Steiner Shallow-Light Trees are Exponentially Better than Spanning Ones *

Michael Elkin[†] Shay Solomon[‡]

Technical Report CS-11-02

December 15, 2010

Abstract

The power of Steiner points was studied in a number of different settings in the context of metric embeddings. Perhaps most notably in the context of *probabilistic tree embeddings* Bartal and Fakcharoenphol et al. [8, 9, 21] used Steiner points to devise near-optimal constructions of such embeddings. However, Konjevod et al. [24] and Gupta [22] demonstrated that Steiner points do not help in this context. Specifically, they showed that any probabilistic tree embedding with distortion D that employs Steiner points can be converted into a probabilistic tree embedding of distortion $O(D)$ that uses only points of the original metric.

Steiner points were also studied in the context of *graph spanners* [3], in the context of *Euclidean spanners* [27, 19, 28], and in the context of *distance preservers* [10]. In all these contexts it is known [3, 27, 19, 28, 10] that Steiner points cannot be used to significantly improve inherent tradeoffs between the involved parameters. The situation is similar in the context of *low-light trees*. Specifically, it is known [17, 19] that essentially the same tradeoff between unweighted diameter and weight that applies to spanning low-light trees applies to Steiner low-light trees too.

These results may lead to a far-reaching conclusion that Steiner points do not help in metric embeddings in general. In this paper we show that this is not the case, and demonstrate that Steiner points do help dramatically in the context of *shallow-light trees*.

Shallow-light trees [6, 7, 23, 12, 13, 14] combine small weight with small distortion with respect to a designated root vertex rt (henceforth, *root-distortion*). Awerbuch et al. [7] and Khuller et al. [23] showed that for any positive real parameter $\epsilon > 0$, one can simultaneously achieve root-distortion $(1 + \epsilon)$ and weight $O(\frac{1}{\epsilon})$ times the weight of the minimum spanning tree $w(MST)$. Moreover, this tradeoff is tight up to constant factors. In this paper we show that by using Steiner points one can simultaneously achieve root-distortion $(1+\epsilon)$ and weight $O(\log\left(\frac{1}{\epsilon}\right)) \cdot w(MST)$. In particular, one can also construct a Steiner tree with weight $O(\log n)\cdot w(MST)$ that *preserves* all distances between rt and other vertices. Furthermore, we show that up to constant factors this tradeoff is tight. These results imply that there is an *exponential separation* between shallow-light trees that use Steiner points and shallow-light trees that do not use them.

Finally, on the way to these results we also address a number of open questions that were posed by Khuller et al. [23]. Specifically, we show that the lower bound on the tradeoff between root-distortion $(1 + \epsilon)$ and weight $\Omega(\frac{1}{\epsilon}) \cdot w(MST)$ of spanning shallow-light trees (1) applies even to 2-dimensional Euclidean metrics rather than to general metrics; (2) applies even if we replace (worst-case) root-distortion by average root-distortion; (3) applies even if the designated root vertex is selected at will.

*Department of Computer Science, Ben-Gurion University of the Negev, POB 653, Beer-Sheva 84105, Israel. E-mail: {elkinm,shayso}@cs.bgu.ac.il
Both authors are partially supported by the Lynn and William Frankel Center for Computer Sciences.

[†]This research has been supported by the BSF grant No. 2008430.

[‡]This research has been supported by the Clore Fellowship grant No. 81265410.

1 Introduction

1.1 Steiner Points in Spanners. Steiner points were studied in various settings in the area of metric spanning trees and spanners. Althofer et al. [3] studied Steiner points in the context of *graph spanners*. Given a possibly weighted undirected graph $G = (V, E, w)$, with $w : E \to \mathbb{R}^+$ being a positive weight function, and an integer parameter $t \geq 1$, a *subgraph t-spanner* of G is a subgraph $G' = (V, H, w)$ of G, i.e., $H \subseteq E$, that approximates all distances of G up to a multiplicative factor of t. The parameter t is called the *stretch* (or *distortion*) of the spanner. A graph $\mathcal{G} = (\mathcal{V}, \mathcal{H}, w')$ with $V \subseteq \mathcal{V}$ is called a *(metric) Steiner subgraph* of G if for every pair $u, v \in V$ of original vertices, the distance $d_{\mathcal{G}}(u, v)$ between u and v in $\mathcal{H}$ is greater or equal to the distance $d_G(u, v)$ between them in G, i.e., $d_{\mathcal{G}}(u, v) \geq d_G(u, v)$ [3]. A (metric) Steiner subgraph $\mathcal{G}$ of G is said to be a *(metric) Steiner t-spanner* of G if it also satisfies that $d_{\mathcal{G}}(u, v) \leq t \cdot d_G(u, v)$, for every pair of original vertices $u, v \in V$.

Althofer et al. [3] observed that for some specific graphs, Steiner spanners can be substantially sparser than subgraph spanners. In particular, they considered the complete n-vertex graph K_n with unit weights. Obviously every subgraph 1-spanner of K_n requires all the $\Omega(n^2)$ edges. On the other hand, the Steiner spanner that uses one additional vertex rt, and connects rt to each original vertex with an edge of weight $\frac{1}{2}$, is a Steiner 1-spanner for K_n with just n edges.

However, Althofer et al. [3] showed that this phenomenon does not hold in general. Specifically, there are n-vertex graphs for which any subgraph $(2t-1)$-spanner requires $\Omega(n^{1+\frac{4}{3t}})$ edges, and any Steiner $(2t-1)$-spanner requires $\Omega(\frac{n^{1+\frac{4}{3t}}}{\log n})$ edges [26, 3]. Moreover, in fact, it is known [3] that a lower bound of $T(n, t)$ on the number of edges required for subgraph t-spanners of general n-vertex graphs implies an analogous lower bound of $\Omega(\frac{T(n,t)}{\log n})$ for Steiner t-spanners. Therefore, while Steiner points can help to produce much sparser spanners for certain specific graphs, they cannot help to significantly improve the bounds on the number of edges requires for spanners of general graphs.

Bollobás et al. [10] studied the impact of Steiner points in the context of *distance preservers*. The notion of distance preserver is closely related to the notion of graph spanner. (See [10] for details.) It was shown in [10] that similarly to the situation with graph spanners, Steiner points cannot help improving the tradeoff between the involved parameters of distance preservers by more than a logarithmic factor.

Rao and Smith [27] analyzed the impact of Steiner points in the context of *Euclidean spanners*, and arrived to similar conclusions. Given a set P of n points in $\mathbb{R}^d$, $d \geq 2$, a graph $G' = (P, H, w), H \subseteq \binom{P}{2}$, with $w(u, v) = \|u - v\|_2 = \|u - v\|$, for every $u, v \in P$, is said to be a *subgraph t-spanner* for P (for $t \geq 1$), if for every $u, v \in P$, the distance $d_{G'}(u, v)$ between u and v in G' is no greater than t times the Euclidean distance $\|u - v\|$ between u and v, i.e., $d_{G'}(u, v) \leq t \cdot \|u - v\|$. A *Euclidean Steiner t-spanner* for P is a graph $\mathcal{G} = (\mathcal{P}, H, w)$, with $P \subseteq \mathcal{P} \subseteq \mathbb{R}^d$ and $w(x, y) = \|x - y\|$ for every $x, y \in \mathcal{P}$, that satisfies $d_{\mathcal{G}}(u, v) \leq t \cdot \|u - v\|$, for every $u, v \in P$. (Note that the Euclidean Steiner spanner notion of [27] is more restrictive than the metric Steiner spanner notion of [3].)

