

Optimal Padded Decomposition For Bounded Treewidth Graphs*

Arnold Filtser¹, Tobias Friedrich², Davis Issac³, Nikhil Kumar⁴, Hung Le⁵,
Nadym Mallek⁶, and Ziena Zeif⁷

¹Bar-Ilan University. Email: `arnold.filtser@biu.ac.il`

^{2,6,7}Hasso Plattner Institute. Email: `firstname.lastname@hpi.de`

⁴University of Waterloo. Email: `nikhil.kumar2@uwaterloo.ca`

³NextSilicon. Email: `davisissac22@gmail.com`

⁵University of Massachusetts Amherst. Email: `hungle@cs.umass.edu`

Abstract

A (β, δ, Δ) -padded decomposition of an edge-weighted graph $G = (V, E, w)$ is a stochastic decomposition into clusters of diameter at most Δ such that for every vertex $v \in V$, the probability that $\mathbf{B}_G(v, \gamma\Delta)$ is entirely contained in the cluster containing v is at least $e^{-\beta\gamma}$ for every $\gamma \in [0, \delta]$. Padded decompositions have been studied for decades and have found numerous applications, including metric embedding, multicommodity flow-cut gap, multicut, and zero extension problems, to name a few. In these applications, parameter β , called the *padding parameter*, is the most important parameter since it decides either the distortion or the approximation ratios. For general graphs with n vertices, $\beta = \Theta(\log n)$.

Klein, Plotkin, and Rao [KPR93] (KPR) showed that K_r -minor-free graphs have padding parameter $\beta = O(r^3)$, which is a significant improvement over general graphs when r is a constant. However, when $r = \Omega(\log n)$, the padding parameter in KPR decomposition can be much worse than $\log n$. A long-standing conjecture is to construct a padded decomposition for K_r -minor-free graphs with padding parameter $\beta = O(\log r)$. Despite decades of research, the best-known result is $\beta = O(r)$, even for graphs with treewidth at most r .

In this work, we make significant progress toward the aforementioned conjecture by showing that graphs with treewidth tw admit a padded decomposition with padding parameter $O(\log \text{tw})$, which is tight. Our padding parameter is strictly better than $O(\log n)$ whenever $\text{tw} = n^{o(1)}$, and is never worse than what is known for general graphs. As corollaries, we obtain an exponential improvement in dependency on treewidth in a host of algorithmic applications: $O(\sqrt{\log n \cdot \log(\text{tw})})$ flow-cut gap, the maxflow-min multicut ratio of $O(\log(\text{tw}))$, an $O(\log(\text{tw}))$ approximation for the 0-extension problem, an $\ell_\infty^{O(\log n)}$ embedding with distortion $O(\log \text{tw})$, and an $O(\log \text{tw})$ bound for integrality gap for the uniform sparsest cut.

*Part of the results in this paper were previously published in conference STOC23 [FIK⁺22].

Contents

1	Introduction	1
2	Preliminaries	4
2.1	From Tree Decomposition to Tree Partition	5
3	Padded Decomposition and Sparse Cover	6
3.1	Proof of Theorem 2	7
3.2	Proof of Theorem 3	8
3.3	Proof of Theorem 1	9
4	Tree-Ordered Nets for Graphs of Bounded Tree-partition Width	11
4.1	The (Semi-)Tree-Ordered Net	13
4.2	The Analysis	17
5	Applications	23
5.1	Flow Sparsifier	23
5.2	Sparse Partition	24
5.3	Universal Steiner Tree and Universal TSP	24
5.4	Steiner Point Removal	25
5.5	Zero Extension	26
5.6	Lipschitz Extension	26
5.7	Embedding into ℓ_p spaces	26
5.8	Stochastic decomposition for minor-free graphs - a reduction from additive stretch embeddings	27

1 Introduction

A basic primitive in designing divide-and-conquer graph algorithms is partitioning a graph into clusters such that there are only a few edges between clusters. This type of primitive has been extensively applied in the design of divide-and-conquer algorithms. Since guaranteeing a few edges crossing different clusters in the worst case could be very expensive, we often seek a good guarantee in a probabilistic sense: the probability that two vertices u and v are placed into two different clusters is proportional to $d_G(u, v)/\Delta$ where $d_G(u, v)$ is the distance between u and v in the input graph G and Δ is the upper bound on the diameter of each cluster. A (stochastic) partition of $V(G)$ with this property is called a *separating decomposition* of G [Fil19a].

In this work, we study a stronger notion of stochastic decomposition, called *padded composition*. More formally, given a weighed graph $G = (V, E, w)$, a partition is Δ -*bounded* if the diameter of every cluster is at most Δ . A distribution \mathcal{D} over partitions is called a (β, δ, Δ) -*padded decomposition*, if every partition is Δ -bounded, and for every vertex $v \in V$ and $\gamma \in [0, \delta]$, we have:

$$\Pr[\mathbf{B}_G(v, \gamma\Delta) \subseteq P(v)] \geq e^{-\beta\gamma} \quad \text{where } P(v) \text{ is the cluster containing } v. \quad (1)$$

That is, the probability that the entire ball $\mathbf{B}_G(v, \gamma\Delta)$ of radius $\gamma\Delta$ around v is clustered together, is at least $e^{-\beta\gamma}$. If G admits a (β, δ, Δ) -padded decomposition for every $\Delta > 0$, we say that G admits (β, δ) -padded decomposition scheme. The parameter β is usually referred to as a padding parameter.

In an influential work, Klein, Plotkin and Rao [KPR93] showed that every K_r minor free graph admits a weak $(O(r^3), \Omega(1))$ -padded decomposition scheme; the padding parameter is $O(r^3)$. This result has found numerous algorithmic applications for solving problems in K_r -minor-free graphs; a few examples are the flow-cut gap of $O(r^3)$ for uniform multicommodity flow [KPR93], extending Lipschitz functions with absolute extendability of $O(r^3)$ [LN05], the maxflow-min multicut ratio of $O(r^3)$ for the multicommodity flow with maximum total commodities [TV93], an $O(r^3 \log \log(n))$ approximation for minimum linear arrangement, minimum containing interval graphs [RR05], an $O(r^3)$ approximation for the 0-extension problem [CKR05], and an $O(r^3 \log n)$ -approximation for the minimum bisection problem [FK02]. An important takeaway is that key parameters quantifying the quality of these applications depend (linearly) on the padding parameter; any improvement to the padding parameter would imply the same improvement in the applications.

Fakcharoenphol and Talwar [FT03] improved the padding parameter of K_r minor free graphs to $O(r^2)$. Abraham, Gavaille, Gupta, Neiman, and Talwar [AGG⁺19] (see also [Fil19a]) improved the padding parameter to $O(r)$. These improvements imply an $O(r)$ dependency on the minor size of all aforementioned applications. The only lower bound is $\Omega(\log r)$ coming from the fact that r -vertex expanders (trivially) exclude K_r as a minor while having padding parameter $\Omega(\log r)$ [Bar96]. Closing the gap between the upper bound of $O(r)$ and the lower bound $\Omega(\log r)$ has been an outstanding problem asked by various authors [FT03, Lee12, AGG⁺19, Fil19a].

Conjecture 1. *There exists a padded decomposition of any K_r -minor-free metric with padding parameter $\beta = O(\log r)$.*

Progress on Conjecture 1 has been made on very special classes of minor-free graphs. Specifically, graphs with *pathwidth* pw admit padding parameter $O(\log pw)$ [AGG⁺19]. This result implies that n -vertex graphs with *treewidth* tw admit padding parameter $O(\log tw + \log \log n)$, since the pathwidth of graphs of treewidth tw is $O(tw \cdot \log(n))$ (also see [KK17]). However, the padding

parameter depends on n . Thus, a significant step towards Conjecture 1 is to show that graphs of treewidth tw admit a padded decomposition with padding parameter $O(\log tw)$. The best-known result (without the dependency on n) for small treewidth graphs is the same as minor-free graphs, implied by the fact that such graphs of treewidth tw exclude K_{tw+2} as a minor. Our first main result is to prove Conjecture 1 for the special case of treewidth- tw graphs:

Theorem 1. *Every weighted graph G with treewidth tw admits a $(O(\log tw), \Omega(1))$ -padded decomposition scheme. Furthermore, such a partition can be sampled efficiently.*

Our Theorem 1 implies that we could replace r^3 with $O(\log tw)$ in the aforementioned problems when the input graphs have treewidth tw : $O(\log(tw))$ for uniform multicommodity flow-cut gap, extending Lipschitz functions with absolute extendability of $O(\log(tw))$, the maxflow-min multicut ratio of $O(\log(tw))$, an $O(\log(tw) \log \log(n))$ approximation for minimum linear arrangement, minimum containing interval graphs, an $O(\log(tw))$ approximation for the 0-extension problem, and an $O(\log(tw) \log n)$ -approximation for the minimum bisection problem. Furthermore, we obtain the first ℓ_1 embedding of treewidth- tw metrics with distortion $\mathcal{O}(\sqrt{\log tw \cdot \log n})$, an $\ell_\infty^{O(\log n)}$ embedding with distortion $O(\log tw)$, and $O(\log tw)$ bound for integrality gap for the uniform sparsest cut. For several of these problems, for example, uniform multicommodity flow-cut gap, maxflow-min multicut ratio, and 0-extension, our results provide the state-of-the-art approximation ratio for an entire range of parameter tw , even when $tw = \Omega(n)$. We refer readers to Section 5 for a more comprehensive discussion of these results.

Here we point out another connection to Conjecture 1. Filtser and Le [FL22] showed that one can embed any K_r -minor-free metric of diameter Δ into a graph with treewidth $O_r(\epsilon^{-2} \cdot (\log \log n)^2)$ and additive distortion $\epsilon \cdot \Delta$. The dependency of $O_r(\cdot)$ on r is currently huge; it is the constant in the Robertson-Seymour decomposition. However, if one could manage to get the same treewidth to be $O(\text{poly}(r/\epsilon))$, then in combination with our Theorem 1, one has a positive answer to Conjecture 1. Even an embedding with a treewidth $O(\text{poly}(r/\epsilon) \text{poly}(\log(n)))$ already implies a padding parameter $O(\log(r) + \log \log(n))$, a significant progress towards Conjecture 1.

Sparse Covers. A related notion to padded decompositions is *sparse cover*. A collection \mathcal{C} of clusters is a (β, s, Δ) -sparse cover if it is Δ -bounded, each ball of radius $\frac{\Delta}{\beta}$ is contained in some cluster, and each vertex belongs to at most s different clusters. A graph admits (β, s) -sparse cover scheme if it admits (β, s, Δ) -sparse cover for every $\Delta > 0$. Sparse covers have been studied for various classes of graphs, such as general graphs [AP90], planar graphs [BLT14], minor-free graphs [KLMN04, BLT14, AGMW10], and doubling metrics [Fil19a].

By simply taking the union of many independently drawn copies of padded decomposition, one can construct a sparse cover. Indeed, given (β, δ, Δ) -padded decomposition, by taking the union of $O(e^{\beta\gamma} \log n)$ partitions (for $\gamma \leq \delta$) one will obtain w.h.p. a $(\gamma, O(e^{\beta\gamma} \log n), \Delta)$ -sparse cover. In particular, using Theorem 1 one can construct $(O(1), O(e^{\log tw} \log n)) = (O(1), tw^{O(1)} \log n)$ -sparse cover scheme. The main question in this context is whether one can construct sparse covers for bounded treewidth graphs with constant cover parameter (β) and sparseness (s) independent from n . If one is willing to sacrifice a (quadratic) dependency on tw on the cover parameter β , then a sparse cover with parameters independent of n is known [KPR93, FT03, AGGM06]. However, in many applications, for example, constructing sparse spanners, it is desirable to have $\beta = O(1)$ as it directly governs the stretch of the spanners.

Theorem 2. *Every graph G with treewidth tw admits a $(6, \text{poly}(tw))$ -sparse cover scheme.*

It is sometimes useful to represent the sparse cover \mathcal{C} as a union of partitions, for example, in metric embeddings [KLMN04], and the construction of ultrametric covers [FL22, Fil23], leading to the notion of *padded partition cover scheme*:

Definition 1 (Padded Partition Cover Scheme). *A collection of partitions $\mathcal{P}_1, \dots, \mathcal{P}_\tau$ is (β, s, Δ) -padded partition cover if (a) $\tau \leq s$, (b) every partition \mathcal{P}_i is Δ -bounded, and (c) for every point x , there is a cluster C in one of the partitions \mathcal{P}_i such that $B(x, \frac{\Delta}{\beta}) \subseteq C$.*

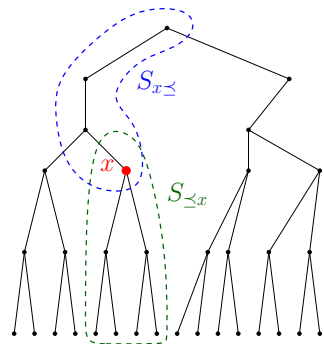
A space (X, d_X) admits a (β, s) -padded partition cover scheme if for every Δ , it admits a (β, s, Δ) -padded partition cover.

While a padded partition cover implies a sparse cover with the same parameters, the reverse direction is not true. For example, graphs with pathwidth pw admit $(10, 5(\text{pw} + 1))$ -sparse cover scheme [Fil20], however, they are only known to admit $(O(\text{pw}^2), 2^{\text{pw}+1})$ -padded partition cover scheme (this is due to K_r -minor free graphs [KPR93, FT03] (see also [KLMN04, Fil20]))¹. That is, the sparseness parameter in the padded partition cover scheme is exponentially worse (in terms of pw) than that of the sparse cover scheme. In this work, we construct a padded partition cover scheme with the same quality as our sparse covers.

Theorem 3. *Every graph G with treewidth tw admits a $(12, \text{poly}(\text{tw}))$ -padded partition cover scheme.*

Tree-Ordered Net. A key new technical insight to all of our aforementioned results is the notion of *tree-ordered net* (Definition 2). A tree order net is analogous to the notion of *nets*, which were used extensively in designing algorithms for metric spaces. More formally, a Δ -net is a set of points N such that every two net points are at a distance at least Δ (i.e. $\min_{x, y \in N} d_X(x, y) \geq \Delta$), and every point has a net point at a distance at most Δ (i.e. $\max_{x \in V} \min_{y \in N} d_X(x, y) \leq \Delta$). Filtser [Fil19a] showed that if there is a Δ -net such that every ball of radius 3Δ contains at most τ net points, that the metric admits a $(O(\log \tau), \Omega(1), O(\Delta))$ -padded decomposition. This result implies that metrics of doubling dimension d have padding parameter $O(d)$ since doubling metrics have sparse nets: $\tau = 2^{O(d)}$. Unfortunately, graphs of small treewidth do not have sparse nets; this holds even in very simple graphs such as star graphs. Nonetheless, we show that small treewidth graphs possess a structure almost as good: a net that is sparse w.r.t. some partial order. We formalize this property via tree-ordered nets. As we will later show, a sparse tree-ordered net is enough to construct the padded decomposition scheme; see Section 3. We believe that the notion of a tree-ordered net is of independent interest.

A tree order \leq of a set V is a partial order (i.e. transitive, reflexive, and anti-symmetric) associated with a rooted tree T and a map $\varphi : V \rightarrow V(T)$ such that $u \leq v$ iff $\varphi(v)$ is an ancestor of $\varphi(u)$ in T . Given a weighted graph $G = (V, E, w)$, a tree order \leq w.r.t. tree T is a *valid order* of G if for every edge $\{u, v\} \in E$, it holds that $u \leq v$ or $v \leq u$ or both (i.e. v is an ancestor of u , or vice versa). A simple consequence of the validity is that every connected subset C in G must contain a maximum element w.r.t. \leq ; see Observation 1. For a vertex $v \in V$, and subset $S \subseteq V$ let $S_{v \leq} = \{u \in S \mid v \leq u\}$ be the ancestors of v w.r.t. T in S . Similarly, let $S_{\leq x} = \{u \in S \mid u \leq x\}$ be all the descendants of x in S . See the illustration on the right.



¹Note that we do not have matching lower bounds, so there is no provable separation.

Definition 2. Given a weighted graph $G = (V, E, w)$ and parameters $\tau, \alpha, \Delta > 0$, a (τ, α, Δ) -tree-ordered net is a triple (N, T, φ) where $N \subseteq V$, and T and φ define a tree order \leq of V such that for every $v \in V$:

- COVERING. There is $x \in N_{v \leq}$ such that $d_{G[V_{\leq x}]}(v, x) \leq \Delta$. That is, there exists an ancestor x of v in N such that the distance from v to x in the subgraph of G induced by descendants of x is at most Δ .
- PACKING. Denote by $N_{v \leq}^{\alpha \Delta} = \{x \in N \cap V_{v \leq} \mid d_{G[V_{\leq x}]}(v, x) \leq \alpha \Delta\}$ the set of ancestor centers of v at distance at most $\alpha \Delta$ from v (w.r.t the subgraphs induced by descendants of the ancestors). Then $|N_{v \leq}^{\alpha \Delta}| \leq \tau$.

While we state our main result in terms of treewidth, it will be more convenient to use the notion of *bounded tree-partition width* [DO96]. We will show that graphs of bounded tree-partition width admit a small tree-ordered net.

