider the grid $\{x \in \mathbb{Z}^2 : \|x\|_\infty \leq \sqrt{n}\}$ on $O(n)$ vertices. In this section, we present the construction of a size-two tree cover $\{T_1, T_2\}$ that covers the grid with distortion $O(1)$. A similar construction (with much detailed presentation) is given in two concurrent works [18, 6].



(a) The recursive construction of T_1 : Tree edges are shown in blue. c is the center of T_1 , c_1 is the center of its northwest subtree, c_2 is the center of the southwest subtree of the c_1 -tree. The path from c to a vertex v in the c_2 tree is shown in orange.

(b) The recursive construction of T_2 : Tree edges are shown in blue. c is the center of T_2 , c_1 is the center of its north subtree, c_2 is the center of the west subtree of the c_1 -tree. The path from c to a vertex v in the c_2 tree is shown in orange. Gray lines form a tilted larger grid.

Figure 6 The construction of T_1 and T_2 .


We construct T_1 as follows. We view T_1 as being rooted at the vertex $(0, 0)$ and consisting of 8 subtrees below c . Four of them are paths: north path $((0, 0), (0, 1), \dots, (0, \sqrt{n}))$, south path: $((0, 0), \dots, (0, -\sqrt{n}))$, the east path $((0, 0), \dots, (\sqrt{n}, 0))$, and the west path: $((0, 0), \dots, (-\sqrt{n}, 0))$. The other remaining are subtrees: the northwest tree, the northeast tree, and the southeast tree, containing all vertices in different orthants (excluding the points on the x and y axis), respectively. The northwest tree, centered at $(-\frac{\sqrt{n}}{2}, \frac{\sqrt{n}}{2})$, is defined recursively (developing 8 subtrees from $(-\frac{\sqrt{n}}{2}, \frac{\sqrt{n}}{2})$, four paths and four subtrees further towards different directions, etc), and is connected to the root $(0, 0)$ by the single edge $((-1, 1), (0, 0))$. See Figure 6 for an illustration.

The recursive construction naturally defines a quad-tree structure of the grid: layer 1 is the whole grid; layer 2 consist of 4 “areas”, corresponding to the vertex sets of the northwest, northeast, southwest, and southeast subtrees (vertices on four paths are excluded from these

regions). Each of these areas are further divided into two “regions” along its diagonal, so the second layer contains a total 8 regions. See Figure 7 for an illustration. Deeper layers and its regions are defined recursively in a similar way. We use the following simple property.

► **Observation 14.** *For each center c , and for each vertex v in the subtree centered at c , the c - v grid distance is preserved up to factor $O(1)$ in tree T_1 . In particular, the grid distances from all vertices to the layer-1 center $(0,0)$ is preserved within factor $O(1)$ in T_1 .*

T_2 is constructed as follows. We start by viewing the original grid as a subset of a larger grid rotated by $\pi/4$ (a “tilted grid”). We then construct a tree T^* recursively on this tilted grid, in the same way as T_1 , and finally restrict this tree onto the original grid (it is easy to verify that T^* induces a connected subtree on vertices of the original grid). This restricted tree is our T_2 . See Figure 6 for an illustration. It can be shown that Observation 14 also holds for T_2 .



■ **Figure 7** Eight regions partitioned from a component, diagonal and orthogonal neighboring.

We now show that the two trees T_1, T_2 constructed above cover the grid with distortion $O(1)$.

Consider a pair u, v of vertices in the grid. Let c be the center of largest-indexed layer in T_1 that has an area containing both u and v . Recall that the c -tree is divided into 8 regions. We distinguish between the following cases.

Case 1. u, v lie in non-adjacent regions

(e.g., regions I and VII in Figure 7). Let θ denote the angle between vectors $\vec{uc}$ and $\vec{vc}$, then $\pi/4 \leq \theta \leq 7\pi/4$. Denote by C the constant hidden in Observation 14, by basic trigonometry,

$$\begin{aligned}
 d(u, v) &= \sqrt{d^2(c, u) + d^2(c, v) - 2 \cos \theta \cdot d(c, u) d(c, v)} \\
 &\geq \sqrt{d^2(c, u) + d^2(c, v) - \sqrt{2} d(c, u) d(c, v)} \\
 &\geq \frac{\sqrt{2 - \sqrt{2}}}{2} \cdot (d(c, u) + d(c, v)).
 \end{aligned}$$

38:16 Lower Bounds on Tree Covers

Denote $C' = \frac{2C}{\sqrt{2}-\sqrt{2}}$. Then the tree distance in T_1 between u and v satisfies

$$\begin{aligned} d_{T_1}(u, v) &= d_{T_1}(u, c) + d_{T_1}(c, v) \leq C \cdot (d(u, c) + d(c, v)) \\ &\leq \frac{2C}{\sqrt{2}-\sqrt{2}} \cdot d(u, v) = C' \cdot d(u, v), \end{aligned}$$

where $d(u, v)$ denotes the grid distance between u and v .

Case 2. u, v lie in adjacent regions

There are two possibilities: diagonally neighboring (e.g., regions I and VIII) or orthogonally neighboring (e.g., regions I and II). In the first case, observe that u and v belong to the same area in the next layer of T_1 , a contradiction to the assumption that c is the center of largest-indexed layer in T_1 that has an area containing both u and v . In the second case, they belong to the same area in the next layer of T_2 . We may then zoom in that layer, further distinguish between the cases where they lie in adjacent or non-adjacent regions, and repeat our analysis.

To sum up, for a pair u, v of vertices: either

- (1) they lie in non-adjacent regions at some layer, in which case by similar analysis in Case 1 we can show that their grid distance is preserved up to factor $O(1)$; or
- (2) they lie in the same area in the last layer of T_1 or T_2 , which is a subtree of $O(1)$ size and therefore preserves grid distances within $O(1)$ factor.