

Terminal Embeddings*

Michael Elkin[†], Arnold Filtser[‡], Ofer Neiman[§]

Department of Computer Science, Ben-Gurion University of the Negev, Beer-Sheva, Israel.
Email: {elkinm,arnoldf,neimano}@cs.bgu.ac.il

Abstract

In this paper we study *terminal embeddings*, in which one is given a finite metric (X, d_X) (or a graph $G = (V, E)$) and a subset $K \subseteq X$ of its points are designated as *terminals*. The objective is to embed the metric into a normed space, while approximately preserving all distances among pairs that contain a terminal. We devise such embeddings in various settings, and conclude that even though we have to preserve $\approx |K| \cdot |X|$ pairs, the distortion depends only on $|K|$, rather than on $|X|$.

We also strengthen this notion, and consider embeddings that approximately preserve the distances between *all* pairs, but provide improved distortion for pairs containing a terminal. Surprisingly, we show that such embeddings exist in many settings, and have optimal distortion bounds both with respect to $X \times X$ and with respect to $K \times X$.

Moreover, our embeddings have implications to the areas of Approximation and Online Algorithms. In particular, [ALN08] devised an $\tilde{O}(\sqrt{\log r})$ -approximation algorithm for sparsest-cut instances with r demands. Building on their framework, we provide an $\tilde{O}(\sqrt{\log |K|})$ -approximation for sparsest-cut instances in which each demand is incident on one of the vertices of K (aka, terminals). Since $|K| \leq r$, our bound generalizes that of [ALN08].

Keywords— embedding, distortion, terminals

*A preliminary version of this paper appeared in APPROX'15.

[†]Supported in part by ISF grant No. (724/15).

[‡]Supported in part by ISF grant No. (523/12), by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n^o303809, and by BSF grant No. 2015813

[§]Supported in part by ISF grant No. (523/12), by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n^o303809, and by BSF grant No. 2015813

Contents

1	Introduction	3
1.1	Algorithmic Applications	4
1.2	Overview of Techniques	5
1.3	Related Work	6
1.4	Subsequent Work	7
2	Preliminaries	8
3	A General Transformation	8
4	Graph Terminal Spanners	12
5	Light Terminal Trees for General Graphs	13
6	Lower Bounds on Distortion-Lightness Tradeoffs for a Single Tree	17
6.1	A lower bound for terminal distortion	17
6.2	A lower bound on the lightness in graphs	18
7	A lower bound on the tradeoff between lightness and stretch for terminal trees in metric spaces	20
8	Probabilistic Embedding into Ultrametrics with Strong Terminal Distortion	25
9	Probabilistic Embedding into Trees with Terminal Congestion	28
10	Strong Terminal Embedding of Graphs into a Distribution of Spanning Trees	31
10.1	Petal decomposition description	31
10.2	Analysis	34
10.3	Expected Distortion	36
11	Applications	48
11.1	Sparsest Cut	48
11.2	Min Bisection	50
11.3	Online Algorithms: Constrained File Migration	51
12	Acknowledgements	52

1 Introduction

Embedding of finite metric spaces is a very successful area of research, due to both its algorithmic applications and its natural geometric appeal. Given two metric spaces (X, d_X) , (Y, d_Y) , we say that X embeds into Y with distortion α if there is a map $f : X \rightarrow Y$ and a constant $c > 0$, such that for all $u, v \in X$,

$$d_X(u, v) \leq c \cdot d_Y(f(u), f(v)) \leq \alpha \cdot d_X(u, v) .$$

Some of the basic results in the field of metric embedding are: a theorem of [Bou85], asserting that any metric space on n points embeds with distortion $O(\log n)$ into Euclidean space (which was shown to be tight by [LLR95]), and probabilistic embedding into a distribution of trees, or ultrametrics,¹ with expected distortion $O(\log n)$ [FRT04], or expected congestion $O(\log n)$ [Rac08] (which are also tight [Bar96]).

In this paper we study a natural variant of embedding, in which the input consists of a finite metric space or a graph, and in addition, a subset of the points are designated as *terminals*. The objective is to embed the metric into a simpler metric (e.g., Euclidean metric), or into a simpler graph (e.g., a tree), while approximately preserving the distances between the terminals to *all other points*. We show that such embeddings, which we call *terminal embeddings*, can have improved parameters compared to embeddings that must preserve all pairwise distances. In particular, the distortion (and the dimension in embedding into normed spaces) depends only on the number of terminals, regardless of the cardinality of the metric space.

We also consider a strengthening of this notion, which we call *strong terminal embedding*. Here we want a distortion bound on *all pairs*, and in addition an improved distortion bound on pairs that contain a terminal. Such strong terminal embeddings enhance the classical embedding results, essentially saying that one can obtain the same distortion for all pairs, with the option to select some of the points, and obtain improved approximation of the distances between any selected point to any other point.

As a possible motivation for studying such embeddings, consider a scenario in which a certain network of clients and servers is given as a weighted graph (where edges correspond to links, weights to communication/travel time). It is conceivable that one only cares about distances between clients and servers, and that there are few servers. We would like to have a simple structure, such as a tree spanning the network, so that the client-server distances in the tree are approximately preserved.

We show that there exists a general phenomenon; essentially any known metric embedding into an ℓ_p space or a graph family can be transformed via a general transformation into a terminal embedding, while paying only a constant blow-up in the distortion. In particular, we obtain a terminal embedding of any finite metric into any ℓ_p space with terminal distortion $O(\log k)$, using only $O(\log k)$ dimensions. We show that a similar general phenomenon² holds also for *strong* terminal embeddings, i.e., that many embeddings can be transformed into strong terminal ones via a general transformation. Many other embeddings into normed spaces, probabilistic embedding into ultrametrics (including capacity preserving ones), and into spanning trees, have their *strong* terminal embedding counterparts, which are constructed directly, that is, not through our general transformation. Our results are tight in most settings.³

¹An ultrametric (U, ρ) is a metric space satisfying $\rho(x, z) \leq \max\{\rho(x, y), \rho(y, z)\}$ for all $x, y, z \in U$.

²Though this transformation is somewhat less general than for the case of ordinary, i.e., not strong, terminal embeddings.

³All our terminal embeddings are tight, except for the probabilistic spanning trees, where they match the state-

It is well known that embedding a graph into a single tree may cause (worst-case) distortion $\Omega(n)$ [RR98]. However, if one only cares about client-server distances, we show that it is possible to obtain distortion $2k - 1$, where k is the number of servers, and that this is tight. Furthermore, we study possible tradeoffs between the distortion and the total weight of the obtained tree. This generalizes the notion of *shallow light trees* [KRY95, ABP92, ES11], which provides a tradeoff between the distortion with respect to a single designated server and the weight of the tree.

We then address probabilistic approximation of metric spaces and graphs by ultrametrics and spanning trees. This line of work started with the results of [AKPW95, Bar96], and culminated in the $O(\log n)$ expected distortion for ultrametrics by [FRT04], and $\tilde{O}(\log n)$ for spanning trees by [AN12]. These embeddings found numerous algorithmic applications, in various settings, see [FRT04, EEST08, AN12] and the references therein for details. In their work on Ramsey partitions, [MN06] implicitly showed that there exists a probabilistic embedding into ultrametrics with expected terminal distortion $O(\log k)$ (see Section 2 for the formal definitions). We generalize this result by showing a *strong* terminal embedding with the same expected $O(\log k)$ distortion guarantee for all pairs containing a terminal, and $O(\log n)$ for all other pairs. We also show a similar result that extends the embedding of [AN12] into spanning trees, with $\tilde{O}(\log k)$ expected distortion for pairs containing a terminal, and $\tilde{O}(\log n)$ for all pairs. A slightly different notion, introduced by [Rö2], is that of trees which approximate the congestion (rather than the distortion), and [Räc08] showed a distribution over trees with expected congestion $O(\log n)$. We provide a strong terminal version of this result, and show expected congestion of $O(\log k)$ for all edges incident on a terminal, and $O(\log n)$ for the rest.

We also consider spanners, with a stretch requirement only for pairs containing a terminal. Our general transformation produces for any $t \geq 1$ a $(4t - 1)$ -terminal stretch spanner with $O(k^{1+1/t} + n)$ edges. The drawback is that this is a metric spanner, not a subgraph of the input graph. We alleviate this issue by constructing a graph spanner with the same stretch and $O(\sqrt{n} \cdot k^{1+1/t} + n)$ edges.⁴ A result of [RTZ05] implicitly provides a terminal graph spanner with $(2t - 1)$ stretch and $O(t \cdot n \cdot k^{1/t})$ edges. Our graph terminal spanner is sparser than that of [RTZ05] as long as $k \leq t \cdot n^{1/2(1+1/t)}$.

1.1 Algorithmic Applications

We overview a few of the applications of our results to approximation and online algorithms. Some of the most striking applications of metric embeddings are to various cut problems, such as the sparsest-cut, min-bisection, and also to several online problems. Our method provides improved guarantees when the input graph has a small set of “important” vertices. Specifically, these vertices can be considered as terminals, and we obtain approximation factors that depend on the cardinality of the terminal set, rather than on the input size. The exact meaning of importance is problem specific; e.g. in the cut problems, we require that the set of important vertices touches every demand pair, or every edge (that is, forms a vertex cover).

For instance, consider the (general) *sparsest-cut* problem [LR99, AR98, LLR95]. We are given a graph $G = (V, E)$ with capacities on the edges $c : E \rightarrow \mathbb{R}_+$, and a collection of pairs $(s_1, t_1), \dots, (s_r, t_r)$ along with their demands D_1, \dots, D_r . The goal is to find a cut $S \subseteq V$ that minimizes the ratio

of-the-art [AN12], and except for our terminal spanners.

⁴Note that the number of edges is linear whenever $k \leq n^{1/(2(1+1/t))}$.

between capacity and demand across the cut:

$$\phi(S) = \frac{\sum_{\{u,v\} \in E} c(u,v) |\mathbf{1}_S(u) - \mathbf{1}_S(v)|}{\sum_{i=1}^r D_i |\mathbf{1}_S(s_i) - \mathbf{1}_S(t_i)|},$$

where $\mathbf{1}_S(\cdot)$ is the indicator for membership in S . Following the breakthrough result of [ARV09], which showed $O(\sqrt{\log n})$ approximation for the uniform demand case, [ALN08] devised an $\tilde{O}(\sqrt{\log r})$ approximation for the general case. Consider the case where there is a set of k important vertices, such that every demand pair contains an important vertex (that is, there is a k size vertex cover of the demand graph). [Gün07] give a $O(\log k)$ approximation for this case. We obtain an $\tilde{O}(\sqrt{\log k})$ approximation using the terminal embedding of negative-type metrics into ℓ_1 . Observe that $k \leq r$, and so this result subsumes the result of [ALN08]. Our bound is particularly useful for instances with many demand pairs but few distinct sources s_i (or few targets t_i).

We also consider other cut problems, and show a similar phenomenon: the $O(\log n)$ approximation for the *min-bisection* problem [Räc08], can be improved to an approximation of only $O(\log k)$, where k is the size of the minimum vertex cover of the input graph. For this result we employ our terminal variant of Räcke’s result [Räc08] on capacity-preserving probabilistic embedding into trees.

We then focus on one application of probabilistic embedding into ultrametrics [Bar96, FRT04], and illustrate the usefulness of our terminal embedding result by the (online) *constrained file migration* problem [BFR95]. Given a graph $G = (V, E)$ representing a network, each node $v \in V$ has a memory capacity m_v , and there is a set of files that reside at the nodes, at most m_v files may be stored at node v at any given time. The cost of accessing a file is the distance in the graph to the node storing it (no copies are allowed). Files can also be migrated from one node to another. This costs D times the distance, for a given parameter $D \geq 1$. When a sequence of file requests from nodes arrives online, the goal is to minimize the cost of serving all requests. [Bar96] showed a algorithm with $O(\log m \cdot \log n)$ competitive ratio for graphs on n nodes, where $m = \sum_{v \in V} m_v$ is the total memory available.⁵ A setting which seems particularly natural is one where there is a small set of nodes who can store files (servers), and the rest of the nodes can only access files but not store them (clients). We employ our probabilistic terminal embedding into ultrametrics to provide a $O(\log m \cdot \log k)$ competitive ratio, for the case where there are k servers. (Note that this ratio is independent of n .)

1.2 Overview of Techniques

The weak variant of our terminal embedding into ℓ_2 maps every terminal x into its image $f(x)$ under an original black-box (e.g., Bourgain’s) embedding of K into ℓ_2 . This embedding is then appended with one additional coordinate. Terminals are assigned 0 value in this coordinate, while each non-terminal point y is mapped to $(f(x), d(x, y))$, where x is the closest terminal to y . It is not hard to see that this embedding guarantees terminal distortion $O(\gamma(k))$, where $\gamma(k)$ is the distortion of the original black-box embedding, i.e., $O(\log k)$ in the case of Bourgain’s embedding. On the other hand, the new embedding employs only $\beta(k) + 1$ dimensions, where $\beta(k)$ is the dimension of the original blackbox embedding (i.e., $O(\log^2 k)$ in the case of Bourgain’s embedding).⁶ This

⁵The original paper shows $O(\log m \cdot \log^2 n)$, the improved factor is obtained by using the embedding of [FRT04].

⁶We can also get dimension $O(\log k)$ for terminal embeddings into ℓ_2 by replacing Bourgain’s embedding with that of [ABN11].

idea easily generalizes to a number of quite general scenarios, and under mild assumptions (see Theorem 2) it can be modified to produce strong terminal embeddings.

This framework, however, does not apply in many important settings, such as embedding into subgraphs, and does not provide strong terminal guarantees in others. Therefore we devise embeddings tailored to each particular setting in a non-black-box manner. For instance, our probabilistic embedding into trees with strong terminal congestion requires an adaptation of a theorem of [AF09], about the equivalence of distance-preserving and capacity-preserving random tree embeddings, to the terminal setting. Perhaps the most technically involved is our probabilistic embedding into spanning trees with strong terminal distortion. This result requires a set of modifications to the recent algorithm of [AN12], which is based on a certain hierarchical decomposition of graphs. We adapt this algorithm by giving preference to the terminals in the decomposition (they are the first to be chosen as cluster centers), and the crux is assuring that the distortion of any pair containing a terminal is essentially not affected by choices made for non-terminals. Furthermore, one has to guarantee that each such pair can be separated in at most $O(\log k)$ levels of the hierarchy.

The basic technical idea that we use for constructing $(4t - 1)$ -terminal subgraph spanners with $O(\sqrt{n}k^{1+1/t} + n)$ edges is the following one. As was mentioned above, our general transformation constructs metric (i.e., non-subgraph) $(4t - 1)$ -terminal spanners with $O(n + k^{1+1/t})$ edges. The latter spanners employ some edges which do not belong to the original graph. We provide these edges as an input to a pairwise preserver. A pairwise preserver [CE05] is a sparse subgraph that preserves exactly all distances between a designated set of vertex pairs. We use these preservers to fill in the gaps in the non-subgraph terminal spanner constructed via our general transformation. As a result we obtain a subgraph terminal spanner which outperforms previously existing terminal spanners of [RTZ05] in a wide range of parameters.

1.3 Related Work

Already in the pioneering work of [LLR95], an embedding that has to provide a distortion guarantee for a subset of the pairs is presented. Specifically, in the context of the sparsest-cut problem, [LLR95] devised a non-expansive embedding of an arbitrary metric into ℓ_1 , with distortion at most $O(\log r)$ for a set of r specified demand pairs.

Terminal distance oracles were studied by [RTZ05], who called them *source restricted* distance oracles. In their paper, [RTZ05] show $(2t - 1)$ -terminal stretch using $O(t \cdot n \cdot k^{1/t})$ space. Implicit in our companion paper [EFN15] is a distance oracle with $(4t - 1)$ -terminal stretch, $O(t \cdot k^{1/t} + n)$ space and $O(1)$ query time. Terminal spanners with additive stretch for unweighted graphs were recently constructed in [KV13]. Specifically, they showed a spanner with $\tilde{O}(n^{5/4} \cdot k^{1/4})$ edges and additive stretch 2 for pairs containing a terminal. Another line of work introduced *distance preservers* [CE05]; these are spanners which preserve exactly distances for a given collection of pairs.

In the context of preserving distances just between the terminals, [Gup01a, CXKR06, EGK⁺14, KKN14] studied embeddings of a graph into a minor over the terminals, while approximately preserving distances. In their work on the requirement cut problem, among other results, [GNR10] obtain for any metric with k specified terminals, a distribution over trees with expected expansion $O(\log k)$ for all pairs, and which is non-contractive for terminal pairs. (Note that this is different from our setting, as the extra guarantee holds for terminals only, not for pairs containing a terminal.)

Another line of research [Moi09, CLLM10, MM10, EGK⁺14] studied cut and vertex sparsifiers. A *cut sparsifier* of a graph $G = (V, E)$ with respect to a subset K of terminals is a graph $H = (K, E_H)$ on just the set of terminals, so that for any subset $A \subset K$, the minimum value of a cut

in G that separates A from $K \setminus A$ is approximately equal to the value of the cut $(A, K \setminus A)$ in H . Note that this notion is substantially different from the notion of terminal congestion-preserving embedding, which we study in the current paper.

1.4 Subsequent Work

In a companion paper [EFN15], we study prioritized metric structures and embeddings. In that setting, along with the input metric (X, d) , a priority ranking of the points of X is given, and the goal is to obtain a data structure (distance oracle, routing scheme) or an embedding with stretch/distortion that depends on the ranking of the points. This has some implications to the terminal setting, since the k terminals can be given as the first k points in the priority ranking. More concretely, implicit in [EFN15] is an embedding into a single (non-subgraph) tree with strong terminal distortion $O(k)$, a probabilistic embedding into ultrametrics with expected strong terminal distortion $O(\log k)$, and embedding into ℓ_p space with strong terminal distortion $\tilde{O}(\log k)$. In the current paper we provide stronger and more general results: our single tree embedding has tight $2k - 1$ stretch, the tree is a subgraph, and it can have low weight as well (at the expense of slightly increased stretch); we obtain probabilistic embedding into *spanning* trees, and in congestion-preserving trees; and our terminal embedding to ℓ_p space has a tight strong terminal distortion $(O(\log k), O(\log n))$ and low dimension. Furthermore, the results of this paper apply to numerous other settings (e.g., embeddings tailored for graphs excluding a fixed minor, negative-type metrics, spanners, etc.).

Following our work, [BFN16] discovered a connection between terminal distortion and coarse partial distortion. First recall the notion of coarse partial distortion, introduced in [KSW09, ABC⁺05]. Let (X, d_X) be a metric space of size $|X| = n$. For $x \in X$ and $\epsilon \in (0, 1)$, let $R(x, \epsilon) = \min\{r : |B(x, r)| \geq \epsilon n\}$. A point y is called ϵ -far from x if $d_X(x, y) \geq R(x, \epsilon)$. We say that an embedding $f : X \rightarrow Y$ has *coarse* $(1 - \epsilon)$ -*partial distortion* λ , if every pair $x, y \in X$ such that both x, y are $\epsilon/2$ -far from each other, has distortion at most λ . The connection between this notion and terminal distortion is roughly as follows.

- If a metric admits a non-contracting embedding (into some target space) with terminal distortion γ for a certain set of k terminals, then this very embedding has coarse $(1 - \frac{8}{k})$ -partial distortion $5 \cdot \gamma$.
- If every metric admits an embedding with coarse $(1 - \frac{1}{k})$ -partial distortion γ , then every metric has embedding with terminal distortion γ (for any set of k terminals).

The terminal embedding results presented here are strictly stronger than those obtainable by using the state-of-the-art coarse partial embeddings with the above transformation. E.g., by [ABC⁺05] it follows that if every n -point metric embeds into ℓ_p with distortion $\alpha(n)$ and dimension $\beta(n)$ (where the embedding needs to fulfill a certain condition), then it embeds into ℓ_p with terminal distortion $\alpha(O(k))$ and dimension $\beta(O(k)) \cdot O(\log n)$. Our generic terminal embedding of Theorem 1 provides dimension independent of n , does not restrict the original embedding, and has improved constants in the distortion. Also our embedding into spanning trees with $(\tilde{O}(\log k), \tilde{O}(\log n))$ -strong terminal distortion improves the embedding obtainable by going through the coarse partial embedding of [ABN15], which would give only $(\tilde{O}(\log^2 k), \tilde{O}(\log^2 n))$ -strong terminal distortion.

2 Preliminaries

Here we provide formal definitions for the notions of terminal distortion. Let (X, d_X) be a finite metric space, with $K \subseteq X$ a set of terminals. Throughout the paper we assume $|K| \leq |X|/2$.

Definition 1. Let (X, d_X) be a metric space, and let $K \subseteq X$ be a subset of terminals. For a target metric (Y, d_Y) , an embedding $f : X \rightarrow Y$ has terminal distortion α if there exists $c > 0$, such that for all $v \in K$ and $u \in X$,⁷

$$d_X(v, u) \leq c \cdot d_Y(f(v), f(u)) \leq \alpha \cdot d_X(v, u) .$$

We say that the embedding has strong terminal distortion (α, β) if it has terminal distortion α , and in addition there exists $c' > 0$, such that for all $u, w \in X$,

$$d_X(u, w) \leq c' \cdot d_Y(f(u), f(w)) \leq \beta \cdot d_X(u, w) .$$

For a metric space (X, d_X) with a terminal set $K \subseteq X$, an α -terminal (metric) spanner is a graph H with X as its vertex set, such that for all $x \in K$ and $y \in X$,

$$d_X(x, y) \leq d_H(x, y) \leq \alpha \cdot d_X(x, y) . \tag{1}$$

If the metric (X, d_X) is induced by the shortest path metric of a weighted graph G , and H is a subgraph of G , we call H a graph spanner of G .

Denote by $\text{diam}(X) = \max_{y, z \in X} \{d_X(y, z)\}$. For any $x \in X$ and $r \geq 0$ let $B_X(x, r) = \{y \in X \mid d_X(x, y) \leq r\}$ (we often omit the subscript when the metric is clear from context). For a point $x \in X$ and a subset $A \subseteq X$, $d_X(x, A) = \min_{a \in A} \{d_X(x, a)\}$. For $K \subseteq X$ we denote by (K, d_K) the metric space where d_K is the induced metric.

For a weighted graph $G = (V, E, w)$ where $w : E \rightarrow \mathbb{R}_+$, given a subgraph H of G , let $w(H) = \sum_{e \in E(H)} w(e)$, and define the lightness of H to be $\Psi(H) = \frac{w(H)}{w(\text{MST}(G))}$, where $w(\text{MST})$ is the weight of a minimum spanning tree of G .

By $\tilde{O}(f(n))$ we mean $O(f(n) \cdot \text{polylog}(f(n)))$.

3 A General Transformation

In this section we present general transformation theorems that create terminal embeddings into normed spaces and graph families from standard ones. We say that a family of graphs \mathcal{G} is leaf-closed, if it is closed under adding leaves. That is, for any $G \in \mathcal{G}$ and $v \in V(G)$, the graph G' obtained by adding a new vertex u and connecting u to v by an edge, belongs to \mathcal{G} . Note that many natural families of graphs are leaf-closed, e.g. trees, planar graphs, minor-free graphs, bounded tree-width graphs, bipartite graphs, general graphs, and many others.

Theorem 1. Let \mathcal{X} be a family of metric spaces. Fix some $(X, d_X) \in \mathcal{X}$, and let $K \subseteq X$ be a set of terminals of size $|K| = k$, such that $(K, d_K) \in \mathcal{X}$. Then the following assertions hold:

⁷In most of our results the embedding has a one-sided guarantee (that is, non-contractive or non-expansive) for all pairs.

- If there are functions $\alpha : \mathbb{N} \rightarrow \mathbb{R}$ and $\gamma : \mathbb{N} \rightarrow \mathbb{N}$, such that every $(Z, d_Z) \in \mathcal{X}$ of size $|Z| = m$ embeds into $\ell_p^{\gamma(m)}$ with distortion $\alpha(m)$, then there is an embedding of X into $\ell_p^{\gamma(k)+1}$ with terminal distortion $2^{(p-1)/p} \cdot ((2\alpha(k))^p + 1)^{1/p}$.⁸
- If \mathcal{G} is a leaf-closed family of graphs, and any $(Z, d_Z) \in \mathcal{X}$ of size $|Z| = m$ embeds into \mathcal{G} with distortion $\alpha(m)$ such that the target graph has at most $\gamma(m)$ edges, then there is an embedding of X into \mathcal{G} with terminal distortion $2\alpha(k) + 1$ and the target graph has at most $\gamma(k) + n - k$ edges.

Remark: The second assertion holds under probabilistic embeddings as well (see Section 8 for formal definition of probabilistic embedding).