Definition 3 (Tree Partition). A tree partition of a graph $G = (V, E)$ is a rooted tree \mathcal{T} whose vertices are bijectively associated with the sets of partition $\{S_1, S_2, \dots, S_m\}$ of V , called bags, such that for each $(u, v) \in E$, there exists a $S_i, S_j \in \mathcal{S}$ such that S_j is a parent of S_i and $\{u, v\} \subseteq S_i \cup S_j$. The width of \mathcal{T} is $\max_{i \in m} \{|S_i|\}$.

Unlike a tree decomposition, bags of a tree partition are disjoint. Therefore, graphs of bounded-tree partition width have a more restricted structure than graphs of bounded treewidth. Indeed, one can show that any graphs of tree-partition width k have treewidth at most $2k - 1$. However, from a metric point of view, graphs of bounded treewidth are the same as graphs of bounded tree-partition width: We could convert a tree decomposition into a tree partition by making copies of vertices; see Lemma 2. We will show in Section 4 that:

Lemma 1. Every weighted graph $G = (V, E, w)$ with a tree-partition width tp admits a $(\text{poly}(\text{tp}), 3, \Delta)$ -tree-ordered net, for every $\Delta > 0$.

In Section 3, we show how to use the tree-ordered net in Lemma 1 to obtain all results stated above.

Follow-up Work. Recently, inspired by our technique, [CCL⁺23] constructed a tree cover with stretch $1 + \varepsilon$ and $2^{(\text{tw}/\varepsilon)^{O(\text{tw})}}$ trees. This result has applications to constructing distance oracles and approximate labeling schemes for graphs of small treewidth. Filtser [Fil24] showed recently that K_r -minor-free graphs admit a $(4 + \varepsilon, O(1/\varepsilon)^r)$ -sparse cover scheme for every $\varepsilon \in (0, 1)$. While the stretch is an absolute constant, the dependency on the minor size is exponential and hence is incomparable to our Theorem 2.

2 Preliminaries

Graphs. We consider connected undirected graphs $G = (V, E)$ with edge weights $w : E \rightarrow \mathbb{R}_{\geq 0}$. We say that vertices v, u are neighbors if $\{v, u\} \in E$. Let d_G denote the shortest path metric in G . $\mathbf{B}_G(v, r) = \{u \in V \mid d_G(v, u) \leq r\}$ is the ball of radius r around v . For a vertex $v \in V$ and a subset $A \subseteq V$, let $d_G(x, A) := \min_{a \in A} d_G(x, a)$, where $d_G(x, \emptyset) = \infty$. For a subset of vertices $A \subseteq V$, let $G[A]$ denote the induced graph on A , and let $G \setminus A := G[V \setminus A]$.

The *diameter* of a graph G is $\text{diam}(G) = \max_{v,u \in V} d_G(v,u)$, i.e. the maximal distance between a pair of vertices. Given a subset $A \subseteq V$, the *weak-diameter* of A is $\text{diam}_G(A) = \max_{v,u \in A} d_G(v,u)$, i.e. the maximal distance between a pair of vertices in A , w.r.t. to original distances d_G . The *strong-diameter* of A is $\text{diam}(G[A])$, the diameter of the graph induced by A . A graph H is a *minor* of a graph G if we can obtain H from G by edge deletions/contractions, and isolated vertex deletions. A graph family \mathcal{G} is *H -minor-free* if no graph $G \in \mathcal{G}$ has H as a minor. We will drop the prefix H in H -minor-free whenever H is not important or clear from the context. Some examples of minor-free graphs are planar graphs (K_5 - and $K_{3,3}$ -minor-free), outer-planar graphs (K_4 - and $K_{3,2}$ -minor-free), series-parallel graphs (K_4 -minor-free) and trees (K_3 -minor-free).

Treewidth. A *tree decomposition* of a graph $G = (V, E)$ is a tree \mathcal{T} where each node $x \in \mathcal{T}$ is associated with a subset S_x of V , called a *bag*, such that: (i) $\cup_{x \in V(\mathcal{T})} S_x = V$, (ii) for every edge $(u, v) \in E$, there exists a bag S_x for some $x \in V(\mathcal{T})$ such that $\{u, v\} \subseteq S_x$, and (iii) for every $u \in V$, the bags containing u induces a connected subtree of \mathcal{T} . The *width* of \mathcal{T} is $\max_{x \in V(\mathcal{T})} \{|S_x| - 1\}$. The *treewidth* of G is the minimum width among all possible tree decompositions of G .

Padded Decompositions. Consider a *partition* \mathcal{P} of V into disjoint clusters. For $v \in V$, we denote by $P(v)$ the cluster $P \in \mathcal{P}$ that contains v . A partition \mathcal{P} is *strongly Δ -bounded* (resp. *weakly Δ -bounded*) if the strong-diameter (resp. weak-diameter) of every $P \in \mathcal{P}$ is bounded by Δ . If the ball $\mathbf{B}_G(v, \gamma\Delta)$ of radius $\gamma\Delta$ around a vertex v is fully contained in $P(v)$, we say that v is *γ -padded* by \mathcal{P} . Otherwise, if $\mathbf{B}_G(v, \gamma\Delta) \not\subseteq P(v)$, we say that the ball is *cut* by the partition.

Definition 4 (Padded Decomposition). *Consider a weighted graph $G = (V, E, w)$. A distribution \mathcal{D} over partitions of G is *strongly (resp. weakly) (β, δ, Δ) -padded decomposition* if every $\mathcal{P} \in \text{supp}(\mathcal{D})$ is *strongly (resp. weakly) Δ -bounded* and for any $0 \leq \gamma \leq \delta$, and $z \in V$,*

$$\Pr[\mathbf{B}_G(z, \gamma\Delta) \subseteq P(z)] \geq e^{-\beta\gamma} .$$

We say that a graph G admits a strong (resp. weak) (β, δ) -padded decomposition scheme, if for every parameter $\Delta > 0$ it admits a strongly (resp. weakly) (β, δ, Δ) -padded decomposition that can be sampled in polynomial time.

2.1 From Tree Decomposition to Tree Partition

We first show that any graph of treewidth tw can be embedded isometrically into a graph of tree-partition width tw by duplicating vertices.

Lemma 2. *Given an edge-weighted graph $G(V, E, w)$ and its tree decomposition of width tw , there is an isometric polynomial time constructible embedding of G into a graph with tree partition of width $\text{tw} + 1$. More formally, there is a graph $H = (X, E_H, w_H)$ with tree partition of width $\text{tw} + 1$ and a map $\phi: V \rightarrow X$ such that $\forall x, y \in V, d_H(\phi(x), \phi(y)) = d_G(x, y)$.*

Proof. Let \mathcal{B} be a tree decomposition of width at most r of $G = (V, E, w)$. We create a graph H and its tree partition as follows (see Figure 1). For each $u \in V$, if u appears in k bags of \mathcal{B} , say B_1, B_2, \dots, B_k , then we make k copies of u , say u_1, u_2, \dots, u_k and replace u in bag B_i with its copies $u_i, i \in [1, k]$. This defines the set of vertices X of H . We then set $\phi(u) = u_1$.

If B_i and B_j are two adjacent bags in \mathcal{B} , we create an edge (u_i, u_j) and assign a weight $w_H(u_i, u_j) = 0$. For each $(u, v) \in E$, there exists at least one bag B_j of the tree decomposition




Figure 1: Converting a tree decomposition in (a) into a tree partition in (b). The thick red edges are the original edges in G , and the thin blue edges are the added edges with weight zero.

of G such that $u, v \in B_j$. We add the edge (u_j, v_j) of weight $w_H(u_j, v_j) = w(u, v)$ to E_H . This completes the construction of H .

The tree partition of H , say \mathcal{T} , has the same structure as the tree decomposition \mathcal{B} of G : for each bag $B \in \mathcal{B}$, there is a corresponding bag $\hat{B} \in \mathcal{T}$ containing copies of vertices of B . As $|B| \leq r+1$, the width of \mathcal{T} is $r+1$. For each $u \in V$, we map it to exactly one copy of u in X . As edges between copies of the same vertex have 0 weight, the distances in G are preserved exactly in H . \square

Using Lemma 2 together with the tree-ordered net in Lemma 1 we will construct (weak) padded decomposition and sparse covers for graphs of small treewidth.

3 Padded Decomposition and Sparse Cover

In this section, we prove three general lemmas constructing a padded decomposition and a sparse cover from a tree-ordered net. These lemmas, together with Lemma 1, imply Theorem 1, Theorem 2, and Theorem 3. We begin with Theorem 2, as its proof is simpler. We observe the following by the definition of a tree order.

Observation 1. *Let \leq be a valid tree order of $G = (V, E, w)$ defined by a tree T . Every subset of vertices C such that $G[C]$ is connected must contain a single maximum element w.r.t \leq .*

Proof. Suppose that there are t maximal elements $u_1, u_2, \dots, u_t \in C$ for $t \geq 2$. Let $A_i = \{v : v \leq u_i\}$ for every $i \in [t]$ and $\mathcal{A} = \{A_1, \dots, A_t\}$. Then \mathcal{A} is a partition of C since $\{u_i\}_{i=1}^t$ are maximal (and T is a tree). As $G[C]$ is connected and $t \geq 2$, there must be some edge $\{x, y\} \in E$ such that $x \in A_i$

and $y \in A_j$ for $i \neq j$. The validity of T implies that either $x \leq y$ or $y \leq x$. However, this means either x is also in A_j or y is also in A_i , contradicting that \mathcal{A} is a partition of C . \square

3.1 Proof of Theorem 2

The following lemma is a reduction from tree-ordered nets to sparse covers.

Lemma 3. *Consider a weighted graph $G = (V, E, w)$ with a (τ, α, Δ) -tree-ordered net (w.r.t. a tree order \leq , associated with a tree T), then G admits a strong $(\frac{4\alpha}{\alpha-1}, \tau, 2\alpha\Delta)$ -sparse cover that can be computed efficiently.*

Proof. Let N be the (τ, α, Δ) -tree-ordered net of G . Let x_1, x_2, \dots be an ordering of the centers in N w.r.t distance from the root in T . Specifically, x_1 is closest to the root, and so on. Note that distances in T are unweighted and unrelated to d_G ; we break ties arbitrarily. For every vertex $x_i \in N$, create a cluster

$$C_i = \mathbf{B}_{G[V_{\leq x_i}]}(x_i, \alpha\Delta),$$

of all the vertices that are at distances at most $\alpha\Delta$ from x_i in the subgraph induced by the descendants of x_i . We now show that $\mathcal{C} = \{C_i\}_i$ is the sparse cover claimed in the lemma.

Clearly, by the triangle inequality, every cluster in \mathcal{C} has a diameter at most $2\alpha\Delta$. Furthermore, observe that every vertex v belongs to at most τ clusters since it has at most τ ancestors in N at distance at most $\alpha\Delta$ (in the respective induced graphs).

Finally, we show that for every vertex v , $\mathbf{B}_G(v, \frac{\Delta}{\beta})$ is fully contained in some cluster in \mathcal{C} , for $\beta = \frac{2}{\alpha-1}$. Let $B = \mathbf{B}_G(v, \frac{\Delta}{\beta})$. Let $v_B \in B$ be the closest vertex to the root w.r.t. T . Observation 1 implies that for every vertex $u \in B$, $u \leq v_B$. By the triangle inequality, for every $u \in B$ it holds that:

$$d_{G[V_{\leq v_B}]}(v_B, u) \leq d_{G[V_{\leq v_B}]}(v_B, v) + d_{G[V_{\leq v_B}]}(v, u) \leq \frac{2\Delta}{\beta} \quad (2)$$

Let $x_B \in N$ be the ancestor of v_B in T that minimizes $d_{G[V_{\leq x}]}(x, v_B)$. Since N is a (τ, α, Δ) -tree-ordered net, it holds that $d_{G[V_{\leq x_B}]}(x_B, v_B) \leq \Delta$. In particular, for every $u \in B$ it holds that:

$$\begin{aligned} d_{G[V_{\leq x_B}]}(x_B, u) &\leq d_{G[V_{\leq x_B}]}(x_B, v_B) + d_{G[V_{\leq x_B}]}(v_B, u) \\ &\leq \Delta + \frac{2\Delta}{\beta} \leq \alpha\Delta \quad (\text{by eq. (2)}) \end{aligned}$$

Thus, B is fully contained in the cluster centered in x_B . As the diameter of each cluster is $2\alpha\Delta$, the padding we obtain is $\frac{2\alpha\Delta}{\frac{\Delta}{\beta}} = \frac{2\alpha}{\frac{\alpha-1}{2}} = \frac{4\alpha}{\alpha-1}$, as required. \square

Now we are ready to show that Theorem 2 follows from Lemma 1, Lemma 2, and Lemma 3. We restate Theorem 2 for convenience.

Theorem 2. *Every graph G with treewidth tw admits a $(6, \text{poly}(tw))$ -sparse cover scheme.*

Proof. Let $\Delta > 0$ be a parameter; we will construct a $(6, \text{poly}(tw), \Delta)$ -sparse cover for G . Let H be the graph of tree-partition width $tw + 1$ in the isometric embedding ϕ of G in Lemma 2. We abuse notation by using v , for each $v \in V(G)$, to denote $\phi(v)$ in H . This means $V(G) \subseteq V(H)$.

Let $\hat{\Delta} = \Delta/6$. By Lemma 1, H admits a $(\text{poly}(tw), 3, \hat{\Delta})$ -tree-ordered net. By Lemma 3, we can construct a $(6, \text{poly}(tw), 6\hat{\Delta} = \Delta)$ -sparse cover of H , denoted by $\hat{\mathcal{C}}$. We then construct

$\mathcal{C} = \{\hat{C} \cap V(G) : \hat{C} \in \hat{\mathcal{C}}\}$. Let C be a cluster in \mathcal{C} and we denote by \hat{C} its corresponding cluster in $\hat{\mathcal{C}}$; that is, $C = \hat{C} \cap V(G)$. We claim that \mathcal{C} is a $(6, \text{poly}(\text{tw}), \Delta)$ -sparse cover for G .

For any $C \in \mathcal{C}$, $\text{diam}(C) \leq \text{diam}(\hat{C}) \leq \Delta$, since ϕ is an isometric embedding. Thus, \mathcal{C} is Δ -bounded. Furthermore, any $v \in V(G)$ belongs to at most $\text{poly}(\text{tw})$ clusters in $\hat{\mathcal{C}}$ and hence it belongs to at most $\text{poly}(\text{tw})$ clusters in \mathcal{C} . Finally, we consider any ball $B_G(v, \Delta/6)$ in G . Observe that $B_G(v, \Delta/6) \subseteq B_H(v, \Delta/6)$. Since $\hat{\mathcal{C}}$ is a $(6, \text{poly}(\text{tw}), \Delta)$ -sparse cover of H , there exists a cluster $\hat{C} \in \hat{\mathcal{C}}$ such that $B_H(v, \Delta/6) \subseteq \hat{C}$, implying $B_G(v, \Delta/6) \subseteq C$. \square

3.2 Proof of Theorem 3

In this subsection, we show that the construction in Lemma 3 can be adapted to obtain a padded partition cover scheme (with a small loss in the padding parameter).

Lemma 4. *Consider a weighted graph $G = (V, E, w)$ with a (τ, α, Δ) -tree-ordered net (w.r.t. a tree order \leq , associated with T) for $\alpha > 2$, then G admits strong $(\frac{4\alpha}{\alpha-2}, \tau, \alpha\Delta)$ -padded partition cover.*

By exactly the same argument in the proof of Theorem 2 in Section 3.1, one can show that Theorem 3 follows directly from Lemma 1, Lemma 2 and Lemma 4: first, we isometrically embed G of treewidth tw to a graph H of tree-partition width $\text{tw} + 1$, and then apply Lemma 4 and Lemma 1 to construct a padded partition cover $\hat{\mathcal{P}}$ for H , which will then be turned into a padded partition cover for G by removing vertices not in G from $\hat{\mathcal{P}}$. Our main focus now is to prove Lemma 4.

Proof of Lemma 4. We say that a set of clusters is a *partial partition* if it is a partition of a subset of vertices; that is, a vertex might not belong to any cluster in a partial partition. We then can turn the partial partitions into partitions of V by adding singleton clusters. Similarly to Lemma 3, we define a set of clusters (note radius $\frac{\alpha}{2} \cdot \Delta$ compared to $\alpha \cdot \Delta$ in Lemma 3):

$$\mathcal{C} = \left\{ C_x = \mathbf{B}_{G[V_{\leq x}]}(x, \frac{\alpha}{2} \cdot \Delta) \right\}_{x \in N}$$

Our partial partitions will consist of clusters in \mathcal{C} only. We construct the partitions greedily: form a partition from a maximal set of disjoint clusters, preferring ones that are closer to the root, and repeat. The pseudocode is given in Algorithm 1.

Algorithm 1: Create Partial Partitions($(\leq, T), N, \mathcal{C}$)

```

1  $A \leftarrow N$ 
2  $i \leftarrow 0$ 
3 while  $A \neq \emptyset$  do
4    $i \leftarrow i + 1$ 
5    $\mathcal{P}_i \leftarrow \emptyset$ 
6   while  $\exists x \in A$  such that  $C_x$  is disjoint from  $\cup \mathcal{P}_i$  do
7     Let  $x \in A$  be a maximal element w.r.t.  $(\leq, T)$  such that  $C_x$  is disjoint from  $\cup_{C \in \mathcal{P}_i} C$ 
8     Remove  $x$  from  $A$ 
9     Add  $C_x$  to  $\mathcal{P}_i$ 
10 return  $\{\mathcal{P}_j\}_{j=1}^i$ 

```

By construction in Line 7, clusters in every partition \mathcal{P}_i are pairwise disjoint. Furthermore, for every $x \in N$, C_x belongs to one of the created partitions due to Line 3. We next argue that the algorithm creates at most τ partial partitions.