Similarly to the case of K_n for general graph spanners, Rao and Smith (see Section 6 of [27]) devised examples of point configurations for which Euclidean Steiner points can drastically improve the weight of spanners, in comparison to subgraph spanners. However, they also demonstrated that "Steiner spanners cannot be too short", i.e., that there are configurations P of n points in $\mathbb{R}^d$ for which any Steiner $(1+\epsilon)$-spanner has weight $\left(\frac{1}{\epsilon}\right)^{\Omega(d)} \cdot w(MST(P))$, where $MST(P)$ stands for the minimum spanning tree of P. (See [27], Theorem 56.) On the other hand, constructions of subgraph $(1+\epsilon)$-spanners of weight $O\left(\left(\frac{1}{\epsilon}\right)^{2d}\right) \cdot w(MST(P))$, for point sets $P \subseteq \mathbb{R}^d$, are known [4, 25].

Other lower bounds of similar flavor were shown for the tradeoffs between weight and unweighted diameter or Euclidean spanners [19], and between number of edges and unweighted diameter of Euclidean spanners [28]. The former lower bound (of [19]) refers to the stronger notion of metric Steiner spanners, whereas the latter bound (of [28]) refers to the notion of Euclidean Steiner ones.

1.2 Steiner Points in Trees. Steiner points were also studied in the context of metric spanning trees. Alon et al. [2] showed that for every graph $G = (V, E, w)$ there exists a probability distribution $\mathcal{D}$ of spanning trees of G such that for every edge $e = (u, v) \in E$, $\mathbb{E}_{T \in \mathcal{D}}\left[\frac{d_T(u,v)}{w(e)}\right] = 2^{O(\sqrt{\log n \log\log n})}$. Such a distribution is called *probabilistic tree embedding* [8], and the value $\max_{e=(u,v)\in E} \mathbb{E}\left[\frac{d_T(u,v)}{w(e)}\right]$ is called the *stretch* of the embedding. Bartal and Fakcharoenphol et al. [8, 9, 21] showed that using metric Steiner points one can drastically improve the bound of [2], and devise probabilistic tree embeddings with stretch $O(\log n)$. The bound $O(\log n)$ is also known to be optimal up to constant factors [2, 8]. However, soon afterwards Konjevod et al. [24] and Gupta [22] demonstrated that the same bounds (up to constant factors) as those of Bartal and Fakcharoenphol et al. [8, 9, 21] can be obtained without Steiner points, i.e., by using only *spanning trees* of the metric induced by G. (Such spanning trees are also said to use *Steiner edges*, rather than Steiner points.) Moreover, more recent studies [18, 1] showed that nearly the same bounds can be obtained by using spanning trees of the original graph G. Therefore, it turns out that neither Steiner points nor Steiner edges can really help to improve probabilistic tree embeddings.

A similar situation is known in the context of *low-light trees*, which combine small weight with small unweighted diameter [17]. In [19] the authors of the current paper showed that Steiner points do not help in this context either, i.e., that any Steiner tree T with a given unweighted diameter can be converted into a spanning tree with the the same (up to constant factors) diameter and weight as those of T.

To summarize, Steiner points were studied in many different settings in the context of trees, graph spanners, Euclidean spanners and distance preservers [3, 27, 24, 22, 10, 19, 28]. In all these settings they are known either not to help much or not to help at all in improving inherent tradeoffs between the involved parameters.

1.3 Steiner points in Shallow-Light Trees. In this paper we study the impact of metric Steiner points on *shallow-light trees.* In a sharp contrast to the situation in spanners, distance preservers, probabilistic tree embeddings and low-light trees (see Sections 1.1 and 1.2), we demonstrate that using metric Steiner points, one can *exponentially* improve the tradeoff between the involved parameters of shallow-light trees.

Shallow-light trees (henceforth, SLTs) were introduced by Awerbuch et al. [6]. Given a graph $G = (V, E, w)$, a designated root vertex $rt \in V$, and parameters $\alpha \geq 1, \beta \geq 1$, a spanning tree T of G is said to be an (α, β)*-SLT* of G if (1) for every vertex $v \in V$, $d_T(rt, v) \leq \alpha \cdot d_G(rt, v)$, and (2) $w(T) \leq \beta \cdot w(MST(G))$, where $MST(G)$ stands for the minimum spanning tree of G. A tree T that satisfies only the first requirement will be referred to as an α*-shortest paths tree*, or shortly, α-SPT, of G. We will say that such a tree T has *root-stretch* at most α. A 1-SPT will also be referred to as an SPT. These definitions can also be extended to metrics rather than to graphs, and to Steiner trees rather than to spanning ones. Awerbuch et al. [7] and Khuller et al. [23] showed that for every graph and every $\epsilon > 0$, a $(1 + \epsilon, O(\frac{1}{\epsilon}))$-SLT exists. Moreover, Khuller et al. [23] showed that there exist graphs G for which any $(1 + \epsilon)$-SPT T has weight[1] $w(T) = \Omega(\frac{1}{\epsilon}) \cdot w(MST(G))$.

In this paper we show that for every graph G there exists a metric Steiner $(1+\epsilon, O(\log \frac{1}{\epsilon}))$-SLT. In other words, we present a construction of Steiner $(1 + \epsilon)$-SPTs whose lightness is *exponentially smaller* than the lightness of the best-known (and, in fact, optimal) spanning $(1+\epsilon)$-SPTs. In view of the lower bound from [23] that shows that spanning $(1 + \epsilon)$-SPTs must have lightness[2] $\Omega(\frac{1}{\epsilon})$, our results demonstrate an *exponential separation* between spanning SLTs and metric Steiner SLTs! We also complement our

[1]For convenience, we will henceforth use a normalized notion of weight, called *lightness*, defined by $\Psi(T) = \frac{w(T)}{w(MST(G))}$.

[2]Strictly speaking, the appropriate notion of lightness for a Steiner tree T should be $\Psi'(T) = \frac{w(T)}{w(SMT(G))}$, where $SMT(G)$ stands for the *minimum Steiner tree* of G. However, since $\frac{1}{2} \cdot w(MST(G)) \leq w(SMT(G)) \leq w(MST(G))$, it follows that $\Psi(T) \leq \Psi'(T) \leq 2 \cdot \Psi(T)$. Hence, up to a factor of 2 these two notions of lightness are equivalent. In this discussion we do not distinguish between them.

construction with matching lower bounds, and show that there are graphs (and, in fact, much stronger than that: we show that there are n-point 2-dimensional Euclidean metrics) for which any metric Steiner $(1+\epsilon)$-SPT must have lightness $\Omega(\log \frac{1}{\epsilon})$.