Proof. We start by proving the first assertion. By the assumption there exists an embedding $f : K \rightarrow \mathbb{R}^{\gamma(k)}$ with distortion $\alpha(k)$ under the ℓ_p norm. We assume w.l.o.g. that f is non-contractive. For each $x \in X$, let $k_x \in K$ be the nearest point to x in K (that is, $d(x, K) = d(x, k_x)$). Define the embedding $\hat{f} : X \rightarrow \mathbb{R}^{\gamma(k)+1}$ by letting for $x \in X$, $\hat{f}(x) = (f(k_x), d(x, k_x))$. Observe that for $t \in K$, $\hat{f}(t) = (f(t), 0)$. Fix any $t \in K$ and $x \in X$. Note that by definition of k_x , $d(x, k_x) \leq d(x, t)$, and by the triangle inequality, $d(t, k_x) \leq d(t, x) + d(x, k_x) \leq 2d(t, x)$, so that,

$$\begin{aligned} \|\hat{f}(t) - \hat{f}(x)\|_p^p &= \|f(t) - f(k_x)\|_p^p + d(x, k_x)^p \\ &\leq (\alpha(k) \cdot d(t, k_x))^p + d(x, k_x)^p \\ &\leq (2\alpha(k) \cdot d(t, x))^p + d(t, x)^p \\ &= d(t, x)^p \cdot ((2\alpha(k))^p + 1) . \end{aligned}$$

On the other hand, since f does not contract distances,

$$\begin{aligned} \|\hat{f}(t) - \hat{f}(x)\|_p^p &= \|f(t) - f(k_x)\|_p^p + d(x, k_x)^p \\ &\geq d(t, k_x)^p + d(x, k_x)^p \\ &\geq (d(t, k_x) + d(x, k_x))^p / 2^{p-1} \\ &\geq d(x, t)^p / 2^{p-1} , \end{aligned}$$

where the second inequality is by the power mean inequality. We conclude that the terminal distortion is at most $2^{(p-1)/p} \cdot ((2\alpha(k))^p + 1)^{1/p}$.

For the second assertion, there is a non-contractive embedding f of K into $G \in \mathcal{G}$ with distortion at most $\alpha(k)$. As above, for each $x \in X \setminus K$ define k_x as the nearest point in K to x . And for each $x \in X$, add to G a new vertex $f(x)$ that is connected by an edge of length $d_G(x, k_x)$ to $f(k_x)$. The resulting graph $G' \in \mathcal{G}$, because it is a leaf-closed family. Fix any $x \in X$ and $t \in K$, then as above $d(t, k_x) \leq 2d(t, x)$, and so

$$\begin{aligned} d_{G'}(f(t), f(x)) &= d_G(f(t), f(k_x)) + d_{G'}(f(x), f(k_x)) \\ &\leq \alpha(k) \cdot d(t, k_x) + d(x, k_x) \\ &\leq d(t, x) \cdot (2\alpha(k) + 1) . \end{aligned}$$

Also note that

$$d_{G'}(f(t), f(x)) = d_G(f(t), f(k_x)) + d(x, k_x) \geq d(t, k_x) + d(x, k_x) \geq d(t, x) ,$$

⁸Note that for any $p, \alpha \geq 1$ we have that $2^{(p-1)/p} \cdot ((2\alpha)^p + 1)^{1/p} \leq 4\alpha$.

so the terminal distortion is indeed $2\alpha(k) + 1$. Since f embeds into a graph with $\gamma(k)$ edges, and we added $n - k$ new edges, the total number of edges is bounded accordingly, which concludes the proof. \square

Next, we study strong terminal embeddings into normed spaces. Fix any metric (X, d) , a set of terminals $K \subseteq X$ and $1 \leq p \leq \infty$. Let $f : K \rightarrow \ell_p$ be a non-expansive embedding. We say that f is *Lipschitz extendable*, if there exists a non-expansive $\hat{f} : X \rightarrow \ell_p$ which is an extension of f (that is, the restriction of \hat{f} to K is exactly f). It is not hard to verify that any Fréchet embedding⁹ is Lipschitz extendable. For example, the embeddings of [Bou85, KLMN05, ALN08] are Fréchet.

Theorem 2. *Let \mathcal{X} be a family of metric spaces such that any $(Z, d_Z) \in \mathcal{X}$ of size $|Z| = m$ embeds into $\ell_p^{\gamma(m)}$ with distortion $\alpha(m)$ by a Lipschitz extendable map. Fix some $(X, d_X) \in \mathcal{X}$, and let $K \subseteq X$ be a set of terminals of size $|K| = k$, such that $(K, d_K) \in \mathcal{X}$. Then there is a (non-expansive) embedding of X into $\ell_p^{\gamma(n) + \gamma(k) + 1}$ with strong terminal distortion $(O(\alpha(k)), O(\alpha(n)))$.*

Remark 1. *In Theorem 2 we assumed that any $(Z, d_Z) \in \mathcal{X}$ of size $|Z| = m$ embeds into $\ell_p^{\gamma(m)}$ with distortion $\alpha(m)$ by a Lipschitz extendable map. Instead, we can have a weaker assumption. We say that $f : K \rightarrow \ell_p$ is *t-Lipschitz extendable* if f is non-expansive and there is $\hat{f} : X \rightarrow \ell_p$ which is extension of f with expansion t . A generalization of Theorem 2 is the following. Suppose that for every $(Z, d_Z) \in \mathcal{X}$ of size $|Z| = m$, there is an embedding $f : Z \rightarrow \ell_p^{\gamma(m)}$ by a t -Lipschitz extendable map, with distortion $\alpha(m)$. Then, a proof similar to Theorem 2 implies that for every $(X, d_X) \in \mathcal{X}$ and $K \subseteq X$ of size $|K| = k$ where $(K, d_K) \in \mathcal{X}$, there is a (non-expansive) embedding of X into $\ell_p^{\gamma(n) + \gamma(k) + 1}$ with strong terminal distortion $(O(t \cdot \alpha(k)), O(\alpha(n)))$.*

Specifically, in the case $p = 2$, every Lipschitz map $f : K \rightarrow \ell_2^{\gamma(k)}$ is $\sqrt{\gamma(k)}$ -Lipschitz extendable. Thus, [JL84] implies that there is a (non-expansive) embedding of X into $\ell_2^{\gamma(n) + O(\log k)}$ with strong terminal distortion $(O(\sqrt{\log k} \cdot \alpha(k)), O(\alpha(n)))$.

Proof. Let $(X, d) \in \mathcal{X}$ be a metric on n points, $K \subseteq X$ of size $|K| = k$. There is a non-expansive embedding $g : X \rightarrow \ell_p^{\gamma(n)}$ with distortion at most $\alpha(n)$, and there exists a Lipschitz extendable embedding $f : K \rightarrow \ell_p^{\gamma(k)}$, which is non-expansive and has distortion $\alpha(k)$. Let \hat{f} be the extension of f to X , note that by definition of Lipschitz extendability, \hat{f} is also non-expansive. Finally, let $h : X \rightarrow \mathbb{R}$ be defined by $h(x) = d(x, K)$. The embedding $F : X \rightarrow \ell_p^{\gamma(n) + \gamma(k) + 1}$ is defined by the concatenation of these maps $F = g \oplus \hat{f} \oplus h$.

Since all the three maps g, \hat{f}, h are non-expansive, it follows that for any $x, y \in X$,

$$\|F(x) - F(y)\|_p^p \leq \|g(x) - g(y)\|_p^p + \|\hat{f}(x) - \hat{f}(y)\|_p^p + |h(x) - h(y)|^p \leq 3d(x, y)^p,$$

so F has expansion at most $3^{1/p}$ for all pairs (which can easily be made 1 without affecting the distortion by more than a constant factor). Also note that

$$\|F(x) - F(y)\|_p \geq \|g(x) - g(y)\|_p \geq \frac{d(x, y)}{\alpha(n)},$$

which implies the distortion bound for all pairs is satisfied. It remains to bound the contraction for all pairs containing a terminal. Let $t \in K$ and $x \in X$, and let $k_x \in K$ be such that $d(x, K) = d(x, k_x)$

⁹In our context, it will be convenient to call an embedding $f : K \rightarrow \ell_p^t$ *Fréchet*, if there are sets $A_1, \dots, A_t \subseteq X$ such that for all $i \in [t]$, and for every $x \in K$, we have $f_i(x) = \frac{d(x, A_i)}{t^{1/p}}$.

(it could be that $k_x = x$). If it is the case that $d(x, t) \leq 3\alpha(k) \cdot d(x, k_x)$ then by the single coordinate of h we get sufficient contribution for this pair:

$$\|F(t) - F(x)\|_p \geq |h(t) - h(x)| = h(x) = d(x, k_x) \geq \frac{d(x, t)}{3\alpha(k)}.$$

The other case is that $d(x, t) > 3\alpha(k) \cdot d(x, k_x)$, here we will get the contribution from \hat{f} . First, observe that by the triangle inequality,

$$d(t, k_x) \geq d(t, x) - d(x, k_x) > d(t, x)(1 - 1/(3\alpha(k))) \geq 2d(t, x)/3. \quad (2)$$

By another application of the triangle inequality, using that \hat{f} is non-expansive, and that f has distortion at most $\alpha(k)$ on the terminals, we get the required bound on the contraction:

$$\begin{aligned} \|F(t) - F(x)\|_p &\geq \|\hat{f}(t) - \hat{f}(x)\|_p \\ &\geq \|\hat{f}(t) - \hat{f}(k_x)\|_p - \|\hat{f}(k_x) - \hat{f}(x)\|_p \\ &\geq \|f(t) - f(k_x)\|_p - d(x, k_x) \\ &\geq \frac{d(t, k_x)}{\alpha(k)} - \frac{d(t, x)}{3\alpha(k)} \\ &\stackrel{(2)}{\geq} \frac{2d(t, x)}{3\alpha(k)} - \frac{d(t, x)}{3\alpha(k)} \\ &= \frac{d(t, x)}{3\alpha(k)}. \end{aligned}$$

□

Remark: The results of Theorem 1 and Theorem 2 hold also if \mathcal{X} is a family of graphs, rather than of metrics, provided that the embedding for this family has the promised guarantees even for graphs with Steiner nodes. (E.g., if $\hat{Z} \in \mathcal{X}$ is a graph and Z is a set of vertices of size m , then there exists a (Lipschitz extendable) embedding of (Z, d_Z) to $\ell_p^{\gamma(m)}$ with distortion $\alpha(m)$, where d_Z is the shortest path metric on \hat{Z} induced on Z .) We note that many embeddings of graph families satisfy this condition, e.g. the embedding of [KLMN05] of planar and minor-free graphs.¹⁰

Corollary 1. *Let (X, d) be a metric space on n points, and $K \subseteq X$ a set of terminals of size $|K| = k$. Then for any $1 \leq p \leq \infty$,*

1. (X, d) can be embedded to $\ell_p^{O(\log k)}$ with terminal distortion $O(\log k)$.
2. If (X, d) is an ℓ_2 metric space, it can be embedded to $\ell_2^{O(\log k)}$ with terminal distortion $O(1)$.
3. For any $t \geq 1$ there exists a $(4t-1)$ -terminal (metric) spanner of X with at most $O(k^{1+1/t}) + n$ edges.

¹⁰We remark that this requirement is needed for those graph families for which the following question is open: given a graph Z in the family with terminals K , is there another graph in the family over the vertex set K , that preserves the shortest-path distances with respect to Z (up to some constant). This question is open, e.g., for planar graphs.

4. If (X, d) is an ℓ_2 metric, for any $t \geq 1$ there exists a $O(t)$ -terminal spanner of X with at most $O(k^{1+1/t^2}) + n$ edges.
5. (X, d) can be embedded to $\ell_p^{O(\log n + \log^2 k)}$ with strong terminal distortion $(O(\log k), O(\log n))$.
6. If (X, d) is a decomposable metric,¹¹ it can be embedded to ℓ_p with strong terminal distortion $(O((\log k)^{\min\{\frac{1}{2}, \frac{1}{p}\}}), O((\log n)^{\min\{\frac{1}{2}, \frac{1}{p}\}}))$.
7. If (X, d) is a negative type metric¹² it can be embedded to ℓ_2 with strong terminal distortion $(\tilde{O}(\sqrt{\log k}), \tilde{O}(\sqrt{\log n}))$. (We note that any ℓ_1 metric is of negative type.)
8. For any $t \geq 1$, (X, d) can be embedded to $\ell_\infty^{O(t \cdot k^{1/t} \cdot \log k)}$ with terminal distortion $O(t)$.

The first two items and the last one use the first assertion of Theorem 1, the next two use its second assertion, and the next three apply Theorem 2. The corollary follows from known embedding results: (1) and (5) are from [Bou85], with improved dimension due to [ABN11], (2) is from [JL84], (3) is from [ADD⁺93] and (4) from [HPIS13], (6) from [KLMN05], (7) from [ALN07, ALN08], and (8) from [Mat96, ABN11].

4 Graph Terminal Spanners

While Theorem 1 provides a general approach to obtain terminal spanners, it cannot provide spanners which are subgraphs of the input graph. We devise a construction of such terminal spanners in this section, while somewhat increasing the number of edges. Specifically, we show the following.

Theorem 3. *For any parameter $t \geq 1$, a graph $G = (V, E)$ on n vertices, and a set of terminals $K \subseteq V$ of size k , there exists a $(4t - 1)$ -terminal graph spanner with at most $O(n + \sqrt{n} \cdot k^{1+1/t})$ edges.*

Remark: Note that the number of edges is linear in n whenever $k \leq n^{1/(2(1+1/t))}$.

We shall use the following result:

Theorem 4 ([CE05]). *Given a weighted graph $G = (V, E)$ on n vertices and a set $P \subseteq \binom{V}{2}$ of size p , then there exists a subgraph G' with $O(n + \sqrt{n} \cdot p)$ edges, such that for all $\{u, v\} \in P$, $d_G(u, v) = d_{G'}(u, v)$.*

Proof of Theorem 3. The construction of the subgraph spanner with terminal distortion will be as follows. Consider the metric induced on the terminals K by the shortest path metric on G . Create a $(2t - 1)$ (metric) spanner H' of this metric, using [ADD⁺93], and let $P \subseteq \binom{K}{2}$ be the set of edges of H' . Note that $p = |P| \leq O(k^{1+1/t})$. Now, apply Theorem 4 on the graph G with the set of pairs P , and obtain a graph G' that for every $\{u, v\} \in P$, has $d_{G'}(u, v) = d_G(u, v)$. This implies that G' is a $(2t - 1)$ -spanner for each pair of vertices $u, w \in K$. Moreover, G' has at most $O(n + \sqrt{n} \cdot p)$ edges. Next, we create the final spanner H out of G' by adding a shortest path tree (actually forest) in G

¹¹See [KLMN05] for a definition of decomposability. We remark that doubling metrics, planar metrics, and more generally metrics arising from graphs excluding a fixed minor, are all decomposable.

¹² (X, d) is called a negative type metric if (X, \sqrt{d}) embeds isometrically to ℓ_2 .

with the set K as its root. This will guarantee that the spanner H will have for each non-terminal, a shortest path to its closest terminal in G . This concludes the construction of H , and now we turn to bounding the distortion. Since H is a subgraph clearly it is non-contractive. Fix any $v \in K$ and $u \in V$, let k_u be the closest terminal to u , then $d_G(k_u, v) \leq d_G(k_u, u) + d_G(u, v) \leq 2d_G(u, v)$, and thus

$$d_H(u, v) \leq d_H(u, k_u) + d_{G'}(k_u, v) \leq d_G(u, v) + (2t - 1)d_G(k_u, v) \leq (4t - 1)d_G(u, v) .$$

Finally observe that the total number of edges in H is at most $O(n + \sqrt{n} \cdot p) = O(n + \sqrt{n} \cdot k^{1+1/t})$. \square

5 Light Terminal Trees for General Graphs

In this section we find a single spanning tree of a given graph, that has both light weight, and approximately preserves distances from a set of specified terminals. Theorem 1 can provide a tree with terminal distortion $2k - 1$ (using that any graph has a tree with distortion $n - 1$), but that tree may not be a subgraph and may have large weight. The result of this section is summarized as follows.

Theorem 5. *For any parameter $\alpha \geq 1$, given a weighted graph $G = (V, E, w)$, and a subset of terminals $K \subseteq V$ of size k , there exists a spanning tree T of G with terminal distortion at most $k \cdot \alpha + (k - 1)\alpha^2$ and lightness at most $2\alpha + 1 + \frac{2}{\alpha - 1}$.*

When substituting $\alpha = 1$ in Theorem 5 we obtain a single tree with terminal distortion exactly $2k - 1$, which is optimal (see Theorem 6), but no guarantee for lightness. More generally, for small $\epsilon > 0$, by setting $\alpha = 1 + \frac{\epsilon}{3k}$ we get terminal distortion $2k - 1 + \epsilon$ and lightness $3 + \frac{6k}{\epsilon} + \frac{2\epsilon}{3k} = O(\frac{k}{\epsilon})$. Also, note that the bound $2\alpha + 1 + \frac{2}{\alpha - 1}$ is minimized by setting $\alpha = 2$, so there is no point in using the theorem with $\alpha > 2$.

Next we describe the algorithm for constructing a spanning tree that satisfies the assertion of Theorem 5. We shall assume w.l.o.g. that all edge weights are different, and every two different paths have different lengths. If it is not the case, then one can break ties in an arbitrary (but consistent) way.

A spanning tree T is an (α, β) -SLT (Shallow Light Tree) (see [KRY95, ABP92, ES11]) with respect to a root $v \in V$, if for all $u \in V$, $d_T(v, u) \leq \alpha \cdot d_G(v, u)$, and T has lightness β . A small modification of an SLT-constructing algorithm produces for any subset $K \subset V$, a forest F , such that every component of F contains exactly one vertex of K .¹³ The forest F has distortion α with respect to K , and *lightness* $1 + \frac{2}{\alpha - 1}$. (Such a forest F is said to have distortion α with respect to K , if for every vertex $u \in V$, $d_F(K, u) \leq \alpha \cdot d_G(K, u)$.)

The algorithm starts by building the aforementioned SLT-forest F from the terminal set K . No two terminals belong to the same connected component of F . Denote $K = \{v_1, \dots, v_k\}$, let V_i be the unique connected component of F containing v_i , and let $T_i \subseteq F$ be the edges of the forest F induced by V_i . It follows that for every $u \in V_i$, $d_F(K, u) = d_{T_i}(v_i, u) \leq \alpha \cdot d_G(K, u)$. Let $G' = (K, E', w')$ be the graph in which two terminals share an edge between them if and only if there is an edge between the components V_i to V_j in G . Formally, $E' = \{\{v_i, v_j\} : \exists u_i \in V_i, u_j \in V_j \text{ such that } \{u_i, u_j\} \in E\}$.

¹³To obtain such a forest F , one should add a new vertex v_K to the graph and connect it to each of the vertices of K with edges of weight zero. Then we compute an (α, β) -SLT with respect to v_K in the modified graph. Finally, we remove v_K from the SLT. The resulting forest is F .

The weight $w'(v_i, v_j)$ is defined to be the length of the shortest path between v_i and v_j which uses *exactly one edge* that does not belong to F . (In other words, among all the paths between v_i and v_j in G which use exactly one edge that does not belong to F , let P be the shortest one. Then $w'(v_i, v_j) = w(P)$.) Note also that $w'(v_i, v_j)$ is given by $w'(v_i, v_j) = \min_{e \in E} \{d_{F \cup \{e\}}(v_i, v_j)\}$. We call the edge $e_{i,j} = \{u_i, u_j\}$ that implements this minimum ($w'(v_i, v_j) = d_{F \cup \{e_{i,j}\}}(v_i, v_j)$) the *representative edge* of $\{v_i, v_j\}$. (Recall that w.l.o.g. the shortest paths, and thus the representative edges, are unique.) Observe that $\{v_i, v_j\} \in E'$ implies that $w'(v_i, v_j) < \infty$. Let T' be the *MST* of G' . Define $R = \left\{ e_{i,j} \mid e_{i,j} \text{ is the representative edge of } e'_{i,j} = (v_i, v_j) \in T' \right\}$. Finally, set $T = F \cup R = \bigcup_{i=1}^k T_i \cup R$. Obviously, T is a spanning tree of G . This concludes the construction, next we turn to the analysis.

As an embedding of a graph into its spanning tree is non-contractive, the tree T will have terminal distortion α if for all $v \in K$, $u \in V$, $d_T(v, u) \leq \alpha \cdot d_G(v, u)$.

The next lemma shows that for every pair of terminals v_i, v_j , there is a path between them in G' in which all edges have weight (with respect to w') at most $\alpha \cdot d_G(v_i, v_j)$.

Lemma 2. [*The bottleneck lemma:*] *For every $v_i, v_j \in K$, there exists a path $P : v_i = z_0, z_1, \dots, z_r = v_j$ in G' such that for every $s = 0, 1, \dots, r-1$, it holds that $\{z_s, z_{s+1}\} \in E'$ and $w'(z_s, z_{s+1}) \leq \alpha \cdot d_G(v_i, v_j)$.*

Proof. Let $P_{i,j} : v_i = u_0, u_1, \dots, u_s = v_j$ be the shortest path from v_i to v_j in G , i.e., $w(P_{i,j}) = d_G(v_i, v_j)$. For each $0 \leq a \leq s$, denote by $V^{(a)}$ the connected component of F that contains u_a , and let $v^{(a)}$ be the unique terminal in that component. Consider the path $P = v^{(0)}, v^{(1)}, \dots, v^{(s)}$. (This path is not necessarily simple. In particular, it might contain self-loops. See Figure 1 for an illustration.) For every index $a < s$,

$$\begin{aligned}
w'(v^{(a)}, v^{(a+1)}) &\stackrel{(1)}{\leq} d_{F \cup \{u_a, u_{a+1}\}}(v^{(a)}, v^{(a+1)}) \\
&= d_F(v^{(a)}, u_a) + d_G(u_a, u_{a+1}) + d_F(u_{a+1}, v^{(a+1)}) \\
&\stackrel{(2)}{\leq} \alpha \cdot d_G(v_i, u_a) + d_G(u_a, u_{a+1}) + \alpha \cdot d_G(v_j, u_{a+1}) \\
&\leq \alpha \cdot (d_G(v_i, u_a) + d_G(u_a, u_{a+1}) + d_G(v_j, u_{a+1})) \\
&\stackrel{(3)}{=} \alpha \cdot d_G(v_i, v_j).
\end{aligned}$$

Note that if for some index a it holds that $v^{(a)} = v^{(a+1)}$ then $w'(v^{(a)}, v^{(a+1)}) = 0$, and the inequality above holds trivially. Otherwise, if $v^{(a)} \neq v^{(a+1)}$, then inequality (1) follows from the assumptions that $\{u_a, u_{a+1}\} \in E$, $u_a \in V^{(a)}$, $u_{a+1} \in V^{(a+1)}$. Inequality (2) follows from the properties of the *SLT* tree T (as $d_F(v^{(a)}, u_a) = d_F(K, u_a) \leq \alpha \cdot d_G(K, u_a) \leq \alpha \cdot d_G(v_i, u_a)$). Equality (3) follows because the edge $\{u_a, u_{a+1}\}$ is on the shortest path from v_i to v_j in G .

In particular, one can remove cycles from P and obtain a simple path with the desired properties. We get a simple path P' such that for every edge v, v' on this path, we have $w'(v, v') \leq \alpha \cdot d_G(v_i, v_j)$, as required. \square

The following is a simple corollary.

Corollary 3. *For $\{v_i, v_j\} \in T'$, we have $w'(v_i, v_j) = d_T(v_i, v_j) \leq \alpha \cdot d_G(v_i, v_j)$.*




Figure 1: An illustration for the bottleneck lemma: v_i and v_j are terminals. The edge $\{u_a, u_{a+1}\}$ belongs to the shortest path from v_i to v_j in G . We conclude that for terminals $v^{(a)}, v^{(a+1)}$ such that $u_a \in V^{(a)}$ and $u_{a+1} \in V^{(a+1)}$, it holds that $w'(v^{(a)}, v^{(a+1)}) \leq \alpha \cdot d_G(v_i, v_j)$.




Figure 2: The two paths P' and $P_{j,i}$ considered in the proof of Lemma 4. The path P' is contained in T' , while all edges of $P_{j,i}$ have weight at most $\alpha \cdot d_G(v_i, v_j)$.

Proof. By Lemma 2, $w'(v_i, v_j) \leq \alpha \cdot d_G(v_i, v_j)$. (Indeed, otherwise the edge $\{v_i, v_j\}$ is strictly the heaviest edge in a cycle in G' , contradiction to the assumption that it belongs to the MST of G' .) Since $\{v_i, v_j\} \in E'$ and the representative edge of $\{v_i, v_j\}$ was taken into T , it follows that $w'(v_i, v_j) = d_T(v_i, v_j)$. \square

We conclude the following lemma, which bounds the distortion of terminal pairs.

Lemma 4. For $v_i, v_j \in K$, we have $d_T(v_i, v_j) \leq d_{T'}(v_i, v_j) \leq \alpha \cdot (k - 1) \cdot d_G(v_i, v_j)$.