Suppose for contradiction that the algorithm does not terminate after creating τ partial partitions. Thus, after τ iterations, there was still an element $x \in A$. In particular, in every iteration i , there exists some vertex $x_i \in N$ such that $x \leq x_i$, and $C_x \cap C_{x_i} \neq \emptyset$. Let y_i be a vertex in $C_x \cap C_{x_i}$. By the triangle inequality:

$$d_G(x, x_i) \leq d_{G[V_{\leq x_i}]}(x, y_i) + d_{G[V_{\leq x_i}]}(y_i, x_i) \leq \alpha \Delta .$$

Note that as every cluster in \mathcal{C} can join only one partial partition, all these centers $\{x_i\}_{i=1}^\tau$ are unique; this contradicts the fact that N is a (τ, α, Δ) -tree-ordered net (as together with x itself, $|N_{x \leq}^{\alpha \Delta}| > \tau$).

It remains to show padding property, which is followed by the same proof as in Lemma 3. Consider a ball $B = \mathbf{B}_G(v, \frac{\Delta}{\beta})$, for $\beta = \frac{4}{\alpha-2}$. Let $v_B \in B$ be the closest vertex to the root w.r.t. T . Following eq. (2), for every $u \in B$, $d_{G[V_{\leq v_B}]}(v_B, u) \leq \frac{2\Delta}{\beta}$. Let $x_B \in N$ be the ancestor of v_B in T that minimizes $d_{G[V_{\leq x}]}(x, v_B)$. Since N is a (τ, α, Δ) -tree-ordered net, it holds that $d_{G[V_{\leq x_B}]}(x_B, v_B) \leq \Delta$. In particular, for every $u \in B$ it holds that

$$\begin{aligned} d_{G[V_{\leq x_B}]}(x_B, u) &\leq d_{G[V_{\leq x_B}]}(x_B, v_B) + d_{G[V_{\leq x_B}]}(v_B, u) \\ &\leq \Delta + \frac{2\Delta}{\beta} = \left(1 + \frac{2 \cdot (\alpha - 2)}{4}\right) \cdot \Delta = \frac{\alpha}{2} \cdot \Delta , \end{aligned}$$

and thus B is fully contained in the cluster centered in x_B a required. \square

3.3 Proof of Theorem 1

In the following lemma, we construct a padded decomposition from a tree-ordered net.

Lemma 5. *Consider a weighted graph $G = (V, E, w)$ with a (τ, α, Δ) -tree-ordered net (w.r.t. a tree order \leq , associated with T), then G admits a weak $\left(16 \cdot \frac{\alpha+1}{\alpha-1} \cdot \ln(2\tau), \frac{\alpha-1}{8 \cdot (\alpha+1)}, (\alpha+1) \cdot \Delta\right)$ -padded decomposition that can be efficiently sampled.*

By exactly the same argument in the proof of Theorem 2 in Section 3.1, one can show that Theorem 1 follows directly from Lemma 1, Lemma 2 and Lemma 5. We will use truncated exponential distribution during the proof of Lemma 5:

Truncated Exponential Distributions. To create padded decompositions, similarly to previous works, we will use truncated exponential distributions. A truncated exponential distribution is an exponential distribution conditioned on the event that the outcome lies in a certain interval. More precisely, the $[\theta_1, \theta_2]$ -truncated exponential distribution with parameter λ , denoted by $\text{Texp}_{[\theta_1, \theta_2]}(\lambda)$, has the density function: $f(y) = \frac{\lambda e^{-\lambda y}}{e^{-\lambda \theta_1} - e^{-\lambda \theta_2}}$, for $y \in [\theta_1, \theta_2]$.

Proof of Lemma 5. Let x_1, x_2, \dots be an ordering of the centers in N w.r.t distances from the root in T . Set $\beta = \frac{\alpha+1}{2}$. For every vertex $x_i \in N$, sample $\delta_i \in [1, \beta]$ according to $\text{Texp}_{[1, \beta]}(\lambda)$, a truncated exponential distribution, with parameter $\lambda = \frac{4}{\alpha-1} \cdot \ln(2\tau)$. Set $R_i = \delta_i \cdot \Delta \in [\Delta, \beta \Delta]$ and create a cluster:

$$C_i = \mathbf{B}_{G[V_{\leq x_i}]}(x, R_i) \setminus \cup_{j < i} C_j .$$

Recall that $\mathbf{B}_{G[S_{\leq x_i}]}(x, R_i)$ is the ball of radius R_i around x_i in the graph induced by all the descendants of x_i . Thus, the cluster C_i of x_i consists of all the points in this ball that did not join the clusters centered at the (proper) ancestors of x_i . Note that C_i might not be connected and that x_i might not even belong to C_i as it could join a previously created cluster. Nonetheless, C_i has a (weak) diameter at most $2R_i \leq 2\beta\Delta = (\alpha + 1) \cdot \Delta$ by the triangle inequality.

We claim that each vertex will eventually be clustered. Indeed, consider a vertex $v \in V$. There exists some vertex $x_i \in N_{v \leq}$ at a distance at most Δ from v in $G[V_{\leq x_i}]$ by the definition of the tree-ordered net. If v did not join any cluster centered at an ancestor of x_i , then v will join C_i because $d_{G[V_{\leq x_i}]}(v, x_i) \leq \Delta \leq R_i$.

It remains to prove the padding property. Consider some vertex $v \in V$ and parameter $\gamma \leq \frac{\alpha-1}{8}$. We argue that the ball $B = \mathbf{B}_G(v, \gamma\Delta)$ is fully contained in $P(v)$ with probability at least $e^{-4\gamma\lambda}$. For $x_i \in N$, denote by \mathcal{F}_i the event that some vertex of B joins the cluster C_i for the first time. That is, $B \cap C_i \neq \emptyset$ and for all $j < i$, $B \cap C_j = \emptyset$. Denote by \mathcal{C}_i the event that \mathcal{F}_i occurred and B is cut by C_i (i.e. $B \not\subseteq C_i$).

Let $v_B \in B$ be the vertex closest to the root of T (w.r.t distances in T). By Observation 1, for every vertex $u \in B$, it holds that $u \leq v_B$. Let $x_B \in N$ be the center that is an ancestor of v_B and minimizes $\mathbf{B}_{G[V_{\leq x}]}(x, v_B)$. Note that v_B will join the cluster of x_B , if it did not join any other cluster. It follows that no descendant of x_B can be the center of the first cluster having a non-trivial intersection with B . This implies that for every center $x_i \notin N_{x_B \leq}$, $\Pr[\mathcal{F}_i] = 0$.

Claim 1. *Let N_B be the set of centers x_i for which $\Pr[\mathcal{F}_i] > 0$. Then $|N_B| \leq \tau$.*

Proof. Observe that \mathcal{F}_i can have non-zero probability only if x_i is an ancestor of x_B and that $d_{G[V_{\leq x_i}]}(x_i, z) \leq R_i \leq \beta\Delta$ for some vertex $z \in B$. As all vertices of B are descendants of x_i , $G[B]$ is a subgraph of $G[V_{\leq x_i}]$. It follows from the triangle inequality that:

$$\begin{aligned} d_{G[V_{\leq x_i}]}(x_i, v_B) &\leq d_{G[V_{\leq x_i}]}(x_i, z) + d_{G[B]}(z, v_B) \\ &\leq (\beta + 2\gamma)\Delta = \left(\frac{\alpha + 1}{2} + 2 \cdot \frac{\alpha - 1}{8} \right) \cdot \Delta < \alpha\Delta \end{aligned} \quad (3)$$

As N is a (τ, α, Δ) -tree-ordered net, there are at most τ centers in N_i satisfying eq. (3), implying the claim. \square

We continue by bounding the probability of a cut by each center.

Claim 2. *For every i , $\Pr[\mathcal{C}_i] \leq (1 - e^{-2\gamma\lambda}) \cdot \left(\Pr[\mathcal{F}_i] + \frac{1}{e^{(\beta-1)\lambda} - 1} \right)$.*

Proof. We assume that by the round i , no vertex in B is clustered, and that $d_{G[V_{\leq x_i}]}(x_i, B) \leq \beta\Delta$, as otherwise $\Pr[\mathcal{C}_i] = \Pr[\mathcal{F}_i] = 0$ and we are done. Let ρ be the minimal value of δ_i such that if $\delta_i \geq \rho$, some vertex of B will join C_i . Formally $\rho = \frac{1}{\Delta} \cdot d_{G[V_{\leq x_i}]}(x_i, B)$. By our assumption, $\rho \leq \beta$. Set $\tilde{\rho} = \max\{\rho, 1\}$. We have:

$$\Pr[\mathcal{F}_i] = \Pr[\delta_i \geq \rho] = \int_{\tilde{\rho}}^{\beta} \frac{\lambda \cdot e^{-\lambda y}}{e^{-\lambda} - e^{-\beta\lambda}} dy = \frac{e^{-\tilde{\rho}\lambda} - e^{-\beta\lambda}}{e^{-\lambda} - e^{-\beta\lambda}}.$$

Let $v_i \in B$ be the closest vertex to x_i w.r.t. $G[V_{\leq x_i}]$. Note that $d_{G[V_{\leq x_i}]}(x_i, v_i) = \rho \cdot \Delta$. Then for every $u \in B$ it holds that

$$d_{G[V_{\leq x_i}]}(u, x_i) \leq d_{G[V_{\leq x_i}]}(v_i, x_i) + 2\gamma\Delta = (\rho + 2\gamma) \cdot \Delta.$$

Therefore, if $\delta_i \geq \rho + 2\gamma$, the entire ball B will be contained in C_i . We conclude that:

$$\begin{aligned}
\Pr[\mathcal{C}_i] &\leq \Pr[\rho \leq \delta_i < \rho + 2\gamma] \\
&= \int_{\tilde{\rho}}^{\min\{\beta, \rho+2\gamma\}} \frac{\lambda \cdot e^{-\lambda y}}{e^{-\lambda} - e^{-\beta\lambda}} dy \\
&\leq \frac{e^{-\tilde{\rho}\lambda} - e^{-(\tilde{\rho}+2\gamma)\lambda}}{e^{-\lambda} - e^{-\beta\lambda}} \\
&= (1 - e^{-2\gamma\lambda}) \cdot \frac{e^{-\tilde{\rho}\lambda}}{e^{-\lambda} - e^{-\beta\lambda}} \\
&= (1 - e^{-2\gamma\lambda}) \cdot \left(\Pr[\mathcal{F}_i] + \frac{e^{-\beta\lambda}}{e^{-\lambda} - e^{-\beta\lambda}} \right) \\
&= (1 - e^{-2\gamma\lambda}) \cdot \left(\Pr[\mathcal{F}_i] + \frac{1}{e^{(\beta-1)\lambda} - 1} \right).
\end{aligned}$$

□

We now bound the probability that the ball B is cut. As the events $\{\mathcal{F}_i\}_{x_i \in N}$ are disjoint, we have that

$$\begin{aligned}
\Pr[\cup_i \mathcal{C}_i] &= \sum_{x_i \in N_B} \Pr[\mathcal{C}_i] \leq (1 - e^{-2\gamma\lambda}) \cdot \sum_{x_i \in N_B} \left(\Pr[\mathcal{F}_i] + \frac{1}{e^{(\beta-1)\lambda} - 1} \right) \\
&\leq (1 - e^{-2\gamma\lambda}) \cdot \left(1 + \frac{\tau}{e^{(\beta-1)\lambda} - 1} \right) \quad (\text{by Claim 1}) \\
&\leq (1 - e^{-2\gamma\lambda}) \cdot (1 + e^{-2\gamma\lambda}) = 1 - e^{-4\gamma\lambda},
\end{aligned}$$

where the last inequality follows as

$$e^{-2\gamma\lambda} = \frac{e^{-2\gamma\lambda} (e^{(\beta-1)\lambda} - 1)}{e^{(\beta-1)\lambda} - 1} \geq \frac{e^{-2\gamma\lambda} \cdot e^{(\beta-1)\lambda} \cdot \frac{1}{2}}{e^{(\beta-1)\lambda} - 1} \stackrel{(*)}{\geq} \frac{\frac{1}{2} \cdot e^{\frac{\alpha-1}{4}\lambda}}{e^{(\beta-1)\lambda} - 1} = \frac{\tau}{e^{(\beta-1)\lambda} - 1}.$$

The inequality $(*)$ holds as $\beta - 1 - 2\gamma \geq \frac{\alpha+1}{2} - 1 - 2 \cdot \frac{\alpha-1}{8} = \frac{\alpha-1}{4}$.

To conclude, we obtain a partition where each cluster has diameter at most $2\beta\Delta \leq (\alpha+1) \cdot \Delta$, and for every $\gamma \leq \frac{\alpha-1}{8}$, every ball of radius $\gamma \cdot \Delta = \frac{\gamma}{\alpha+1} \cdot (\alpha+1) \cdot \Delta$ is cut with probability at most $e^{-4\gamma\lambda} = e^{-\frac{\gamma}{\alpha+1} \cdot (\alpha+1) \cdot 4\lambda}$. Thus the padding parameter is $(\alpha+1) \cdot 4 \cdot \lambda = \frac{\alpha+1}{\alpha-1} \cdot 16 \cdot \ln(2\tau)$, while the guarantee holds for balls of radius up to $\frac{1}{\alpha+1} \cdot \frac{\alpha-1}{8}$. The lemma now follows. □

4 Tree-Ordered Nets for Graphs of Bounded Tree-partition Width

In this section, we show that graphs of bounded tree-partition width have small tree-ordered nets as claimed in Lemma 1, which we restate here for convenience.

Lemma 1. *Every weighted graph $G = (V, E, w)$ with a tree-partition width tp admits a $(\text{poly}(\text{tp}), 3, \Delta)$ -tree-ordered net, for every $\Delta > 0$.*

Herein, let \mathcal{T} be a tree partition of G of width tp . Recall that in the definition of a tree order \leq realized by a tree T and a map $\varphi : V \rightarrow V(T)$, \leq is a partial order, which means \leq is anti-symmetric. This means φ is injective, i.e., two different vertices in V will be mapped to two distinct vertices in $V(T)$. In our construction presented below, it is convenient to drop the anti-symmetric property to allow two distinct vertices to be mapped to the same vertex in $V(T)$. We call a tree order \leq without the anti-symmetric property a *semi-tree order*.

We then can extend the notion of a tree-ordered net to a semi-tree-ordered net in a natural way: a triple (N, T, φ) is a (τ, α, Δ) -semi-tree-ordered net if T and φ define a semi-tree order on V , and N satisfies the covering and packing properties as in Definition 2. Specifically:

- **Covering:** for every $v \in V$ there is $x \in N$ such that $\phi(v) \leq \phi(x)$, and $d_{G[V_{\leq x}]}(v, x) \leq \Delta$ for the set $V_{\leq x} = \{u \mid \phi(u) \leq \phi(x)\}$.
- **Packing:** for a vertex $v \in V$, denote by $N_{v \leq}^{\alpha \Delta} = \{x \in N \cap V_{v \leq} \mid d_{G[V_{\leq x}]}(v, x) \leq \alpha \Delta\}$ the set of ancestor centers of v at distance at most $\alpha \Delta$ (here $V_{v \leq} = \{u \mid \phi(v) \leq \phi(u)\}$). Then $|N_{v \leq}^{\alpha \Delta}| \leq \tau$.

In Section 4.1, we show that graphs of tree-partition width tp admit a good semi-tree-ordered net.

Lemma 6. *Every weighted graph $G = (V, E, w)$ with a tree-partition width tp admits a $(\text{poly}(\text{tp}), 3, \Delta)$ -semi-tree-ordered net, for every $\Delta > 0$.*

Given Lemma 6, we now show how to construct a good tree-ordered net as claimed in Lemma 1.

Proof of Lemma 1. Let $(N, \hat{T}, \hat{\varphi})$ be a $(\text{poly}(\text{tp}), 3, \Delta)$ -semi-tree-ordered net of G . For each vertex $\hat{x} \in \hat{T}$, let $\hat{\varphi}^{-1}(\hat{x}) \subseteq V$ be the set of vertices in G that are mapped to \hat{x} . That is, $\hat{\varphi}^{-1}(x) = \{v \in V : \hat{\varphi}(v) = \hat{x}\}$. Roughly speaking, to construct a tree order T for G , we simply replace \hat{x} in \hat{T} by a path, say $P_{\hat{x}}$, composed of vertices in $\hat{\varphi}^{-1}(\hat{x})$. The packing property remains the same, but the covering property might not hold if vertices in $\hat{\varphi}^{-1}(\hat{x})$ are ordered arbitrarily along $P_{\hat{x}}$. Our idea to guarantee the packing property is to place net points in $\hat{\varphi}^{-1}(\hat{x})$ closer to the root of the tree.