One particularly interesting point on our tradeoff curve is $\epsilon = 0$. Observe that there are (simple) n-vertex graphs for which any spanning shortest paths tree from a designated root vertex $rt \in V$ has lightness $\Omega(n)$. Consider, for example, a path P_{n-1} of $n-1$ vertices $(v_1, v_2, \ldots, v_{n-1})$, with each v_i connected to v_{i+1} via a unit-weight edge, for $i \in \{1, 2, \ldots, n-2\}$. Add also the root vertex rt, and connect it to all vertices of P_{n-1} via edges of weight $\frac{n-1}{2}$. Call the resulting graph G. Obviously, any MST of G uses a single edge that connects rt to P_{n-1}, and has weight $O(n)$. On the other hand, the SPT of G rooted at rt contains all the $n-1$ edes (rt, v_i), $i \in \{1, 2, \ldots, n-1\}$, and has weight $\Omega(n^2)$, i.e., lightness $\Omega(n)$. (The same bound of $\Omega(n)$ on the lightness of a spanning SPT can also be obtained for simple Euclidean 2-dimensional point sets, such as the following one. Spread $n-1$ points uniformly on the boundary of a unit circle, and add the root point rt at the center of the circle. Denote this point set by $\tilde{C}_n$.) On the other hand, we show that by using metric Steiner points one can construct a shortest paths tree with lightness $O(\log n)$ for any n-vertex graph! Moreover, as was mentioned above, up to constant factors this logarithmic upper bound is tight, even for Euclidean 2-dimensional metrics. In addition, our constructions of Steiner $(1+\epsilon, O(\log \frac{1}{\epsilon}))$-SLTs and of Steiner SPTs with lightness $O(\log n)$ can both be implemented within $O(n^2)$ time.

We remark, however, that the Steiner points that we use are *metric* ones, even if the original metric is a Euclidean one. (See Section 1.1 for the distinction between metric and Euclidean Steiner points.) On the other hand, it is easy to see that any Euclidean Steiner SPT for $\tilde{C}_n$ rooted at rt has lightness $\Omega(n)$. More generally, we show that for any $\epsilon > 0$, any Euclidean Steiner $(1+\epsilon)$-SPT for $\tilde{C}_n$ rooted at rt has lightness $\Omega(\sqrt{\frac{1}{\epsilon}})$. Therefore, our results imply also an *exponential separation* between the lightness of metric Steiner SLTs (which is $\Theta(\log \frac{1}{\epsilon})$) and the lightness of Euclidean Steiner SLTs (which is $\Omega(\sqrt{\frac{1}{\epsilon}})$).

1.4 Additional Results.

Euclidean Shallow-Light Trees. We also proved a number of results concerning *spanning* SLTs (rather than Steiner SLTs). In the Conclusions section of their paper [23] Khuller et al. asked four questions. In this paper we address three of them.

First, Khuller et al. [23] showed that there are graphs for which any $(1+\epsilon)$-SPT has lightness $1 + \frac{2}{\epsilon}$. They asked whether the same lower bound applies to Euclidean metrics. Up to constant factors we answer this question in the affirmative, and show that there exist configurations of n points in the Euclidean plane for which any $(1+\epsilon)$-SPT rooted at a designated root-vertex rt has lightness $\Omega(\frac{1}{\epsilon})$. We also show that there exist configurations for which any α-SPT has lightness at least $1 + \Omega(\frac{1}{\alpha})$, for any $\alpha > 1$.

Second, Khuller et al. [23] suggested to relax the notion of root-stretch, and replace it with the notion of *average root-stretch*. A spanning tree T for a graph G is said to have average root-stretch at most $1+\epsilon$ with respect to a designated root vertex rt if $\frac{\sum_{v \in V \setminus \{rt\}} d_T(rt,v)}{\sum_{v \in V \setminus \{rt\}} d_G(rt,v)} \leq 1 + \epsilon$. (Observe that the average root-stretch of a tree is bounded above by its root-stretch.) They asked whether their lower bound applies when replacing root-stretch by average root-stretch. We show that there exist configurations of n points in the Euclidean plane for which any spanning tree with average root-stretch at most $1+\epsilon$ (respectively, α) with respect to a certain designated root vertex rt has lightness $\Omega(\frac{1}{\epsilon})$ (resp., at least $1+\Omega(\frac{1}{\alpha})$).

Third, Khuller et al. [23] asked whether their lower bound applies if one is allowed to select the root vertex at will. We show that there are configurations of n points in the Euclidean plane for which any spanning tree with average root-stretch at most $(1+\epsilon)$ (respectively, α) with respect to *any vertex* rt has lightness $\Omega(\frac{1}{\epsilon})$ (resp., at least $1+\Omega(\frac{1}{\alpha})$).

In other words, up to constant factors we settle three out of the four open questions posed by Khuller et al. [23]. The fourth question concerns approximation algorithms for SLTs; this question is left open.

We also show that Steiner *edges* do not help in the context of SLTs. In other words, we show that for any graph $G = (V, E, w)$ and any designated root vertex $rt \in V$, every spanning tree T of the metric induced by G (i.e., T may include Steiner edges but not Steiner vertices) can be converted into a spanning tree T' of G whose weight and root-stretch are bounded by the weight and root-stretch of T, respectively. (Hence there exist metrics (and not just graphs) for which any $(1+\epsilon)$-SPT has lightness $1 + \frac{2}{\epsilon}$.)

Euclidean spanners. Another result that we prove in this paper concerns the tradeoff between the stretch and the maximum degree of Euclidean spanners. Das and Narasimhan [16] (see also [5, 4, 11]) showed that for any configuration of n points in the d-dimensional Euclidean space $\mathbb{R}^d$, and any $\epsilon > 0$, there exists a $(1+\epsilon)$-spanner with maximum degree $O(\epsilon^{-d+1})$. We show that this bound on the degree is tight up to constant factors, i.e., that there exist configurations of n points in $\mathbb{R}^d$ for which any $(1+\epsilon)$-spanner has maximum degree $\Omega(\epsilon^{-d+1})$.

1.5 Our Techniques. The most technically involved result in this paper is that for any n-vertex graph G there exists a metric Steiner SPT with lightness $\mathrm{O}(\log n)$, and that this bound is tight up to constant factors. Other results are deduced (in a non-trivial) way either from this result, or from the methods developed on our way to this result.

Our method in deriving this result is to identify a *core* example, i.e., a simple example that manages to encapsulate the inherent complexity of the problem. Our core example is $\tilde{C}_n$ (see Section 1.3). We start with devising a construction of Steiner SPTs with respect to the root vertex rt that lies at the center of the circle with lightness $O(\log n)$. Then we proceed to extending it to all possible graphs. This extension is not at all simple, but nevertheless, the construction for the core example provides us with certain guidelines as to how to proceed to the general case.

Note also that it is sufficient to prove the lower bound only for the core example. The fact that already this simple Euclidean metric acheives the worst (up to constant factors) possible tradeoff is, in our opinion, interesting on its own right. We remark that a similar in spirit approach was taken in [17] for the analysis of low-light trees. There a different core example was identified. Specificially, it was a path P_n of n points lying on the same line, with unit distance between consecutive points. It was shown in [17] that essentially tight lower bounds can be proved already for this elementary metric, and moreover, that constructions for P_n extend naturally to general metrics. However, the techniques that are used to construct spanning low-light trees in [17] are fundamentally different from the techniques that we use here to construct metric Steiner shallow-light trees. Also, the proofs of lower bounds for the different core examples here and in [17] have nothing in common. In particular, the proof of lower bound in [17] relies on analyzing a linear program from [20] for the minimum linear arrangement problem. On the other hand, here we empoy direct combinatorial arguments, and methods from the area of low-distortion embeddings for working with Steiner points. (See, e.g., [24, 22, 8].)

1.6 Structure of the Paper. In Section 2 we present our construction of metric Steiner SLTs. Therein we start (Section 2.1) with the construction of metric Steiner SPTs with logarithmic lightness, and proceed (Section 2.2) with its generalization for metric Steiner SLTs with root-stretch at most $1+\epsilon$ and lightness $O(\log \frac{1}{\epsilon})$. The matching lower bound on the lightness of metric Steiner SLTs is given in Section 3. In Sections 4 and 5 we provide our lower bounds for Euclidean spanning SLTs and spanners, respectively. Finally, the argument that shows that Steiner edges do not help appears in Section 6.