Proof. Let $P' : v_i = v^{(0)}, v^{(1)}, \dots, v^{(h)} = v_j$ be the (unique) path in T' between v_i and v_j . Since T' is a spanning tree of the k -vertex graph G' , it follows that $h \leq k - 1$. Observe also that for every index $a \in [h - 1]$, by Corollary 3, the edge $w'(v^{(a)}, v^{(a+1)}) = d_T(v^{(a)}, v^{(a+1)})$. Also, we next argue that $w'(v^{(a)}, v^{(a+1)}) \leq \alpha \cdot d_G(v_i, v_j)$. Indeed, suppose for contradiction that $w'(v^{(a)}, v^{(a+1)}) >$

$\alpha \cdot d_G(v_i, v_j)$. Let $P_{j,i}$ be a path between v_j and v_i in G' such that all its edges have weight at most $\alpha \cdot d_G(v_i, v_j)$. The existence of this path is guaranteed by Lemma 2. In particular, since $w'(v^{(a)}, v^{(a+1)}) > \alpha \cdot d_G(v_i, v_j)$, it follows that $\{v^{(a)}, v^{(a+1)}\} \notin P_{j,i}$. Consider the cycle $P' \circ P_{j,i}$ in G' . It is not necessarily a simple cycle, but since $\{v^{(a)}, v^{(a+1)}\} \notin P_{j,i}$, the edge $\{v^{(a)}, v^{(a+1)}\}$ belongs to a simple cycle C contained in $P' \circ P_{j,i}$. The heaviest edge of C clearly does not belong to $P_{j,i}$, because the edge $\{v^{(a)}, v^{(a+1)}\}$ is heavier than each of them. Hence the heaviest edge belongs to P' , but $P' \subseteq T'$. This is a contradiction to the assumption that T' is an MST of G' . (See Figure 2 for an illustration). Hence $d_T(v^{(a)}, v^{(a+1)}) = w'(v^{(a)}, v^{(a+1)}) \leq \alpha \cdot d_G(v_i, v_j)$. Finally,

$$\begin{aligned} d_T(v_i, v_j) &\leq \sum_{a=0}^{h-1} d_T(v^{(a)}, v^{(a+1)}) = \sum_{a=0}^{h-1} w'(v^{(a)}, v^{(a+1)}) \\ &\leq \sum_{a=0}^{h-1} \alpha \cdot d_G(v_i, v_j) \leq h \cdot \alpha \cdot d_G(v_i, v_j) \leq \alpha \cdot (k-1) \cdot d_G(v_i, v_j). \end{aligned}$$

□

Next, we analyze the terminal distortion of T .

Lemma 5. *The terminal distortion of T is at most $k \cdot \alpha + (k-1)\alpha^2$.*

Proof. For each terminal $v_i \in K$ and any vertex $u \in V_j$, for some $j \in \{1, 2, \dots, k\}$, it holds that

$$\begin{aligned} d_T(v_i, u) &\leq d_T(v_i, v_j) + d_T(v_j, u) \leq \alpha \cdot (k-1) \cdot d_G(v_i, v_j) + \alpha \cdot d_G(v_i, u) \\ &\leq \alpha \cdot (k-1) \cdot (d_G(v_i, u) + d_G(u, v_j)) + \alpha \cdot d_G(v_i, u) \\ &\leq \alpha \cdot (k-1) \cdot (d_G(v_i, u) + \alpha \cdot d_G(v_i, u)) + \alpha \cdot d_G(v_i, u) \\ &= (k \cdot \alpha + (k-1)\alpha^2) \cdot d_G(v_i, u). \end{aligned}$$

The last inequality is because $d_G(v_j, u) \leq d_F(v_j, u) = d_F(K, u) \leq \alpha \cdot d_G(K, u) \leq \alpha \cdot d_G(v_i, u)$.

□

Next, we analyze the lightness of T . A tree $T = (S, E', w')$ is called a *Steiner tree* for a graph $G = (V, E, w)$ if (1) $V \subseteq S$, (2) for any edge $e \in E \cap E'$, the edge has the same weight in both G and T , i.e. $w(e) = w'(e)$, and (3) for any pair of vertices $u, v \in V$ it holds that $d_T(u, v) \geq d_G(u, v)$. The *minimum Steiner tree* T of G , denoted $SMT(G)$, is a Steiner tree of G with minimum weight. It is well-known that for any graph G , $w(SMT(G)) \geq \frac{1}{2}w(MST(G))$. (See, e.g., [GP68], Section 10.) The next lemma bounds the lightness of the tree T .

Lemma 6. *The lightness of T is bounded by $\Psi(T) \leq 2\alpha + 1 + \frac{2}{\alpha-1}$.*

Proof. The main challenge is to bound $w(R)$. (Recall that R is the set of the representative edges of T' , where T' is the MST of the graph G' .) Consider an edge $\{v_i, v_j\} \in T'$, its representative edge $e_{i,j}$ is part of a path from v_i to v_j of weight $w'(v_i, v_j)$. In particular $w(e_{i,j}) \leq w'(v_i, v_j)$. Therefore $w(R) = \sum_{e_{i,j} \in R} w(e_{i,j}) \leq \sum_{\{v_i, v_j\} \in T'} w'(v_i, v_j) = w'(T')$. Next we provide an upper bound for $w'(T')$. Define the graph \tilde{G} as the complete graph on the vertex set K , with weights \tilde{w} induced by d_G (the shortest path distances in G). Let \tilde{T} be the *MST* of \tilde{G} . We build a new tree \hat{T} by the following process:

1. Let $\hat{T} \leftarrow \tilde{T}$;
2. For each $\{v_i, v_j\} = \tilde{e} \in \tilde{T}$:
 - (a) Let $P_{\tilde{e}}$ be a path from v_i to v_j which consists of edges in E' , such that for each edge e in $P_{\tilde{e}}$, $w'(e) \leq \alpha \cdot d_G(v_i, v_j) = \alpha \cdot \tilde{w}(\tilde{e})$; (By Lemma 2, such a path exists);
 - (b) Let $e' \in P_{\tilde{e}}$ be an edge such that $(\hat{T} \setminus \{\tilde{e}\}) \cup \{e'\}$ is connected;
 - (c) Set $\hat{T} \leftarrow (\hat{T} \setminus \{\tilde{e}\}) \cup \{e'\}$;

In each step in the loop we replace an edge $\tilde{e} = \{v_i, v_j\}$ from \tilde{T} by an edge e' from G' of weight $w'(e') \leq \alpha \cdot \tilde{w}(\tilde{e})$. Hence, the resulting tree \hat{T} is a spanning tree of G' , and $w'(\hat{T}) \leq \alpha \cdot \tilde{w}(\tilde{T})$. Since T' is the MST of G' , it follows that $w'(T') \leq w'(\hat{T})$. The MST of G is a Steiner tree for \tilde{G} , so that $\tilde{w}(SMT(\tilde{G})) \leq w(MST(G))$. Also, $\tilde{w}(MST(\tilde{G})) = \tilde{w}(\tilde{T}) \leq 2 \cdot \tilde{w}(SMT(\tilde{G})) \leq 2 \cdot w(MST(G))$. We obtain that

$$w(R) \leq w'(T') \leq w'(\hat{T}) \leq \alpha \cdot \tilde{w}(\tilde{T}) \leq 2 \cdot \alpha \cdot w(MST(G)).$$

Since $w(F) \leq \left(1 + \frac{2}{\alpha-1}\right) \cdot w(MST(G))$, we conclude that

$$w(T) = w(R \cup F) = w(R) + w(F) \leq \left(2\alpha + 1 + \frac{2}{\alpha-1}\right) \cdot w(MST(G)).$$

□

6 Lower Bounds on Distortion-Lightness Tradeoffs for a Single Tree

In this section we provide lower bounds on the terminal distortion and lightness of k -terminal trees. For each lower bound we start with showing a lower bound for graphs and proceed to showing a lower bound for metric spaces. (Observe that the latter is more general.) The lower bounds exhibit similar tradeoffs in both cases, while the analysis is significantly simpler for graphs than for metric spaces.

6.1 A lower bound for terminal distortion

Recall that Theorem 5 implies that every graph contains a spanning tree with terminal distortion $2k - 1$. Here is a matching lower bound.

Theorem 6. *For any k there is a weighted graph G with $n = 2k$ vertices and k terminals such that any spanning tree has terminal distortion at least $2k - 1$.*

Proof. Consider the cycle graph C_{2k} such that the terminals and the non-terminal vertices alternate (as usual V' stands for the set of terminals). Any spanning tree T of C_{2k} is obtained by removing

a single edge $\{v_i, v_{i+1}\}$. Observe that either v_i or v_{i+1} is a terminal. Hence the terminal distortion is at least

$$\frac{d_{C_{2k} \setminus \{v_i, v_{i+1}\}}(v_i, v_{i+1})}{d_{C_{2k}}(v_i, v_{i+1})} = 2k - 1.$$

□

Remark 2. *The same result for any $n > 2k$ follows by adding $n - 2k$ non-terminal vertices to the graph, and connecting them all to an arbitrary vertex of C_{2k} . In the resulting graph every spanning tree has terminal distortion at least $2k - 1$.*

Next we extend Theorem 6 to metric spaces.

Theorem 7. *For any k there is a metric space (M, d) with $n = 2k$ vertices and k terminals such that any spanning tree for (M, d) has terminal distortion at least $2k - 1$.*

Proof. Let M be the metric generated by the cycle graph C_{2k} , where there are k terminals and the terminals and the non-terminal vertices alternate. In [Gup01b], Lemma 7.1, Gupta showed that the distortion of every spanning tree of M is at least $n - 1 = 2k - 1$. Moreover, the maximal distortion is achieved by an original edge e of C_{2k} (because if we had distortion t on every original edge, the triangle inequality implies distortion at most t on every pair). One of the two endpoints of e is a terminal. Hence the terminal distortion of any spanning tree for M is at least $2k - 1$. □

Remark 3. *One can extend Theorem 7 to $n > 2k$ by adding $n - 2k$ non-terminal vertices to M at a large distance from all the terminals. For a spanning tree T of M , if some shortest path between a terminal v and a vertex u which belongs to C_{2k} uses at least one of the new vertices then the terminal distortion is too high. Hence if the terminal distortion is small, the tree restricted to C_{2k} is connected. Hence by Theorem 7 it follows that the terminal distortion is at least $2k - 1$.*

6.2 A lower bound on the lightness in graphs

In this section we prove a lower bound on the tradeoff between terminal distortion and lightness of k -terminal trees.

Theorem 8. *For any positive integer parameters k, n and $\epsilon > 0$ such that $k \leq \frac{\epsilon}{2}n$, there exists a graph G with $N = (n + 1)k$ vertices, such that any spanning tree T for G with terminal distortion at most $(2k - 1) \left(1 + \frac{\epsilon}{k^2}\right)$ has lightness at least $\Omega\left(\frac{1}{\epsilon}\right)$.*

Proof. Consider the following graph G on $N = k \cdot (n + 1)$ vertices. There are k terminals $V' = \{v_0, v_1, \dots, v_{k-1}\}$ in G . For every index $i \in [0, k - 1]$, there is a n -vertex path P_i . All edges in these paths have unit weight. Also, for each index $i \in [0, k - 1]$, both v_i and $v_{i+1 \pmod k}$ are connected to every vertex in P_i by edges of weight w , where $w > 1$ is a parameter that will be determined later. (To simplify the notation we will henceforward write v_{i+1} instead $v_{i+1 \pmod k}$). Generally, all the arithmetic operations on indexes of vertices v_0, \dots, v_{k-1} and paths P_0, \dots, P_{k-1} are performed modulo k .) See Figure 3 for an illustration.

Each spanning tree of G contains $N - 1 = k + kn - 1$ edges. There are $k \cdot (n - 1)$ edges of unit weight, and all the other edges have weight w . Hence the weight of the MST is at least $k(n - 1) \cdot 1 + (2k - 1) \cdot w$. It is easy to verify that there actually exists a spanning tree of that weight.




Figure 3: An illustration of the graph used in the proof of Theorem 8. The k terminals are depicted by the big dots. The vertices v_i and v_{i+1} are connected to each vertex of an n -vertex path P_i by edges of weight w .

We will show that for any $\beta < \frac{1}{w(2k-1)}$, every tree with terminal distortion at most $(2k-1)(1+\beta)$ has weight at least $\Omega(nw)$.

Let T be a spanning tree for G with terminal distortion at most $(2k-1)(1+\beta)$, for some $\beta > 0$. There exists an index i such that the path between v_i and v_{i+1} in T does not use vertices from the path P_i . (Otherwise there is a cycle in T that passes through v_0, \dots, v_{k-1} .) Without loss of generality assume that $i = 0$. Therefore, $d_T(v_0, v_1) \geq (k-1) \cdot 2w$.

Claim 7. For every vertex u in P_0 , if $\beta < \frac{1}{(2k-1)w}$ then either (v_0, u) or (v_1, u) is an edge of T .

Proof. Without loss of generality the shortest path from v_0 to u goes through v_1 . Assume for contradiction that the edge (v_1, u) does not belong to T . Then the terminal distortion is at least

$$\begin{aligned} \frac{d_T(v_0, u)}{d_G(v_0, u)} &= \frac{d_T(v_0, v_1) + d_T(v_1, u)}{d_G(v_0, u)} \geq \frac{(k-1)2w + w + 1}{w} \\ &= 2k - 1 + \frac{1}{w} = (2k-1) \left(1 + \frac{1}{(2k-1)w} \right) \\ &> (2k-1)(1+\beta), \end{aligned}$$

contradiction. □

By Claim 7, $\beta < \frac{1}{(2k-1)w}$ implies that $w(T) \geq n \cdot w$. Hence for every k -terminal tree T with terminal distortion at most $(2k-1)(1+\beta)$, with $\beta < \frac{1}{(2k-1)w}$, it holds that $\Psi(T) = \frac{w(T)}{w(MST)} \geq \frac{nw}{k(n-1) + (2k-1)w}$. We set $w = \frac{k}{\epsilon}$. Then the condition $\beta < \frac{1}{(2k-1)w}$ translates to $\beta < \frac{\epsilon}{(2k-1)k}$. This condition implies that

$$\Psi(T) = \frac{n \frac{k}{\epsilon}}{k(n-1) + (2k-1) \frac{k}{\epsilon}} = \frac{nk}{\epsilon k(n-1) + (2k-1)k} \geq \frac{nk}{\epsilon nk + 2k^2} = \frac{n}{\epsilon n + 2k}.$$

As $k \leq \frac{\epsilon}{2}n$ we obtain $\Psi(T) \geq \frac{1}{2\epsilon}$. □

Our algorithm from Theorem 5, for $\alpha = 1 + O(\epsilon)$, guarantees terminal distortion $(2k - 1)(1 + O(\epsilon))$ and lightness $O(\frac{1}{\epsilon})$. In the graph G from the above proof lightness smaller than $\frac{1}{\epsilon}$ implies terminal distortion at least $(2k - 1)\left(1 + \frac{\epsilon}{(2k-1)k}\right)$. Hence our bounds are tight for $k = O(1)$, but generally there is a gap of $O(k^2)$ between the error term in the upper and lower bounds.

In Appendix 7 we extend this lower bound to metric spaces.

7 A lower bound on the tradeoff between lightness and stretch for terminal trees in metric spaces

In this section we extend our lower bound from Section 6.2 to the metric case.

The metric space case is similar to the graph case. We will use the metric closure of the graph from the previous section. The main difficulty is however to show that the non-graph edges do not help at all.

For a positive integer parameter k and a k -sequence n_0, n_1, \dots, n_{k-1} of positive integer numbers, we define a graph $G_{k, n_0, n_1, \dots, n_{k-1}}$. The graph has $n = k + \sum_i n_i$ vertices and k terminals. The k terminals are $V' = \{v_0, \dots, v_{k-1}\}$. For every index $i \in [0, k - 1]$ there is an n_i -vertex path P_i . (We will use P_i to denote both the path, and the set of vertices in the path.) All edges in these paths have unit weight. Also, for each $i \in [0, k - 1]$, both v_i and v_{i+1} (the index arithmetic is modulo k) are connected to every vertex in P_i by edges of weight w , for a parameter $w > 1$. Observe that the graph G from Section 6.2 satisfies $G = G_{k, n_0, n_1, \dots, n_{k-1}}$ with $n_0 = n_1 = \dots = n_{k-1} = n$. Let $\overline{G}_{k, n_0, \dots, n_{k-1}}$ denote the metric closure of $G_{k, n_0, \dots, n_{k-1}}$. We also write $G = G_{k, n_0, \dots, n_{k-1}}$ and $\overline{G} = \overline{G}_{k, n_0, \dots, n_{k-1}}$.

Lemma 8. *For any spanning tree T of $\overline{G}_{k, n_0, \dots, n_{k-1}}$ there exists an index i such that for any vertex $u \in P_i$ either $\frac{d_T(v_i, u)}{d_{\overline{G}_{k, n_0, \dots, n_{k-1}}}(v_i, u)} \geq 2k - 1$ or $\frac{d_T(v_{i+1}, u)}{d_{\overline{G}_{k, n_0, \dots, n_{k-1}}}(v_{i+1}, u)} \geq 2k - 1$.*

Remark 4. *Observe that for a graph spanning tree T this lemma follows directly from the observation that there exists an index i such that the path in T from v_i to v_{i+1} does not contain vertices of P_i . Indeed, for this index i and a vertex $u \in P_i$, either $(v_i, u) \notin T$ or $(v_{i+1}, u) \notin T$. In the former case $\frac{d_T(v_i, u)}{d_{G_{k, n_0, \dots, n_{k-1}}}(v_i, u)} \geq 2k - 1$ and in the latter $\frac{d_T(v_{i+1}, u)}{d_{G_{k, n_0, \dots, n_{k-1}}}(v_{i+1}, u)} \geq 2k - 1$. The lemma proves this statement in a much greater generality, specifically, for T being a spanning tree of the metric closure $\overline{G}_{k, n_0, \dots, n_{k-1}}$ of $G_{k, n_0, \dots, n_{k-1}}$.*

Proof. For a spanning tree T of $\overline{G}_{k, n_0, \dots, n_{k-1}}$ and an index $i \in [0, k - 1]$, denote by $t_{T, i}$ the number of vertices u in P_i such that $\frac{d_T(v_i, u)}{d_{\overline{G}_{k, n_0, \dots, n_{k-1}}}(v_i, u)} \geq 2k - 1$ or $\frac{d_T(v_{i+1}, u)}{d_{\overline{G}_{k, n_0, \dots, n_{k-1}}}(v_{i+1}, u)} \geq 2k - 1$. We will show that for any spanning tree T there exists an index i such that $t_{T, i} = n_i$. For a tree T , define $\gamma(T) = \min_i \{n_i - t_{T, i}\}$. Observe that $\gamma(T) \geq 0$. It suffices to prove that for every tree T , $\gamma(T) = 0$. Also, let $\mu = \max_T \{\gamma(T)\}$, where the maximum is taken over all spanning trees of $\overline{G}_{k, n_0, \dots, n_{k-1}}$. It is enough to show that $\mu = 0$, i.e., that for every spanning tree T , $\gamma(T) = 0$.

For each vertex, we define the right and the left *hemisphere* with respect to this vertex. Consider a graph, where we replace each path P_i by a node p_i . We obtain the $2k$ -cycle C_{2k} . The *right hemisphere* of v_i consists of all the vertices between v_i and its antipodal vertex (in the path $p_i, v_{i+1}, p_{i+1}, \dots$), while the *left hemisphere* of v_i consists of all the vertices in the other shortest




Figure 4: An illustration of the partition of the graph to the right and the left hemispheres with respect to a terminal v_i . This example is for odd k . Note that the path $P_{i+\frac{k-1}{2}}$ is both in the right and the left hemispheres.

path from v_i to its antipodal vertex $(p_{i-1}, v_{i-1}, p_{i-2}, \dots)$. Formally, for a terminal v_i , if k is even, let $R(v_i) = \cup \{P_i, \{v_{i+1}\}, P_{i+1}, \dots, \{v_{i+\frac{k}{2}}\}\}$ and $L(v_i) = \cup \{\{v_{i+\frac{k}{2}}\}, P_{i+\frac{k}{2}}, \{v_{i+\frac{k}{2}+1}\}, \dots, P_{i-1}\}$ denote the right and the left hemispheres with respect to v_i , respectively. Similarly, if k is odd then $R(v_i) = \cup \{P_i, \{v_{i+1}\}, \dots, P_{i+\frac{k-1}{2}}\}$ and $L(v_i) = \cup \{P_{i+\frac{k-1}{2}}, \{v_{i+\frac{k+1}{2}}\}, \dots, P_{i-1}\}$. In addition we define hemispheres for non-terminal vertices. For an index i , all the vertices in P_i have the same hemispheres. For a vertex $u \in P_i$, if k is even, then $R(u) = \cup \{\{v_{i+1}\}, P_{i+1}, \{v_{i+2}\}, \dots, P_{i+\frac{k}{2}}\}$ and $L(u) = \cup \{P_{i+\frac{k}{2}}, \{v_{i+\frac{k}{2}+1}\}, P_{i+\frac{k}{2}+1}, \dots, \{v_i\}\}$. If k is odd then $R(u) = \cup \{\{v_{i+1}\}, P_{i+1}, \{v_{i+2}\}, \dots, \{v_{i+\frac{k+1}{2}}\}\}$ and $L(u) = \cup \{\{v_{i+\frac{k+1}{2}}\}, P_{i+\frac{k+1}{2}}, \{v_{i+\frac{k+1}{2}+1}\}, \dots, \{v_i\}\}$. (See Figure 4 for an illustration.)

For a vertex $u \in V$ and an index $i \in [0, k-1]$, we say that all the edges from u to P_i , i.e., $\{u, z \mid z \in P_i\}$, are of the *same type*. In addition, the edge from u to v_i has a unique type. (Note that each non-terminal vertex might have edges of $2k$ different types, while a terminal vertex can be incident to edges of $2k-1$ different types.) An edge $e = \{u, z\}$ such that u, z are in the same path P_i called a *path-internal edge*. A vertex u is called a *same-path neighbor* of a, b in the tree T , if there are edges $\{u, a\}, \{u, b\}$ in T and a, b belong to the same path P_i . For a simple path π in $\overline{G}_{k, n_0, \dots, n_{k-1}}$ we say that π is a *one-sided path* if for any internal vertex x in π , the two edges $(x, y_1), (x, y_2)$ which are incident on x in π , connect x to different hemispheres, i.e., e.g., y_1 is in the left hemisphere with respect to x , and y_2 is in the right one.

The proof (that $\mu = 0$) is by induction on $\sum_{i=1}^k n_i = A$. The base case where $\sum_{i=1}^k n_i = k$ (i.e., for all $i \in [0, k-1]$, $n_i = 1$) follows by Theorem 7.

The induction step: assume that the claim is true for A and we will prove it for $A+1$. Let T be some spanning tree of $\overline{G}_{k, n_0, \dots, n_{k-1}}$ with minimal weight among all the trees with $\gamma(T) = \mu$. By our assumption $\sum_{i=1}^k n_i = A+1$.

Claim 9. For any vertex u (either terminal or a non-terminal one), if there exist two edges




Figure 5: An illustration of the special edges that we use in the proof of Lemma 8.




Figure 6: An illustration of deleting a path-internal edge $\{u, z\}$ in the tree T .

$\{u, a\}, \{u, b\}$ in T that connect u to two vertices $a, b \in R(u)$, then the two vertices a and b belong to the same path P_i . The same is true for the left hemisphere $L(u)$ of u as well.

Proof. Suppose for contradiction that there exist two edges $\{u, a\}, \{u, b\}$ in T as above (i.e., with $a, b \in R(u)$) and such that these two edges have different type. (In other words, $d_{\overline{G}}(u, a) \neq d_{\overline{G}}(u, b)$.) Without loss of generality $d_{\overline{G}}(u, a) < d_{\overline{G}}(u, b)$. We construct a new tree T' by replacing $\{u, b\}$ by $\{a, b\}$. This change decreases the weight of the tree. Note that any path π in T that uses the edge $\{u, b\}$ can be replaced by a similar path that uses the edges $\{u, a\}$ and $\{a, b\}$ instead of $\{u, b\}$. Clearly, the length of the path does not change. Hence for every index i , the value $t_{T,i}$ does not increase. Hence $\gamma(T) = \min_i \{n - t_{T,i}\}$ does not decrease, and since T is a tree with $\gamma(T) = \mu = \max_{T''} \{\gamma(T'')\}$, it follows that $\gamma(T') = \gamma(T) = \mu$. This is a contradiction to the minimality of the weight of T among trees with $\gamma()$ value equal to μ . Obviously, the same argument also applies for $a, b \in L(u)$. \square

We now continue proving Lemma 8. The rest of the proof splits into a number of cases which are characterized by existence of certain edges in the tree T .