We now formally describe the construction of T . Let $N_{\hat{x}} = N \cap \hat{\varphi}^{-1}(x)$. $N_{\hat{x}}$ is the set of net points that are mapped to \hat{x} by $\hat{\varphi}$. Let $P_{\hat{x}}$ be a *rooted* path created from vertices in $\hat{\varphi}^{-1}(x)$ where all vertices in $\hat{\varphi}^{-1}(x) \setminus N_{\hat{x}}$ are descendants of every vertex in $N_{\hat{x}}$. Note that if $N_{\hat{x}} \neq \emptyset$, then the root of $P_{\hat{x}}$ is a vertex in $N_{\hat{x}}$. Vertices in $N_{\hat{x}}$ are ordered arbitrarily in $P_{\hat{x}}$ and vertices in $\hat{\varphi}^{-1}(x) \setminus N_{\hat{x}}$ are also ordered arbitrarily. (If $\hat{\varphi}^{-1}(x) = \emptyset$, then $P_{\hat{x}}$ will be a single vertex that does not correspond to any vertex in $V(G)$.) Then we create the tree T by connecting all the paths $\{P_{\hat{x}} : \hat{x} \in V(\hat{T})\}$ in the following way: if (\hat{x}, \hat{y}) is an edge in \hat{T} such that x is the parent of y , then we connect the root of $P_{\hat{y}}$ to the (only) leaf of $P_{\hat{x}}$. The bijection between $P_{\hat{x}}$ and $\hat{\varphi}^{-1}(x)$, unless when $\hat{\varphi}^{-1}(x) = \emptyset$, naturally induce an injective map $\varphi : V(G) \rightarrow T$.

Let \leq_T ($\leq_{\hat{T}}$) be the tree order (semi-tree order resp.) induced by T and φ (\hat{T} and $\hat{\varphi}$ resp.). To show that T and φ induce a tree order, it remains to show the validity: for every edge $\{u, v\} \in E$, either $u \leq_T v$ or $v \leq_T u$. Since $\leq_{\hat{T}}$ is a semi-tree ordered, w.l.o.g., we can assume that $\hat{\varphi}(u)$ is an ancestor of $\hat{\varphi}(v)$; it is possible that $\hat{\varphi}(u) = \hat{\varphi}(v)$. If $\hat{\varphi}(u) \neq \hat{\varphi}(v)$, then every vertex in $P_{\hat{\varphi}(u)}$ will be an ancestor of every vertex in $P_{\hat{\varphi}(v)}$, implying that $\varphi(u)$ is an ancestor of $\varphi(v)$ and hence $u \leq_T v$. If $\hat{\varphi}(u) = \hat{\varphi}(v)$, then u and v belong to the same path $P_{\hat{\varphi}(u)}$ and hence either $u \leq_T v$ or $v \leq_T u$. Thus, the validity follows.

Finally, we show that (N, T, φ) is a $(\text{poly}(\text{tp}), 3, \Delta)$ -tree-ordered net. Observe by the construction of T that:

Observation 2. *if $u \leq_T v$ then $u \leq_{\hat{T}} v$.*

The converse of Observation 2 might not be true, specifically in the case when u and v are mapped to the same vertex in \hat{T} by $\hat{\varphi}$. By Observation 2, for any vertex $x \in V$, $V_{\leq_T x} \subseteq V_{\leq_{\hat{T}} x}$. Thus, for any v and x such that $v \leq_T x$, if $d_G[V_{\leq_T x}](v, x) \leq 3\Delta$, then $d_G[V_{\leq_{\hat{T}} x}](v, x) \leq 3\Delta$. Therefore, $N_{v \leq_T}^{3\Delta} \subseteq N_{v \leq_{\hat{T}}}^{3\Delta}$ and hence $|N_{v \leq_T}^{3\Delta}| \leq |N_{v \leq_{\hat{T}}}^{3\Delta}| = \text{poly}(\text{tp})$, giving the packing property.

Next, we show the covering property. By the covering property of $(N, \hat{T}, \hat{\varphi})$, for every $v \in V$, there exists $x \in N_{v \leq_{\hat{T}}}$ such that $d_G[V_{\leq_{\hat{T}} x}](x, v) \leq \Delta$. Let Q be a shortest path in $G[V_{\leq_{\hat{T}} x}]$ from v to x . Note that Q might contain vertices that are not in $G[V_{\leq_T x}]$.

Claim 3. *If Q contains $y \notin V_{\leq_T x}$, then $y \in N$ and $x \leq_T y$.*

Proof. Since $y \leq_{\hat{T}} x$, as $y \notin V_{\leq_T x}$, it must be the case that $\hat{\varphi}(y) = \hat{\varphi}(x)$ by the construction of T . Let $\hat{x} = \hat{\varphi}(x)$. Since net points in $\hat{\varphi}^{-1}(\hat{x})$ are placed closer to the root in $P_{\hat{x}}$ and $y \notin V_{\leq_T x}$, y must also be a net point and $x \leq_T y$. \square

Let $z \in N \cap Q$ such that z has the highest order in T . By Claim 3, every $y \in Q[v, z]$ satisfies $y \in V_{\leq_T z}$. This implies $z \in N_{\leq_T v}$ and $G[V_{\leq_T z}](v, z) \leq w(Q[x, z]) \leq w(Q) \leq \Delta$, giving the covering property. \square

4.1 The (Semi-)Tree-Ordered Net

For a bag B in \mathcal{T} , we define \mathcal{T}_B to be the subtree of \mathcal{T} rooted at B . For a subtree \mathcal{T}' of \mathcal{T} , we define $V[\mathcal{T}']$ to be the union of all bags in \mathcal{T}' ; that is, $V[\mathcal{T}'] = \cup_{B \in \mathcal{T}'} B$.

Constructing Cores. We construct a set \mathcal{R} of subsets of the vertex set V whose union covers V ; see Algorithm 2. Each set in \mathcal{R} is called a *core*. We then construct a tree ordering of the vertices of G using these sets.

We describe the notation used in Algorithm 2. For a subset $U \subseteq V$ of vertices, we define $\mathbf{B}_G(U, \Delta) = \cup_{u \in U} \mathbf{B}_G(u, \Delta)$ to be the ball of radius Δ centered at U . We say that a vertex is *covered* if it is part of at least one core constructed so far. For any subset $X \subseteq V$ of vertices, we use $\text{Uncov}(X)$ to denote the set of uncovered vertices of X . We say that a bag \mathcal{T} is *covered* if all the vertices in the bag are covered; a bag is *uncovered* if at least one vertex in the bag is uncovered.

The algorithm proceeds in *rounds*—the while loop in line 5—covers vertices in each round and continues until all vertices are covered. During each round, we process each connected component \mathcal{T}' of the forest induced on \mathcal{T} by the uncovered bags of \mathcal{T} . Note that \mathcal{T} is rooted and \mathcal{T}' is a rooted subtree of \mathcal{T} . The algorithm then works on \mathcal{T}' in a top-down manner. The basic idea is to take an unvisited bag B of \mathcal{T}' (initially all bags in \mathcal{T}' are marked unvisited) closest to the root (line 9), carve a ball of radius Δ centered at the *uncovered vertices* of B in a subgraph H (defined in line 10) of G as a new core R (line 11), mark all bags intersecting with R as visited, and repeat. We call the bag B in line 9 the *center bag* of R , and the uncovered vertices in B are called the *centers* of the core R . Graph H is the subgraph of G induced by uncovered vertices in the subtree of \mathcal{T}' rooted at B and *the attachments* of the bags of this subtree—the set $A(X)$ associated with each bag X . We will clarify the role of the attachments below. See Figure 2 for an illustration.

For each bag B of \mathcal{T} a subset of vertices, we associated with B a set of vertices $V[\mathcal{T}_B]$, which are contained in bags of the subtree \mathcal{T}_B , called an *attachment* of B , denoted by $A(B)$. The attachments

allowed the cores created in a round to contain vertices in the cores formed in previous rounds. The general idea is that $A(B)$ contains all the cores (in previous rounds) whose center bags are children of B . However, $A(B)$ is also subjected to changes, as described next.

For a subtree \mathcal{T}' of \mathcal{T} , we use $A[\mathcal{T}']$ to denote the union of attachments of all bags of \mathcal{T}' . Formally, $A[\mathcal{T}'] = \cup_{B \in \mathcal{T}'} A(B)$. When forming a core from a center bag B in a connected component \mathcal{T}' , we construct a graph H , which is a subgraph G induced by uncovered vertices in the subtree \mathcal{T}'_B rooted at B and the attachment $A[\mathcal{T}'_B]$ (line 10). The core R is a ball of radius at most Δ from uncovered vertices of B in H (line 11). We call H the support graph of R . Note that R may contain vertices in the attachments of the bags in \mathcal{T}'_B , and for each such bag, we remove vertices in R from its attachment (line 14).

Once we construct a core R from the center bag B , we have to update the attachment of the parent bag of B to contain R (line 18), unless B is the root of \mathcal{T}' . If B is the root of \mathcal{T}' , either B has no parent bag— B is also the root of \mathcal{T} in this case—or the parent bag of B is covered, and hence will not belong to any connected component of uncovered bags. In both cases, the parent of B will not be directly involved in any subsequent rounds; they could be involved indirectly via the attachments. It is useful to keep in mind that if X is already a covered bag, the attachment update in Line 18 by adding a core centered at child bag B has *no effect* on subsequent rounds as we only consider uncovered bags.

Algorithm 2: CONSTRUCTCORES(\mathcal{T}, G)

```

1  $\mathcal{R} \leftarrow \emptyset$ 
2  $A(B) \leftarrow \emptyset$  for every bag  $B \in \mathcal{T}$  // the attachment.
3 mark every vertex of  $V$  uncovered
4  $\text{round} \leftarrow 1$ 
5 while there is an uncovered vertex do
6   for each tree  $\mathcal{T}'$  of the forest induced by uncovered bags of  $\mathcal{T}$  do
7     mark all bags of  $\mathcal{T}'$  unvisited
8     while there is an unvisited bag in  $\mathcal{T}'$  do
9       pick an unvisited bag  $B$  of  $\mathcal{T}'$  closest to the root (breaking ties arbitrarily)
10       $H \leftarrow$  the graph induced in  $G$  by  $\text{Uncov}(V[\mathcal{T}'_B]) \cup A[\mathcal{T}'_B]$ 
11      add  $\mathbf{B}_H(\text{Uncov}(B), \Delta)$  as a new core  $R$  into  $\mathcal{R}$ 
12      mark uncovered vertices of  $R$  as covered
13      mark all the bags of  $\mathcal{T}'$  that intersects  $R$  as visited
14      for every bag  $X \in \mathcal{T}'_B$  s.t.  $A(X) \cap R \neq \emptyset$  do
15         $A(X) \leftarrow A(X) \setminus R$  // remove vertices of  $R$  from the attachment
16      if  $B$  is not the root of  $\mathcal{T}'$  then
17         $X \leftarrow$  parent bag of  $B$  in  $\mathcal{T}'$ 
18         $A(X) \leftarrow A(X) \cup R$  // update the attachment of  $B$ 's parent.
19    $\text{round} \leftarrow \text{round} + 1$ 
20 return  $\mathcal{R}$ 

```

As alluded above, cores in a specific round could contain vertices of cores in previous rounds via attachments. However, as we will show in the following claim that cores in the same round are

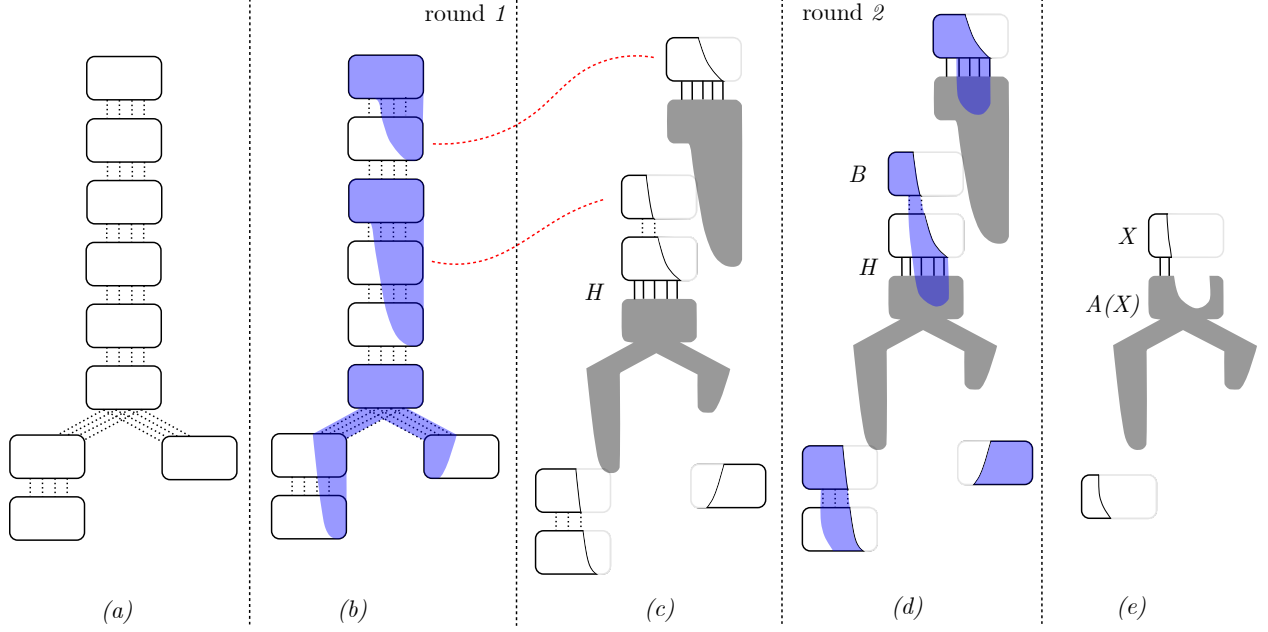


Figure 2: Illustrating two rounds of Algorithm 2. (b) In the first round, the algorithm creates (blue) cores from \mathcal{T} ; the attachment of every bag is \emptyset . (c) After the first round, uncovered bags of \mathcal{T} form a forest; these are the white bags. Some bags now have non-empty attachments, illustrated by the gray-shaded regions. (d) In the second round, the algorithm considers each connected component \mathcal{T}' of uncovered bags and creates cores by carving balls $\mathbf{B}_H(\text{Uncov}(B), \Delta)$ in graph H from uncovered vertices of B . The attachment allows a core of the 2nd round to grow within a core of the 1st round. (e) The remaining uncovered bags and their attachments after round 2. A core in round 2 carves out a portion of the attachment of X .

vertex-disjoint.

Claim 4. *The cores and attachments satisfy the following properties:*

1. $\cup_{R \in \mathcal{R}} R = V$.
2. For any bag X in a connected component of unvisited bags \mathcal{T}' of \mathcal{T} considered in line 6, $A(X)$ contains vertices in the descendant bags of X that are currently not in \mathcal{T}' . As a corollary, at any point of the algorithm, $A(X)$ only contains vertices in descendant bags of X , excluding X .
3. For every core $R \in \mathcal{R}$ centered at a bag B , vertices in R are in descendant bags of B . (B is considered its own descendant.)
4. For every core $R \in \mathcal{R}$ centered at a bag B , if a vertex $u \in R$ was covered before R is created, then u is in the attachments of descendant bags of B .
5. Let R_1 and R_2 be two different cores created from the same connected component \mathcal{T}' in the same round. Then $V(R_1) \cap V(R_2) = \emptyset$.

Proof. We observe that the algorithm only terminates when every vertex is covered. Furthermore, the algorithm only marks an uncovered vertex as covered when it is contained in a new core (line 12), implying item 1.

Item 2 follows from the fact that whenever we update the attachment of an uncovered bag X in line 18 by a core R , the center bag B of R become covered and hence will be disconnected from the connected component containing X in the next round.

For item 3, we observe that the support graph H of R in line 10 contains uncovered vertices in descendant bags of B and their attachments. By item 2, the attachment of a bag X only contains vertices in the descendant bags of X . Thus, vertices in H are in descendant bags of B , as claimed.

For item 4, observe that the only way for a covered vertex to be considered in subsequent rounds is via the attachments. Thus, item 4 follows from item 3.

For the last item 5, suppose otherwise: there is a vertex v in a bag Z such that $v \in R_1 \cap R_2$. W.l.o.g, we assume that R_1 is created before R_2 . Let B_1 (B_2 , resp.) be the center bag of R_1 (R_2 , resp.). Then B_1 is the ancestor of B_2 by construction, and furthermore, by item 3, Z are descendants of both B_1 and B_2 .

If $Z \in \mathcal{T}'$, then when R_1 was created, all the bags on the path from B_1 to Z will be marked visited, and hence B_2 will also be marked visited, contradicting that R_2 was created from B_2 . Otherwise, $Z \notin \mathcal{T}'$ and that means v belongs to the attachment $A(Y)$ of some bag $Y \in \mathcal{T}'$. Since all bags on the path from B_1 to Y are marked when R_1 was created, B_2 is not an ancestor of Y . But this means v cannot belong to a descendant bag of B_2 , contradicting item 3. \square

A crucial property of the core construction algorithm used in our analysis is that it has a natural hierarchy of clusters associated with it. To define this hierarchy, we need some new notation. Let \mathcal{T}' be a connected component of uncovered bags in a specific round. Let $C_{\mathcal{T}'} = \text{Uncov}(V[\mathcal{T}']) \cup A[\mathcal{T}']$ be a set of vertices, called the *cluster* associated with \mathcal{T}' . Note that $V(H) \subseteq C_{\mathcal{T}'}$. The following lemma implies that clusters over different rounds form a hierarchy.