1.7 Preliminaries. Let $T = (T, rt)$ be either a spanning tree or a Steiner tree of a graph $G = (V, E, w)$ rooted at some designated point rt. The *stretch* between a pair u, v of vertices in T is defined as $Str_T(u, v) = \frac{d_T(u,v)}{d_G(u,v)}$. The *root-stretch* and *average root-stretch* of (T, rt) are defined as $rtStr(T, rt) = \max\{Str_T(rt, v) \mid v \in V \setminus \{rt\}\}$ and $AvgStr(T, rt) = \frac{\sum_{v \in V \setminus \{rt\}} Str_T(rt,v)}{|V|-1}$, respectively. We remark that this definition of average root-stretch is slightly different than the one used by Khuller et al. [23]. (See Section 1.4.) Nevertheless, all bounds on average root-stretch presented in this paper apply with respect to both definitions. For a pair of non-negative integers k, n, $k \leq n$, we denote the sets $\{k, k+1, \ldots, n\}$ and $\{1, 2, \ldots, n\}$ by $[k, n]$ and $[n]$, respectively.

2 Upper Bounds for Steiner SLTs

This section is devoted to upper bound constructions for Steiner SLTs.

2.1 Steiner SPTs with Logarithmic Lightness

In this section we devise a construction of Steiner SPTs for general metrics with logarithmic lightness. We harness this construction in Section 2.2 to produce a construction of Steiner SLTs.

Let $M = (V, dist)$ be an n-point metric, let rt be a designated (root) point in V, and let $M' = (V \setminus \{rt\}, dist)$ be the $(n-1)$-point metric induced by the point set of $V \setminus \{rt\}$.

Consider a Hamiltonian path H' of M'. In what follows we construct a binary Steiner tree $T' = T'(H')$ for M' rooted at a Steiner point rt' of weight $O(\log n) \cdot w(H')$. The tree T' will also satisfy the following property. For any vertex x in T', there exists a number $\rho(x) \geq 0$, such that for any point v in $V \setminus \{rt\}$ that belongs to the subtree T'_x of T' rooted at x,

$$dist(rt, v) - d_{T'}(x, v) \;=\; \rho(x). \tag{1}$$

We show (Corollary 2.8) that the rooted tree (T, rt), $T = T(H')$, obtained from T' by adding to it an edge (rt, rt') of weight $\rho(rt')$, is a Steiner SPT for M with weight at most $O(\log n) \cdot w(H') + \rho(rt') \leq O(\log n) \cdot w(H') + w(MST(M))$. (See Figure 1 for an illustration.) In particular, if we take H' to be

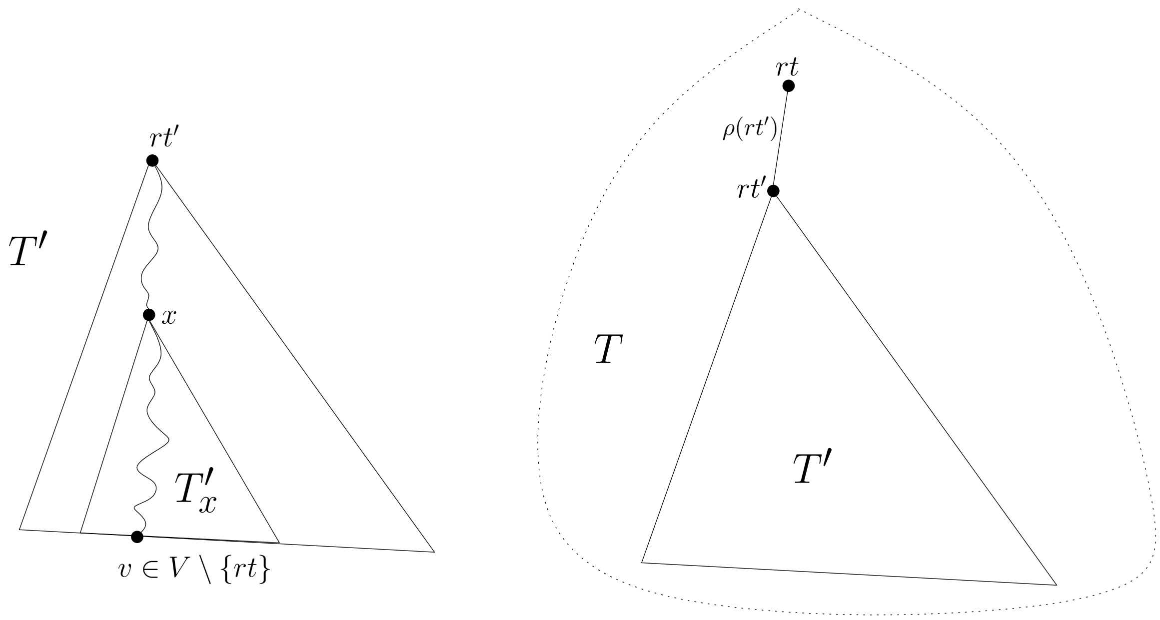

Figure 1: The trees T' and T.

a Hamiltonian path for M' with weight $O(w(MST(M'))) = O(w(MST(M)))$ (e.g., if H' is a TSP for M'), then the lightness of T is bounded by $O(\log n)$.

Denote the points of the Hamiltonian path H', from left to right, by $p_1, p_2, \ldots, p_{n-1}$. To construct T', we start by building a skeleton of a full balanced binary tree rooted at rt' with N leaves, where N is the smallest integer power of 2 greater or equal to $n-1$. (Note that $N < 2n-2$ and $\log N = \lceil \log(n-1) \rceil$.) All inner vertices of T' are Steiner points. Denote the N leaves of this tree, from left to right, by $\ell_1, \ell_2, \ldots, \ell_N$. For each index $i \in [n-2]$, the leaf ℓ_i corresponds to the point p_i, and for each index $i \in [n-1, N]$, the leaf ℓ_i corresponds to the point p_{n-1}, i.e., $p_i = p_{n-1}$ for each $i \in [n-1, N]$. (Thus, the last point p_{n-1} of H' is duplicated $N - n + 1$ times.)

Before we describe the weight assignment of edges in T', we need to introduce some notation.
For a vertex x in T', denote its left child by $L(x)$, its right child by $R(x)$, and the set of leaves in the subtree T'_x of T' rooted at x by $Leaves(x)$. (For a leaf x, we have $L(x) = R(x) = NULL$ and $Leaves(x) = \{x\}$.) Observe that for an inner vertex x, $Leaves(x) = Leaves(L(x)) \cup Leaves(R(x))$ and $Leaves(rt') = \{p_1, p_2, \ldots, p_N\} = \{p_1, p_2, \ldots, p_{n-1}\} = V \setminus \{rt\}$. The weight assignment of edges in the tree is computed recursively bottom-up, so that the weights $w_L(x)$ and $w_R(x)$ of the two edges $(x, L(x))$ and $(x, R(x))$ connecting an inner vertex x with its two children are computed only after all other edge weights in the subtree T'_x have been computed. We hold for each vertex x in T' three variables Δ_x, δ_x and $\rho(x)$, and use them to compute the weights $w_L(x)$ and $w_R(x)$ in the following way.
If x is a leaf, we set $\delta_x = \Delta_x = 0$ and $\rho(x) = dist(rt, x)$, and define $w_L(x) = w_R(x) = 0$.
For an inner vertex x, we set $\delta_x = \rho(L(x)) - \rho(R(x))$. Also, consider a pair of leaves $x_L \in Leaves(L(x))$ and $x_R \in Leaves(R(x))$. Let $\Delta(x_L, x_R) = dist(x_L, x_R) - (d_{T'}(L(x), x_L) + d_{T'}(R(x), x_R))$. Recall that the distance between x_L and x_R in the tree T' needs to be at least $dist(x_L, x_R)$. Also, the distance between these two vertices in T' is realized by the path

$$P'(x_L, x_R) = P'(x_L, L(x)) \circ (L(x), x) \circ (x, R(x)) \circ P'(R(x), x_R),$$

where $\circ$ stands for the concatenation, and $P'(x_L, L(x))$ (respectively, $P'(R(x), x_R)$) stands for the path connecting x_L with $L(x)$ (resp., $R(x)$ with x_R) in the tree T'. (See Figure 2 for an illustration.) In other