- Case 1: There is a path-internal edge in T . Assume there is a path-internal edge between vertices in a path P_j , for some index $j \in [0, k - 1]$. Note that if all the vertices which have path-internal edges in P_j have more than one path-internal edge incident on them in T , then the graph induced by those vertices contains a cycle. However, this graph is a subgraph of T , contradiction. Hence there is a vertex $u \in P_j$ such that u has exactly one path-internal edge $\{u, z\}$ incident on u in T . Assume without loss of generality that $j = 0$. Delete the vertex u and replace each edge $\{u, a\}$ incident on u in T by an edge $\{z, a\}$. (Note that u has only one

path-internal edge in T . Thus is the edge $\{u, z\}$. Therefore vertex a as above do not belong to $P_j = P_0$.) Denote the resulting tree by T' . Note that T' is a spanning tree for $\overline{G}_{k, n_0-1, n_1, \dots, n_{k-1}}$. (See Figure 6 for an illustration.) (T' is obviously connected and we reduce the number of edges by exactly 1, hence it is a spanning tree of $\overline{G}_{k, n_0-1, n_1, \dots, n_{k-1}}$.) Note that for every two vertices a, b in $\overline{G}_{k, n_0-1, n_1, \dots, n_{k-1}}$, we have $d_{T'}(a, b) \leq d_T(a, b)$, to see this consider the shortest path π from a to b in T . If the vertex u is not taking part in π then the claim is trivial, hence we assume that u is taking part. If the path π include the edge $\{u, z\}$, then obviously $\pi \setminus \{\{u, z\}\}$ is a path from a to b in T' and therefore $d_{T'}(a, b) \leq d_T(a, b)$. Otherwise, the path π does not include the edge $\{u, z\}$, set $\pi = aw_1 \dots w_r u w_{r+1} \dots w_m b$ where $w_r, w_{r+1} \notin P_0$ because u has only one path-internal edge to z , therefore $\pi' = aw_1 \dots w_r z w_{r+1} \dots w_m b$ is path from a to b in T' of length smaller or equal then π in T ($w_r, w_{r+1} \notin P_0$ implies $d_T(w_r, u) = d_{T'}(w_r, z)$ and $d_T(u, w_{r+1}) = d_{T'}(z, w_{r+1})$, and all the other edges remain in T with smaller or equal weight.) Moreover, for each $q \in P_i$, $d_{\overline{G}_{k, n_0, \dots, n_{k-1}}}(v_i, q) = d_{\overline{G}_{k, n_0, \dots, n_{k-1}}}(v_{i+1}, q) = d_{\overline{G}_{k, n_0-1, n_1, \dots, n_{k-1}}}(v_i, q) = d_{\overline{G}_{k, n_0-1, n_1, \dots, n_{k-1}}}(v_{i+1}, q) = w$, therefore $d_T(v_i, q) \geq d_{T'}(v_i, q)$ and $d_T(v_{i+1}, q) \geq d_{T'}(v_{i+1}, q)$ implies $\max \left\{ \frac{d_{T'}(v_i, q)}{\overline{G}_{k, n_0, \dots, n_{k-1}}(v_i, q)}, \frac{d_{T'}(v_{i+1}, q)}{\overline{G}_{k, n_0, \dots, n_{k-1}}(v_{i+1}, q)} \right\} \geq \max \left\{ \frac{d_{T'}(v_i, q)}{\overline{G}_{k, n_0-1, n_1, \dots, n_{k-1}}(v_i, q)}, \frac{d_{T'}(v_{i+1}, q)}{\overline{G}_{k, n_0-1, n_1, \dots, n_{k-1}}(v_{i+1}, q)} \right\}$. Denote $n'_0 = n_0 - 1$, $n'_i = n_i$, for every $i \in [1, k-1]$. By the induction hypothesis, there is an index i such that $t_{T', i} = n'_i$. If $i \neq 0$ then $t_{T', i} = n'_i$, and in particular by the previews argument $t_{T, i} \geq t_{T', i}$. Since, by definition, $t_{T, i} \leq n_i$, it follows that $t_{T, i} = n_i$, as required. If $i = 0$, then $t_{T', 0} = n'_0 = n_0 - 1$. A vertex $a \in P_0 \setminus \{u, z\}$, contributes 1 to $t_{T', 0}$, and hence contributes 1 also to $t_{T, 0}$. Moreover, $d_T(v_0, u), d_T(v_0, z) \geq d_{T'}(v_0, z)$, and analogously $d_T(v_1, u), d_T(v_1, z) \geq d_{T'}(v_1, z)$. Since z contributes 1 to $t_{T', 0}$ it follows that u and z each contribute 1 to $t_{T, 0}$ and so $t_{T, 0} = n_0$, completing the proof of case 1.

- Case 2: There are no path-internal edges in T , and there is a path P_i , with two vertices $u_1, u_2 \in P_i$ which have a same-path neighbor $z \notin P_i$. Without loss of generality $i = 0$. Similarly to the previous case, delete u_2 and replace any edge of the form $\{u_2, a\}$ by edge of the form $\{u_1, a\}$. Denote the resulting tree by T' . Note that T' is a spanning tree for $\overline{G}_{k, n_0-1, n_1, \dots, n_{k-1}}$. (T' is obviously connected and we reduce the number of edges by exactly 1.) For each two vertices a, b in $\overline{G}_{k, n_0-1, n_1, \dots, n_{k-1}}$, $d_{T'}(a, b) \leq d_T(a, b)$, to see this, consider the shortest path π from a to b in T . If the vertex u_2 does not taking part in π then the claim is trivial. Otherwise, set $\pi = aw_1 \dots w_r u_2 w_{r+1} \dots w_m b$, therefore $\pi' = aw_1 \dots w_r u_1 w_{r+1} \dots w_m b$ is path from a to b in T' of equal length to the path π in T (There are no path-internal edges hence $d_T(w_r, u_1) = d_{T'}(w_r, u_2)$ and $d_T(u_1, w_{r+1}) = d_{T'}(u_2, w_{r+1})$, and all the other edges remain with equal weight.). As in previews case, this argument implies for every $q \in P_i$, $\max \left\{ \frac{d_{T'}(v_i, q)}{\overline{G}_{k, n_0, \dots, n_{k-1}}(v_i, q)}, \frac{d_{T'}(v_{i+1}, q)}{\overline{G}_{k, n_0, \dots, n_{k-1}}(v_{i+1}, q)} \right\} \geq \max \left\{ \frac{d_{T'}(v_i, q)}{\overline{G}_{k, n_0-1, n_1, \dots, n_{k-1}}(v_i, q)}, \frac{d_{T'}(v_{i+1}, q)}{\overline{G}_{k, n_0-1, n_1, \dots, n_{k-1}}(v_{i+1}, q)} \right\}$. Denote $n'_0 = n_0 - 1$, $n'_i = n_i$, for every $i \in [1, k-1]$. By the induction hypothesis, there is a path P'_i such that $t_{T', i} = n'_i$. If $i \neq 0$ then obviously $t_{T, i} = n'_i$. Otherwise, $i = 0$, vertex $a \in P_0 \setminus \{u_2\}$, contributes 1 to $t_{T', 0}$, and hence contributes 1 also to $t_{T, 0}$. In particular $\max \{d_T(u_2, v_0), d_T(u_2, v_{k-1})\} \geq \max \{d_{T'}(u_1, v_0), d_{T'}(u_1, v_{k-1})\}$, hence u_2 also contributes 1 to $t_{T', 0}$, we get $t_{T, 0} = t_{T', 0} + 1 =$

$n'_0 + 1 = n_0$ as required.

- Case 3: There are no path-internal edges, and no same-path neighbors in T . In particular, any vertex has at most two edges: one edge to his right and left hemispheres (There is at most one type in each hemisphere, because two edges of the same type implies same-path neighbor.) Hence any path π in T is one-sided. There have to be two terminals v_i, v_{i+1} such that there is no path between them of length $2w$ (otherwise, there is a cycle). Hence for all $u \in P_i$ there are no edges to both v_i, v_{i+1} . If without loss of generality $\{u, v_i\} \notin T$, then the shortest one-sided path between u and v_i is of length at least $2k - 1$, implying that the terminal distortion is at least $2k - 1$, hence $t_{T,i} = 0$ as required!

□

Theorem 9. For any positive integer parameters k, n and $\epsilon > 0$ such that $k \leq \frac{\epsilon}{2}n$, any spanning tree T of $\overline{G}_{k,n,n,\dots,n}$ where $w = \frac{k}{\epsilon}$, with terminal distortion at most $(2k - 1) \left(1 + \frac{\epsilon}{k^2}\right)$ has lightness at least $\Omega\left(\frac{1}{\epsilon}\right)$.

Proof. By Lemma 8 there is a path P_i such that any $u \in P_i$ has distortion at least $2k - 1$ from v_i or v_{i+1} . We will show that each $u \in P_i$ is included in non path-internal edge of weight at least w . Some observations before we prove it:

- Each $u \in P_i$ is at distance exactly $(2k - 1) \cdot w$ from v_i or v_{i+1} . Otherwise the distance is at least $(2k - 1)w + 1$ which implies

$$\begin{aligned} \frac{d_T(u, v_i)}{d_{\overline{G}_{k,n,\dots,n}}(u, v_i)} &= \frac{(2k - 1)w + 1}{w} = (2k - 1) \left(1 + \frac{1}{(2k - 1)w}\right) \\ &= (2k - 1) \left(1 + \frac{\epsilon}{(2k - 1)k}\right) > (2k - 1) \left(1 + \frac{\epsilon}{k^2}\right) \end{aligned}$$

in contradiction to the $(2k - 1) \left(1 + \frac{\epsilon}{k^2}\right)$ terminal distortion.

- For two neighbors u_1, u_2 in a tree, and any other vertex v . The shortest path from u_1 to v goes through u_2 or the shortest path from u_2 to v goes through u_1 . Otherwise there is a cycle.

Assume by contradiction, that there is $u \in P_i$ with no outgoing edge of weight at least w . Let p be the first vertex on the shortest path from u to v_i , and q be the first vertex on the shortest path from u to v_{i+1} . W.l.o.g $d_T(v_i, u) = (2k - 1)w$. Hence $d_T(v_i, p) < (2k - 1)w$ which implies $d_T(v_{i+1}, p) = (2k - 1)w$. The shortest path from p to v_{i+1} goes through u , as otherwise $d_T(v_{i+1}, u) > (2k - 1)w$, contradiction. Hence $p \neq q$. Necessarily the shortest path from q to v_i goes through u , which implies $d_T(v_i, q) > d_T(v_i, u) = (2k - 1)w$, contradiction.

We conclude that $w(T) > n \cdot w$. Hence,

$$\Psi(T) = \frac{w(T)}{w(MST)} \geq \frac{n \cdot w}{(2k - 1)w + (n - 1)k} = \frac{n \cdot \frac{k}{\epsilon}}{(2k - 1)\frac{k}{\epsilon} + (n - 1)k} > \frac{n}{\epsilon n + 2k}.$$

As $k \leq \frac{\epsilon}{2}n$ we obtain $\Psi(T) \geq \frac{1}{2\epsilon}$ as required.

□




Figure 7: A visualization of the contradiction during the proof of Theorem 9. We assume that u does not have outgoing edge of weight w . From the assumption that the distance from u to v_i is $(2k - 1)w$ we conclude that the distance from q to v_i is larger than $(2k - 1)w$, a contradiction.

8 Probabilistic Embedding into Ultrametrics with Strong Terminal Distortion

In this section we show our terminal variant of the result of [FRT04]. We note that in a companion paper [EFN15] a more general result is shown, and we give the full details here for completeness.

Recall that an ultrametric (U, d) is a metric space satisfying a strong form of the triangle inequality, that is, for all $x, y, z \in U$, $d(x, z) \leq \max\{d(x, y), d(y, z)\}$. The following definition is known to be an equivalent one (see [BLMN05]).

Definition 2. A finite ultrametric \mathcal{U} is a metric space (U, d) whose elements are the leaves of a rooted labeled tree T . Each node $z \in T$ is associated with a label $\Phi(z) \geq 0$ such that if $q \in T$ is a descendant of z then $\Phi(q) \leq \Phi(z)$ and $\Phi(q) = 0$ iff q is a leaf. The distance between leaves $z, q \in U$ is defined as $d_T(z, q) = \Phi(\text{lca}(z, q))$ where $\text{lca}(z, q)$ is the least common ancestor of z and q in T .

We now define probabilistic embeddings with terminal distortion. For a class of metrics \mathcal{Y} and a fixed metric space X , a distribution \mathcal{D} over embeddings $f_Y : X \rightarrow Y$ with $Y \in \mathcal{Y}$ has *expected terminal distortion* α if each f_Y is non-contractive (that is, for all $u, w \in X$ and $Y \in \text{supp}(\mathcal{D})$, it holds that $d_X(u, w) \leq d_Y(f_Y(u), f_Y(w))$), and for all $v \in K$ and $u \in X$,

$$\mathbb{E}_{Y \sim \mathcal{D}}[d_Y(f_Y(v), f_Y(u))] \leq \alpha \cdot d_X(v, u) .$$

The notion of strong terminal distortion translates to this setting in the obvious manner.

Let (X, d) be a metric on n points, and $K \subseteq X$ the set of terminals of size $|K| = k$. Theorem 1 combined with [FRT04] gives an embedding to a distribution over ultrametrics with expected terminal distortion $O(\log k)$. However, this does not give any guarantee on the distortion of two non-terminal points.¹⁴ In order to obtain *strong* terminal distortion, we modify the FRT algorithm, so that it gives “preference” to the terminals. In particular, in the heart of the FRT algorithm is a construction of a random partitioning scheme, based on [CKR04], whose purpose is to decompose

¹⁴In fact, it is not hard to verify that this approach may lead to unbounded distortion.

the metric to bounded diameter pieces, such that the probability of separating pairs is “sufficiently small”. The partition is created by choosing a random permutation of X , a random radius in some specified interval, and then creating clusters as balls of the chosen radius, in the order given by the permutation. To obtain improved expected distortion guarantees for the terminals, we enforce them to be the first k points of the permutation. We show that this restriction improves dramatically the expected distortion of *any pair* that contains a terminal, while all the other pairs suffer only a factor of 2 in the expected distortion bound.

Theorem 10. *Given a metric space (X, d) of size $|X| = n$ and a subset of terminals $K \subseteq X$ of size $|K| = k$, there exists a distribution over embeddings of X into ultrametrics with strong terminal distortion $(O(\log k), O(\log n))$.*

Proof. Let Δ be the diameter of X . We assume w.l.o.g. that the minimal distance in X is 1, and let δ be the minimal integer so that $\Delta \leq 2^\delta$. We shall create a hierarchical laminar partition, where for each $i \in \{0, 1, \dots, \delta\}$, the clusters of level i have diameter at most 2^i , and each of them is contained in some level $i + 1$ cluster. Recall Definition 2. The ultrametric is built in the natural manner, the root corresponds to the level δ cluster which is X , and each cluster in level i corresponds to an inner node of the ultrametric with label 2^i , whose children correspond to the level $i - 1$ clusters contained in it. The leaves correspond to singletons, that is, to the elements of X . Clearly, the ultrametric will dominate (X, d) .¹⁵

In order to define the partition, we sample a permutation uniformly at random from the set of permutations that have the terminals as the preface (there are $k! \cdot (n - k)!$ such permutations), and sample a number $\beta \in [1, 2]$ randomly according to distribution with the probability density function $p(x) = \frac{1}{x \ln 2}$. In each step, each vertex x chooses the vertex u with minimal value according to π among the vertices of distance at most $\beta_i = \beta \cdot 2^{i-2}$ from x , and joins the cluster of u . Note that a vertex may not belong to the cluster associated with it, and some clusters may be empty (which we can discard). The description of the hierarchical partition appears in Algorithm 1.

Algorithm 1 Modified FRT(X, K)

- 1: Choose a random permutation π of X such that $\pi^{-1}(K) = [k]$.
 - 2: Choose $\beta \in [1, 2]$ randomly by the distribution with the following probability density function $p(x) = \frac{1}{x \ln 2}$.
 - 3: Let $D_\delta = X$; $i \leftarrow \delta - 1$.
 - 4: **while** D_{i+1} has non-singleton clusters **do**
 - 5: Set $\beta_i \leftarrow \beta \cdot 2^{i-2}$.
 - 6: **for** $l = 1, \dots, n$ **do**
 - 7: **for** every cluster S in D_{i+1} **do**
 - 8: Create a new cluster in D_i , consisting of all unassigned vertices in S closer than β_i to $\pi(l)$.
 - 9: **end for**
 - 10: **end for**
 - 11: $i \leftarrow i - 1$.
 - 12: **end while**
-

¹⁵A metric d' on X dominates d if for all $x, y \in X$, $d(x, y) \leq d'(x, y)$.

Let T denote the (random) ultrametric created by the hierarchical partition of Algorithm 1, and $d_T(u, v)$ the distance between u and v in T . Consider the clustering step at some level i , where clusters in D_{i+1} are picked for partitioning. In each iteration l , all unassigned vertices z such that $d(z, \pi(l)) \leq \beta_i$, assign themselves to the cluster of $\pi(l)$. $\pi(l)$ is called the *center* of the new cluster. Fix an arbitrary pair $\{v, u\}$. We say that center w *settles* the pair $\{v, u\}$ at level i , if it is the first center so that at least one of u and v gets assigned to its cluster. Note that exactly one center w settles any pair $\{v, u\}$ at any particular level. Further, we say that a center w *cuts* the pair $\{v, u\}$ at level i , if it settles them at this level, and exactly one of u and v is assigned to the cluster of w at level i . Whenever w cuts a pair $\{v, u\}$ at level i , $d_T(v, u)$ is set to be $2^{i+1} \leq 8\beta_i$. We charge this length to the vertex w and define $d_T^w(v, u)$ to be $\sum_i \mathbf{1}(w \text{ cuts } \{v, u\} \text{ at level } i) \cdot 8\beta_i$ (where $\mathbf{1}(\cdot)$ denotes an indicator function). Clearly, $d_T(v, u) \leq \sum_{w \in X} d_T^w(v, u)$.

We now arrange the vertices of K in non-decreasing order of their distance in X from the pair $\{v, u\}$ (breaking ties arbitrarily). Consider the s th terminal w_s in this sequence. We now upper bound the expected value of $d_T^{w_s}(v, u)$ for an arbitrary w_s . W.l.o.g assume that $d_X(w_s, v) \leq d_X(w_s, u)$. For a center w_s to cut $\{v, u\}$, it must be the case that:

1. $d_X(w_s, v) \leq \beta_i < d_X(w_s, u)$ for some i .
2. w_s settles $\{v, u\}$ at level i .

Note that for each $x \in [d_X(w_s, v), d_X(w_s, u)]$, the probability that $\beta_i \in [x, x + dx)$ is at most $\frac{dx}{x \cdot \ln 2}$. Conditioning on β_i taking such a value x , any one of w_1, \dots, w_s can settle $\{v, u\}$, and the probability that w_s is first in the permutation among w_1, \dots, w_s is $\frac{1}{s}$. Thus we obtain,

$$\mathbb{E}[d_T^{w_s}(v, u)] \leq \int_{d_X(w_s, v)}^{d_X(w_s, u)} 8x \cdot \frac{dx}{x \ln 2} \cdot \frac{1}{s} = \frac{8(d_X(w_s, u) - d_X(w_s, v))}{s \cdot \ln 2} \leq \frac{16}{s} \cdot d_X(v, u). \quad (3)$$

The crucial observation is that if $y \in X \setminus K$, and at least one of v, u is a terminal, w.l.o.g. $v \in K$, then y cannot settle $\{v, u\}$. The reason is that v always appears before y in π , so v will surely be assigned to a cluster when it is the turn of y to create a cluster. This leads to the conclusion that for all $v \in K$ and $u \in X$

$$\mathbb{E}[d_T(v, u)] \leq \sum_{s=1}^k \mathbb{E}[d_T^{w_s}(v, u)] \stackrel{(3)}{\leq} 16d_X(v, u) \sum_{s=1}^k \frac{1}{s} = O(\log k) \cdot d_X(v, u).$$

It remains to bound the expected distortion of non-terminal pairs, so fix some $v, u \in X \setminus K$. Arrange the vertices y_1, \dots, y_{n-k} of $X \setminus K$ according to their distance from $\{v, u\}$, in non-decreasing order. A similar reasoning as above gives the bound of (3) for the s th vertex y_s in this ordering to cut $\{v, u\}$ (recall that the permutation over the non-terminals is uniformly distributed). In fact, the bound is even slightly loose, because we disregard the possibility that a terminal will settle $\{v, u\}$. We conclude that

$$\begin{aligned} \mathbb{E}[d_T(v, u)] &\leq \sum_{s=1}^k \mathbb{E}[d_T^{w_s}(v, u)] + \sum_{s=1}^{n-k} \mathbb{E}[d_T^{y_s}(v, u)] \\ &\leq 16d_X(v, u) \left(\sum_{s=1}^k \frac{1}{s} + \sum_{s=1}^{n-k} \frac{1}{s} \right) \\ &= O(\log k + \log n) \cdot d_X(v, u) \\ &= O(\log n) \cdot d_X(v, u). \end{aligned}$$

□

9 Probabilistic Embedding into Trees with Terminal Congestion

In this section we focus on embeddings into trees that approximate capacities of cuts, rather than distances between vertices. This framework was introduced by Räcke [R02] (for a single tree), and in [Rac08] he showed how to obtain a capacity preserving probabilistic embedding from a distance preserving one, such as the ones given by [FRT04]. Later, [AF09] showed a complete equivalence between these notions in random tree embeddings. Here we show our terminal variant of these results. Informally, we construct a distribution over *capacity-dominating* trees (each cut in each tree is at least as large as the corresponding cut in the original graph), and for each edge, its expected congestion is bounded accordingly, with an improved bound for edges *containing a terminal*.

We next elaborate on the notions of capacity and congestion, and their relation to distance and distortion, following the notation of [AF09]. Given a graph $G = (V, E)$, let \mathcal{P} be a collection of multisets of edges in G . A map $M : E \rightarrow \mathcal{P}$, where $M(e)$ is a path between the endpoints of e , is called a *path mapping* (the path is not necessarily simple). Denote by $M_{e'}(e)$ the number of appearances of e' in $M(e)$.

The path mapping relevant to the rest of this section is constructed as follows: given a graph G and a tree $T = (V, E_T)$ (not necessarily a subgraph of G , but on the same vertex set), for each edge $e' \in E_T$ let $P_G(e')$ be a shortest path between the endpoints of e' in G (breaking ties arbitrarily), and similarly for $e \in E$, let $P_T(e)$ be the unique path between the endpoints of e in T . Then for an edge $e \in E$, where $P_T(e) = e'_1 e'_2 \dots e'_r$, the path $M(e)$ is defined as $M(e) = P_G(e'_1) \circ P_G(e'_2) \circ \dots \circ P_G(e'_r)$ (where \circ denotes concatenation). In what follows fix a tree T , and let M be the path mapping of T .

Fix a weight function $w : E \rightarrow \mathbb{R}_+$, and a capacity function $c : E \rightarrow \mathbb{R}_+$. For an edge $e \in E$, $\text{dist}_T(e) = \sum_{e' \in E} M_{e'}(e) \cdot w(e')$ is the weight of the path $M(e)$, and $\text{load}_T(e) = \sum_{e' \in E} M_{e'}(e) \cdot c(e')$ is the sum (with multiplicities) of the capacities of all the edges whose path is using e . Define $\text{distortion}_T(e) = \frac{\text{dist}_T(e)}{w(e)}$ to be the distortion of e in T , and $\text{cong}_T(e) = \frac{\text{load}_T(e)}{c(e)}$ is the congestion of e . Note that if T is a subgraph of G , then $\text{dist}_T(e)$ is the length of the unique path between the endpoints of e , while $\text{load}_T(e)$ is the total capacity of all the edges of E that are in the cut obtained by deleting e from T (for $e \notin T$, $\text{load}_T(e) = 0$).

Definition 3. Let $K \subseteq V$ be a set of terminals of size k , and let $E_K \subseteq E$ be the set of edges that contain a terminal. We say that a distribution \mathcal{D} over trees has strong terminal congestion (α, β) if for every $e \in E_K$.

$$\mathbb{E}_{T \sim \mathcal{D}}[\text{cong}_T(e)] \leq \alpha,$$

and for any $e \in E$, $\mathbb{E}_{T \sim \mathcal{D}}[\text{cong}_T(e)] \leq \beta$.

A tight connection between distance preserving and capacity preserving mappings was shown in [AF09]. We generalize their theorem to the terminal setting in the following manner.

Theorem 11. *The following statements are equivalent for a graph G :*

- For every possible weight assignment G admits a probabilistic embedding into trees with strong terminal distortion (α, β) .

- For every possible capacity assignment G admits a probabilistic embedding into trees with strong terminal congestion (α, β) .