Lemma 7. *Let \mathcal{T}_1 and \mathcal{T}_2 be two different connected components of uncovered bags. If \mathcal{T}_1 and \mathcal{T}_2 belong to the same round, then $C_{\mathcal{T}_1} \cap C_{\mathcal{T}_2} = \emptyset$. If \mathcal{T}_1 and \mathcal{T}_2 belong to two consecutive rounds such that \mathcal{T}_2 is a subtree of \mathcal{T}_1 , then $C_{\mathcal{T}_2} \subseteq C_{\mathcal{T}_1}$.*

Proof. We prove the lemma by induction. The base case holds since in round 1, we only have a single tree T whose associated cluster $C_{\mathcal{T}} = V(G)$. Let \mathcal{T}' be a connected component of uncovered bags at round i , and $\mathcal{T}_1, \mathcal{T}_2$ be two different components of uncovered bags that are subtrees of \mathcal{T}' in round $i + 1$. It suffices to show that $C_{\mathcal{T}_i} \subseteq C_{\mathcal{T}'}$ for any $i = 1, 2$, and $C_{\mathcal{T}_1} \cap C_{\mathcal{T}_2} = \emptyset$.

Observe that in a round i , the attachment of a bag X either shrinks (due to the removal in line 15) or grows by the addition of the cores. As any core R_1 created from \mathcal{T}' is a subset of the cluster $C_{\mathcal{T}'}$, $A(X)$ remains a subset of $C_{\mathcal{T}'}$ in both cases. This means all attachments of bags in \mathcal{T}_1 are subsets of $C_{\mathcal{T}'}$, implying that $C_{\mathcal{T}_1} \subseteq C_{\mathcal{T}'}$. The same argument gives $C_{\mathcal{T}_2} \subseteq C_{\mathcal{T}'}$.

Let B_1 and B_2 be the root bag of \mathcal{T}_1 and \mathcal{T}_2 , respectively. Assume w.l.o.g that B_1 is an ancestor of B_2 . Let X be the (only) leaf of B_1 that is the ancestor of B_2 . By item 4 of Claim 4, only the attachment of X could possibly have a non-empty intersection $C_{\mathcal{T}_2}$. Suppose that there exists a vertex $v \in A(X) \cap C_{\mathcal{T}_2}$. Let B be the child bag of X that is an ancestor of B_2 . Then $B \in \mathcal{T}'$ and B the center bag of a core R created in round i . Observe that $v \in R$. When R was created, all the uncovered vertices in \mathcal{T}' in R were marked covered. This means $R \cap \text{Uncov}(V[\mathcal{T}_2]) = \emptyset$. Thus, $v \in R \cap A(Y)$ for some bag $Y \in \mathcal{T}'$ in round $i + 1$. However, by line 15, v will be removed from

$A(Y)$ in round i and hence will not be present in $A(Y)$ in round $i + 1$, a contradiction. Thus, $C_{\mathcal{T}_1} \cap C_{\mathcal{T}_2} = \emptyset$. \square

The Semi-tree-ordered Net. We now construct a tree ordering of the vertices using \mathcal{R} . We define the rank of a core R , denoted by $\text{rank}(R)$, to be the round number when R is constructed. That is, $\text{rank}(R) = i$ if R was constructed in the i -th round. Lastly, for a vertex v , we define $Q(v)$ to be the center bag of the *smallest-rank core* containing v . We observe that:

Observation 3. *The core R covering v for the first time is the smallest-rank core containing v .*

We now define the tree ordering of vertices, and the tree-ordered net as follows.

TREEORDERING

- Let T be the rooted tree that is *isomorphic* to the tree partition \mathcal{T} of G ; each node $x \in T$ corresponds to a bag B_x in \mathcal{T} .
- The map $\varphi : V \rightarrow V(T)$ maps each vertex v to a node $\varphi(v)$ such that its corresponding bag $B_{\varphi(v)}$ is exactly $Q(v)$. This map naturally induces a partial ordering of vertices in V : $u \leq v$ if and only if $Q(u)$ is an ancestor of $Q(v)$ in the tree partition \mathcal{T} .
- The tree-ordered net N is the union of all the centers of the cores in \mathcal{R} . Recall that the centers of a core is the set of uncovered vertices of the center bag in line 11.

The semi-tree-ordered net for G is (N, T, φ) .

The following lemma follows directly from the description of the algorithm and the discussion above.

Lemma 8. *Given the tree partition \mathcal{T} of G , the semi-tree-ordered net (N, T, φ) can be constructed in polynomial time.*

We can see that the tree ordered induced by T and φ is a semi-tree-order since it does not have the anti-symmetric property: there could be two different vertices u and v such that $Q(u) = Q(v)$ and hence they will be mapped to the same vertex of T via φ .

4.2 The Analysis

We now show all properties of the tree-ordered net (N, T, φ) as stated in Lemma 1. Specifically, we will show the covering property in Lemma 9 and the packing property in Lemma 14. We start with the covering property.

Lemma 9. *Let (N, T, φ) be the tree-ordered net defined in TREEORDERING. For every vertex $v \in V$ there is an $x \in N_{v \leq}$ such that $d_{G[V_{\leq x}]}(v, x) \leq \Delta$.*

Proof. By item 1 of Claim 4, v is in at least one core of \mathcal{R} . Let $R \in \mathcal{R}$ be the core of the smallest rank containing v . Let B be the center bag of R . Note that $Q(v) = B$ by definition.

Let $U \subseteq R$ be the set of uncovered vertices in R when R is created; all vertices in U will be marked as covered in line 12 after R is created. We observe that:

1. All vertices in U are equivalent under the tree ordering \leq . This is because $Q(u) = B$ for every $u \in U$ by Observation 3 and definition of $Q(\cdot)$.
2. For any two vertices $z \in R \setminus U$ and $u \in U$, $z < u$. To see this, observe that, by item 4 in Claim 4, z is in the attachment of some bag Y , which is the descendant of the center bag B . By item 2 and the definition of $Q(\cdot)$, $Q(z)$ is a (proper) descendant of Y . Therefore, $Q(z)$ is a descendant of B , which is $Q(u)$, implying that $z < u$.

Let H be the support graph (in line 10) of R . By the construction of R (line 11), $d_{H[R]}(v, U) \leq \Delta$. This means there exists an $x \in U$ such that $d_{H[R]}(v, x) \leq \Delta$. The two observations above imply that $H[R] \subseteq G[V_{\leq x}]$. Thus, $d_{G[V_{\leq x}]}(v, x) \leq \Delta$ as desired. \square

The packing property is substantially more difficult to prove. Our argument goes roughly as follows. First, we show that cores satisfy various properties, and one of them is that, for every bag B in \mathcal{T} , there are at most $O(\text{tp}^2)$ cores intersecting B , and moreover, for every vertex v , there are at most $O(\text{tp})$ cores that contain v (Corollary 1). This allows us to bound $|N_{v \leq}^{2\Delta}|$ via bounding the number of cores that contain at least one vertex in $N_{v \leq}^{2\Delta}$. This set of cores can be partitioned into two sets: those that contain v —there are only $O(\text{tp})$ of them—and those that do not contain v . To bound the size of the latter set, we basically construct a sequence of cores of strictly increasing ranks $R_1^*, R_2^*, \dots, R_\ell^*$ for $\ell \leq \text{tp}$, and for each R_j^* , show that there are only $O(\text{tp}^2)$ ancestral cores of R_j^* that are not ancestors of lower ranked cores in the sequence. This implies a bound of $O(\text{tp}^3)$ on the set of cores, giving the packing property.

We begin by analyzing several properties of the cores.

Lemma 10. *The cores of the same rank are vertex-disjoint.*

Proof. Let R_1 and R_2 be two cores of the same rank, say r . Let B_k be the center bags of R_k for $k = 1, 2$. If B_1 and B_2 belong to two different connected components of uncovered bags in round r , then by Lemma 7, $R_1 \cap R_2 = \emptyset$. Thus, we only consider the complementary case where B_1 and B_2 belong to the same connected component.

Suppose for contradiction that there exists $v \in R_1 \cap R_2$. By item 3 in Claim 4, R_k only contains vertices in descendant bags of B_k . As $R_1 \cap R_2 \neq \emptyset$, either B_1 is an ancestor of B_2 or B_2 is ancestor of B_1 . W.l.o.g., we assume that B_1 is an ancestor of B_2 . Let Y be the bag containing v . Then Y is the descendant of both B_1 and B_2 . As endpoints of edges of G are either in the same bag or in two adjacent bags of the tree partition, $R_1 \cap Y \neq \emptyset$ implies that $R_1 \cap B_2 \neq \emptyset$. Thus, B_2 will be marked as visited in line 13 of the algorithm, and hence R_2 will not have the same rank as R_1 , a contradiction. \square

Lemma 11. *The algorithm has at most tp iterations. Therefore, $\text{rank}(R) \leq \text{tp}$ for every $R \in \mathcal{R}$.*

Proof. Let B' be any bag of \mathcal{T} . We claim that in every iteration, at least one uncovered vertex of B' gets covered. Thus, after tp iterations, every vertex of B' is covered, and hence the algorithm will terminate.

Consider an arbitrary iteration where B' remains uncovered; that is, at least some vertex of B' is uncovered. Let \mathcal{T}' be the connected component of uncovered bags of \mathcal{T} containing B' . If B' got picked in line 11, then all uncovered vertices of B' are marked as covered (in line 12), and hence the claim holds. Otherwise, B' is marked visited in line 13 when a core $R \in \mathcal{R}$ is created. Let H be the support graph of R . By item 2 in Claim 4, the attachment of every bag in \mathcal{T}' does not contain

any vertex of B' . Thus, H only contains uncovered vertices of B' . As $R \cap B' \neq \emptyset$, R contains at least one uncovered vertex of B' , which will be marked as covered in line 12; the claim holds. \square

We obtain the following corollary of Lemma 10 and Lemma 11.

Corollary 1. *Each vertex is contained in at most tp cores. Furthermore, the number of cores intersecting any bag is at most tp^2 .*

Proof. By Lemma 10, cores of the same rank are vertex-disjoint. Thus, each vertex belongs to at most one core of a given rank. As there are tp different ranks by Lemma 11, each vertex belongs to at most tp cores.

Let B be any bag in \mathcal{T} . As the cores of the same rank are vertex-disjoint by Lemma 10, there is at most tp cores of a given rank intersecting B . As there are at most tp different ranks by Lemma 11, the total number of cores intersecting any bag is at most tp^2 . \square

We define the *rank* of a vertex v , denoted by $\text{rank}(v)$, to be the lowest rank among all the cores containing it: $\text{rank}(v) = \min_{R \in \mathcal{R} \wedge v \in R} \text{rank}(R)$. The following also follows from the algorithm.

Lemma 12. *Let B be the center bag of a core R , and U be its center. If R has rank i , then every vertex $v \in B \setminus U$ has rank strictly smaller than i .*

Proof. When R is constructed (in line 11), all uncovered vertices of B are in U and will be marked as covered afterward. Thus, before R is constructed, vertices in $B \setminus U$ must be marked covered, and furthermore, B is not marked as visited. This means vertices in $B \setminus U$ are marked in previous rounds, and hence have ranks strictly smaller than i . \square

Next, we introduce central concepts in the proof of the packing property. We say that a core R_1 is an ancestor (descendant resp.) of R_2 if the center-bag of R_1 is an ancestor (descendant resp.) of R_2 . For any bag of \mathcal{T} , we define its *level* to be the hop distance from the root in \mathcal{T} . Also, we define the *graph rooted at B* to be the subgraph of G induced by the subtree of \mathcal{T} rooted at B . For any core R , we define a graph $G_{\mathcal{T}}[R]$ to be the subgraph of G induced by vertices in the bags of the subtrees of \mathcal{T} rooted at the center-bag of R ; see Figure 3(a).

Let $\mathcal{A}^{\prec}(R)$ be the set of cores that are ancestors of R and have rank strictly less than R . We define the *shadow domain* of a core R to be the graph obtained from $G_{\mathcal{T}}[R]$ by removing the vertices that are contained in at least one core in $\mathcal{A}^{\prec}(R)$. The *shadow* of a core R , denoted by $\text{shadow}(R)$, is defined as the ball of radius Δ centered around R in the shadow domain of R . The *strict shadow* of R is defined as its shadow minus itself, i.e., $\text{shadow}(R) \setminus R$. To bound the size of $N_{v \leq \Delta}^{2\Delta}$, for a vertex v we are interested in the cores and their shadows where v is located. See Figure 3(b).

We observe the following properties of ranks and shadow domains.

Lemma 13. *Let R_1 and R_2 be two cores such that $\text{rank}(R_1) \geq \text{rank}(R_2)$ and R_1 is an ancestor of R_2 . For each $i \in \{1, 2\}$, let D_i , B_i , and Y_i be the shadow domain, the center bag, and the center of R_i , respectively. (Note that $Y_i \subseteq B_i$.)*

1. *If $V(D_1) \cap (B_2 \setminus Y_2) \neq \emptyset$, then there exists a core R_3 that is an ancestor of R_2 and descendant of R_1 such that $\text{rank}(R_3) < \text{rank}(R_2)$.*

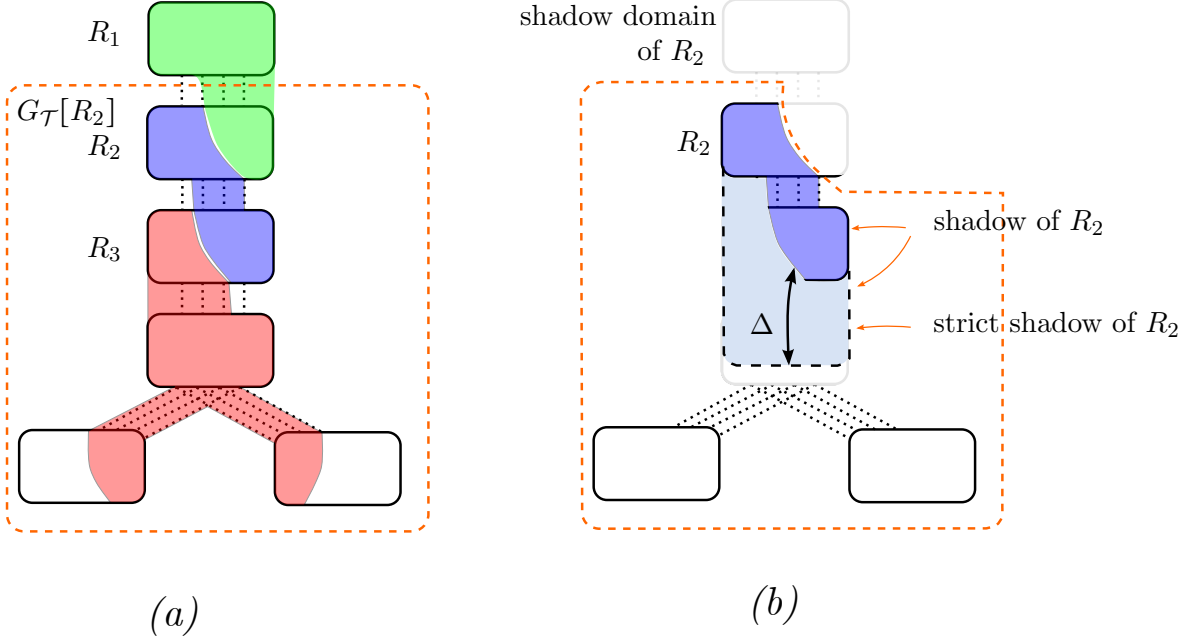


Figure 3: (a) Core R_2 and graph $G_{\mathcal{T}}[R_2]$ rooted at its center bag, (b) the shadow domain, shadow and strict shadow of core R_2 .

<ul style="list-style-type: none"> ◦ Rank of R: the round number when R is constructed. ◦ Central bag of R - the bag from line 9, from which we grew the core. ◦ Core R_1 is an ancestor of R_2 if the center-bag of R_1 is an ancestor of R_2 (independent of round). ◦ The level of a bag B is the hop distance from the root in \mathcal{T}. ◦ The graph rooted at B is the subgraph of G induced by the subtree of \mathcal{T} rooted at B. 	<ul style="list-style-type: none"> ◦ $G_{\mathcal{T}}[R]$ denotes to be the graph rooted at the center-bag of R (independent of rank). ◦ $\mathcal{A}^{\leftarrow}(R)$ - cores that are ancestors of R and have rank strictly less than R. ◦ Shadow domain of R is the graph obtained from $G_{\mathcal{T}}[R]$ by removing the vertices that are contained in at least one core in $\mathcal{A}^{\leftarrow}(R)$. ◦ shadow($R$): Shadow of R is the ball of radius Δ centered around R in the shadow domain of R. ◦ Strict shadow of R: the shadow minus R: $\text{shadow}(R) \setminus R$.
--	--

Figure 4: Key definitions used during the proof of Lemma 8.

2. If $\text{rank}(R_1) > \text{rank}(R_2)$ and there are no cores of rank smaller than R_2 whose center bag is in the path between B_1 and B_2 (in \mathcal{T}), then $V(G_{\mathcal{T}}[R_2]) \cap V(D_1) \subseteq V(D_2)$. In other words, every vertex in the shadow domain of R_1 in $G_{\mathcal{T}}[R_2]$ is in the shadow domain of R_2 .