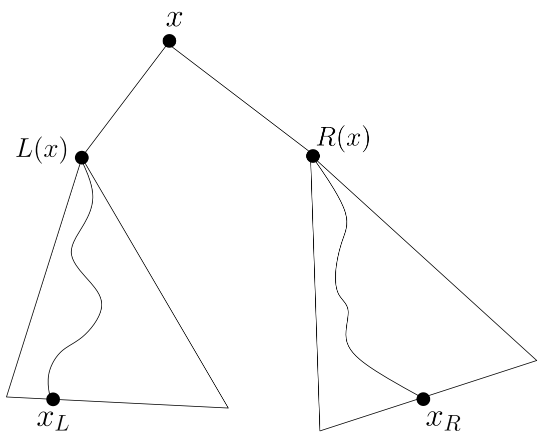

Figure 2: The path $P'(x_L, x_R)$ that connects x_L with x_R in T'.

words, it must hold that

$$dist(x_L, x_R) \;\le\; d_{T'}(x_L, L(x)) + w_L(x) + w_R(x) + d_{T'}(R(x), x_R),$$

and so

$$w_L(x) + w_R(x) \;\ge\; dist(x_L, x_R) - (d_{T'}(x_L, L(x)) + d_{T'}(R(x), x_R)) \;=\; \Delta(x_L, x_R).$$

Therefore, $\Delta(x_L, x_R)$ is a lower bound on the sum of the weights that we need to assign to the edges $(L(x), x)$ and $(x, R(x))$. We will call $\Delta(x_L, x_R)$ the *distance surplus* of the pair (x_L, x_R). Finally, the *distance surplus* of the vertex x is defined as the maximum distance surplus over all pairs (x_L, x_R) with $x_L \in Leaves(L(x))$ and $x_R \in Leaves(R(x))$, i.e.,

$$\Delta_x \;=\; \max\{\Delta(x_L, x_R) \mid x_L \in Leaves(L(x)), x_R \in Leaves(R(x))\}. \tag{2}$$

Note that the choice of the weights $w_L(x)$ and $w_R(x)$ for the edges $(L(x), x)$ and $(x, R(x))$, respectively, needs to satisfy $w_L(x) + w_R(x) \geq \Delta_x$. This inequality motivates the definition (2).

Given the values δ_x and Δ_x determined as above, we set the weights $w_L(x)$ and $w_R(x)$ as follows. If $|\delta_x| \leq \Delta_x$, we set $w_L(x) = \frac{\Delta_x+\delta_x}{2}$, $w_R(x) = \frac{\Delta_x-\delta_x}{2}$. Otherwise, we set $w_L(x) = \max\{\delta_x, 0\}$, $w_R(x) = \max\{-\delta_x, 0\}$. (In the latter case, either $w_L(x)$ or $w_R(x)$ is equal to zero, and the other parameter is equal to $|\delta_x|$.) Finally, having computed the weight assignment for the entire subtree T'_x, we pick an arbitrary vertex v in $Leaves(x)$, and set

$$\rho(x) = dist(rt, v) - d_{T'}(x, v).$$

We start the analysis of our construction with the following observation.

Observation 2.1 *For any vertex x in T', (1) $w_L(x), w_R(x) \geq 0$, (2) $w_L(x) + w_R(x) = \max\{\Delta_x, |\delta_x|\}$, and (3) $w_L(x) - w_R(x) = \delta_x$, or equivalently, $\rho(L(x)) - w_L(x) = \rho(R(x)) - w_R(x)$.*

The third statement of Observation 2.1 demonstrates the intuitive meaning of the variable $\delta(x)$. It is the difference between $w_L(x)$ and $w_R(x)$.

Next, we show that if the numbers $\rho(x)$ are set as was described above, then they satisfy Equation (1).

Lemma 2.2 *For any vertex x in T' and any vertex v in $Leaves(x)$, $dist(rt, v) - d_{T'}(x, v) = \rho(x)$.*

Proof: The proof is by induction on the depth $h = h(T'_x)$ of the subtree T'_x. The basis $h = 0$ is trivial.
Induction Step: We assume that the statement holds for the two children $L(x)$ and $R(x)$ of x, and prove it for x. Consider an arbitrary pair v, w of vertices in $Leaves(x)$.
Next, we show that $dist(rt, v) - d_{T'}(x, v) = dist(rt, w) - d_{T'}(x, w)$, which suffices. (Indeed, $\rho(x)$ was set as $dist(rt, u) - d_{T'}(x, u)$, for an arbitrary leaf $u \in Leaves(x)$. Hence, it suffices to show that for any leaf $u \in Leaves(x)$, the expression $dist(rt, u) - d_{T'}(x, u)$ is equal to the same value.)
If both v and w belong to $T'_{L(x)}$ or if they both belong to $T'_{R(x)}$, then the result follows easily from the induction hypothesis. Specifically, if $v, w \in T'_{L(x)}$ then $dist(rt, v) - d_{T'}(L(x), v) = dist(rt, w) - d_{T'}(L(x), w)$. However, $dist(rt, v) - d_{T'}(x, v) = (dist(rt, v) - d_{T'}(L(x), v)) - w_L(x)$, and similarly, $dist(rt, w) - d_{T'}(x, w) = (dist(rt, w) - d_{T'}(L(x), w)) - w_L(x)$. Hence $dist(rt, v) - d_{T'}(x, v) = dist(rt, w) - d_{T'}(x, w)$, as required. (See Figure 3 for an illustration.) The case when $v, w \in T'_{R(x)}$ is analogous.
We may henceforth suppose without loss of generality that $v \in Leaves(L(x))$ and $w \in Leaves(R(x))$. By construction, $d_{T'}(x, v) = d_{T'}(L(x), v) + w_L(x)$ and $d_{T'}(x, w) = d_{T'}(R(x), w) + w_R(x)$. By the third statement of Observation 2.1 and the induction hypothesis,

$$\begin{aligned} dist(rt, v) - d_{T'}(x, v) &= dist(rt, v) - d_{T'}(L(x), v) - w_L(x) \;=\; \rho(L(x)) - w_L(x) \\ &= \rho(R(x)) - w_R(x) = dist(rt, w) - d_{T'}(R(x), w) - w_R(x) \\ &= dist(rt, w) - d_{T'}(x, w). \qquad \blacksquare \end{aligned}$$

Lemma 2.2 shows that for an inner vertex x and a leaf $v \in Leaves(x)$, $dist(rt, v) = d_{T'}(x, v) + \rho(x)$. The next lemma shows that $\rho(x) \geq 0$. Hence $dist(rt, v) \geq d_{T'}(x, v)$. Eventually we will need to guarantee $dist(rt, v) = d_T(rt, v)$, where $T = T' \cup \{(rt, rt')\}$. Intuitively, the meaning of the value $\rho(x)$ is that if the vertex x were the root of T', i.e., if $rt' = x$, then the edge $(rt, x) = (rt, rt')$ in T would need to be of weight $\rho(x) = \rho(rt')$. (See Figure 1 for an illustration.)