A corollary of Theorem 11, achieved by applying Theorem 10 using the algorithmic framework described in [AF09], is:¹⁶

Corollary 10. *For any graph $G = (V, E)$ on n vertices, a set $K \subseteq V$ of k terminals, and any capacity function, there exists a distribution over trees with strong terminal congestion $(O(\log k), O(\log n))$. Moreover, there is an efficient procedure to sample from the distribution.*

Proof of Theorem 11. Assuming the first item holds we prove the second. Let $\kappa(e) = \begin{cases} 1/\alpha & e \in E_K \\ 1/\beta & \text{otherwise} \end{cases}$.

Given any capacity function $c : E \rightarrow \mathbb{R}_+$, we would like to show that there exists a distribution \mathcal{D}' such that for any $e \in E$, $\mathbb{E}_{T \sim \mathcal{D}'}[\kappa(e) \cdot \text{cong}_T(e)] \leq 1$. By applying the minimax principle (as in [AKPW95]), it suffices to show that for any coefficients $\{\lambda_e\}_{e \in E}$ with $\lambda_e \geq 0$ and $\sum_{e \in E} \lambda_e = 1$, there exists a single tree T such that

$$\sum_{e \in E} \lambda_e \cdot \kappa(e) \cdot \text{cong}_T(e) \leq 1 . \quad (4)$$

To this end, define the weights $w(e) = \kappa(e) \cdot \frac{\lambda_e}{c(e)}$, and by the first assertion there exists a distribution \mathcal{D} over trees such that for any $e \in E$,

$$\mathbb{E}_{T \sim \mathcal{D}}[\kappa(e) \cdot \text{distortion}_T(e)] \leq 1 .$$

By applying the minimax again, there exists a single tree T such that

$$\sum_{e \in E} \lambda_e \cdot \kappa(e) \cdot \text{distortion}_T(e) \leq 1 .$$

Now,

$$\begin{aligned} 1 &\geq \sum_{e \in E} \lambda_e \cdot \kappa(e) \cdot \text{distortion}_T(e) \\ &= \sum_{e \in E} \lambda_e \cdot \kappa(e) \cdot \frac{\sum_{e' \in E} M_{e'}(e) \cdot w(e')}{w(e)} \\ &= \sum_{e \in E} \lambda_e \cdot \kappa(e) \cdot \frac{\sum_{e' \in E} M_{e'}(e) \cdot \kappa(e') \cdot \lambda_{e'}/c(e')}{\kappa(e) \cdot \lambda_e/c(e)} \\ &= \sum_{e' \in E} \lambda_{e'} \cdot \kappa(e') \cdot \frac{\sum_{e \in E} M_{e'}(e) \cdot c(e)}{c(e')} \\ &= \sum_{e' \in E} \lambda_{e'} \cdot \kappa(e') \cdot \text{cong}_T(e') , \end{aligned}$$

which concludes the proof of (4). The second direction is symmetric. \square

¹⁶Even though the embedding of Theorem 10 is into ultrametrics, which contain Steiner vertices, these can be removed while increasing the distortion of each pair by at most a factor of 8 [Gup01a].

Capacity Domination Property. As [Räc08, AF09] showed, under the natural capacity assignment, any tree T supported by the distribution of Theorem 11 has the following property: Any multi-commodity flow in G can be routed in T with no larger congestion. We would like to show this explicitly, using the language of cuts, as this will be useful for the algorithmic applications.

Fix some tree $T = (V, E_T)$, and for any edge $e' \in E_T$ let $S_{T,e'} \subseteq V$ be the cut obtained by deleting e' from T . Define the capacities $C_T : E_T \rightarrow \mathbb{R}_+$ by

$$C_T(e') = \sum_{e \in E(S_{T,e'}, \bar{S}_{T,e'})} c(e) ,$$

where $E(S, \bar{S})$ denotes the set of edges in the graph crossing the cut S . (Observe that for spanning trees, $C_T(e) = \text{load}_T(e)$.)

Lemma 11. *For any graph $G = (V, E)$ and tree $T = (V, E_T)$ with capacities as defined above, for any set $S \subseteq V$ it holds that*

$$\sum_{e \in E(S, \bar{S})} c(e) \leq \sum_{e' \in E_T(S, \bar{S})} C_T(e') \leq \sum_{e \in E(S, \bar{S})} \text{load}_T(e) . \quad (5)$$

Proof. We begin with the left inequality. For any graph edge $e \in E(S, \bar{S})$, there exists a tree edge $e' \in E_T(S, \bar{S})$ such that $e' \in P_T(e)$, because the path $P_T(e)$ must cross the cut. Since removing e' from T separates the endpoints of e , $C_T(e')$ will contain the term $c(e)$. We conclude that

$$\sum_{e \in E(S, \bar{S})} c(e) \leq \sum_{e' \in E_T(S, \bar{S})} C_T(e') .$$

For the right inequality, consider an edge $e \in E$, and note that for any tree edge $e' \in E_T$ such that $e \in P_G(e')$, every edge $e'' \in E(S_{T,e'}, \bar{S}_{T,e'})$ will have $e \in M(e'')$ and thus contribute to $\text{load}_T(e)$ (perhaps multiple times, due to different e'). This implies that

$$\text{load}_T(e) = \sum_{\{e' \in E_T : e \in P_G(e')\}} C_T(e') . \quad (6)$$

Next, observe that any tree edge $e' \in E_T(S, \bar{S})$ must have at least one graph edge $e \in E(S, \bar{S})$ such that $e \in P_G(e')$. This suggests that

$$\begin{aligned} \sum_{e \in E(S, \bar{S})} \text{load}_T(e) &\stackrel{(6)}{=} \sum_{e \in E(S, \bar{S})} \sum_{\{e' \in E_T : e \in P_G(e')\}} C_T(e') \\ &= \sum_{e' \in E_T} |E(S, \bar{S}) \cap P_G(e')| \cdot C_T(e') \\ &\geq \sum_{e' \in E_T(S, \bar{S})} C_T(e') . \end{aligned}$$

□

10 Strong Terminal Embedding of Graphs into a Distribution of Spanning Trees

In this section we consider strong terminal embedding of graphs into a distribution over their spanning trees. We will show $(\tilde{O}(\log k), \tilde{O}(\log n))$ -strong terminal distortion as promised in the introduction.

Theorem 12. *Given a weighted graph $G = (V, E, w)$ on n vertices, and a subset of k terminals $K \subseteq V$, there exists a distribution over spanning trees of G with strong terminal distortion $(\tilde{O}(\log k), \tilde{O}(\log n))$.*

We will use the framework of [AN12], in particular their petal decomposition structure in order to obtain the distribution over spanning trees. Roughly speaking, it is an iterative method to build a spanning tree. In each level, the current graph is partitioned into smaller diameter pieces, called petals, and a single central piece, which are then connected by edges in a tree structure. Each of the petals is a ball in a certain metric. The main advantage of this framework, is that it produces a spanning tree whose diameter is proportional to the diameter of the graph, but allows very large freedom for the choice of radii of the petals. Specifically, if the graph diameter is Δ , each radius can be chosen in an interval of length $\approx \Delta$. Intuitively, if a pair is separated by the petal decomposition, then its distance in the tree will be $O(\Delta)$. So we would like a method to choose the radii, that will give the appropriate bounds both on pairs containing a terminal, and for all other pairs, simultaneously. We note that in the star-partition framework of [EEST08] (used also in [ABN08]), the tree radius introduces a factor to the distortion that inherently depends on n , which does not seem to allow terminal distortion that depends on k alone.

In order to “take care” of pairs containing a terminal, we will need to somehow give the terminals a “preference” in the decomposition. Since the [CKR04, FRT04] style partitioning scheme cannot work in a graph setting (because it does not produce strong diameter clusters, and the cluster’s center may not be contained in it), we turn to the partitions based on truncated exponential distribution as in [Bar96, Bar04, ABN08]. To implement the “terminal preference” idea, we first build petals with the terminals as centers, and only then build petals for the remaining points. There are few technical subtleties needed to assure that pairs containing a terminal suffer small expected distortion. By a careful choice for the petal’s center and the interval from which to choose the radius, we guarantee that in a level where no terminals are separated, none of the relevant balls around terminals are in danger of being cut. Using that the radius of clusters decreases geometrically, we conclude that each ball is “at risk” in at most $O(\log k)$ levels.

10.1 Petal decomposition description

In this section we present the petal decomposition algorithm, and quote some of its properties. We do not provide full proofs of these properties, these can be found in [AN12]. The petals are built using an alternative metric on the graph. Recall that for a subset $X \subseteq V$, d_X denotes the shortest path metric in the graph induced by X .

Definition 4 (Cone metric). *Given a graph $G = (V, E, w)$, a subset $X \subseteq V$ and points $x, y \in X$, define the cone-metric $\rho = \rho(X, x, y) : X^2 \rightarrow \mathbb{R}^+$ as $\rho(u, v) = |(d_X(x, u) - d_X(y, u)) - (d_X(x, v) - d_X(y, v))|$.*

This is actually a pseudo metric. For a cone metric $\rho = \rho(X, x, y)$,

$$\rho(y, u) = |(d_X(x, u) - d_X(y, u)) - (d_X(x, y) - d_X(y, y))| = d_X(x, y) + d_X(y, u) - d_X(x, u)$$

is the difference between the shortest path from x to u in $G[X]$ and the shortest path from x to u in $G[X]$ that goes through y .¹⁷ Therefore, the ball $B_{(X,\rho)}(y, r)$ in the cone metric $\rho = \rho(X, x, y)$ centered in y with radius r , contains all such vertices u . Denote by $P_{xy}(G)$ the shortest¹⁸ path between x and y in G . If Y is a subset of the vertices of G , we denote by $P_{xy}(Y)$ the shortest path from x to y in $G[Y]$.

A petal is a union of balls in the cone metric. In the `create-petal` algorithm, while working in a subgraph $G[Y]$ with two specified vertices: a center x_0 and a target t , we define $W_r(Y, x_0, t) = \bigcup_{p \in P_{x_0 t}: d_Y(p, t) \leq r} B_{(Y, \rho(Y, x_0, p))}(p, \frac{r - d_Y(p, t)}{2})$ which is union of balls in the cone metric, where any vertex p in the shortest path from x_0 to t of distance at most r from t is a center of a ball with radius $\frac{r - d_Y(p, t)}{2}$. We will often omit the parameters and write just W_r if Y, x_0 and t are clear from the context. The next claim, which is implicit in [AN12], states that W_r is monotone (in r) and that W_{r+4l} contains the ball with radius l around W_r (where the ball is taken in the shortest path metric in $G[Y]$).

Claim 12. For $W_r(Y, x_0, t) = \bigcup_{p \in P_{x_0 t}: d_Y(p, t) \leq r} B_{(Y, \rho(Y, x_0, p))}(p, (r - d_Y(p, t))/2)$

1. W_r is monotone in r , i.e., for $r \leq r'$ it holds that $W_r \subseteq W_{r'}$.
2. For every $y \in W_r$ and l , the ball $B_Y(y, l)$ is contained in W_{r+4l} , i.e., $\forall z, d_Y(y, z) \leq l \Rightarrow z \in W_{r+4l}$

Next, we quote some properties of the petal decomposition algorithm proved in [AN12]. Their proofs remain valid even after our modification, because the algorithm allows for arbitrary choice of targets after the first petal is built, and allows any radius in the appropriate interval. For a subset $X \subseteq V$, let Δ_{X, x_0} denote the radius of $G[X]$ with respect to x_0 (the maximal distance from x_0 in $G[X]$). We often omit the subscript x_0 if it is clear from the context.

Fact 1. For a graph G , a vertex set X , and some vertices $x_0, t \in X$, `petal-decomposition` ($G[X], x_0, t, K$) returns as output the sequence $(X_0, \dots, X_s, \{y_1, x_1\}, \dots, \{y_s, x_s\}, t_0, \dots, t_s)$ such that for each $j \leq s$, $\Delta_{X_j, x_j} \leq 7\Delta_{X, x_0}/8$.

Fact 2. The `hierarchical-petal-decomposition` algorithm returns a tree.

Fact 3. Every edge $e \in E$ can have its weight multiplied by $\frac{1}{2}$ at most once throughout the execution of the algorithm.¹⁹

Fact 4. For a graph G and tree T created by the `hierarchical-petal-decomposition`,

$$\Delta_T \leq 4\Delta_G.$$

Fact 5. Let $1 \leq j \leq s$ be an integer and $z \in Y_j$,²⁰ then $P_{x_0 z}(X) \subseteq G[Y_j]$.

¹⁷For $X \subseteq V$, $G[X]$ is the induced graph on the vertices of X .

¹⁸For simplicity, we assume that for each pair of vertices there is a unique shortest path in G .

¹⁹This multiplication may happen only for edges in the special paths P_{x_j, t_j} for a petal X_j which is not the first special petal. Note that after the creation of X_j , t_j will be the target for the first special petal in the decomposition of X_j which will ensure that the weights of the edges in P_{x_j, t_j} won't decrease again.

²⁰ Y_j is defined in the `petal-decomposition` procedure.

From Fact 4 we can deduce that if the radius of G was Δ (with respect to some vertex x_0) then the diameter of the tree created by the hierarchical petal decomposition is bounded by 8Δ . In lines 15 and 24 of the `petal-decomposition` procedure we divide the weight of some edges by 2. By Fact 3 it can happen to any edge at most once, so in the analysis of the distortion we will ignore this factor (for simplicity).

For completeness, we present the full petal decomposition algorithm. We illustrate below the main changes compared to the version presented in [AN12].

1. The order of choosing targets in `petal-decomposition`. In the original algorithm the targets were chosen arbitrarily, while in our version we first choose terminals, and only when there are no terminals left sufficiently far from x_0 , we choose the other vertices. The purpose of this ordering, is to ensure that only the terminal petals may cut certain balls around the other terminals.
2. The segments sent to create petal procedure. In line 22 in `petal-decomposition` we send an interval of a different size from the original algorithm. The purpose of this change is to ensure that petals which are created in line 22 are far away from any non-covered terminal.
3. We stop constructing petals when all points outside $B_X(x_0, 7\Delta/8)$ have been clustered (rather than $B_X(x_0, 3\Delta/4)$). The purpose is to ensure that small balls around terminals in X_0 are not cut.
4. The choice of the radius in the `create-petal` procedure is done by sampling according to a truncated exponential distribution with parameter λ . The parameter is computed using the density of terminals or other vertices around the target depending on whether the petal will contain a terminal.
5. The interval in which we choose the radius is determined using the `Reform-Interval` procedure. The purpose of this is to ensure that certain balls around terminals may be cut only at levels when some terminals are separated (note that there can be at most k such levels). The new interval returned is guaranteed to be contained in the original one.

Algorithm 2 $T = \text{hierarchical-petal-decomposition}(G[X], x_0, t, K)$

```

1: if  $|X| = 1$  then
2:   return  $G[X]$ .
3: end if
4: Let  $(X_0, \dots, X_s, (y_1, x_1), \dots, (y_s, x_s), t_0, \dots, t_s) = \text{petal-decomposition}(G[X], x_0, t, K)$ ;
5: for each  $j \in [0, \dots, s]$  do
6:    $T_j = \text{hierarchical-petal-decomposition}(G[X_j], x_j, t_j, K \cap X_j)$ ;
7: end for
8: Let  $T$  be the tree formed by connecting  $T_0, \dots, T_s$  using the edges  $\{y_1, x_1\}, \dots, \{y_s, x_s\}$ ;

```

Algorithm 3 $(X_0, \dots, X_s, \{y_1, x_1\}, \dots, \{y_s, x_s\}, t_0, \dots, t_s) = \text{petal-decomposition}(G[X], x_0, t, K)$

```

1: Let  $\Delta = \Delta_{X, x_0}$ ; Let  $Y_0 = X$ ; Set  $j = 1$ ;
2: if  $d_X(x_0, t) \geq 5\Delta/8$  then
3:   Let  $(X_1, x_1) = \text{create-petal}(X, t, x_0, [d_X(x_0, t) - 5\Delta/8, d_X(x_0, t) - \Delta/2], K)$ ;
4:   Let  $Y_1 = Y_0 \setminus X_1$ ;
5:   Let  $y_1$  be the neighbor of  $x_1$  on  $P_{x_0 t}$  (the one closer to  $x_0$ );
6:   Set  $t_0 = y_1, t_1 = t; j = 2$ ;
7: else
8:   set  $t_0 = t$ .
9: end if
10: Creating the terminal petals:
11: while  $K \cap Y_{j-1} \setminus B_X(x_0, 3\Delta/4) \neq \emptyset$  do
12:   Let  $t_j \in Y_{j-1} \cap K$  be an arbitrary terminal satisfying  $d_X(x_0, t_j) > 3\Delta/4$ ;
13:   Let  $(X_j, x_j) = \text{create-petal}(Y_{j-1}, t_j, x_0, [0, \Delta/8], K)$ ;
14:    $Y_j = Y_{j-1} \setminus X_j$ ;
15:   For each edge  $e \in P_{x_j t_j}$ , set its weight to be  $w(e)/2$ ;
16:   Let  $y_j$  be the neighbor of  $x_j$  on  $P_{x_0 t_j}$  (the one closer to  $x_0$ );
17:   Let  $j = j + 1$ ;
18: end while
19: Creating the non-terminal petals:
20: while  $Y_{j-1} \setminus B_X(x_0, 7\Delta/8) \neq \emptyset$  do
21:   Let  $t_j \in Y_{j-1}$  be an arbitrary vertex satisfying  $d_X(x_0, t_j) > 7\Delta/8$ ;
22:   Let  $(X_j, x_j) = \text{create-petal}(Y_{j-1}, t_j, x_0, [0, \Delta/32], \emptyset)$ ;
23:    $Y_j = Y_{j-1} \setminus X_j$ ;
24:   For each edge  $e \in P_{x_j t_j}$ , set its weight to be  $w(e)/2$ ;
25:   Let  $y_j$  be the neighbor of  $x_j$  on  $P_{x_0 t_j}$  (the one closer to  $x_0$ );
26:   Let  $j = j + 1$ ;
27: end while
28: Let  $s = j - 1$ ;
29: Creating the stigma  $X_0$ :
30: Let  $X_0 = Y_s$ ;

```

10.2 Analysis

In what follows, fix a cluster X with center x_0 , a target t , and the set of terminals K contained in X . Let $(X_0, \dots, X_s, \{y_1, x_1\}, \dots, \{y_s, x_s\}, t_0, \dots, t_s)$ be the result of applying $\text{petal-decomposition}(G[X], x_0, t, K)$. The following claim will be useful for bounding the number of levels in which a relevant ball around a terminal is in danger of being cut.

Claim 13. *The following assertions hold true:*

- If X_1 is a petal created at line 3 of $\text{petal-decomposition}$ and $K \cap X_1 = \emptyset$, then for all $v \in K$, $B_X(v, \Delta/160) \cap X_1 = \emptyset$.
- If X_j is a petal created at lines 3 or 13 of $\text{petal-decomposition}$ and $K \subseteq X_j$, then for all $v \in K$, $B_X(v, \Delta/160) \subseteq X_j$.

Algorithm 4 $(W, x) = \text{create-petal}(X, Y, t, x_0, [lo, hi], K)$

- 1: $W_r = \bigcup_{p \in P_{x_0 t}: d_Y(p, t) \leq r} B_{(Y, \rho(Y, x_0, p))}(p, (r - d_Y(p, t))/2)$;
 - 2: **if** $K \neq \emptyset$ **then**
 - 3: $[lo, hi] = \text{Reform-Interval}(Y, t, x_0, [lo, hi], K)$;
 - 4: $L_k = \lceil \log \log k \rceil$; $R = lo - hi$;
 - 5: Let $1 \leq q \leq L_k$ be the minimal integer satisfying $|W_{lo+qR/L_k}|_k \leq \frac{2|X|_k}{2^{\log^{1-q}/L_k}}$; Set $a = lo + (q - 1)R/L_k$, $b' = a + R/2L_k$, $\chi = \frac{|X|_k + 1}{|W_a|_k}$, $\hat{\chi} = \max\{\chi, e\}$;
 - 6: Set $\lambda = (2 \ln \hat{\chi}) / (b' - a) = (4L_k \ln \hat{\chi}) / R = (160L_k \ln \hat{\chi}) / \Delta$;
 - 7: **else**
 - 8: $L_n = \lceil \log \log n \rceil$; $R = lo - hi$;
 - 9: Let $1 \leq q \leq L_n$ be the minimal integer satisfying $|W_{lo+qR/L}| \leq \frac{2|X|_n}{2^{\log^{1-q}/L_n}}$; Set $a = lo + (q - 1)R/L_n$, $b' = a + R/2L_n$, $\chi = \frac{|X|_n + 1}{|W_a|}$, $\hat{\chi} = \max\{\chi, e\}$;
 - 10: Set $\lambda = (2 \ln \hat{\chi}) / (b' - a) = (4L_n \ln \hat{\chi}) / R = (2^7 L_n \ln \hat{\chi}) / \Delta$;
 - 11: **end if**
 - 12: Sample $r \in [a, b']$ according to the truncated exponential distribution with parameter λ , which has the following density function: $f(r) = \frac{\lambda \cdot e^{-\lambda r}}{e^{-\lambda \cdot a} - e^{-\lambda \cdot b'}}$;
 - 13: Let $r' \leq r$ be the maximum value such that there exists a point $p_{r'}$ of distance r' from t on $P_{x_0 t}$;
 - 14: **return** $(W_r, p_{r'})$;
-

- If X_j is a petal created at line 22 of *petal-decomposition*, then for all $v \in K \cap Y_{j-1}$, $B_X(v, \Delta/16) \cap X_j = \emptyset$.

Proof. Let W_r be as defined in line 1 of *create-petal* when forming the petal X_j . For the first assertion, note that the condition of line 6 in *reform-interval* must be satisfied, no terminals are in $W_{lo+2\alpha/5}$, and thus hi' is set to $lo + \alpha/5$ (we use α as in *reform-interval*, the size of the interval sent by *create-petal*.) Now, for every $u \in X_1 \subseteq W_{hi'}$ and $v \in K$, $d_X(u, v) \geq \alpha/20 = \Delta/160$, as otherwise, Claim 12 implies that $v \in W_{hi'+4\alpha/20} = W_{lo+2\alpha/5}$, contradiction.

For the second assertion, observe that $j \in \{1, 2\}$, because X_2 certainly contains a terminal since we are at the “creating terminal petals” stage. As $K \subseteq X_j$, it must be that the condition of line 3 in *reform-interval* is satisfied, and in line 4 we set lo' to be $lo + 4\alpha/5$. We conclude that if $u \in B_X(v, \Delta/160)$ for some $v \in K$, then $d_{Y_{j-1}}(u, v) \leq \Delta/160$ (because by the first assertion, this ball around v was not cut by the first petal). By Claim 12, $u \in W_{lo+3\alpha/5+\Delta/40} = W_{lo+4\alpha/5} \subseteq X_j$.

Finally we prove the third assertion. By the termination condition in line 11 of *petal-decomposition*, it must be that $v \in B_X(x_0, 3\Delta/4)$. By Fact 5, all distances in Y_{j-1} to x_0 remain the same as in X . Note that the parameter r of the petal $X_j = W_r$ is at most $\Delta/32$, and that $d_X(x_0, t_j) \geq 7\Delta/8$. By the definitions of a petal and of the cone-metric, for any $p \in P_{x_0 t}$ with $d_X(p, t) \leq \Delta/32$ we have that if $u \in B_{(Y_{j-1}, \rho(Y_{j-1}, x_0, p))}(p, (r - d_{Y_{j-1}}(p, t))/2)$ then

$$\begin{aligned} d_X(x_0, u) &\geq d_X(x_0, p) + d_{Y_{j-1}}(p, u) - (r - d_{Y_{j-1}}(p, t))/2 > d_X(x_0, p) - r \\ &\geq d_X(x_0, t) - d_X(p, t) - \Delta/32 \geq 7\Delta/8 - 2\Delta/32 \geq 13\Delta/16. \end{aligned}$$

By the triangle inequality we obtain that $d_X(v, u) \geq d_X(x_0, u) - d_X(x_0, v) > \Delta/16$, which concludes the proof. \square

Algorithm 5 $[lo', hi'] = \text{Reform-Interval}(Y, t, x_0, [lo, hi], K)$

```

1: Let  $W_r = \bigcup_{p \in P_{x_0 t}: d_Y(p, t) \leq r} B_{(Y, \rho(Y, x_0, p))}(p, (r - d_Y(p, t))/2)$ ;
2: Set  $\alpha \leftarrow hi - lo$ ;
3: if  $K \subseteq W_{lo+3\alpha/5}$  then
4:   return  $[lo + 4\alpha/5, hi]$ ;
5: end if
6: if  $K \cap W_{lo+2\alpha/5} = \emptyset$  then
7:   return  $[lo, lo + \alpha/5]$ ;
8: end if
9: return  $[lo + 2\alpha/5, lo + 3\alpha/5]$ ;

```

By Claim 13, given a terminal v in cluster X of radius Δ , and a ball $B = B_X(v, l)$ of radius at most $l \leq \Delta/160$, the ball B might be cut in the petal decomposition of X only if the terminal v is separated from some other terminal. `petal-decomposition` picks terminal targets arbitrary but deterministically. Therefore we will be more specific, and decide that `petal-decomposition` pick terminal targets by increasing order of the remaining terminals. If the first petal is built in line 3, then the first target is given, and by the `reform-interval` procedure, it is deterministically determined whether the first petal will contain all, part, or none of the terminals. If there are terminals remaining, and the first special petal (if exists) contains none of the terminals, then the second target is once again picked deterministically, and again `reform-interval` deterministically determines whether the petal will contain all or only part of the terminals. In any case, the issue of whether all the terminals will be in single petal or not, is resolved deterministically. Hence, given a cluster X , root x_0 and target t , it is known if there is some danger that B may be cut. Moreover, there are at most k levels in which some terminals are separated, and hence involving a danger of cutting balls around terminals. (If $l > \Delta/160$ then we do not care whether B is cut or not, because the distortion induced on v and points outside of B will be a constant.)