Proof. We first show item 1. Let x be a vertex in $V(D_1) \cap (B_2 \setminus Y_2)$. By Lemma 12, $\text{rank}(x) < \text{rank}(R_2)$, implying that x is contained in a core R_3 such that $\text{rank}(R_3) < \text{rank}(R_2)$. Furthermore, R_3 is a descendant of R_1 as otherwise, by the definition of shadow domain, $R_3 \cap V(D_1) = \emptyset$ and hence $x \notin V(D_1)$, a contradiction. Also, R_3 is an ancestor of R_2 as the center-bag of R_2 contains a vertex of R_3 , which is x .

We show item 2 by contrapositive. Let x be a vertex in $G_{\mathcal{T}}[R_2]$ such that $x \notin V(D_2)$. We show

that $x \notin V(D_1)$. Since x is not in $V(D_2)$, there exists an ancestor core R_3 of R_2 such that $x \in R_3$ and $\text{rank}(R_3) < \text{rank}(R_2)$. By the assumption in item 2, R_3 is also an ancestor of R_1 and hence $R_3 \in \mathcal{A}^<(R_1)$. Thus, $x \notin V(D_1)$ as claimed. \square

Equipped with the lemmas above, we are finally ready to prove the packing property with $\alpha = 2$.

Lemma 14. *Let (N, T, φ) be the tree-ordered net defined in TREEORDERING. For every vertex $v \in V$, $|N_{v \leq}^{2\Delta}| \leq \text{tp}^4 + \text{tp}^2$.*

Proof. Consider a vertex $v \in V$. Recall that the net N is the set of centers of all cores in \mathcal{R} . Let $\mathfrak{C}_{\text{a11}}$ be the set of cores that have a non-empty intersection with the vertices in $N_{v \leq}^{2\Delta}$. Instead of bounding $|N_{v \leq}^{2\Delta}|$ directly, we bound $|\mathfrak{C}_{\text{a11}}|$. Note that every vertex in $N_{v \leq}^{2\Delta}$ is in the center of one of the cores in $\mathfrak{C}_{\text{a11}}$.

Let $\mathfrak{C}^\Delta \subseteq \mathfrak{C}_{\text{a11}}$ be the cores that contain v and let $\mathfrak{C} := \mathfrak{C}_{\text{a11}} \setminus \mathfrak{C}^\Delta$. By Corollary 1, $|\mathfrak{C}^\Delta| \leq \text{tp}$. The bulk of our proof below is to show that $|\mathfrak{C}| \leq \text{tp}^3$. Since each core has at most tp centers, we have $|N_{v \leq}^{2\Delta}| \leq \text{tp} \cdot |\mathfrak{C}_{\text{a11}}| \leq \text{tp}^4 + \text{tp}^2$, as claimed.

We now focus on proving that $|\mathfrak{C}| \leq \text{tp}^3$. Let B_v be the bag containing v . Let \mathcal{P} be the path in \mathcal{T} from the root bag to B_v . We claim that every core in \mathfrak{C} has a center bag on \mathcal{P} . This is because every core R in \mathfrak{C} contains a vertex, say x , in $N_{v \leq}^{2\Delta}$, whose bag is an ancestor of B_v . By construction in line 11, every vertex in R is in a descendant of the center bag of R . This means B_v is a descendant of the center bag of R , implying the claim.

The rest of our proof goes as follows:

- Let r_1 be the lowest rank among all cores in \mathfrak{C} . We will show in Claim 5 below that there is only one core in \mathfrak{C} having rank r_1 . Let this unique core in \mathfrak{C} with rank r_1 be R_1^* . Let \mathfrak{C}_1 be the set of all cores in \mathfrak{C} that are ancestors of R_1^* , including R_1^* . We then show in the next Claim 6 that each core in \mathfrak{C}_1 intersects the center of R_1^* . Since the center is contained in a bag, Corollary 1 implies that $|\mathfrak{C}_1| \leq \text{tp}^2$.
- Next, let $\bar{\mathfrak{C}}_1 = \mathfrak{C} \setminus \mathfrak{C}_1$. If $\bar{\mathfrak{C}}_1$ is non-empty, we define r_2 be the lowest rank among cores in $\bar{\mathfrak{C}}_1$. Note that $r_2 > r_1$. Claim 5 below again implies that there is only one core in $\bar{\mathfrak{C}}_1$ having rank r_2 . Let this unique core in $\bar{\mathfrak{C}}_1$ with rank r_2 be R_2^* . Let \mathfrak{C}_2 be the set of all ancestor cores R_2^* in $\bar{\mathfrak{C}}_1$ including R_2^* . Then Claim 6 implies that each core in \mathfrak{C}_2 intersects the center of R_2^* . Since the center is contained in a bag, Corollary 1 gives that $|\mathfrak{C}_2| \leq \text{tp}^2$.
- Inductively, we define the sequence of sets of cores $\mathfrak{C}_1, \mathfrak{C}_2, \mathfrak{C}_3, \dots, \mathfrak{C}_\ell$ and $r_1 < r_2 < \dots < r_\ell$ until $\bar{\mathfrak{C}}_\ell = \bar{\mathfrak{C}}_{\ell-1} \setminus \mathfrak{C}_\ell$ is empty. Here, r_j is defined as the lowest rank among cores in $\bar{\mathfrak{C}}_{j-1}$, and $\bar{\mathfrak{C}}_0 = \mathfrak{C}$. For each $j \in [\ell]$, Claim 5 implies that there exist a unique core R_j^* in $\bar{\mathfrak{C}}_{j-1}$ with rank r_j . Let \mathfrak{C}_j be the set of all cores in $\bar{\mathfrak{C}}_j$ that are ancestors of R_j^* , including R_j^* . Claim 6 then implies that all cores in \mathfrak{C}_j intersect the center of R_j^* in. Thus, $|\mathfrak{C}_j| \leq \text{tp}^2$ for each $j \in [\ell]$ by Corollary 1. By Lemma 11, $r_1 < r_2 < \dots < r_\ell \leq \text{tp}$, implying that $|\mathfrak{C}| \leq \text{tp}^3$ as desired.

For the rest of the proof, we prove two claims.

Claim 5. *There is only one core in $\bar{\mathfrak{C}}_{j-1}$ having rank r_j for each $j \in [\ell]$.*

Proof. Suppose otherwise. Then there are two cores R_1 and R_2 in $\bar{\mathfrak{C}}_{j-1}$ having rank r_j . Assume w.l.o.g. that R_1 is an ancestor of R_2 . By Lemma 10, R_1 and R_2 are disjoint as they have the same rank. Also, since $R_1, R_2 \in \mathfrak{C}$, the center-bags of R_1 and R_2 lie on the path \mathcal{P} .

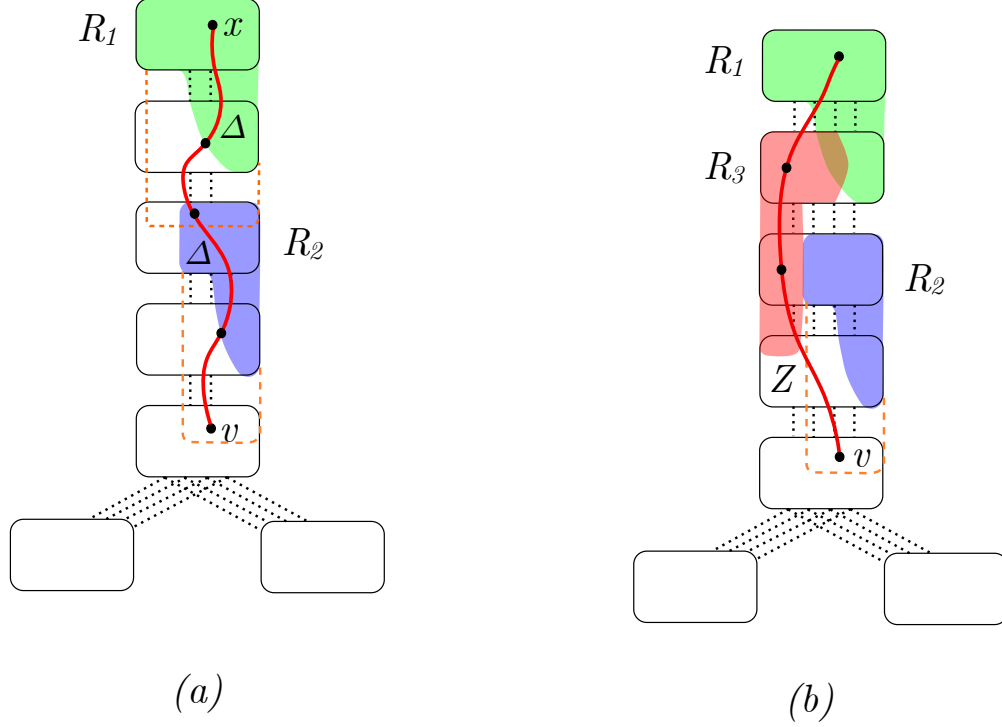


Figure 5: Illustration for the proof of Claim 5.

Since R_1 contains a point of $N_{v \leq}^{2\Delta}$, v is in the shadow of R_1 . Thus, there is a path Z of length at most 2Δ that goes from the center of R_1 to v in the shadow domain of R_1 . This path Z has to intersect the center bag of R_2 to get to v .

Suppose this intersection occurs at a vertex in the center of R_2 ; see Figure 5(a). Then the path Z goes from the center of R_1 to outside R_1 and then into the center of R_2 and then to outside of R_2 . It has to go outside of R_2 as v is in the strict shadow of R_2 . Also, it has to go outside of R_1 before entering R_2 as R_1 and R_2 are vertex disjoint (however, it is possible that there is an edge from R_1 to R_2). This means Z has a length of more than $\Delta + \Delta = 2\Delta$, a contradiction.

Now, suppose this intersection occurs at a vertex in the center-bag of R_2 that is not in the center of R_2 ; see Figure 5(b). Then by item 1 (of Lemma 13) it follows that there is at least one core that is ancestor of R_2 and descendant of R_1 , and having rank lower than r_j . Let R_3 be the one among such cores whose center-bag has the smallest level.

We claim that the path Z intersects the center of R_3 . Suppose otherwise. However, the path has to intersect the center-bag of R_3 to get to v . Then, by item 1 (of Lemma 13) it follows that there is a core that is ancestor of R_3 and descendant of R_1 , and having rank lower than r_j , a contradiction to the selection of R_3 .

By definition of shadow domain, R_3 is disjoint from the shadow domain of R_2 . Hence v is not in R_3 . However, v is contained in the shadow of R_3 as the part of the path Z from the center of R_3 to v is contained in the shadow domain of R_3 and has length at most 2Δ . Thus, v is contained in the strict shadow of R_3 . Since R_3 is a descendant of R_1 , we have that $R_3 \notin \mathfrak{C}_1 \cup \mathfrak{C}_2 \cup \dots \cup \mathfrak{C}_{j-1}$ and hence $R_3 \in \bar{\mathfrak{C}}_{j-1}$. Thus, there is a core in $\bar{\mathfrak{C}}_{j-1}$ that has rank lower than r_j , a contradiction to the choice of r_j . \square

Claim 6. *Each core in \mathfrak{C}_j intersects the center of R_j^* for each $j \in [\ell]$.*

Proof. Suppose this is not true, and let j be the minimum index for which the claim does not hold. Then there exists a core $R \in \mathfrak{C}_j$ that is disjoint from the center of R_j^* . Note that R_j^* is a descendant of R by definition of \mathfrak{C}_j and the rank of R_j^* is strictly less than the rank of R by the definition of R_j^* . Since v is in the shadow of R there is a path Z of length at most 2Δ that goes from the center of R to v in the shadow domain of R . This path Z has to intersect the center-bag of R_j^* to get to v .

Suppose this intersection occurs at a vertex in the center of R_j^* . Then the path Z goes from the center of R outside R and then into the center of R_j^* and then to outside of R_j^* . It has to go outside of R_j^* as v is in the strict shadow of R_j^* . Also, it has to go outside of R before entering R_j^* as R and the center of R_j^* are vertex disjoint (this is what we assumed for the sake of contradiction). This means Z has a length of more than $\Delta + \Delta = 2\Delta$, a contradiction.

Now, suppose this intersection occurs at a vertex in the center-bag of R_j^* that is not in the center of R_j^* . Then, by item 1 (of Lemma 13), there exist at least one core that is an ancestor of R_j^* and a descendant of R and having rank lower than r_j . Let R' be the one among such cores whose center-bag has the smallest level. The path Z has to intersect the center of R' as otherwise there is a core that is an ancestor of R' and a descendant of R and having rank lower than r_j , contradicting the selection of R' .

Note that v is not in R' as R' is disjoint from the shadow domain of R_j^* by definition of shadow domain. However, v is contained in the shadow of R' as the part of the path Z from center of R' to v is contained in the shadow domain of R' (by item 2 of Lemma 13) and has length at most 2Δ . Thus, v is contained in the strict shadow of R' . Since j is the minimum index for which the statement of the claim does not hold, if $R' \in \mathfrak{C}_1 \cup \dots \cup \mathfrak{C}_{j-1}$, then $R' \notin \mathfrak{C}_j$. This implies that $R' \in \mathfrak{C}_j$. Thus we have a core in \mathfrak{C}_j having rank strictly smaller than r_j , a contradiction. \square

The discussion in the proof outline with Claim 5 and Claim 6 proves the lemma. \square

5 Applications

In this section, we gave a more detailed exposition of the applications of Theorem 1 and Theorem 2 mentioned in Section 1. The list of applications here is not meant to be exhaustive, and we believe that our result will find further applications.

5.1 Flow Sparsifier

Given an edge-capacitated graph $G = (V, E, c)$ and a set $K \subseteq V$ of terminals, a K -flow is a flow where all the endpoints are terminals. A flow-sparsifier with quality $\rho \geq 1$ is another capacitated graph $H = (K, E_H, c_H)$ such that (a) any feasible K -flow in G can be feasibly routed in H , and (b) any feasible K -flow in H can be routed in G with congestion ρ (see [EGK⁺14] for formal definitions).

Englert *et al.* [EGK⁺14] showed that given a graph G which admits a $(\beta, \frac{1}{\beta})$ -padded decomposition scheme, then for any subset K , one can efficiently compute a flow-sparsifier with quality $O(\beta)$. Using their result, we obtain the following corollary of Theorem 1.

Corollary 2. *Given an edge-capacitated graph $G = (V, E, c)$ with treewidth tw , and a subset of terminals $K \subseteq V$, one can efficiently compute a flow sparsifier with quality $O(\log \text{tw})$.*

The best previously known result had quality $O(\text{tw})$ approximation [EGK⁺14, AGG⁺19]. Thus, our result improves the dependency on tw exponentially. For further reading on flow sparsifiers, see [Moi09, MT10, MM10, CLLM10, Chu12, AGK14].

We note that, though the graph G has treewidth tw , the flow-sparsifier H in Corollary 2 can have arbitrarily large treewidth. Having low treewidth is a very desirable property of a graph, and naturally, we would like to have a sparsifier of small treewidth. A flow-sparsifier $H = (K, E_H, c_H)$ of G is called *minor-based* if H is a minor of G . That is, H can be obtained from G by deleting/contracting edges, and deleting vertices. Englert *et al.* [EGK⁺14] showed that given a graph G which admits a $(\beta, \frac{1}{\beta})$ -padded decomposition scheme, then for any subset K , one can efficiently compute a minor-based flow-sparsifier with quality $O(\beta \log \beta)$. By Theorem 1, we obtain:

Corollary 3. *Given an edge-capacitated graph $G = (V, E, c)$ with treewidth tw , and a subset of terminals $K \subseteq V$, one can efficiently compute a minor-based flow-sparsifier $H = (K, E_H, c_H)$ with quality $O(\log \text{tw} \log \log \text{tw})$. In particular, H also has treewidth tw .*

5.2 Sparse Partition

Given a metric space (X, d_X) , a (α, τ, Δ) -sparse partition is a partition \mathcal{C} of X such that:

- **Low Diameter:** $\forall C \in \mathcal{C}$, the set X has diameter at most Δ ;
- **Sparsity:** $\forall x \in X$, the ball $B_X(x, \frac{\Delta}{\alpha})$ intersects at most τ clusters from \mathcal{C} .

We say that the metric (X, d_X) admits a (α, τ) -sparse partition scheme if for every $\Delta > 0$, X admits a (α, τ, Δ) -sparse partition.

Jia *et al.* [JLN⁺05] implicitly proved (see [Fil20] for an explicit proof) that if a space admits a (β, s) -sparse cover scheme, then it admits a (β, s) -sparse partition scheme. Therefore, by g Theorem 2 we obtain:

Corollary 4. *Every graph G with treewidth tw admits a $(O(1), \text{poly}(\text{tw}))$ -sparse partition scheme.*

Previously, it was only known that graphs with treewidth tw admit $(O(\text{tw}^2), 2^{\text{tw}})$ -sparse partition scheme [FT03, Fil20]. Thus, Corollary 4 implies an exponential improvement in the dependency on tw . For further reading on sparse partitions, see [JLN⁺05, BDR⁺12, Fil20, CJF⁺23, Fil24].

5.3 Universal Steiner Tree and Universal TSP

We consider the problem of designing a network that allows a server to broadcast a message to a single set of clients. If sending a message over a link incurs some cost, then designing the best broadcast network is classically modeled as the Steiner tree problem [HR92]. However, if the server has to solve this problem repeatedly with different client sets, it is desirable to construct a single network that will optimize the cost of the broadcast for every possible subset of clients. This setting motivates the *Universal Steiner Tree* (UST) problem.