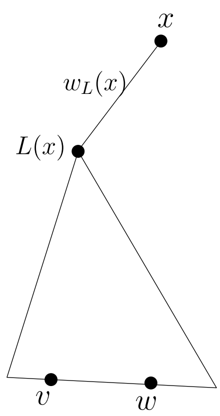

Figure 3: The case $v, w \in Leaves(L(x))$.

Lemma 2.3 *For any vertex x in T', $\rho(x) \geq 0$.*

Proof: The proof is by induction on the depth $h = h(T'_x)$ of the subtree T'_x. The basis $h = 0$ is trivial.
Induction Step: We assume that the statement holds for the two children $L(x)$ and $R(x)$ of x, and prove it for x. Suppose without loss of generality that $w_L(x) \leq w_R(x)$. (The complementary case $w_L(x) > w_R(x)$ is symmetric.) By the third statement of Observation 2.1, $\delta_x \leq 0$, and so $|\delta_x| = -\delta_x$.

Suppose first that $|\delta_x| \leq \Delta_x$. By Observation 2.1, $w_L(x) + w_R(x) = \Delta_x$ and $w_L(x) - w_R(x) = \delta_x$. Hence $2 \cdot w_L(x) = \Delta_x + \delta_x$. Let $x_L \in Leaves(L(x))$ and $x_R \in Leaves(R(x))$ be two vertices for which $\Delta_x = \Delta(x_L, x_R) = dist(x_L, x_R) - d_{T'}(L(x), x_L) - d_{T'}(R(x), x_R)$. By Lemma 2.2, $\rho(L(x)) = dist(rt, x_L) - d_{T'}(L(x), x_L)$ and $\rho(R(x)) = dist(rt, x_R) - d_{T'}(R(x), x_R)$. Altogether

$$\begin{aligned} 2 \cdot w_L(x) &= \Delta_x + \delta_x \;=\; \Delta_x + \rho(L(x)) - \rho(R(x)) \\ &= dist(x_L, x_R) + dist(rt, x_L) - dist(rt, x_R) - 2 \cdot d_{T'}(L(x), x_L) \\ &\leq 2 \cdot dist(rt, x_L) - 2 \cdot d_{T'}(L(x), x_L). \end{aligned}$$

(The last inequality holds by the triangle inequality.) Hence $w_L(x) \leq dist(rt, x_L) - d_{T'}(L(x), x_L)$.

Otherwise, $|\delta_x| > \Delta_x$. In this case $w_L(x) = 0$. By Lemma 2.2 and the induction hypothesis, $dist(rt, x_L) - d_{T'}(L(x), x_L) = \rho(L(x)) \geq 0$, and so $w_L(x) \leq dist(rt, x_L) - d_{T'}(L(x), x_L)$.

We have shown that in both cases, $w_L(x) \leq dist(rt, x_L) - d_{T'}(L(x), x_L)$. Also, by Lemma 2.2, $\rho(x) = dist(rt, x_L) - d_{T'}(x, x_L)$. It follows that

$$\rho(x) \;=\; dist(rt, x_L) - d_{T'}(x, x_L) \;=\; dist(rt, x_L) - d_{T'}(L(x), x_L) - w_L(x) \;\geq\; 0. \qquad \blacksquare$$

The following lemma shows that the tree T' dominates the metric M'.

Lemma 2.4 *For any vertex x in T' and any pair v, w of vertices in $Leaves(x)$, $d_{T'}(v, w) \geq dist(v, w)$.*

Proof: The proof is by induction on the depth $h = h(T'_x)$ of the subtree T'_x. The basis $h = 0$ holds vacuously.
Induction Step: We assume that the statement holds for the two children $L(x)$ and $R(x)$ of x, and prove it for x.
If both v and w belong to $T'_{L(x)}$ or if they both belong to $T'_{R(x)}$, then the result follows immediately from the induction hypothesis.
We may henceforth suppose without loss of generality that $v \in Leaves(L(x))$ and $w \in Leaves(R(x))$. By

construction, we have $d_{T'}(v,w) = d_{T'}(x,v) + d_{T'}(x,w)$, $d_{T'}(x,v) = d_{T'}(L(x),v) + w_L(x)$ and $d_{T'}(x,w) = d_{T'}(R(x),w) + w_R(x)$. By definition, $\Delta_x \geq \Delta(v,w) = dist(v,w) - (d_{T'}(L(x),v) + d_{T'}(R(x),w))$. By the second statement of Observation 2.1, $w_L(x) + w_R(x) \geq \Delta_x$. Altogether,

$$\begin{aligned} d_{T'}(v,w) &= d_{T'}(L(x),v) + w_L(x) + d_{T'}(R(x),w) + w_R(x) \\ &\geq d_{T'}(L(x),v) + d_{T'}(R(x),w) + \Delta_x \ \geq \ dist(v,w). \quad \blacksquare \end{aligned}$$

Next, we analyze the weight of the tree T'.

For a vertex x in T', let $f(x)$ and $l(x)$, $f(x) \leq l(x)$, be the indices in $[N]$ for which $Leaves(x) = \{p_{f(x)}, p_{f(x)+1}, \ldots, p_{l(x)}\}$. For a pair i,j of indices in $[N]$, $i \leq j$, let $Wt(i,j) = \sum_{k=i}^{j-1} dist(p_k, p_{k+1})$ denote the sum of all edge weights along the subpath $(p_i, p_{i+1}, \ldots, p_j)$ of the path $\tilde{H} = (p_1, p_2, \ldots, p_N)$. Since all points $p_n, p_{n+1}, \ldots, p_N$ are copies of the same point p_{n-1}, we have $Wt(1,N) = Wt(1,n-1)$. Hence,

$$w(\tilde{H}) \ = \ Wt(1,N) \ = \ Wt(1,n-1) \ = \ w(H').$$

We use the following claim to prove Lemma 2.6, which is, in turn, used to obtain an upper bound on the weight of the tree T'.

Claim 2.5 *For any vertex x in T', there exist indices i and j in $[f(x), l(x)]$, such that $d_{T'}(x,p_i) \leq Wt(i, l(x))$ and $d_{T'}(x,p_j) \leq Wt(f(x), j)$.*

Proof: The proof is by induction on the depth $h = h(T'_x)$ of the subtree T'_x.
Basis: $h = 0$. In this case x is a leaf, and there exists an index k in $[N]$, such that $x = p_k$. Thus $f(x) = l(x) = k$, and for $i = j = k$, we have $d_{T'}(x,p_i) = d_{T'}(x,p_j) = Wt(i,l(x)) = Wt(f(x),j) = 0$.
Induction Step: We assume that the statement holds for the two children $L(x)$ and $R(x)$ of x, and prove it for x.