10.3 Expected Distortion

When dealing with weighted graphs, it could be that a tiny ball participates in many recursive levels, and thus is “threatened” many times. In order to avoid such situations, we shall change the algorithm slightly, and when performing a petal decomposition on a cluster X of radius Δ , we shall contract for each vertex $u \in X$, the ball of radius Δ/n^3 around u . In addition, if terminals are separated (which is determined deterministically) then for each terminal $v \in K$, we contract the ball of radius Δ/k^3 around v . These contractions are done sequentially; when a vertex contracts a ball of radius l it becomes a “supernode” of all the vertices within distance l , and for any edge leaving a vertex in the ball to some vertex z , we will have a corresponding edge with the same weight from the supernode to z . These contractions yield a cluster X' , on which we run the `petal-decomposition` algorithm (note that X' may contains supernodes - vertices that correspond to several original vertices, and that X' induces a multi-graph). After the partition of X' to X_0, \dots, X_s is determined, we expand back the contracted balls, so that each vertex belongs to the cluster of its supernode.

This guarantees that a ball of radius l around each terminal can participate only at partitions of radii in the range $[l, k^3 \cdot l]$, and by Fact 1, the radii decrease by a factor of at least $7/8$ every level, so there are at most $16 \log k$ levels in which each such ball participates. Similarly, a ball of radius l' around some vertex $u \in X$, can participate only at partitions of radii in the range $[l', n^3 \cdot l']$, so

there are at most $16 \log n$ levels in which each such ball participates. This contraction, while saving small balls from being cut, may have an effect on the radius of the tree when we expand back the vertices. We claim that this will be a minor increase. To see this, note that in a particular level, expanding back the balls around terminals can increase distances by at most $2\Delta/k^2$ (because every contracted ball has diameter at most $2\Delta/k^3$, and there can be at most k contracted balls). Similarly, expanding back all the other contracted balls may increase distances by at most an additional term of $2\Delta/n^2$. Now, since there are at most k iterations in which terminals are separated (only in such levels the balls around them are contracted), even if the radius is increased by a factor of $1 + 2/k^2$ in every one of them, the total increase is at most $(1 + 2/k^2)^k < e^{2/k}$. Similarly, there are at most n iterations all in all (since $\emptyset \subsetneq X_0 \subsetneq X$), and even if the radius is increased by a factor of $1 + 2/n^2$ in every one of them, the total increase is at most $(1 + 2/n^2)^n < e^{2/n}$. Henceforth we shall ignore this minor increase, as it affects the distortion of every pair by at most a factor of $e^{2/k+2/n}$.

We will show that the distribution generated by the **hierarchical-petal-decomposition**(G, x_0, x_0, K) algorithm has strong terminal distortion ($\tilde{O}(\log k), \tilde{O}(\log n)$) (where x_0 is an arbitrary vertex). The hierarchical partition naturally corresponds to a laminar family of clusters, that are arranged in a hierarchical tree structure denoted by \mathcal{T} (with V as the root, and each cluster X which is partitioned by **petal-decomposition** to X_0, \dots, X_s , has them as its children in \mathcal{T}). The *level* of a cluster is its distance in \mathcal{T} to the root. Note that there might be several trees that correspond to a single hierarchical partition (this depends on the precise edges connecting the different petals). For a pair of vertices x, y and hierarchical tree \mathcal{T} , let $d_{\mathcal{T}}(x, y)$ be the maximal distance between x and y in a tree T that corresponds to the hierarchical tree \mathcal{T} . Note that if x is separated from y in the hierarchical tree \mathcal{T} in cluster X of radius Δ_X , then by Fact 4, $d_{\mathcal{T}}(x, y) \leq 8\Delta_X$.

Fix any two points $x, y \in V$ and denote by $B_x = B_G(x, d(x, y))$. For $j \in [s]$, we say that B_x is *cut* by X_j if $B_x \cap X_j \notin \{\emptyset, B_x\}$. Further, B_x is cut in X if there exists some petal X_j that cuts B_x . For $i \geq 0$, let $S_{X,i}$ be the event that $B_x \subseteq X$ and X is a cluster in level i . When calling **petal-decomposition** on the cluster X , a sequence of clusters X_1, \dots, X_s, X_0 is generated, with Y_j as defined in the algorithm. For $Y \subseteq X$ and integer $j \geq 1$, let $S_{X,i,Y,j}$ denote the event that $S_{X,i}$ holds, $B_x \subseteq Y$, and $Y = Y_{j-1}$. Define the following events

$$\begin{aligned} \mathcal{C}_{X,i,j} &= \{B_x \cap X_j \notin \{\emptyset, B_x\}\} && \text{“}B_x \text{ cut in level } i \text{ at iteration } j\text{”}, \\ \mathcal{F}_{X,i,j} &= \{B_x \cap X_j = \emptyset\} && \text{“}B_x \text{ is not cut in level } i \text{ iteration } j\text{”}. \end{aligned}$$

Denote by $\mathcal{F}_{X,i,(<j)}$ the event $\bigwedge_{0 < h < j} \mathcal{F}_{X,i,h}$, by $\bar{\mathcal{F}}_{X,i,j} = \{B_x \cap X_j \neq \emptyset\}$ the event that B_x is either cut or contained in X_j , and let

$$\mathcal{E}_{X,i,j} = \{\mathcal{C}_{X,i,j} \wedge \mathcal{F}_{X,i,(<j)}\} \quad \text{“}B_x \text{ is cut in level } i \text{ iteration } j, \text{ but not before”}.$$

Finally, define $\mathcal{E}_{X,i} = \left\{ \mathcal{S}_{X,i} \wedge \bigcup_j \mathcal{E}_{X,i,j} \right\}$, as the event that B_x is cut for the first time in level i . With our new notations, we can bound the expected distance, as follows:

$$\begin{aligned}
\mathbb{E}[d_{\mathcal{T}}(x, y)] &\leq \sum_{\mathcal{T}} \Pr[\mathcal{T}] \cdot d_{\mathcal{T}}(x, y) \\
&= \sum_{\mathcal{T}} \sum_{i \in \mathbb{N}} \sum_{X \subseteq V} \Pr[\mathcal{E}_{X,i}] \cdot \Pr[\mathcal{T} \mid \mathcal{E}_{X,i}] \cdot d_{\mathcal{T}}(x, y) \\
&= \sum_{X,i} \Pr[\mathcal{E}_{X,i}] \sum_{\mathcal{T}} \Pr[\mathcal{T} \mid \mathcal{E}_{X,i}] \cdot d_{\mathcal{T}}(x, y) \\
&\leq 8 \sum_{X,i} \Pr[\mathcal{E}_{X,i}] \cdot \Delta_X \cdot \sum_{\mathcal{T}} \Pr[\mathcal{T} \mid \mathcal{E}_{X,i}] \\
&= 8 \sum_{X,i} \Pr[\mathcal{E}_{X,i}] \cdot \Delta_X. \tag{7}
\end{aligned}$$

Recall the constants defined in `create-petal`: $L_k = \lceil \log \log k \rceil$ and $L_n = \lceil \log \log n \rceil$. For $Y \subseteq X \subseteq V$ define $\varphi_k(X, Y, x) = \max \left\{ \frac{|X|_k}{|B_Y(x, \Delta_X / (2^{8L_k}))|_k}, e \right\}$, and similarly $\varphi_n(X, Y, x) = \max \left\{ \frac{|X|}{|B_Y(x, \Delta_X / (2^{8L_n}))|}, e \right\}$. The following lemma bounds the probability that a ball around a terminal or an arbitrary vertex, is cut. We defer its proof to Section 10.3.1.

Lemma 14. *For any integer i , and a cluster X with radius Δ_X (from some root vertex), it holds that if x is a terminal ($x \in K$) then (where x, y are as defined earlier)*

$$\Pr[\mathcal{E}_{X,i}] \Delta_X \leq 2^{14} d(x, y) L_k \sum_{j \leq k+1} \sum_{Y \subseteq X} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \overline{\mathcal{F}}_{X,i,j}] \ln(\varphi_k(X, Y, x)). \tag{8}$$

Moreover, without dependence whether x is a terminal, it holds that

$$\begin{aligned}
\Pr[\mathcal{E}_{X,i}] \Delta_X &\leq 2^{14} d(x, y) L_k \sum_{j \leq k+1} \sum_{Y \subseteq X} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \overline{\mathcal{F}}_{X,i,j}] \ln(\varphi_k(X, Y, x)) \\
&\quad + 2^{14} d(x, y) L_n \sum_j \sum_{Y \subseteq X} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \overline{\mathcal{F}}_{X,i,j}] \ln(\varphi_n(X, Y, x)). \tag{9}
\end{aligned}$$

The algorithm returns us a hierarchical tree \mathcal{T} . We refer to the base level (where there are only one cluster that contains all the vertices) as level 0, to the next level as level 1 and so on. For a hierarchical tree \mathcal{T} , let $A_{\mathcal{T}} \subseteq \mathbb{N}$ be the set of levels i , in which B_x is included in a cluster X with radius Δ_X , such that $d(x, y)$ is larger than Δ_X/k^3 . Note that if x is a terminal, due to the contractions, if $d(x, y) \leq \Delta_X/k^3$, then the ball B_x could not be cut, hence $\mathcal{E}_{X,i} = \emptyset$. By Fact 1, $A_{\mathcal{T}} \subseteq \mathbb{N}$ is of size at most $\log_{8/7}(2 \cdot k^3) + 1 < 22 \log k$. (The radius has to be in the range $[d(u, v)/2, k^3 d(u, v)]$, once the radius is less than $d(u, v)/2$, B_x will surely be cut.) For any level i , denote by $A_{\mathcal{T},i} = \{i' \in A_{\mathcal{T}} : i' \geq i\}$. Similarly, let $Q_{\mathcal{T}}$ be the set of levels i , in which B_x is included in a cluster X with diameter at most $\Delta_X \leq n^3 \cdot d(x, y)$. For every vertex x due to the contraction, if $d(x, y) \leq \Delta_X/n^3$, $\mathcal{E}_{X,i} = \emptyset$. By Fact 1, $Q_{\mathcal{T}} \subseteq \mathbb{N}$ is of size at most $22 \log n$. For any level i , denote by $Q_{\mathcal{T},i} = \{i' \in Q_{\mathcal{T}} : i' \geq i\}$.

Define $l_k = \lceil \log_{8/7}(2^9 L_k) \rceil$, and $l_n = \lceil \log_{8/7}(2^9 L_n) \rceil$. Consider the partition of the indices of all the levels in the hierarchical tree to the sets $I_m = \{i \in \mathbb{N} : i = m \pmod{l_k}\}$ for $m \in \{0, 1, \dots, l_k - 1\}$. Similarly, for $g \in \{0, 1, \dots, l_n - 1\}$, $J_g = \{i \in \mathbb{N} : i = g \pmod{l_n}\}$. The next lemma, combined with Lemma 14, is used to bound the cut probability.

Claim 15. For every positive integers m, g and t , it holds that:

$$\sum_{i \in I_m: i \geq t} \sum_{X, j, Y} \Pr [\mathcal{S}_{X, i, Y, j} \wedge \bar{\mathcal{F}}_{X, i, j}] \ln(\varphi_k(X, Y, x)) \leq \sum_{X, \mathcal{T}} \Pr[\mathcal{S}_{X, t} \wedge \mathcal{T}] (\ln |X|_k + |A_{\mathcal{T}, t}|) \quad (10)$$

$$\sum_{i \in J_g: i \geq t} \sum_{X, j, Y} \Pr [\mathcal{S}_{X, i, Y, j} \wedge \bar{\mathcal{F}}_{X, i, j}] \ln(\varphi_n(X, Y, x)) \leq \sum_{X, \mathcal{T}} \Pr[\mathcal{S}_{X, t} \wedge \mathcal{T}] (\ln |X| + |Q_{\mathcal{T}, t}|) \quad (11)$$

Proof. Fix any such m , we will show (10) by (reverse) induction on $t \in I_m$.

For the base case, note that when t is sufficiently large, x and y must have been separated. E.g. if $t = \log_{8/7}(2\text{diam}(G)/d(x, y))$, at levels $i \geq t$ the radius of any cluster will be less than $d(x, y)/2$ (using that the radius drops by at least $7/8$ at every level). We get that $\Pr[\mathcal{E}_{X, i}] = 0$ for any possible X and $i \geq t$. Hence $\sum_{i \in I_m: i \geq t} \sum_{X, j, Y} \Pr [\mathcal{S}_{X, i, Y, j} \wedge \bar{\mathcal{F}}_{X, i, j}] = 0$ and (10) holds.

Assume (10) holds for $t + l_k \in I_m$, and prove for t . Observe that if Z is a cluster in level $t + l_k$ such that event $\mathcal{S}_{Z, t+l_k}$ holds, then there must be a cluster X at level t (the ancestor of Z) such that $\mathcal{S}_{X, t}$ holds, and there are also Y and j such that $Y = Y_{j-1}$, event $\bar{\mathcal{F}}_{X, i, j}$ holds, and $Z \subseteq Y \subseteq X$ (X_j is the ancestor of Z at level $t + 1$). Also note that $\Delta_X \geq (8/7)^{l_k} \Delta_Z = 2^9 L_k \cdot \Delta_Z$, therefore $Z \subseteq B_Z(x, 2\Delta_Z) \subseteq B_Y(x, \Delta_X/(2^8 L_k))$. It follows that

$$\begin{aligned} & \sum_{Z, \mathcal{T}} \Pr[\mathcal{S}_{Z, t+l_k} \wedge \mathcal{T}] (\ln |Z|_k + |A_{\mathcal{T}, t+l_k}|) \\ &= \sum_{\mathcal{T}, X, Y, Z, j: Z \subseteq Y \subseteq X} \Pr[\mathcal{S}_{X, t, Y, j} \wedge \bar{\mathcal{F}}_{X, t, j} \wedge \mathcal{T}] \cdot \Pr[\mathcal{S}_{Z, t+l_k} \mid \mathcal{S}_{X, t, Y, j} \wedge \bar{\mathcal{F}}_{X, t, j} \wedge \mathcal{T}] \cdot \\ & \quad \cdot (\ln |Z|_k + |A_{\mathcal{T}, t+l_k}|) \\ &\leq \sum_{\mathcal{T}, X, Y, Z, j: Z \subseteq Y \subseteq X} \Pr[\mathcal{S}_{X, t, Y, j} \wedge \bar{\mathcal{F}}_{X, t, j} \wedge \mathcal{T}] \cdot \\ & \quad \cdot \Pr[\mathcal{S}_{Z, t+l_k} \mid \mathcal{S}_{X, t, Y, j} \wedge \bar{\mathcal{F}}_{X, t, j} \wedge \mathcal{T}] \left(\ln |B_Y(x, \frac{\Delta_X}{2^8 L_k})|_k + |A_{\mathcal{T}, t+l_k}| \right) \\ &\leq \sum_{\mathcal{T}, X, Y, j: Y \subseteq X} \Pr[\mathcal{S}_{X, t, Y, j} \wedge \bar{\mathcal{F}}_{X, t, j} \wedge \mathcal{T}] \left(\ln |B_Y(x, \frac{\Delta_X}{2^8 L_k})|_k + |A_{\mathcal{T}, t+l_k}| \right). \end{aligned} \quad (12)$$

Where the last inequality is by changing the order of summation and the fact that $\sum_{Z \subseteq Y} \Pr[\mathcal{S}_{Z, t+l_k} \mid \mathcal{S}_{X, t, Y, j} \wedge \bar{\mathcal{F}}_{X, t, j} \wedge \mathcal{T}] \leq 1$. Now, if in the hierarchical tree \mathcal{T} , $t \in A_{\mathcal{T}}$, then

$|A_{\mathcal{T},t}| \geq |A_{\mathcal{T},t+l_k}| + 1$, so that

$$\begin{aligned}
& \sum_{i \in I_m, i \geq t} \sum_{X,j,Y} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \bar{\mathcal{F}}_{X,i,j}] \ln(\varphi_k(X, Y, x)) \\
&= \sum_{i \in I_m, i \geq t+l_k} \sum_{X,j,Y} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \bar{\mathcal{F}}_{X,i,j}] \ln(\varphi_k(X, Y, x)) \\
&\quad + \sum_{X,j,Y} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j}] \ln(\varphi_k(X, Y, x)) \\
&\stackrel{(10)}{\leq} \sum_{X,\mathcal{T}} \Pr[\mathcal{S}_{X,t} \wedge \mathcal{T}] (\ln |X|_k + |A_{\mathcal{T},t+l_k}|) + \sum_{X,j,Y_{j-1}} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j}] \ln(\varphi_k(X, Y_{j-1}, x)) \\
&\stackrel{(12)}{\leq} \sum_{\mathcal{T}, X, Y, j: Y \subseteq X} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \wedge \mathcal{T}] \left(\ln |B_Y(x, \frac{\Delta_X}{2^8 L_k})|_k + |A_{\mathcal{T},t+l_k}| \right) \\
&\quad + \sum_{\mathcal{T}, X, Y, j} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \wedge \mathcal{T}] \ln \varphi_k(X, Y, x) \\
&\leq \sum_{\mathcal{T}, X, Y, j} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \wedge \mathcal{T}] (\ln |X|_k + |A_{\mathcal{T},t}|) . \tag{13}
\end{aligned}$$

In the last inequality we used that if $\varphi_k(X, Y, x) = \frac{|X|_k}{|B_Y(x, \Delta_X/2^8 L_k)|_k}$, then

$$\ln |B_Y(x, \Delta_X/2^8 L_k)|_k + |A_{\mathcal{T},t+l}| + \ln \varphi_k(X, Y, x) = \ln |X|_k + |A_{\mathcal{T},t+l}| \leq \ln |X|_k + |A_{\mathcal{T},t}| ,$$

and if $\varphi_k(X, Y, x) = e$, then

$$\begin{aligned}
\ln |B_Y(x, \Delta_X/2^8 L)|_k + |A_{\mathcal{T},t+l}| + \ln \varphi_k(X, Y, x) &= \ln |B_Y(x, \Delta_X/2^8 L)|_k + |A_{\mathcal{T},t+l}| + 1 \\
&\leq \ln |X|_k + |A_t| .
\end{aligned}$$

In addition note that

$$\begin{aligned}
& \sum_{\mathcal{T}, X, Y, j} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \wedge \mathcal{T}] (\ln |X|_k + |A_{\mathcal{T},t}|) \\
&= \sum_{\mathcal{T}, X, Y, j} \Pr[\mathcal{S}_{X,t} \wedge \mathcal{T}] \cdot \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \mid \mathcal{S}_{X,t} \wedge \mathcal{T}] (\ln |X|_k + |A_{\mathcal{T},t}|) \\
&= \sum_{\mathcal{T}, X} \Pr[\mathcal{S}_{X,t} \wedge \mathcal{T}] (\ln |X|_k + |A_{\mathcal{T},t}|) \sum_{Y,j} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \mid \mathcal{S}_{X,t} \wedge \mathcal{T}] \\
&= \sum_{\mathcal{T}, X} \Pr[\mathcal{S}_{X,t} \wedge \mathcal{T}] (\ln |X|_k + |A_{\mathcal{T},t}|) \tag{14}
\end{aligned}$$

where the last equality follows by $\sum_{Y,j} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \mid \mathcal{S}_{X,t} \wedge \mathcal{T}] = 1$, which holds because for every different Y, j , the events $\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j}$ are disjoint. This concludes the proof of (10).

The proof of (11) is fully symmetric. Fix some g , we will show (11) by (reverse) induction on $t \in J_g$. The base case is trivial (for every X and $t \geq \log_{8/7}(2 \text{diam}(G)/d(x, y))$, $\Pr[\mathcal{S}_{X,i}] = 0$ and therefore $\sum_{i \in J_g: i \geq t} \sum_{X,j,Y} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \bar{\mathcal{F}}_{X,i,j}] = 0$.) Assume (11) holds for $t+l_n \in J_g$, and prove for t . For cluster Z such that event $\mathcal{S}_{Z,t+l_n}$ holds, there are unique clusters X, Y and an index

j , such that the events $\mathcal{S}_{X,t}, \bar{\mathcal{F}}_{X,i,j}$ holds, $Y = X_j$, and $Z \subseteq B_Z(x, 2\Delta_Z) \subseteq B_Y(x, \Delta_X/(2^8 L_n))$. Similarly to equation (12) we get

$$\begin{aligned}
& \sum_{Z, \mathcal{T}} \Pr[\mathcal{S}_{Z,t+l_n} \wedge \mathcal{T}] (\ln |Z| + |Q_{\mathcal{T},t+l_n}|) \\
&= \sum_{\mathcal{T}, X, Y, Z, j : Z \subseteq Y \subseteq X} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \wedge \mathcal{T}] \cdot \Pr[\mathcal{S}_{Z,t+l_n} \mid \mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \wedge \mathcal{T}] \cdot \\
&\quad \cdot (\ln |Z| + |Q_{\mathcal{T},t+l_n}|) \\
&\leq \sum_{\mathcal{T}, X, Y, j : Y \subseteq X} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \wedge \mathcal{T}] \left(\ln |B_Y(x, \frac{\Delta_X}{2^8 L_n})| + |Q_{\mathcal{T},t+l_n}| \right). \tag{15}
\end{aligned}$$

Now we make the induction step:

$$\begin{aligned}
& \sum_{i \in J_g, i \geq t} \sum_X \sum_{j, Y} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \bar{\mathcal{F}}_{X,i,j}] \ln(\varphi_n(X, Y, x)) \\
&= \sum_{i \in J_g, i \geq t+l_n} \sum_X \sum_{j, Y} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \bar{\mathcal{F}}_{X,i,j}] \ln(\varphi_n(X, Y, x)) \\
&\quad + \sum_X \sum_{j, Y} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j}] \ln(\varphi_n(X, Y, x)) \\
&\stackrel{(11)}{\leq} \sum_{X, \mathcal{T}} \Pr[\mathcal{S}_{X,t+l_n} \wedge \mathcal{T}] (\ln |X| + |Q_{\mathcal{T},t+l_n}|) \\
&\quad + \sum_X \sum_{j, Y} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j}] \ln(\varphi_n(X, Y, x)) \\
&\stackrel{(15)}{\leq} \sum_{\mathcal{T}, X, Y, j} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \wedge \mathcal{T}] \left(\ln |B_Y(x, \frac{\Delta_X}{2^8 L_k})| + |Q_{\mathcal{T},t+l_n}| \right) \\
&\quad + \sum_{\mathcal{T}, X, Y, j} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \wedge \mathcal{T}] \ln(\varphi_n(X, Y_{j-1}, x)) \\
&\leq \sum_{\mathcal{T}, X, Y, j} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \wedge \mathcal{T}] (\ln |X| + |Q_{\mathcal{T},t}|) .
\end{aligned}$$

By the same calculation as in equation (14) we have $\sum_{\mathcal{T}, X, Y, j} \Pr[\mathcal{S}_{X,t,Y,j} \wedge \bar{\mathcal{F}}_{X,t,j} \wedge \mathcal{T}] (\ln |X| + |Q_{\mathcal{T},t}|) = \sum_{X, \mathcal{T}} \Pr[\mathcal{S}_{X,t} \wedge \mathcal{T}] (\ln |X| + |Q_{\mathcal{T},t}|)$, Which concludes the proof of (11). \square

Proof of Theorem 12. The proof of Theorem 12 follows from Lemma 14, Claim 15 and equation (7). We start by proving the first assertion of Theorem 12. Recall that for every hierarchical tree \mathcal{T} , $|A_{\mathcal{T}}| \leq 22 \log k$. Note that $|X|_k \leq k$.