Given a metric space (X, d_X) and root $r \in X$, a ρ -approximate UST is a weighted tree T over X such that for every $S \subseteq X$ containing r , we have

$$w(T\{S\}) \leq \rho \cdot \text{OPT}_S$$

where $T\{S\} \subseteq T$ is the minimal subtree of T connecting S , and OPT_S is the minimum weight Steiner tree connecting S in X .

A closely related problem to UST is the *Universal Traveling Salesman Problem* (UTSP). Consider a postman providing post service for a set X of clients with n different locations (with distance measure d_X). Each morning, the postman receives a subset $S \subseteq X$ of the required deliveries for the day. In order to minimize the total tour length, one solution may be to compute each morning an (approximation of an) Optimal TSP tour for the set S . An alternative solution will be to compute a *Universal TSP* (UTSP) tour R containing all the points X . Given a subset S , $R\{S\}$ is the tour visiting all the points in S w.r.t. the order induced by R . Given a tour T , denote its length by $|T|$. The *stretch* of R is the maximum ratio among all subsets $S \subseteq X$ between the length of $R\{S\}$ and the length of the optimal TSP tour on S , $\max_{S \subseteq X} \frac{|R\{S\}|}{|\text{Opt}(S)|}$.

Jia *et al.* [JLN⁺05] showed that for every n -point metric space that admits (σ, τ) -sparse partition scheme, there is a polynomial time algorithm that given a root $\text{rt} \in V$ computes a UST with stretch $O(\tau\sigma^2 \log_\tau n)$. In addition, Jia *et al.* [JLN⁺05] also showed that such a metric admit a UTSP with stretch $O(\tau\sigma^2 \log_\tau n)$. Using our Corollary 4, we conclude:

Corollary 5. *Consider an n -point graph with treewidth tw . Then its shortest path metric admits a solution to both universal Steiner tree and universal TSP with stretch $\text{poly}(\text{tw}) \cdot \log n$.*

The best-known previous result for both problems had a stretch of $\exp(\text{tw}) \cdot \log n$ [Fil20], which is exponentially larger than ours in terms of the tw dependency. For further reading on the UTSP and TSP problems, see [PBI89, BG89, JLN⁺05, GHR06, HKL06, SS08, GKSS10, BCK11, BDR⁺12, BLT14, Fil20, Fil24]. Interestingly, our solution to the UST problem in Corollary 5 produces a solution which is not a subgraph of the low treewidth input graph G . If one requires that the UST will be a subgraph of G , the current state of the art is by Busch *et al.* [BCF⁺23] who obtained stretch $O(\log^7 n)$. To date, this is also the best known upper bound for graph with bounded treewidth.

5.4 Steiner Point Removal

Given a graph $G = (V, E, w)$, and a subset of terminals $K \subseteq V$, in the Steiner point removal problem, we are looking for a graph $H = (K, E_H, w_H)$, which is a minor of G . We say that H has stretch t , if for every $u, v \in K$, $d_G(u, v) \leq d_H(u, v) \leq t \cdot d_G(u, v)$. Englert *et al.* [EGK⁺14] showed that given a graph G which admits a $(\beta, \frac{1}{\beta})$ -padded decomposition scheme, one can compute a distribution \mathcal{D} over minors $H = (K, E_H, w_H)$, such that for every $u, v \in K$, and $H \in \text{supp}(\mathcal{D})$, $d_G(u, v) \leq d_H(u, v)$, and $\mathbb{E}_{H \sim \mathcal{D}}[d_H(u, v)] \leq O(\beta \cdot \log \beta) \cdot d_G(u, v)$. That is, the minor has expected stretch $O(\beta \log(\beta))$. Thus, by Theorem 1, we obtain:

Corollary 6. *Given a weighted graph $G = (V, E, w)$ with treewidth tw , and a subset of terminals $K \subseteq V$, one can efficiently sample a minor $H = (K, E_H, c_H)$ of G such that for every $u, v \in K$, $d_G(u, v) \leq d_H(u, v)$, and $\mathbb{E}[d_H(u, v)] \leq O(\log \text{tw} \log \log \text{tw}) \cdot d_G(u, v)$.*

The best previously known result had expected stretch $O(\text{tw} \log \text{tw})$ [EGK⁺14, AGG⁺19]. For the classic Steiner point removal, where we are looking for a single minor, the stretch is $2^{O(\text{tw} \log \text{tw})}$ [CCL⁺23]. For further reading on the Steiner point removal problem, see [Gup01, CXKR06, KKN15, Che18, Fil19b, FKT19, Fil20, HL22, CCL⁺23].

5.5 Zero Extension

In the 0-Extension problem, the input is a set X , a terminal set $K \subseteq X$, a metric d_K over the terminals, and an arbitrary cost function $c: \binom{X}{2} \rightarrow \mathbb{R}_+$. The goal is to find a *retraction* $f: X \rightarrow K$ that minimizes $\sum_{\{x,y\} \in \binom{X}{2}} c(x,y) \cdot d_K(f(x), f(y))$. A retraction is a surjective function $f: X \rightarrow K$ that satisfies $f(x) = x$ for all $x \in K$. The 0-Extension problem, first proposed by Karzanov [Kar98], generalizes the Multiway Cut problem [DJP⁺92] by allowing d_K to be any discrete metric (instead of a uniform metric).

For the case where the metric (K, d_K) on the terminals admits a $(\beta, \frac{1}{\beta})$ -padded-decomposition, Lee and Naor [LN05] (see also [AFH⁺04, CKR04]) showed an $O(\beta)$ upper bound. By Theorem 1, we get:

Corollary 7. *Consider an instance of the 0-Extension problem $(K \subseteq X, d_K, c: \binom{X}{2} \rightarrow \mathbb{R}_+)$, where the metric (K, d_K) is a sub-metric of a shortest path metric of a graph with treewidth tw . Then, one can efficiently find a solution with cost at most $O(\log \text{tw})$ times the cost of the optimal. In particular, there is a $O(\log \text{tw})$ -approximation algorithm for the multi-way cut problem for graphs of treewidth tw .*

The best previously known result was $O(\text{tw})$ approximation [CKR04, AGG⁺19]. For further reading on the 0-Extension problem, see [CKR04, FHRT03, AFH⁺04, LN05, EGK⁺14, FKT19].

5.6 Lipschitz Extension

For a function $f: X \rightarrow Y$ between two metric spaces $(X, d_X), (Y, d_Y)$, set $\|f\|_{\text{Lip}} = \sup_{x,y \in X} \frac{d_Y(f(x), f(y))}{d_X(x,y)}$ to be the Lipschitz parameter of the function. In the Lipschitz extension problem, we are given a map $f: Z \rightarrow Y$ from a subset Z of X . The goal is to extend f to a function \tilde{f} over the entire space X , while minimizing $\|\tilde{f}\|_{\text{Lip}}$ as a function of $\|f\|_{\text{Lip}}$. Lee and Naor [LN05], proved that if a space admits a $(\beta, \frac{1}{\beta})$ -padded decomposition scheme, then given a function f from a subset $Z \subseteq X$ into a closed convex set in some Banach space, then one can extend f into $\tilde{f}: X \rightarrow C$ such that $\|\tilde{f}\|_{\text{Lip}} \leq O(\beta) \cdot \|f\|_{\text{Lip}}$. By Theorem 1, we obtain:

Corollary 8 (Lipschitz Extension). *Consider a graph $G = (V, E, w)$ with treewidth tw , and let $f: V' \rightarrow C$ be a map from a subset $V' \subseteq V$ into C , where C is a convex closed set of some Banach space. Then there is an extension $\tilde{f}: X \rightarrow C$ such that $\|\tilde{f}\|_{\text{Lip}} \leq O(\log \text{tw}) \cdot \|f\|_{\text{Lip}}$.*

5.7 Embedding into ℓ_p spaces

Metric embedding is a map between two metric spaces that preserves all pairwise distances up to a small stretch. We say that embedding $f: X \rightarrow Y$ between (X, d_X) and (Y, d_Y) has distortion t if for every $x, y \in X$ it holds that $d_X(x, y) \leq d_Y(f(x), f(y)) \leq t \cdot d_X(x, y)$. Krauthgamer, Lee, Mendel and Naor [KLMN04] (improving over Rao [Rao99]) showed that every n -point metric space that admits $(\beta, \frac{1}{\beta})$ -padded decomposition scheme can be embedded into an ℓ_p space with distortion $O(\beta^{1-\frac{1}{p}} \cdot (\log n)^{\frac{1}{p}})$. Previously, it was known that the shortest path metric of an n point graph with treewidth tw (or more generally K_{tw} -minor free) embeds into ℓ_p space with distortion $O((\text{tw})^{1-\frac{1}{p}} \cdot (\log n)^{\frac{1}{p}})$ [KLMN04, AGG⁺19]. By Theorem 1, we conclude:

Corollary 9. *Let G be an n -point weighted graph with treewidth tw . Then there exist embedding of G into ℓ_p ($1 \leq p \leq \infty$) with distortion $O((\log \text{tw})^{1-\frac{1}{p}} \cdot (\log n)^{\frac{1}{p}})$.*

Since every finite subset of the Euclidean space ℓ_2 embed isometrically into ℓ_p (for $1 \leq p \leq \infty$), it follows that for $1 \leq p \leq \infty$ such G can be embedded into ℓ_p space (in particular ℓ_1) with distortion $O(\sqrt{\log \text{tw} \cdot \log n})$.

A norm space of special interest is ℓ_∞ , as every finite metric space embeds into ℓ_∞ isometrically (i.e. with distortion 1, this is the so called Fréchet embedding). However the dimension of such embedding is very large: $\Omega(n)$. Krauthgamer *et al.* [KLMN04] proved that every graph with treewidth tw (or more generally K_{tw} -minor free) embeds into ℓ_∞^d with dimension $d = \tilde{O}(3^{\text{tw}}) \cdot \log n$ and distortion $O(\text{tw}^2)$. This was recently improved by Filtser [Fil24] who showed that every graph with treewidth tw (or more generally K_{tw} -minor free) embeds into ℓ_∞^d with dimension $d = \tilde{O}(\text{tw}^2) \cdot \log n$ and distortion $O(\text{tw})$ (alternatively distortion $3+\varepsilon$ and dimension $d = \tilde{O}(\frac{1}{\varepsilon})^{\text{tw}+1} \cdot \log n$). More generally, Filtser [Fil24] proved that if a graph G admits a (β, s) -padded partition cover scheme, then G embeds into ℓ_∞^d with distortion $(1+\varepsilon) \cdot 2\beta$ and dimension $d = O\left(\frac{\tau}{\varepsilon} \cdot \log \frac{\beta}{\varepsilon} \cdot \log\left(\frac{n\beta}{\varepsilon}\right)\right)$. Using our Theorem 3, we conclude:

Corollary 10. *Let G be an n -point weighted graph with treewidth tw . Then there exist embedding of G into ℓ_∞^d with distortion $O(1)$ and dimension $d = \text{poly}(\text{tw}) \cdot \log n$.*

Note that Corollary 10 provides an exponential improvement in the dependence on tw compared to the previous best known embedding into ℓ_∞ with constant distortion. For further reading about metric embedding into ℓ_p spaces see [OS81, Bou85, LLR95, Rao99, KLMN04, GNRS04, LR10, AFGN22, KLR19, Fil19b, Kum22, Fil24].

5.8 Stochastic decomposition for minor-free graphs - a reduction from additive stretch embeddings

A major open question is to determine the padding parameter of K_r -minor-free graphs. The current state of the art is $O(r)$ [AGG⁺19] (see also [Fil19a]), while the natural conjecture is exponentially smaller $O(\log r)$. A somewhat weaker guarantee, we call (t, p, Δ) -stochastic decomposition, is a distribution over partitions with diameter Δ such that every pair of vertices at distance at most $\frac{\Delta}{t}$ is clustered together with probability at least p . Compared to padded decomposition, the guarantee here is over pairs (instead of balls), and there is a threshold distance for the guarantee (instead of a linear dependence as in padded decomposition). Nonetheless, for many applications, a stochastic decomposition is good enough. Moreover, in most cases, the parameters of padded and stochastic decompositions are the same. The only case that we are aware of where they are different is high-dimensional Euclidean spaces. In particular, for minor-free graphs, the best-known result is a $(O(r), \frac{1}{2}, \Delta)$ -stochastic decomposition [AGG⁺19]. Here, we argue that if the parameters of recently studied stochastic embeddings with additive distortion are improved, then our Theorem 1 will imply much better stochastic decompositions for minor-free graphs.

Definition 5 (Stochastic embedding with additive distortion). *Consider an n -vertex weighted graph $G = (V, E, w)$ with diameter D . A stochastic additive embedding of G into a distribution over embeddings f into graphs H with treewidth τ and expected additive distortion $\varepsilon \cdot D$, is a distribution \mathcal{D} over maps f from G into graphs H , such that*

- *Small treewidth: Any graph H in the support has treewidth at most τ .*

- *Dominating:* for every $u, v \in V$, and $(f, H) \in \text{supp}(\mathcal{D})$, $d_G(u, v) \leq d_H(f(u), f(v))$.
- *Expected additive distortion:* for every $u, v \in V$, $\mathbb{E}_{(f, H) \in \text{supp}(\mathcal{D})} [d_H(f(u), f(v))] \leq d_G(u, v) + \varepsilon \cdot D$.

Recently, Filtser and Le [FL22] (see also [FKS19, CFKL20, FL21]) constructed a stochastic embedding of every n -vertex K_r -minor free graph G with diameter D , into a distribution over graphs with treewidth $O(\frac{(\log \log n)^2}{\varepsilon^2}) \cdot \chi(r)$ and expected additive distortion $\varepsilon \cdot D$. Here $\chi(r)$ is an extremely fast-growing function of r (an outcome of the minor-structure theorem [RS03]). If one would construct such stochastic embedding into treewidth $\text{poly}(\frac{r}{\varepsilon})$, then Lemma 15 below will imply stochastic decomposition with parameter $O(\log r)$. Furthermore, even allowing treewidth $\text{poly}(\frac{r \cdot \log n}{\varepsilon})$ will imply stochastic decomposition with parameter $O(\log r + \log \log n)$ which is yet unknown as well.

Lemma 15. *Suppose that there is a coordinate monotone function $\pi(\varepsilon, r, n)$, such that every n -vertex K_r -minor free graph with diameter D admits a stochastic embedding into a distribution over graphs with treewidth $\pi(\varepsilon, r, n)$ and expected additive distortion $\varepsilon \cdot D$. Then every for every $\Delta > 0$, every n -vertex K_r -minor free admits a $(t, \frac{1}{2}, \Delta)$ -stochastic decomposition, for $t = O(\log \pi(\Omega(\frac{1}{r}), r, n))$.*

Proof. Consider an n -vertex, K_r -minor-free graph $G = (V, E, w)$. Fix $\Delta > 0$. We will construct a distribution \mathcal{D} over partitions of G such that every $\mathcal{P} \in \text{supp}(\mathcal{D})$ will be weakly Δ -bounded and for every pair u, v at distance $\frac{\Delta}{t}$, u, v will be clustered together with probability at least $\frac{1}{2}$.

The partition will be constructed as follows:

1. Use Filtser's [Fil19a] strong padded decomposition scheme with diameter parameter $\Delta' = O(\frac{r}{t}) \cdot \Delta$ to sample a partition \mathcal{P}' of V . Note that every cluster $C \in \mathcal{P}'$ is a K_r -minor free graph of diameter Δ' .
2. Fix $\varepsilon = \frac{\Delta}{5 \cdot t \cdot \Delta'} = \Omega(\frac{1}{r})$. For every cluster $C \in \mathcal{P}'$, sample an embedding f_C of $G[C]$ into a graph H_C with treewidth $\pi(\varepsilon, r, |C|) \leq \pi(\Omega(\frac{1}{r}), r, n)$ and expected additive distortion $\varepsilon \cdot \Delta'$. We will abuse notation and use $v \in C$ to denote both v and $f_C(v)$.
3. For every $C \in \mathcal{P}'$, use Theorem 1 to sample a $(O(\log(\pi(\Omega(\frac{1}{r}), r, n))), \Omega(1), \Delta)$ -padded decomposition \mathcal{P}_C of H_C .

The final partition we return is the union of all the created partitions $\mathcal{P} = \cup_{C \in \mathcal{P}'} (\mathcal{P}_C \cap V)$. We claim that \mathcal{P} has a weak diameter at most Δ . Consider a cluster $P \in \mathcal{P}$ and $u, v \in P$. There is some cluster $C \in \mathcal{P}'$ containing both u and v . Then, it holds that $d_G(u, v) \leq d_{G[C]}(u, v) \leq d_{H_C}(u, v) \leq \Delta$.

Consider a pair of vertices u, v at distance $\leq \frac{\Delta}{t}$. Denote by Ψ_1 the event that the ball of radius $\frac{\Delta}{t}$ around v belongs to a single cluster in \mathcal{P}' . Denote by Ψ_2 the event that conditioned on Ψ_1 occurring, $d_{H_C}(u, v) \leq d_G(u, v) + 5 \cdot \varepsilon \cdot \Delta'$. Denote by Ψ_3 the event that conditioned on Ψ_2 occurring, u and v belong to the same cluster of \mathcal{P}_C .