Suppose first that $|\delta_x| \leq \Delta_x$. By the first two statements of Observation 2.1, $w_L(x), w_R(x) \leq \max\{\Delta_x, |\delta_x|\} = \Delta_x$. Let $p_i \in Leaves(L(x))$ and $p_j \in Leaves(R(x))$ be two points for which $\Delta_x = \Delta(p_i,p_j) = dist(p_i,p_j) - (d_{T'}(L(x),p_i) + d_{T'}(R(x),p_j))$. Clearly both i and j are indices in $[f(x), l(x)]$. Observe that

$$\begin{aligned} d_{T'}(x,p_i) &= w_L(x) + d_{T'}(L(x),p_i) \ \leq \ \Delta_x + d_{T'}(L(x),p_i) \\ &= dist(p_i,p_j) - d_{T'}(R(x),p_j) \ \leq \ dist(p_i,p_j) \ \leq \ Wt(i,j) \ \leq \ Wt(i,l(x)). \end{aligned}$$

(The one before last inequality holds by the triangle inequality.) Similarly, we get that

$$\begin{aligned} d_{T'}(x,p_j) &= w_R(x) + d_{T'}(R(x),p_j) \ \leq \ \Delta_x + d_{T'}(R(x),p_j) \\ &= dist(p_i,p_j) - d_{T'}(L(x),p_i) \ \leq \ dist(p_i,p_j) \ \leq \ Wt(i,j) \ \leq \ Wt(f(x),j). \end{aligned}$$

Otherwise, $|\delta_x| > \Delta_x$. Suppose without loss of generality that $w_L(x) \leq w_R(x)$. In this case, we have $w_L(x) = 0$. By the induction hypothesis, there exist indices i and j in $[f(L(x)), l(L(x))]$, such that $d_{T'}(L(x),p_i) \leq Wt(i, l(L(x)))$ and $d_{T'}(L(x),p_j) \leq Wt(f(L(x)), j)$. Since $w_L(x) = 0$, $d_{T'}(x,p_i) = d_{T'}(L(x),p_i)$ and $d_{T'}(x,p_j) = d_{T'}(L(x),p_j)$. Also, $[f(L(x)), l(L(x))] \subset [f(x), l(x)]$. Consequently, i and j serve as two indices in $[f(x), l(x)]$ for which $d_{T'}(x,p_i) \leq Wt(i, l(L(x))) \leq Wt(i, l(x))$ and $d_{T'}(x,p_j) \leq Wt(f(L(x)), j) \leq Wt(f(x), j)$. $\blacksquare$

The next lemma shows that for every inner vertex x in T', the sum of weights of the two edges that descend from x is no greater than the length of the sub-path of the Hamiltonian path H that traverses all vertices from $Leaves(x)$.

Lemma 2.6 *For any vertex x in T', $w_L(x) + w_R(x) \leq Wt(f(x), l(x))$.*

Proof: The proof is by induction on the depth $h = h(T'_x)$ of the subtree T'_x. The basis $h = 0$ is trivial. *Induction Step:* We assume that the statement holds for the two children $L(x)$ and $R(x)$ of x, and prove it for x.

Suppose first that $|\delta_x| \leq \Delta_x$. By Observation 2.1, $w_L(x) + w_R(x) = \Delta_x$. Let $p_i \in Leaves(L(x))$ and $p_j \in Leaves(R(x))$ be two points for which $\Delta_x = \Delta(p_i, p_j) = dist(p_i, p_j) - (d_{T'}(L(x), p_i) + d_{T'}(R(x), p_j))$. It follows that

$$w_L(x) + w_R(x) \;=\; \Delta_x \;\leq\; dist(p_i, p_j) \;\leq\; Wt(i, j) \;\leq\; Wt(f(x), l(x)).$$

(The one before last inequality holds by the triangle inequality.)

Otherwise, $|\delta_x| > \Delta_x$. Suppose without loss of generality that $w_L(x) \leq w_R(x)$. In this case, we have $|\delta_x| = -\delta_x$, and so $w_L(x) + w_R(x) = -\delta_x = \rho(R(x)) - \rho(L(x))$. By Claim 2.5, there exists an index a in $[f(L(x)), l(L(x))]$, such that $d_{T'}(L(x), p_a) \leq Wt(f(L(x)), a)$. By Lemma 2.2, $\rho(L(x)) = dist(rt, p_a) - d_{T'}(L(x), p_a)$ and $\rho(R(x)) = dist(rt, p_b) - d_{T'}(R(x), p_b)$, for an arbitrary index b in $[f(R(x)), l(R(x))]$. It follows that

$$\begin{aligned}
w_L(x) + w_R(x) &= \rho(R(x)) - \rho(L(x)) \\
&= dist(rt, p_b) - d_{T'}(R(x), p_b) - dist(rt, p_a) + d_{T'}(L(x), p_a) \\
&\leq dist(rt, p_b) - dist(rt, p_a) + d_{T'}(L(x), p_a) \\
&\leq dist(p_a, p_b) + Wt(f(L(x)), a) \;\leq\; Wt(f(L(x)), b) \;\leq\; Wt(f(x), l(x)).
\end{aligned}$$

(The second and third inequalities hold by the triangle inequality. See Figure 4 for an illustration.) ∎

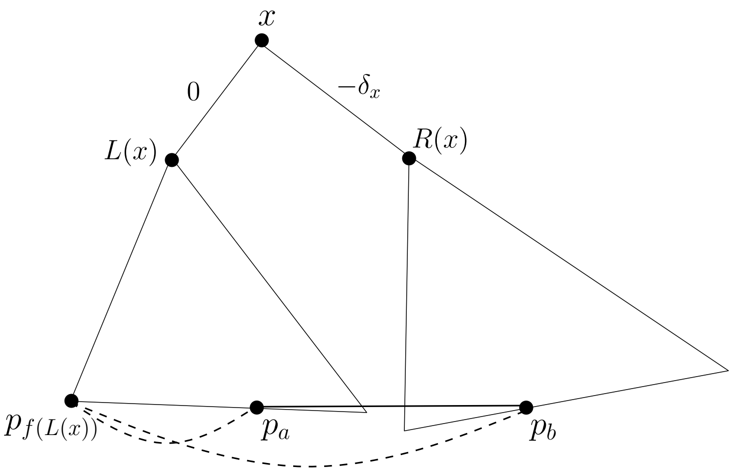

Figure 4: An illustration for the inequality $dist(p_a, p_b) + Wt(f(L(x)), a) \leq Wt(f(L(x)), b)$.

Notice that the depth of T' is $\log N = \lceil \log(n-1) \rceil$. The *level* of a vertex in T' is defined as its unweighted distance from rt. We denote by V_i the set of all vertices in T' of level i, for each index $i \in [0, \log N]$. Also, denote by E_i the set of all edges in T' that connect a vertex in V_i with a vertex in V_{i+1}, for each index $i \in [0, \log N - 1]$, and denote by W_i the sum of all edge weights in E_i.

The following lemma implies that the weight $w(T')$ of T' is at most $\lceil \log(n-1) \rceil \cdot w(H')$.