$$\begin{aligned}
\mathbb{E}[d_T(x, y)] &\stackrel{(7)}{\leq} 8 \sum_{X,i} \Pr[\mathcal{E}_{X,i}] \cdot \Delta_X \\
&= 8 \sum_{m=0}^{l_k-1} \sum_{i \in I_m} \sum_X \Pr[\mathcal{E}_{X,i}] \cdot \Delta_X \\
&\stackrel{(8)}{\leq} 2^{17} L_k \cdot d(x, y) \sum_{m=0}^{l_k-1} \sum_{i \in I_m} \sum_{X, j \leq k+1, Y} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \bar{\mathcal{F}}_{X,i,j}] \ln(\varphi_k(X, Y, x)) \\
&\stackrel{(10)}{\leq} 2^{17} d(x, y) L_k \sum_{m=0}^{l_k-1} \sum_{X, \mathcal{T}} \Pr[\mathcal{S}_{X,0} \wedge \mathcal{T}] (\ln |X|_k + |A_{\mathcal{T},0}|) \\
&\leq 2^{17} L_k \cdot d(x, y) (\ln k + 22 \log k) l_k \sum_{\mathcal{T}} \Pr[\mathcal{S}_{V,0} \wedge \mathcal{T}] \\
&\leq 2^{29} \log k \log \log k \log \log \log k \cdot d(x, y) = \tilde{O}(\log k) d(x, y) .
\end{aligned}$$

Similarly, for the second assertion of Theorem 12, recall that for every hierarchical tree \mathcal{T} , $|Q_{\mathcal{T}}| \leq 22 \log n$. Note that $|X| \leq n$.

$$\begin{aligned}
\mathbb{E}[d_T(x, y)] &\stackrel{(7)}{\leq} 8 \sum_{X,i} \Pr[\mathcal{E}_{X,i}] \cdot \Delta_X \\
&\stackrel{(9)}{\leq} 8 \sum_{X,i} 2^{14} d(x, y) L_k \sum_{j \leq k} \sum_Y \Pr[\mathcal{S}_{X,i,Y,j} \wedge \bar{\mathcal{F}}_{X,i,j}] \ln(\varphi_k(X, Y, x)) \\
&\quad + 8 \sum_{X,i} 2^{14} d(x, y) L_n \sum_j \sum_Y \Pr[\mathcal{S}_{X,i,Y,j} \wedge \bar{\mathcal{F}}_{X,i,j}] \ln(\varphi_n(X, Y, x)) \\
&\leq 2^{17} L_k \cdot d(x, y) \sum_{m=0}^{l_k-1} \sum_{i \in I_m} \sum_{X, j \leq k, Y} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \bar{\mathcal{F}}_{X,i,j}] \ln(\varphi_k(X, Y, x)) \\
&\quad + 2^{17} d(x, y) L_n \sum_{g=0}^{l_n-1} \sum_{i \in J_g} \sum_{X, j, Y} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \bar{\mathcal{F}}_{X,i,j}] \ln(\varphi_n(X, Y, x)) \\
&\stackrel{(10) \wedge (11)}{\leq} 2^{17} d(x, y) L_k \sum_{m=0}^{l_k-1} \sum_{i \in I_m} \sum_{X, \mathcal{T}} \Pr[\mathcal{S}_{X,0} \wedge \mathcal{T}] (\ln |X|_k + |A_{\mathcal{T},0}|) \\
&\quad + 2^{17} d(x, y) L_n \sum_{g=0}^{l_n-1} \sum_{i \in J_g} \sum_{X, \mathcal{T}} \Pr[\mathcal{S}_{X,0} \wedge \mathcal{T}] (\ln |X| + |Q_{\mathcal{T},0}|) \\
&\leq (2^{29} \log k \log \log k \log \log \log k + 2^{29} \log n \log \log n \log \log \log n) d(x, y) \\
&= \tilde{O}(\log n) d(x, y) .
\end{aligned}$$

□

10.3.1 Proof of Lemma 14

Let X be the vertex set of the graph in the current level i of the petal decomposition, with arbitrary center x_0 , target t and radius Δ (with respect to x_0). Recall the vertices $x, y \in X$, and the ball $B_x = B_G(x, d(x, y))$. Set $\gamma = d(x, y)/\Delta$. As X and i are fixed, and all the probabilistic events in the statement of the lemma are contained in $\mathcal{S}_{X,i}$, we will implicitly condition all the probabilistic events during the proof on $\mathcal{S}_{X,i}$ (i.e. our sample space is restricted to $\mathcal{S}_{X,i}$). We shall also omit X and i from the subscript of the probabilistic events (i.e., we will write $\mathcal{C}_j, \mathcal{F}_j, \mathcal{F}_{(<j)}, \mathcal{E}_j, \mathcal{E}$ and we will write $\mathcal{Z}_{Y,j}$ instead of $\mathcal{S}_{X,i,Y,j}$).

The petal decomposition algorithm returns a partition $(X_0, \dots, X_s, (y_1, x_1), \dots, (y_s, x_s), t_0, \dots, t_s)$. We make a small change in the numbering of the created petals: let t_1, \dots, t_{k+1} be the terminal targets chosen in lines 6 or 12 of **petal-decomposition**, and t_{k+2}, \dots, t_{n+1} be the non-terminal targets chosen in line 21. Observe that in that notation we always have exactly $n+1$ petals, while there might be an index j such that $X_j = \emptyset$, in that case we say that X_j is an imaginary petal. Note that there are at most $k+1$ terminal targets because there are just k terminals (in addition to the first special petal whose target is not necessarily a terminal).

Recall the definitions from **create-petal** procedure, some of them depend on the type of petal (terminal or non-terminal petal), and some of them are actually random variables (which depend on the previously created petals). We will write these with an index j to clarify which petal they correspond to. In the terminal petal case ($j \leq k+1$), $L_k = \lceil \log \log k \rceil$, $R_j = hi - lo = \Delta/40$ (after reformation), $a_j = lo + (q_j - 1)R_j/L_k$, $b'_j = a_j + R_j/(2L_k)$, $b_j = a_j + R_j/L_k$, where q_j is chosen in such a way that

$$\frac{2|X|_k}{2^{\log^{1-(q-1)/L_k} k}} \leq |W_{a_j}|_k \leq |W_{b_j}|_k \leq \frac{2|X|_k}{2^{\log^{1-q/L_k} k}}. \quad (16)$$

Also, $\hat{\chi}_j = \max\{\frac{|X|_k+1}{|W_{a_j}|_k}, e\}$, $\lambda_j = \frac{2 \ln \hat{\chi}_j}{b'_j - a_j} = \frac{4L_k \ln \hat{\chi}_j}{R_j} = \frac{160L_k \ln \hat{\chi}_j}{\Delta}$. The radius r_j of the petal $X_j = W_{r_j}$, is chosen from $[a_j, b'_j]$ with density function $f(r) = \frac{\lambda_j \cdot e^{-\lambda_j r}}{e^{-\lambda_j \cdot a_j} - e^{-\lambda_j \cdot b'_j}}$. Analogously, in the non-terminal case ($j > k+1$): $L_n = \lceil \log \log n \rceil$, $R_j = hi - lo = \Delta/32$ (no reformation), $a_j = lo + (q_j - 1)R_j/L_n$, $b'_j = a_j + R_j/(2L_n)$, $b_j = a_j + R_j/L_n$, where q_j is chosen in such a way that $\frac{2|X|}{2^{\log^{1-(q_j-1)/L_n} n}} \leq |W_{a_j}| \leq |W_{b_j}| \leq \frac{2|X|}{2^{\log^{1-q_j/L_n} n}}$. Also, $\hat{\chi}_j = \max\{\frac{|X|+1}{|W_{a_j}|}, e\}$, $\lambda_j = \frac{2 \ln \hat{\chi}_j}{b'_j - a_j} = \frac{4L_n \ln \hat{\chi}_j}{R_j} = \frac{128L_n \ln \hat{\chi}_j}{\Delta}$. The radius r_j of the petal $X_j = W_{r_j}$, is chosen from $[a_j, b'_j]$ with density function $f(r) = \frac{\lambda_j \cdot e^{-\lambda_j r}}{e^{-\lambda_j \cdot a_j} - e^{-\lambda_j \cdot b'_j}}$.

Let $\delta_j = e^{-8\lambda_j \gamma \Delta}$. Towards the proof, we assume that $\gamma \leq \frac{1}{2^{10} L_k}$ as otherwise the assertions of the lemma are trivial (as $\Pr[\mathcal{E}_{X,i}] \leq \sum_j \sum_{Y \subseteq X} \Pr[\mathcal{S}_{X,i,Y,j} \wedge \overline{\mathcal{F}}_{X,i,j}]$).

Claim 16. For every $1 \leq j \leq n+1$, $\Pr[C_j | \mathcal{Z}_{Y,j}] \leq (1 - \delta_j) \left(\Pr[\overline{\mathcal{F}}_j | \mathcal{Z}_{Y,j}] + \frac{2}{\hat{\chi}_j^2} \right)$.

Proof. For ease of notation we write simply λ, δ, χ, a and b' instead of $\lambda_j, \delta_j, \chi_j, a_j, b'_j$, as the proof is the same for both of the cases.²¹ Let ρ be the minimal number greater than a such that $W_\rho \cap B_x \neq \emptyset$. If $\rho \geq b'$, then trivially $\Pr[\overline{\mathcal{F}}_j | \mathcal{Z}_{Y,j}] = 1$ and hence $\mathcal{P}_r[C_j | \mathcal{Z}_{Y,j}] = 0$ and we are

²¹Note that if $X_j = \emptyset$ is imaginary petal created only for clarity in notation, then $\Pr[C_j | \mathcal{Z}_{Y,j}] = 0$ and the claim is obvious.

done. Otherwise, $\rho \in [a, b']$, the probability that B_x intersects X_j is

$$\Pr[\overline{\mathcal{F}}_j \mid \mathcal{Z}_{Y,j}] = \int_{\rho}^{b'} f(r) dr = \int_{\rho}^{b'} \frac{\lambda \cdot e^{-\lambda r}}{e^{-\lambda \cdot a} - e^{-\lambda \cdot b'}} dr = \frac{e^{-\lambda \rho} - e^{-\lambda b'}}{e^{-\lambda \cdot a} - e^{-\lambda \cdot b'}}.$$

The ball B_x is of diameter at most $2\gamma\Delta$, therefore by Claim 12, $B_x \subseteq W_{\rho+8\gamma\Delta}$. Hence

$$\begin{aligned} \Pr[C_j \mid \mathcal{Z}_{Y,j}] &\leq \int_{\rho}^{\rho+8\gamma\Delta} f(r) dr = \int_{\rho}^{\rho+8\gamma\Delta} \frac{\lambda \cdot e^{-\lambda r}}{e^{-\lambda \cdot a} - e^{-\lambda \cdot b'}} dr = \frac{e^{-\lambda \rho} - e^{-\lambda(\rho+8\gamma\Delta)}}{e^{-\lambda \cdot a} - e^{-\lambda \cdot b'}} \\ &= \frac{e^{-\lambda \rho}}{e^{-\lambda \cdot a} - e^{-\lambda \cdot b'}} \left(1 - e^{-\lambda 8\gamma\Delta}\right) = (1 - \delta) \frac{e^{-\lambda \rho} - e^{-\lambda \cdot b'} + e^{-\lambda \cdot b'}}{e^{-\lambda \cdot a} - e^{-\lambda \cdot b'}} \\ &= (1 - \delta) \left(\frac{e^{-\lambda \rho} - e^{-\lambda \cdot b'}}{e^{-\lambda \cdot a} - e^{-\lambda \cdot b'}} + \frac{e^{-\lambda \cdot b'}}{e^{-\lambda \cdot a} - e^{-\lambda \cdot b'}} \right) \\ &= (1 - \delta) \left(\Pr[\overline{\mathcal{F}}_j \mid \mathcal{Z}_{Y,j}] + \frac{1}{e^{\lambda \cdot (b' - a)} - 1} \right) \\ &\leq (1 - \delta) \left(\Pr[\overline{\mathcal{F}}_j \mid \mathcal{Z}_{Y,j}] + \frac{2}{\hat{\chi}^2} \right). \end{aligned}$$

Where the last inequality follows since $\hat{\chi} \geq 2$. □

We are now ready to bound the probability that one of the terminal petals cuts B_x .

$$\begin{aligned} \Pr\left[\bigvee_{j \leq k+1} C_j\right] &= \sum_{j \leq k+1} \sum_{Y \subseteq X} \Pr[C_j \mid \mathcal{Z}_{Y,j}] \cdot \Pr[\mathcal{Z}_{Y,j}] \tag{17} \\ &\leq \sum_{j \leq k+1} \sum_Y (1 - \delta_j) \left(\Pr[\overline{\mathcal{F}}_j \mid \mathcal{Z}_{Y,j}] + \frac{2}{\hat{\chi}_j^2} \right) \cdot \Pr[\mathcal{Z}_{Y,j}] \\ &= \sum_{j \leq k+1} \sum_Y (1 - \delta_j) \Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] + \sum_{j \leq k+1} \sum_Y (1 - \delta_j) \frac{2}{\hat{\chi}_j^2} \cdot \Pr[\mathcal{Z}_{Y,j}]. \end{aligned}$$

To see how the first equality follows, note that the probability that B_x is cut by the first $k+1$ petals ($\Pr\left[\bigvee_{j \leq k+1} C_j\right]$) is equal to the sum of the probabilities that B_x cut at the first time by petal j ($\sum_{j \leq k+1} \Pr\left[C_j \wedge \bigwedge_{l < j} \overline{C}_l\right]$). While for each j , the probability that B_x is cut at the first time by petal j is equal to the probability that B_x is active at iteration j (i.e. $\mathcal{F}_{(< j)}$), and is indeed cut by petal j ($\sum_{Y \subseteq X} \Pr[C_j \mid \mathcal{Z}_{Y,j}] \cdot \Pr[\mathcal{Z}_{Y,j}]$).

Note that for $j \leq k+1$, $1 - \delta_j = 1 - e^{-8\lambda_j \gamma \Delta} = 1 - e^{-8 \cdot 160 \gamma L_k \ln \hat{\chi}_j} < 2^{11} \gamma L_k \ln \hat{\chi}_j$. For a set $Y = Y_{j-1}$ such that $\Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] \neq 0$, necessarily a target t_j and q_j are chosen so that $B_x \cap W_{b'_j} \neq \emptyset$ (otherwise it is not possible that X_j intersects B_x). As B_x is a ball of radius $\gamma\Delta \leq \Delta/(2^{10} L_k)$, increasing the petal radius by $4\gamma\Delta$ we have $x \in W_{b'_j + 4\gamma\Delta}$, and moreover $B\left(x, \frac{\Delta}{2^9 L_k}\right) \subseteq W_{b'_j + 4\gamma\Delta + \frac{\Delta}{2^7 L_k}} \subseteq W_{b'_j + \frac{\Delta}{2^8 L_k} + \frac{\Delta}{2^7 L_k}} \subseteq W_{b'_j + \frac{R_j}{2L_k}} = W_{b'_j}$.

We will show that for such a Y , (for which $\Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] \neq 0$), $\log(\hat{\chi}_j) \leq 6 \log(\varphi_k(X, Y, x))$.²² We may assume that $\hat{\chi}_j = \frac{|X|_k + 1}{|W_{a_j}|_k}$, as otherwise ($\hat{\chi}_j = e$) the bound is trivial. Using that $\log^{1/L_k} k \leq 3$ we get:

$$\begin{aligned}
\log \hat{\chi}_j &= \log \left(\frac{|X|_k + 1}{|W_{a_j}|_k} \right) \\
&\leq \log \left(\frac{2|X|_k}{|W_{a_j}|_k} \right) \\
&\stackrel{(16)}{\leq} \log^{1-(q-1)/L_k} k \\
&\leq 3 \log^{1-q/L_k} k \\
&\stackrel{(16)}{\leq} 3 \log \left(\frac{2|X|_k}{|W_{b_j}|_k} \right) \\
&\leq 3 \log \left(\frac{2|X|_k}{|B(x, \Delta/2^9 L_k)|_k} \right) \\
&= 3 (\log(\varphi_k(X, Y, x)) + 1) \\
&\leq 6 \log(\varphi_k(X, Y, x)) .
\end{aligned}$$

Note that in particular it is true that $\ln \hat{\chi}_j \leq 6 \ln(\varphi_k(X, Y, x))$. Hence we can bound the first component of the summation in (17):

$$\begin{aligned}
\sum_{j \leq k+1} \sum_Y (1 - \delta_j) \Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] &\leq 2^{11} \gamma L_k \sum_{j \leq k+1} \sum_Y \ln(\hat{\chi}_j) \Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] \quad (18) \\
&\leq 6 \cdot 2^{11} \gamma L_k \sum_{j \leq k+1} \sum_Y \ln(\varphi_k(X, Y, x)) \Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}]
\end{aligned}$$

For bounding the second component ($\sum_{j \leq k+1} \sum_Y (1 - \delta_j) \frac{2}{\hat{\chi}_j^2} \cdot \Pr[\mathcal{Z}_{Y,j}]$), note that $\frac{1}{\hat{\chi}_j} = \min \left\{ \frac{|W_{a_j}|_k}{|X|_k + 1}, \frac{1}{e} \right\} \leq \frac{|W_{a_j}|_k}{|X|_k}$. In addition, the probability of the event $S_{Y,j}$ is equal to the sum of probabilities of sequences X_0, \dots, X_s , over all sequences for which $Y = (X_j \cup \dots \cup X_s \cup X_0)$ (while we abuse notation, $\Pr[X_0, \dots, X_s]$ is the probability that the petal decomposition returned the partition X_0, \dots, X_s). Note also that for each such partition, X_0, \dots, X_s , $|W_{a_j}|_k \leq |X_j|_k$ and $\sum_{j \leq k+1} \frac{|X_j|_k}{|X|_k} \leq 1$, because all the petals, X_1, \dots, X_{k+1} , are pairwise disjoint. We get

²² Recall that $\varphi_k(X, Y, x) = \max \left\{ \frac{|X|_k}{|B_Y(x, \Delta_X/(2^8 L_k))|_k}, e \right\}$.

$$\begin{aligned}
& \sum_{j \leq k+1} \sum_Y \frac{\Pr[S_{Y,j}]}{\hat{\chi}_j} \\
&= \sum_{j \leq k+1} \sum_{Y: B_x \subseteq Y} \sum_{(X_0, \dots, X_s): Y = (X_j \cup \dots \cup X_s \cup X_0)} \Pr[X_0, \dots, X_s] \frac{1}{\hat{\chi}_j} \\
&\leq \sum_{j \leq k+1} \sum_{Y: B_x \subseteq Y} \sum_{(X_0, \dots, X_s): Y = (X_j \cup \dots \cup X_s \cup X_0)} \Pr[X_0, \dots, X_s] \frac{|W_{a_j}|_k}{|X|_k} \\
&\leq \sum_{j \leq k+1} \sum_{(X_0, \dots, X_s): Y: B_x \subseteq Y, Y = (X_j \cup \dots \cup X_s \cup X_0)} \Pr[X_0, \dots, X_s] \frac{|X_j|_k}{|X|_k} \\
&\leq \sum_{j \leq k+1} \sum_{(X_0, \dots, X_s)} \Pr[X_0, \dots, X_s] \frac{|X_j|_k}{|X|_k} \\
&= \sum_{(X_0, \dots, X_s)} \Pr[X_0, \dots, X_s] \sum_{j \leq k+1} \frac{|X_j|_k}{|X|_k} \\
&\leq \sum_{(X_0, \dots, X_s)} \Pr[X_0, \dots, X_s] = 1.
\end{aligned}$$

Where the third inequality follows by the fact that for a fixed j and a particular partition X_0, \dots, X_s of X there might be only a single set Y such that $B_x \subseteq Y = X_j \cup \dots \cup X_s \cup X_0$. As $\frac{\ln \hat{\chi}_j}{\hat{\chi}_j} \leq 1$ we can bound the second component of the summation in (17):

$$\begin{aligned}
\sum_{j \leq k+1} \sum_Y (1 - \delta_j) \frac{2}{\hat{\chi}_j^2} \cdot \Pr[\mathcal{Z}_{Y,j}] &\leq 2^{11} \gamma L_k \sum_{j \leq k+1} \sum_Y \ln \hat{\chi}_j \frac{2}{\hat{\chi}_j^2} \cdot \Pr[\mathcal{Z}_{Y,j}] \\
&\leq 2^{12} \gamma L_k \sum_{j \leq k+1} \sum_Y \frac{\Pr[\mathcal{Z}_{Y,j}]}{\hat{\chi}_j} \\
&\leq 2^{12} \gamma L_k \\
&\leq 2^{12} \gamma L_k \sum_{Y,j} \Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] \ln \varphi_k(X, Y, x). \tag{19}
\end{aligned}$$

Where the last inequality follows by $\sum_{Y,j} \Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] \ln \varphi_k(X, Y, x) \geq \sum_{Y,j} \Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] = 1$ (recall that all the probabilities are implicitly conditioned on $\mathcal{S}_{X,i}$, and $\overline{\mathcal{F}}_j$ has to hold for some j).

We conclude:

$$\begin{aligned}
& \Pr \left[\bigvee_{j \leq k+1} C_j \right] \\
& \stackrel{(17)}{\leq} \sum_{j \leq k+1} \sum_Y (1 - \delta_j) \Pr [\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] + \sum_{j \leq k+1} \sum_Y (1 - \delta_j) \frac{2}{\hat{\chi}_j^2} \cdot \Pr[\mathcal{Z}_{Y,j}] \\
& \stackrel{(18) \wedge (19)}{\leq} (6 \cdot 2^{11} + 2^{12}) \gamma L_k \sum_{j \leq k+1} \sum_Y \Pr [\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] \ln(\varphi_k(X, Y, v)) \\
& = 2^{14} \gamma L_k \sum_{j \leq k+1} \sum_Y \Pr [\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] \ln(\varphi_k(X, Y, v)) .
\end{aligned}$$

Claim 17. *If $x \in K$, then for $j > k + 1$, it holds that $\Pr[C_j \mid \mathcal{Z}_{Y,j}] = 0$.*

Proof. Since we condition on $\mathcal{Z}_{Y,j}$, it implies that $B_x \subseteq Y_{j-1}$. As $\gamma < 1/16$, the third assertion of Claim 13 implies that $X_j \cap B_x = \emptyset$. \square

We are ready to prove the first statement of Lemma 14. For terminal $v \in K$, Using Claim 17,

$$\begin{aligned}
\Pr[\mathcal{E}] &= \Pr \left[\bigvee_{j \leq n+1} C_j \right] = \Pr \left[\bigvee_{j \leq k+1} C_j \right] \\
&\leq 2^{14} \gamma L_k \sum_{j \leq k+1} \sum_Y \Pr [\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] \ln(\varphi_k(X, Y, v)) .
\end{aligned}$$

Using a symmetric argument, we can also show the second assertion of Lemma 14. Set $\mu = \Pr \left(\bigvee_{j \leq k+1} C_j \right)$ the probability that B_x is cut by terminal petal. Let $I_n = k + 2, \dots, n + 1$ be all the indices of the non-terminal petals. Note that for $j \in I_n$, $1 - \delta_j = 1 - e^{-8\lambda_j \gamma \Delta} = 1 - e^{-8 \cdot 128 \gamma L_n \ln \hat{\chi}_j} \leq 2^{10} \gamma L_n \ln \hat{\chi}_j$. Hence for every vertex $x \in X$,

$$\begin{aligned}
\Pr[\mathcal{E}] &= \Pr \left[\bigvee_{j \leq n+1} C_j \right] \\
&= \sum_{j \leq n+1} \sum_{Y \subseteq X} \Pr[C_j \mid \mathcal{Z}_{Y,j}] \cdot \Pr[\mathcal{Z}_{Y,j}] \\
&= \mu + \sum_{j \in I_n} \sum_{Y \subseteq X} \Pr[C_j \mid \mathcal{Z}_{Y,j}] \cdot \Pr[\mathcal{Z}_{Y,j}] \\
&\leq \mu + \sum_{j \in I_n} \sum_Y (1 - \delta_j) \left(\Pr[\overline{\mathcal{F}}_j \mid \mathcal{Z}_{Y,j}] + \frac{2}{\hat{\chi}_j^2} \right) \cdot \Pr[\mathcal{Z}_{Y,j}] \\
&\leq \mu + 2^{10} \gamma L_n \left(\sum_{j \in I_n} \sum_Y \ln \hat{\chi}_j \cdot \Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] + \sum_{j \in I_n} \sum_Y \frac{2}{\hat{\chi}_j} \cdot \Pr[\mathcal{Z}_{Y,j}] \right)
\end{aligned}$$

where the first inequality is by Claim 16. For Y such that $\Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] \neq 0$ it holds that $B_x \cap W_{b'_j} \neq \emptyset$. As $\gamma \leq \frac{1}{2^{10} L_k} \leq \frac{1}{2^{10} L_n}$, by Claim 12 we have that $B_x \left(x, \frac{\Delta}{2^9 L_n} \right) \subseteq$

$W_{b'_j+4\gamma\Delta+\frac{\Delta}{2^7L_n}} \subseteq W_{b'_j+\frac{\Delta}{2^8L_n}+\frac{\Delta}{2^7L_n}} \subseteq W_{b'_j+\frac{R_j}{2L_n}} = W_{b_j}$. We will show that $\ln \hat{\chi}_j \leq 6 \ln \varphi_n(X, Y, x)$. If $\hat{\chi}_j = e$ this is trivial, hence we will assume that $\hat{\chi}_j = \frac{|X|+1}{|W_{a_j}|}$. By the maximality of q_j ($\frac{2|X|}{2^{\log^{1-(q-1)/L_n} n}} \leq |W_{a_j}| \leq |W_{b_j}| \leq \frac{2|X|}{2^{\log^{1-q/L_n} n}}$) we get $\log \hat{\chi}_j = \log \left(\frac{|X|+1}{|W_{a_j}|} \right) \leq \log \left(\frac{2|X|}{|W_{a_j}|} \right) \leq \log^{1-(q-1)/L_n} n \leq 3 \log^{1-q/L_n} n \leq 3 \log \left(\frac{2|X|}{|W_{b_j}|} \right) \leq 3 \log \left(\frac{2|X|}{|B(v, \Delta/2^9 L_n)|} \right) \leq 3 (\log(\varphi_n(X, Y, x)) + 1) \leq 6 \log(\varphi_n(X, Y, x))$.