Let C denote the cluster containing v in \mathcal{P}' . By choosing the constant in Δ' large enough, it holds that

$$\Pr[\Psi_1] = \Pr\left[\mathbf{B}_G(v, \frac{\Delta}{t}) \subseteq C\right] = \Pr\left[\mathbf{B}_G(v, \frac{\Delta}{t \Delta'} \cdot \Delta') \subseteq C\right] \geq e^{-O(r) \cdot \frac{\Delta}{t \Delta'}} \geq \frac{4}{5}.$$

Note that if $\mathbf{B}_G(v, \frac{\Delta}{t}) \subseteq C$, that the entire shortest path from u to v is contained in C , and hence $d_{G[C]}(u, v) = d_G(u, v)$. Using Markov, the probability that the additive distortion by f_C is too large is bounded by:

$$\begin{aligned} \Pr[\overline{\Psi}_2 \mid \Psi_1] &= \Pr[d_{H_C}(u, v) > d_{G[C]}(u, v) + 5 \cdot \epsilon \cdot \Delta'] \\ &\leq \frac{\mathbb{E}[d_{H_C}(u, v) - d_{G[C]}(u, v)]}{5 \cdot \epsilon \cdot \Delta'} \leq \frac{1}{5}. \end{aligned}$$

We conclude $\Pr[\Psi_2 \mid \Psi_1] \geq \frac{4}{5}$. Note that if Ψ_2 indeed occurred, then $d_{H_C}(u, v) \leq d_{G[C]}(u, v) + 5 \cdot \epsilon \cdot \Delta' = d_{G[C]}(u, v) + \frac{\Delta}{t} \leq \frac{2\Delta}{t}$. Finally in step 3 we sample a padded decomposition $C_{\mathcal{P}}$ of H_C . As H_C has treewidth $\pi(\Omega(\frac{1}{r}), r, n)$, the probability that u and v are clustered together is bounded by:

$$\Pr[\Psi_3 \mid \Psi_1 \wedge \Psi_2] \geq \Pr\left[\mathbf{B}_{H_C}(v, \frac{2\Delta}{t}) \subseteq \mathcal{P}_C(c)\right] \geq e^{-O(\log \pi(\Omega(\frac{1}{r}), r, n)) \cdot \frac{2}{t}} \geq \frac{4}{5},$$

where the last inequality holds for a large enough constant in the definition of t . We conclude

$$\Pr[\mathcal{P}(v) = \mathcal{P}(u)] \geq \Pr[\Psi_1 \wedge \Psi_2 \wedge \Psi_3] \geq \left(\frac{4}{5}\right)^3 > \frac{1}{2}.$$

□

Acknowledgements. Arnold Filtser was supported by the Israel Science Foundation (grant No. 1042/22). Hung Le was supported by the NSF CAREER Award No. CCF-223728, an NSF Grant No. CCF-2121952, and a Google Research Scholar Award.

References

- [AFGN22] Ittai Abraham, Arnold Filtser, Anupam Gupta, and Ofer Neiman. Metric embedding via shortest path decompositions. *SIAM J. Comput.*, 51(2):290–314, 2022.
- [AFH⁺04] Aaron Archer, Jittat Fakcharoenphol, Chris Harrelson, Robert Krauthgamer, Kunal Talwar, and Éva Tardos. Approximate classification via earthmover metrics. In J. Ian Munro, editor, *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004, New Orleans, Louisiana, USA, January 11-14, 2004*, pages 1079–1087. SIAM, 2004.
- [AGG⁺19] Ittai Abraham, Cyril Gavoille, Anupam Gupta, Ofer Neiman, and Kunal Talwar. Cops, robbers, and threatening skeletons: Padded decomposition for minor-free graphs. *SIAM J. Comput.*, 48(3):1120–1145, 2019.
- [AGGM06] Ittai Abraham, Cyril Gavoille, Andrew V. Goldberg, and Dahlia Malkhi. Routing in networks with low doubling dimension. In *26th IEEE International Conference on Distributed Computing Systems (ICDCS 2006), 4-7 July 2006, Lisboa, Portugal*, page 75, 2006.
- [AGK14] Alexandr Andoni, Anupam Gupta, and Robert Krauthgamer. Towards $(1 + \epsilon)$ -approximate flow sparsifiers. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 279–293, 2014.
- [AGMW10] Ittai Abraham, Cyril Gavoille, Dahlia Malkhi, and Udi Wieder. Strong-diameter decompositions of minor free graphs. *Theory Comput. Syst.*, 47(4):837–855, 2010.

- [AP90] Baruch Awerbuch and David Peleg. Sparse partitions (extended abstract). In *31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22-24, 1990, Volume II*, pages 503–513, 1990.
- [Bar96] Yair Bartal. Probabilistic approximations of metric spaces and its algorithmic applications. In *37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, Vermont, USA, 14-16 October, 1996*, pages 184–193, 1996.
- [BCF⁺23] Costas Busch, Da Qi Chen, Arnold Filtser, Daniel Hathcock, D. Ellis Hershkowitz, and Rajmohan Rajaraman. One tree to rule them all: Poly-logarithmic universal steiner tree. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023*, pages 60–76. IEEE, 2023.
- [BCK11] Anand Bhalgat, Deeparnab Chakrabarty, and Sanjeev Khanna. Optimal lower bounds for universal and differentially private steiner trees and tsps. In *APPROX-RANDOM*, pages 75–86. Springer, 2011.
- [BDR⁺12] Costas Busch, Chinmoy Dutta, Jaikumar Radhakrishnan, Rajmohan Rajaraman, and Srinivasagopalan Srivathsan. Split and join: Strong partitions and universal steiner trees for graphs. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 81–90, 2012.
- [BG89] Dimitris Bertsimas and Michelangelo Grigni. Worst-case examples for the spacefilling curve heuristic for the euclidean traveling salesman problem. *Operations Research Letters*, 8(5):241–244, 1989.
- [BLT14] Costas Busch, Ryan LaFortune, and Srikanta Tirthapura. Sparse covers for planar graphs and graphs that exclude a fixed minor. *Algorithmica*, 69(3):658–684, 2014.
- [Bou85] J. Bourgain. On lipschitz embedding of finite metric spaces in hilbert space. *Israel Journal of Mathematics*, 52(1):46–52, Mar 1985.
- [CCL⁺23] H. Chang, J. Conroy, H. Le, L. Milenkovic, S. Solomon, and C. Than. Covering planar metrics (and beyond): $O(1)$ trees suffice. In *64th IEEE Symposium on Foundations of Computer Science, FOCS '23*, pages 2231–2261, 2023.
- [CFKL20] Vincent Cohen-Addad, Arnold Filtser, Philip N. Klein, and Hung Le. On light spanners, low-treewidth embeddings and efficient traversing in minor-free graphs. In Sandy Irani, editor, *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 589–600. IEEE, 2020.
- [Che18] Yun Kuen Cheung. Steiner point removal - distant terminals don't (really) bother. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, Louisiana, USA, January 7-10, 2018*, 2018.
- [Chu12] Julia Chuzhoy. On vertex sparsifiers with steiner nodes. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 673–688, 2012.
- [CJF⁺23] Artur Czumaj, Shaofeng H.-C. Jiang, Arnold Filtser, Robert Krauthgamer, Pavel Veselý, and Mingwei Yang. Streaming facility location in high dimension via new geometric hashing. *CoRR*, abs/2204.02095v5, 2023.
- [CKR04] Gruia Călinescu, Howard J. Karloff, and Yuval Rabani. Approximation algorithms for the 0-extension problem. *SIAM J. Comput.*, 34(2):358–372, 2004.
- [CKR05] Gruia Calinescu, Howard Karloff, and Yuval Rabani. Approximation algorithms for the 0-extension problem. *SIAM Journal on Computing*, 34(2):358–372, 2005.

- [CLLM10] Moses Charikar, Tom Leighton, Shi Li, and Ankur Moitra. Vertex sparsifiers and abstract rounding algorithms. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 265–274, 2010.
- [CXKR06] T.-H. Chan, Donglin Xia, Goran Konjevod, and Andrea Richa. A tight lower bound for the steiner point removal problem on trees. In *Proceedings of the 9th International Conference on Approximation Algorithms for Combinatorial Optimization Problems, and 10th International Conference on Randomization and Computation, APPROX'06/RANDOM'06*, pages 70–81, Berlin, Heidelberg, 2006. Springer-Verlag.
- [DJP⁺92] Elias Dahlhaus, David S. Johnson, Christos H. Papadimitriou, Paul D. Seymour, and Mihalis Yannakakis. The complexity of multiway cuts (extended abstract). In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing, May 4-6, 1992, Victoria, British Columbia, Canada*, pages 241–251, 1992.
- [DO96] Guoli Ding and Bogdan Oporowski. On tree-partitions of graphs. *Discret. Math.*, 149(1-3):45–58, 1996.
- [EGK⁺14] Matthias Englert, Anupam Gupta, Robert Krauthgamer, Harald Räcke, Inbal Talgam-Cohen, and Kunal Talwar. Vertex sparsifiers: New results from old techniques. *SIAM J. Comput.*, 43(4):1239–1262, 2014.
- [FHRT03] Jittat Fakcharoenphol, Chris Harrelson, Satish Rao, and Kunal Talwar. An improved approximation algorithm for the 0-extension problem. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 12-14, 2003, Baltimore, Maryland, USA.*, pages 257–265, 2003.
- [FIK⁺22] Tobias Friedrich, Davis Issac, Nikhil Kumar, Nadym Mallek, and Ziena Zeif. Approximate max-flow min-multicut theorem for graphs of bounded treewidth. *CoRR*, abs/2211.06267, 2022.
- [Fil19a] Arnold Filtser. On strong diameter padded decompositions. In Dimitris Achlioptas and László A. Végh, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2019, September 20-22, 2019, Massachusetts Institute of Technology, Cambridge, MA, USA*, volume 145 of *LIPICs*, pages 6:1–6:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [Fil19b] Arnold Filtser. Steiner point removal with distortion $o(\log k)$ using the relaxed-voronoi algorithm. *SIAM J. Comput.*, 48(2):249–278, 2019.
- [Fil20] Arnold Filtser. Scattering and sparse partitions, and their applications. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPICs*, pages 47:1–47:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [Fil23] Arnold Filtser. Labeled Nearest Neighbor Search and Metric Spanners via Locality Sensitive Orderings. In Erin W. Chambers and Joachim Gudmundsson, editors, *39th International Symposium on Computational Geometry (SoCG 2023)*, volume 258 of *Leibniz International Proceedings in Informatics (LIPICs)*, pages 33:1–33:18, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [Fil24] Arnold Filtser. On sparse covers of minor free graphs, low dimensional metric embeddings, and other applications, 2024.
- [FK02] Uriel Feige and Robert Krauthgamer. A polylogarithmic approximation of the minimum bisection. *SIAM Journal on Computing*, 31(4):1090–1118, 2002.

- [FKS19] Eli Fox-Epstein, Philip N. Klein, and Aaron Schild. Embedding planar graphs into low-treewidth graphs with applications to efficient approximation schemes for metric problems. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1069–1088. SIAM, 2019.
- [FKT19] Arnold Filtser, Robert Krauthgamer, and Ohad Trabelsi. Relaxed voronoi: A simple framework for terminal-clustering problems. In *2nd Symposium on Simplicity in Algorithms, SOSA@SODA 2019, January 8-9, 2019 - San Diego, CA, USA*, pages 10:1–10:14, 2019.
- [FL21] Arnold Filtser and Hung Le. Clan embeddings into trees, and low treewidth graphs. In Samir Khuller and Virginia Vassilevska Williams, editors, *STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021*, pages 342–355. ACM, 2021.
- [FL22] Arnold Filtser and Hung Le. Low treewidth embeddings of planar and minor-free metrics. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 1081–1092. IEEE, 2022.
- [FT03] Jittat Fakcharoenphol and Kunal Talwar. An improved decomposition theorem for graphs excluding a fixed minor. In *Approximation, Randomization, and Combinatorial Optimization: Algorithms and Techniques, 6th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX 2003 and 7th International Workshop on Randomization and Approximation Techniques in Computer Science, RANDOM 2003, Princeton, NJ, USA, August 24-26, 2003, Proceedings*, pages 36–46, 2003.
- [GHR06] Anupam Gupta, Mohammad Taghi Hajiaghayi, and Harald Räcke. Oblivious network design. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*, pages 970–979. ACM Press, 2006.
- [GKSS10] Igor Gorodezky, Robert D Kleinberg, David B Shmoys, and Gwen Spencer. Improved lower bounds for the universal and a priori tsp. In *APPROX-RANDOM*, pages 178–191. Springer, 2010.
- [GNRS04] Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Cuts, trees and l_1 -embeddings of graphs. *Comb.*, 24(2):233–269, 2004.
- [Gup01] Anupam Gupta. Steiner points in tree metrics don’t (really) help. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '01*, pages 220–227, Philadelphia, PA, USA, 2001. Society for Industrial and Applied Mathematics.
- [HKL06] Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, and Frank Thomson Leighton. Improved lower and upper bounds for universal TSP in planar metrics. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 649–658, 2006. <https://dl.acm.org/citation.cfm?id=1109557.1109628>.
- [HL22] D. Ellis Hershkowitz and Jason Li. $O(1)$ steiner point removal in series-parallel graphs. In Shiri Chechik, Gonzalo Navarro, Eva Rotenberg, and Grzegorz Herman, editors, *30th Annual European Symposium on Algorithms, ESA 2022, September 5-9, 2022, Berlin/Potsdam, Germany*, volume 244 of *LIPICs*, pages 66:1–66:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [HR92] Frank K. Hwang and Dana S. Richards. Steiner tree problems. *Networks*, 22(1):55–89, 1992.
- [JLN⁺05] Lujun Jia, Guolong Lin, Guevara Noubir, Rajmohan Rajaraman, and Ravi Sundaram. Universal approximations for tsp, steiner tree, and set cover. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 386–395, 2005.

- [Kar98] Alexander V. Karzanov. Minimum 0-extensions of graph metrics. *Eur. J. Comb.*, 19(1):71–101, 1998.
- [KK17] Lior Kamma and Robert Krauthgamer. Metric decompositions of path-separable graphs. *Algorithmica*, 79(3):645–653, 2017.
- [KKN15] Lior Kamma, Robert Krauthgamer, and Huy L. Nguyen. Cutting corners cheaply, or how to remove steiner points. *SIAM J. Comput.*, 44(4):975–995, 2015.
- [KLMN04] Robert Krauthgamer, James R. Lee, Manor Mendel, and Assaf Naor. Measured descent: A new embedding method for finite metrics. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*, pages 434–443, 2004.
- [KLR19] Robert Krauthgamer, James R. Lee, and Havana Rika. Flow-cut gaps and face covers in planar graphs. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 525–534. SIAM, 2019.
- [KPR93] Philip N. Klein, Serge A. Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing, May 16-18, 1993, San Diego, CA, USA*, pages 682–690, 1993.
- [Kum22] Nikhil Kumar. An approximate generalization of the okamura-seymour theorem. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 1093–1101. IEEE, 2022.
- [Lee12] James Lee. A simpler proof of the KPR theorem. <https://tcsmath.wordpress.com/2012/01/11/a-simpler-proof-of-the-kpr-theorem/>, 2012. Accessed: 2023-22-05.
- [LLR95] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- [LN05] James R. Lee and Assaf Naor. Extending lipschitz functions via random metric partitions. *Inventiones mathematicae*, 160(1):59–95, 2005.
- [LR10] James R. Lee and Prasad Raghavendra. Coarse differentiation and multi-flows in planar graphs. *Discret. Comput. Geom.*, 43(2):346–362, 2010.
- [MM10] Konstantin Makarychev and Yury Makarychev. Metric extension operators, vertex sparsifiers and lipschitz extendability. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 255–264, 2010.
- [Moi09] Ankur Moitra. Approximation algorithms for multicommodity-type problems with guarantees independent of the graph size. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia, USA*, pages 3–12, 2009.
- [MT10] Robin A. Moser and Gábor Tardos. A constructive proof of the general lovász local lemma. *J. ACM*, 57(2):11:1–11:15, 2010.
- [OS81] Haruko Okamura and P.D. Seymour. Multicommodity flows in planar graphs. *Journal of Combinatorial Theory, Series B*, 31(1):75–81, 1981.
- [PBI89] Loren K Platzman and John J Bartholdi III. Spacefilling curves and the planar travelling salesman problem. *Journal of the ACM (JACM)*, 36(4):719–737, 1989.
- [Rao99] Satish Rao. Small distortion and volume preserving embeddings for planar and euclidean metrics. In *Proceedings of the fifteenth annual symposium on Computational geometry, SCG '99*, pages 300–306, New York, NY, USA, 1999. ACM.
- [RR05] Satish Rao and Andréa W. Richa. New approximation techniques for some linear ordering problems. *SIAM Journal on Computing*, 34(2):388–404, 2005.

- [RS03] Neil Robertson and Paul D. Seymour. Graph minors. XVI. excluding a non-planar graph. *J. Comb. Theory, Ser. B*, 89(1):43–76, 2003.
- [SS08] Frans Schalekamp and David B Shmoys. Algorithms for the universal and a priori tsp. *Operations Research Letters*, 36(1):1–3, 2008.
- [TV93] Éva Tardos and Vijay V. Vazirani. Improved bounds for the max-flow min-multicut ratio for planar and k-free graphs. *Information Processing Letters*, 47(2):77–80, 1993.