By similar arguments (and definitions) to (19) we get

$$\begin{aligned} \sum_{j \in I_n} \sum_Y \frac{\Pr[S_{Y,j}]}{\hat{\chi}_j} &\leq \sum_{j \in I_n} \sum_{Y: B_x \subseteq Y(X_0, \dots, X_s) : Y = (X_j \cup \dots \cup X_s \cup X_0)} \Pr[X_0, \dots, X_s] \frac{|W_{a_j}|}{|X|} \\ &\leq \sum_{j \in I_n} \sum_{(X_0, \dots, X_s)} \Pr[X_0, \dots, X_s] \frac{|X_j|}{|X|} \\ &= \sum_{(X_0, \dots, X_s)} \Pr[X_0, \dots, X_s] \sum_{j \leq k+1} \frac{|X_j|}{|X|} \leq 1. \end{aligned}$$

As $1 \leq \sum_{Y,j} \Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] \ln \varphi_n(X, Y, x)$ we get

$$\begin{aligned} \Pr[\mathcal{E}] &\leq \mu + 2^{10} \gamma L_n \left(6 \sum_{j \in I_n} \sum_Y \ln(\varphi_k(X, Y, x)) \cdot \Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}] + 2 \right) \\ &\leq \mu + 2^{13} \gamma L_n \sum_{j \in I_n} \sum_Y \ln(\varphi_n(X, Y, x)) \cdot \Pr[\overline{\mathcal{F}}_j \wedge \mathcal{Z}_{Y,j}]. \end{aligned}$$

Hence the second assertion of Lemma 14 follows as well.

11 Applications

In this section we illustrate several algorithmic applications of our techniques. Some of our applications are suitable for graphs with a small vertex cover. Recall that for a graph $G = (V, E)$, a set $A \subseteq V$ is a vertex cover of G , if for any edge $e \in E$, at least one of its endpoints is in A . A polynomial time 2-approximation algorithm to this problem is folklore.

11.1 Sparsest Cut

In the sparsest-cut problem we are given a graph $G = (V, E)$ with capacities on the edges $c : E \rightarrow \mathbb{R}_+$, and a collection of pairs $(s_1, t_1), \dots, (s_r, t_r)$ along with their demands D_1, \dots, D_r . The goal is to find a cut $S \subseteq V$ that minimizes the ratio between capacity and demand across the cut:

$$\phi(S) = \frac{\sum_{\{u,v\} \in E} c(u,v) |\mathbf{1}_S(u) - \mathbf{1}_S(v)|}{\sum_{i=1}^r D_i |\mathbf{1}_S(s_i) - \mathbf{1}_S(t_i)|},$$

where $\mathbf{1}_S(\cdot)$ is the indicator for membership in S . Arora et. al. [ALN08] present an $\tilde{O}(\sqrt{\log r})$ approximation algorithm to this problem. Our contribution is the following.

Theorem 13. *If there exists a set $K \subseteq V$ of size k such that any demand pair contains a vertex of K , then there exists a $\tilde{O}(\sqrt{\log k})$ approximation algorithm for the sparsest-cut problem.*

The key ingredient of the algorithm of [ALN08] is a non-expansive embedding from ℓ_2^2 (negative-type metrics) into ℓ_1 , which has $\tilde{O}(\sqrt{\log r})$ contraction for all demand pairs. We will use the strong terminal embedding for negative type metrics given in item (7) of Corollary 1 to improve the distortion to $\tilde{O}(\sqrt{\log k})$.

We now elaborate on how to use the embedding of ℓ_2^2 into ℓ_1 to obtain an approximation algorithm for the sparsest-cut, all the details can be found in [LLR95, ARV09, ALN08], and we provide them just for completeness. First, write down the following SDP relaxation with triangle inequalities:

Algorithm 6 Sparsest Cut SDP Relaxation

$$\begin{aligned} \min \quad & \sum_{\{u,v\} \in E} c(u,v) \cdot \|\bar{u} - \bar{v}\|_2^2 \\ \text{s.t.} \quad & \sum_{i=1}^r D_i \cdot \|\bar{s}_i - \bar{t}_i\|_2^2 = 1 \\ & \text{For all } u, v, w \in V, \|\bar{u} - \bar{v}\|_2^2 + \|\bar{v} - \bar{w}\|_2^2 \geq \|\bar{u} - \bar{w}\|_2^2 \\ & \text{For all } u \in V, \bar{u} \in \mathbb{R}^n \end{aligned}$$

Note that this is indeed a relaxation: if S is the optimal cut, set $\rho = \sum_{i=1}^r D_i \cdot |\mathbf{1}_S(s_i) - \mathbf{1}_S(t_i)|$; for $u \in S$ set $\bar{u} = (\frac{1}{\rho}, 0, \dots, 0)$, and for $u \notin S$, set $\bar{u} = (0, \dots, 0)$. It can be checked to be a feasible solution of value equal to that of the cut S .

Let $K \subseteq V$ be a vertex cover of the demand graph $(V, \{\{s_i, t_i\}_{i=1}^r\})$ of size at most $2k$ (recall that we can find such a cover in polynomial time). Let $X = \{\bar{v} \in \mathbb{R}^n \mid v \in V\}$ be an optimal solution to the SDP (it can be computed in polynomial time), which is in particular an ℓ_2^2 (pseudo) metric. By Corollary 1 there exists a non-expansive embedding $f : X \rightarrow \ell_1$ with terminal distortion $\tilde{O}(\sqrt{\log k})$ (where K is the terminal set).²³ This implies that for any $u, v \in V$ and any $1 \leq i \leq r$,

$$\begin{aligned} \|\bar{u} - \bar{v}\|_2^2 &\geq \|f(\bar{v}) - f(\bar{u})\|_1 \\ \|\bar{s}_i - \bar{t}_i\|_2^2 &\leq \tilde{O}(\sqrt{\log k}) \cdot \|f(\bar{s}_i) - f(\bar{t}_i)\|_1. \end{aligned} \tag{20}$$

Let $\|f(\bar{v}) - f(\bar{u})\|_1 = \sum_{S \subseteq V} \alpha_S |\mathbf{1}_S(v) - \mathbf{1}_S(u)|$ be a representation of the ℓ_1 metric as a nonnegative linear combination of cut metrics (it is well known that there is such a representation with

²³The embedding of Corollary 1 is in fact into ℓ_2 , but there is an efficient randomized algorithm to embed any finite subset of ℓ_2 into ℓ_1 with constant distortion [FLM77].

polynomially many cuts S having $\alpha_S > 0$). We conclude

$$\begin{aligned}
\text{opt}(\text{SDP}) &= \sum_{\{u,v\} \in E} c(u,v) \cdot \|\bar{u} - \bar{v}\|_2^2 \\
&= \frac{\sum_{\{u,v\} \in E} c(u,v) \cdot \|\bar{u} - \bar{v}\|_2^2}{\sum_{i=1}^r D_i \cdot \|\bar{s}_i - \bar{t}_i\|_2^2} \\
&\stackrel{(20)}{\geq} \frac{\sum_{\{u,v\} \in E} c(u,v) \cdot \|f(\bar{v}) - f(\bar{u})\|_1}{\sum_{i=1}^r D_i \cdot \tilde{O}(\sqrt{\log k}) \cdot \|f(\bar{s}_i) - f(\bar{t}_i)\|_1} \\
&= \frac{1}{\tilde{O}(\sqrt{\log k})} \cdot \frac{\sum_{\{u,v\} \in E} c(u,v) \cdot \sum_{S \subseteq V} \alpha_S |\mathbf{1}_S(v) - \mathbf{1}_S(u)|}{\sum_{i=1}^r D_i \cdot \sum_{S \subseteq V} \alpha_S |\mathbf{1}_S(s_i) - \mathbf{1}_S(t_i)|} \\
&\geq \frac{1}{\tilde{O}(\sqrt{\log k})} \min_{S: \alpha_S > 0} \frac{\sum_{\{u,v\} \in E} c(u,v) \cdot |\mathbf{1}_S(v) - \mathbf{1}_S(u)|}{\sum_{i=1}^r D_i \cdot |\mathbf{1}_S(s_i) - \mathbf{1}_S(t_i)|} \\
&= \min_{S: \alpha_S > 0} \frac{\phi(S)}{\tilde{O}(\sqrt{\log k})}.
\end{aligned}$$

In particular, among the polynomially many sets $S \subseteq V$ with $\alpha_S > 0$, there exists one which has sparsity at most $\tilde{O}(\sqrt{\log k})$ times larger than the optimal one.

11.2 Min Bisection

In the min-bisection problem, we are given a graph $G = (V, E)$ on an even number n of vertices, with capacities $c : E \rightarrow \mathbb{R}_+$. The purpose is to find a partition of V into two equal parts $S \subseteq V$ and $\bar{S} = V \setminus S$, that minimizes $\sum_{e \in E(S, \bar{S})} c(e)$. This problem is NP-hard, and the best known approximation is $O(\log n)$ by [Räc08]. We obtain the following generalization.

Theorem 14. *There exists a $O(\log k)$ approximation algorithm for min-bisection, where k is the size of a minimal vertex cover of the input graph.*

Proof. Our algorithm follows closely the algorithm of [Räc08], the major difference is that we use our embedding into trees with terminal congestion. The algorithm starts by efficiently finding a vertex cover $K \subseteq V$ of size at most $2k$. We will call the vertices in K terminals. Let \mathcal{D} a distribution over trees with strong terminal congestion ($O(\log k), O(\log n)$) given by Corollary 10. The algorithm will sample a tree $T = (V, E_T)$ from \mathcal{D} , find an optimal bisection in T and return it. We refer the reader to Section 9 for details on notation and on the definition of capacities $C_T : E_T \rightarrow \mathbb{R}_+$ for T . We note that there is polynomial time algorithm (by dynamic programming) to find a min-bisection in trees.

It remains to analyze the algorithm. Let $S \subseteq V$ be the optimal solution in G , and S_T be the optimal bisection for the tree T . The expected cost of using S_T in G can be bounded using

Lemma 11 as follows

$$\begin{aligned}
\sum_{T \in \text{supp}(\mathcal{D})} \Pr[T] \sum_{e \in E(S_T, \bar{S}_T)} c(e) &\stackrel{(5)}{\leq} \sum_T \Pr[T] \sum_{e' \in E_T(S_T, \bar{S}_T)} C_T(e') \\
&\leq \sum_T \Pr[T] \sum_{e' \in E_T(S, \bar{S})} C_T(e') \\
&\stackrel{(5)}{\leq} \sum_T \Pr[T] \sum_{e \in E(S, \bar{S})} \text{load}_T(e) \\
&= \sum_{e \in E(S, \bar{S})} \mathbb{E}_{T \sim \mathcal{D}} [\text{load}_T(e)] \\
&\leq \sum_{e \in E(S, \bar{S})} O(\log k) \cdot c(e) \\
&= O(\log k) \cdot \text{opt}(G) \quad ,
\end{aligned}$$

where the last inequality uses that every edge touches a terminal, so its expected congestion is $O(\log k)$. □

11.3 Online Algorithms: Constrained File Migration

We illustrate the usefulness of our probabilistic terminal embedding into ultrametric via the constrained file migration problem. This is an online problem, in which we are given a graph $G = (V, E)$ representing a network, each node $v \in V$ has a memory capacity m_v , and a parameter $D \geq 1$. There is some set of files that reside at the nodes, at most m_v files may be stored at node v in any given time. The cost of accessing a file that currently lies at v from node u is $d_G(u, v)$ (no copies of files are allowed). Files can also be migrated from one node to another, this costs D times the distance. When a sequence of file requests arrives online, the goal is to minimize the cost of serving all requests. The *competitive ratio* of an online algorithm is the maximal ratio between its cost to the cost of an optimal (offline) solution. For randomized algorithms the expected cost is used.

We consider the case where there exists a small set of vertices which are allowed to store files (i.e. $m_v > 0$). One may think about these vertices as servers who store files, while allowing file requests from all end users. Let $K \subseteq V$ be the set of terminal vertices that are allowed to store files, with $|K| = k$. Our result is captured by the following theorem.

Theorem 15. *There is a randomized algorithm for the constrained file migration problem with competitive ratio $O(\log m \cdot \log k)$, where k vertices can store files and m is the total memory available.*

This theorem generalizes a result of [Bar96], who showed an algorithm with competitive ratio $O(\log m \cdot \log n)$ for arbitrary graphs on n nodes. Both results are based on the following theorem. (Recall that a 2-HST is an ultrametric (see Definition 2) such that the ratio between the label of a node to any of its children's label is at least 2.)

Theorem 16 ([Bar96]). *For any 2-HST, there is a randomized algorithm with competitive ratio $O(\log m)$ for constrained file migration with total memory m .*

By Theorem 10 there is a distribution \mathcal{D} over embeddings of G into ultrametrics with expected terminal distortion $O(\log k)$, but in fact every tree in that distribution is a 2-HST. Assume that in the optimal (offline) solution there are s_{uv} times a file residing on v was accessed by u , and t_{uv} files were migrated from v to u . Let $c_{uv} = s_{uv} + D \cdot t_{uv}$ be the total cost of file traffic from v to u . Note that as $m_v = 0$ for any $v \notin K$, then for any $u \in V$ we have $c_{uv} = 0$. Using the fact that the terminal distortion guarantee of \mathcal{D} applies to all of the relevant distances, we obtain that

$$\begin{aligned} \text{opt}_G &= \sum_{u \in V, v \in K} c_{uv} \cdot d_G(u, v) \\ &\geq \frac{1}{O(\log k)} \cdot \sum_{u \in V, v \in K} c_{uv} \cdot \mathbb{E}_{T \sim \mathcal{D}}[d_T(u, v)] \\ &= \frac{1}{O(\log k)} \cdot \mathbb{E}_{T \sim \mathcal{D}} \left[\sum_{u \in V, v \in K} c_{uv} \cdot d_T(u, v) \right]. \end{aligned} \tag{21}$$

Observe that for any tree $T \in \text{supp}(\mathcal{D})$ we could have served the request sequence in the same manner as the optimal algorithm, which would have the cost $\sum_{u \in V, v \in K} c_{uv} \cdot d_T(u, v)$. In particular, the optimal solution opt_T for the same requests with the input graph T cannot be larger than that, i.e.

$$\sum_{u \in V, v \in K} c_{uv} \cdot d_T(u, v) \geq \text{opt}_T. \tag{22}$$

Our algorithm will operate as follows: Pick a random tree according to the distribution \mathcal{D} , pick a random strategy S for transmitting files in T according to the distribution $\mathcal{S}(T)$ guaranteed to exist by Theorem 16, and serve the requests according to S . Denote by $\text{cost}_H(S)$ the cost of applying strategy S with distances taken in the graph H . For any possible $T \in \text{supp}(\mathcal{D})$ it holds that

$$\text{opt}_T \geq \frac{\mathbb{E}_{S \sim \mathcal{S}(T)}[\text{cost}_T(S)]}{O(\log m)} \geq \frac{\mathbb{E}_{S \sim \mathcal{S}(T)}[\text{cost}_G(S)]}{O(\log m)}, \tag{23}$$

where the last inequality holds since T dominates G (i.e. $d_T(u, v) \geq d_G(u, v)$ for all $u, v \in V$). Combining equations (21), (22) and (23) we get that

$$\text{opt}_G \geq \frac{\mathbb{E}_{T \sim \mathcal{D}} \mathbb{E}_{S \sim \mathcal{S}(T)}[\text{cost}_G(S)]}{O(\log k \log m)}.$$

Hence our randomized algorithm has $O(\log m \log k)$ competitive ratio, as promised.

12 Acknowledgements

We would like to thank Robert Krauthgamer, Yair Bartal and Manor Mendel for fruitful discussions, an anonymous referee for an idea leading to Theorem 2, and another additional anonymous referee for Remark 1.

References

- [ABC⁺05] Ittai Abraham, Yair Bartal, T-H. Hubert Chan, Kedar Dhamdhere Dhamdhere, Anupam Gupta, Jon Kleinberg, Ofer Neiman, and Aleksandrs Slivkins. Metric embeddings with relaxed guarantees. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '05, pages 83–100, Washington, DC, USA, 2005. IEEE Computer Society.
- [ABN08] Ittai Abraham, Yair Bartal, and Ofer Neiman. Nearly tight low stretch spanning trees. In *FOCS '08: Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 781–790, Washington, DC, USA, 2008. IEEE Computer Society.
- [ABN11] Ittai Abraham, Yair Bartal, and Ofer Neiman. Advances in metric embedding theory. *Advances in Mathematics*, 228(6):3026 – 3126, 2011.
- [ABN15] Ittai Abraham, Yair Bartal, and Ofer Neiman. Embedding metrics into ultrametrics and graphs into spanning trees with constant average distortion. *SIAM J. Comput.*, 44(1):160–192, 2015.
- [ABP92] B. Awerbuch, A. Baratz, and D. Peleg. Efficient broadcast and light-weight spanners. Technical Report CS92-22, The Weizmann Institute of Science, Rehovot, Israel., 1992.
- [ADD⁺93] I. Althöfer, G. Das, D. P. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993.
- [AF09] Reid Andersen and Uriel Feige. Interchanging distance and capacity in probabilistic mappings. Technical report, 2009.
- [AKPW95] Noga Alon, Richard M. Karp, David Peleg, and Douglas West. A graph-theoretic game and its application to the k -server problem. *SIAM J. Comput.*, 24(1):78–100, 1995.
- [ALN07] Sanjeev Arora, James R. Lee, and Assaf Naor. Fréchet embeddings of negative type metrics. *Discrete & Computational Geometry*, 38(4):726–739, 2007.
- [ALN08] S. Arora, J. R. Lee, and A. Naor. Euclidean distortion and the sparsest cut. *Journal of the American Mathematical Society* 21, 1:1–21, 2008.
- [AN12] Ittai Abraham and Ofer Neiman. Using petal-decompositions to build a low stretch spanning tree. In *STOC*, pages 395–406, 2012.
- [AR98] Yonatan Aumann and Yuval Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM Journal on Computing*, 27(1):291–301, 1998.
- [ARV09] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2), 2009.
- [Bar96] Yair Bartal. Probabilistic approximations of metric spaces and its algorithmic applications. In *FOCS*, pages 184–193, 1996.
- [Bar04] Yair Bartal. Graph decomposition lemmas and their role in metric embedding methods. In *ESA*, pages 89–97, 2004.
- [BFN16] Yair Bartal, Arnold Filtser, and Ofer Neiman. Constructing almost minimum spanning trees with constant average distortion. In *SODA*, 2016.
- [BFR95] Yair Bartal, Amos Fiat, and Yuval Rabani. Competitive algorithms for distributed data management. *J. Comput. Syst. Sci.*, 51(3):341–358, 1995.
- [BLMN05] Yair Bartal, Nathan Linial, Manor Mendel, and Assaf Naor. Some low distortion metric ramsey problems. *Discrete & Computational Geometry*, 33(1):27–41, 2005.
- [Bou85] J. Bourgain. On lipschitz embedding of finite metric spaces in hilbert space. *Israel Journal of Mathematics*, 52(1-2):46–52, 1985.

- [CE05] Don Coppersmith and Michael Elkin. Sparse source-wise and pair-wise distance preservers. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, pages 660–669, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.
- [CKR04] Gruia Călinescu, Howard J. Karloff, and Yuval Rabani. Approximation algorithms for the 0-extension problem. *SIAM Journal on Computing*, 34(2):358–372, 2004.
- [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.
- [EEST08] Michael Elkin, Yuval Emek, Daniel A. Spielman, and Shang-Hua Teng. Lower-stretch spanning trees. *SIAM Journal on Computing*, 38(2):608–628, 2008.
- [EFN15] Michael Elkin, Arnold Filtser, and Ofer Neiman. Prioritized metric structures and embedding. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 489–498, 2015.
- [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.
- [ES11] Michael Elkin and Shay Solomon. Steiner shallow-light trees are exponentially lighter than spanning ones. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 373–382, 2011.
- [FLM77] T. Figiel, J. Lindenstrauss, and V.D. Milman. The dimension of almost spherical sections of convex bodies. *Acta Mathematica*, 139(1):53–94, 1977.
- [FRT04] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci.*, 69(3):485–497, 2004.
- [GNR10] Anupam Gupta, Viswanath Nagarajan, and R. Ravi. An improved approximation algorithm for requirement cut. *Oper. Res. Lett.*, 38(4):322–325, 2010.
- [GP68] E. N. Gilbert and H. O. Pollak. Steiner minimal trees. *SIAM Journal on Applied Mathematics*, 16(1):1–29, Jan 1968.
- [Gün07] Oktay Günlük. A new min-cut max-flow ratio for multicommodity flows. *SIAM J. Discrete Math.*, 21(1):1–15, 2007.
- [Gup01a] 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.
- [Gup01b] Anupam Gupta. Steiner points in tree metrics don't (really) help. In *SODA*, pages 220–227, 2001.
- [HPIS13] Sariel Har-Peled, Piotr Indyk, and Anastasios Sidiropoulos. Euclidean spanners in high dimensions. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '13, pages 804–809. SIAM, 2013.
- [JL84] William Johnson and Joram Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. *Contemporary Mathematics*, 26:189206, 1984.

- [KKN14] Lior Kamma, Robert Krauthgamer, and Huy L. Nguyen. Cutting corners cheaply, or how to remove steiner points. In *SODA*, pages 1029–1040, 2014.
- [KLMN05] Robert Krauthgamer, James R. Lee, Manor Mendel, and Assaf Naor. Measured descent: a new embedding method for finite metrics. *Geometric and Functional Analysis*, 15(4):839–858, 2005.
- [KRY95] Samir Khuller, Balaji Raghavachari, and Neal E. Young. Balancing minimum spanning trees and shortest-path trees. *Algorithmica*, 14(4):305–321, 1995.
- [KSW09] Jon Kleinberg, Aleksandrs Slivkins, and Tom Wexler. Triangulation and embedding using small sets of beacons. *J. ACM*, 56(6):32:1–32:37, September 2009.
- [KV13] Telikepalli Kavitha and Nithin M. Varma. Small stretch pairwise spanners. In *Proceedings of the 40th International Conference on Automata, Languages, and Programming - Volume Part I, ICALP’13*, pages 601–612, Berlin, Heidelberg, 2013. Springer-Verlag.
- [LLR95] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.
- [LR99] Frank Thomson Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46(6):787–832, 1999.
- [Mat96] Jiri Matoušek. On the distortion required for embedding finite metric spaces into normed spaces. volume 93, pages 333–344, 1996.
- [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.
- [MN06] Manor Mendel and Assaf Naor. Ramsey partitions and proximity data structures. In *FOCS*, pages 109–118, 2006.
- [Moi09] Ankur Moitra. Approximation algorithms for multicommodity-type problems with guarantees independent of the graph size. In *FOCS*, pages 3–12, 2009.
- [R02] Harald Räcke. Minimizing congestion in general networks. In *Proceedings of the 43rd Symposium on Foundations of Computer Science, FOCS ’02*, pages 43–52, Washington, DC, USA, 2002. IEEE Computer Society.
- [Rac08] Harald Räcke. Optimal hierarchical decompositions for congestion minimization in networks. In *STOC*, pages 255–264, 2008.
- [RR98] Yuri Rabinovich and Ran Raz. Lower bounds on the distortion of embedding finite metric spaces in graphs. *Discrete & Computational Geometry*, 19(1):79–94, 1998.
- [RTZ05] Liam Roditty, Mikkel Thorup, and Uri Zwick. Deterministic constructions of approximate distance oracles and spanners. In *Proceedings of the 32Nd International Conference on Automata, Languages and Programming, ICALP’05*, pages 261–272, Berlin, Heidelberg, 2005. Springer-Verlag.