Lemma 2.7 *For each index $i \in [0, \log N - 1]$, $W_i \leq w(H')$.*

Proof: Fix an arbitrary index $i \in [0, \log N - 1]$. By definition, we have

$$W_i \;=\; \sum_{e \in E_i} w(e) \;=\; \sum_{x \in V_i} (w_L(x) + w_R(x)).$$

By construction, for any pair x, x' of distinct vertices in V_i, $Leaves(x)$ and $Leaves(x')$ are disjoint, and so either $f(x) \le l(x) < f(x') \le l(x')$ or $f(x') \le l(x') < f(x) \le l(x)$ must hold. It follows that $\sum_{x \in V_i} Wt(f(x), l(x)) \;\le\; Wt(1, N) \;=\; Wt(1, n-1) \;=\; w(H')$. Lemma 2.6 implies that $w_L(x) + w_R(x) \;\le\; Wt(f(x), l(x))$, for any $x \in V_i$. Altogether,

$$W_i \;=\; \sum_{x \in V_i} w_L(x) + w_R(x) \;\le\; \sum_{x \in V_i} Wt(f(x), l(x)) \;\le\; w(H'). \qquad \blacksquare$$

Note that the tree T' consists of $2N - 1 = O(n)$ vertices. Also, it is easy to verify that it can be constructed in $O(n^2)$ time, disregarding the time needed to compute the Hamiltonian path H'.
Define $W^*(M, rt) = \min\{dist(rt, v) \mid v \in V \setminus \{rt\}\}$. Clearly, $W^*(M, rt) \le w(MST(M))$.
Lemmas 2.2, 2.3, 2.4 and 2.7 yield the following corollary.

Corollary 2.8 *The rooted tree (T, rt), $T = T(H')$, obtained from T' by adding to it an edge (rt, rt') of weight $\rho(rt')$, is a Steiner SPT for M with weight at most $\lceil \log(n-1) \rceil \cdot w(H') + W^*(M, rt)$ and $O(n)$ vertices. Moreover, T can be constructed in $O(n^2)$ time, disregarding the time needed to compute the Hamiltonian path H'.*

Proof: First, we argue that all edge weights in T are non-negative. Indeed, the first statement of Observation 2.1 implies that all edge weights in T' are non-negative. The only edge of T that does not belong to T' is (rt, rt'), and its weight $\rho(rt')$ is non-negative by Lemma 2.3.

By Lemma 2.4, for any two points v and w in $V \setminus \{rt\}$, $d_T(v, w) = d_{T'}(v, w) \ge dist(v, w)$. Also, by Lemma 2.2, for any point v in $V \setminus \{rt\}$, we have $d_T(rt, v) = d_{T'}(rt', v) + \rho(rt') = dist(rt, v)$. It follows that (T, rt) is a Steiner shortest paths tree for M.

To bound the weight of the tree, first note that for any point $v \in V \setminus \{rt\}$, $\rho(rt') = dist(rt, v) - d_{T'}(rt', v)$, and so $\rho(rt') \le W^*(M, rt)$. By Lemma 2.7, $w(T') \;\le\; \log N \cdot w(H') = \lceil \log(n-1) \rceil \cdot w(H')$, implying that $w(T) \;=\; w(T') + \rho(rt') \le \lceil \log(n-1) \rceil \cdot w(H') + W^*(M, rt)$.

Given the bound on the construction time of T', we conclude that T can also be constructed in $O(n^2)$ time, disregarding the time needed to compute the Hamiltonian path H'. $\blacksquare$

Finally, we remark that it is possible to construct within $O(n^2)$ time a Hamiltonian path $L(M')$ for M' with weight at most $2 \cdot w(MST(M')) = O(w(MST(M)))$. The following claim is provided for completeness.

Claim 2.9 *For any n-point metric X, one can construct in $O(n^2)$ time a Hamiltonian path $L(X)$ with weight at most $2 \cdot w(MST(X))$.*

Proof: Let $T(X)$ be an MST of X rooted at an arbitrary designated point $rt \in X$. Let $D(X)$ be an Euler tour of $T(X)$, starting at rt. For every vertex $v \in X$, remove from $D(X)$ all occurrences of v except for the first one, and denote by $L(X) = (v_1 = rt, v_2, \ldots, v_n)$ the resulting Hamiltonian path of X. It is easy to verify that $w(L(X)) \le 2 \cdot w(T(X)) = 2 \cdot w(MST(X))$.

Since the metric X contains at most $O(n^2)$ edges, its MST can be computed within $O(n^2)$ time (cf. [15], chapter 23). The construction of the Euler tour $D(X)$, as well as the subsequent removal of all non-first occurrences of elements from it, can be performed in $O(n)$ time in the straight-forward way. The claim follows. $\blacksquare$

To optimize the bounds on the weight and construction time of the tree T in Corollary 2.8, we take H' to be the Hamiltonian path $L(M')$ that is guaranteed by Claim 2.9. Specifically, the weight bound of T is reduced to $O(\log n) \cdot w(MST(M))$, and the overall construction time of T is reduced to $O(n^2)$. We derive the main result of this section.

Theorem 2.10 *The rooted tree (T, rt) returned by our construction is a Steiner SPT with lightness $O(\log n)$ and $O(n)$ vertices for the input n-point metric M. Moreover, our construction can be implemented within $O(n^2)$ time.*

Observe that one can easily get rid of zero weight edges that are present in our construction. This can be done just by iteratively contracting all such edges.

2.2 Steiner SLTs

In Section 2.1 we showed that for any n-point metric M and a designated point rt there exists a Steiner tree T which preserves the distances between rt and all other points of M (i.e., an SPT of M with respect to rt) and has lightness at most $O(\log n)$. In this section we generalize this result and show that for any $\epsilon > 0$ there exists a Steiner SLT that provides a $(1+\epsilon)$-approximation to the distances between rt and all other points and has lightness $O(\log \frac{1}{\epsilon})$.

This generalization is based on the following ideas. In [7] Awerbuch et al. devised a construction of *spanning* SLTs with lightness $O\left(\frac{1}{\epsilon}\right)$. Their construction identifies a set $\mathcal{B}$ of special points, called *break-points*, and connects each of the points $B \in \mathcal{B}$ to rt via shortest paths. Our construction replaces these shortest paths by the Steiner SPT for the set $\mathcal{B}$ rooted at rt, which was constructed in Section 2.1. It is pretty obvious that the resulting tree satisfies the desired distance properties. Also, by Theorem 2.10, its lightness is $O(\log |\mathcal{B}|)$. If we could show that $|\mathcal{B}| = O\left(\frac{1}{\epsilon}\right)$, this would finish the proof. However, it is easy to see that this is generally not the case. For example, if the metric M is the unit clique and ϵ is some small constant, every non-root point will be identified as a breakpoint, and we will thus get $|\mathcal{B}| = n-1$. To overcome this obstacle we refine the bound on $w(T)$ from Theorem 2.10, and express it in terms of the sum of all root-distances $\sum_{v \in V \setminus \{rt\}} dist(rt, v)$, rather than in terms of the number n ot points in M. We then use this refined bound to analyze the weight of our shallow-light trees.

We start with the following simple observation.

Claim 2.11 $\sum_{v \in V \setminus \{rt\}} d_{T'}(rt', v) \;=\; \sum_{i=0}^{\log N - 1} 2^{\log N - (i+1)} \cdot W_i$.

Proof: Fix an arbitrary index $i \in [0, \log N - 1]$, and consider an edge $e = (v_i, v_{i+1}) \in E_i$, $v_i \in V_i$, $v_{i+1} \in V_{i+1}$. For each vertex $v \in Leaves(v_{i+1})$, the edge e belongs to the path connecting rt' with v in T'. Moreover, the edge e does not belong to paths that connect rt' to other vertices $z \in V \setminus \{rt\} \setminus Leaves(v_{i+1})$. Hence,

$$\sum_{v \in V \setminus \{rt\}} d_{T'}(rt', v) \;=\; \sum_{e=(v_i, v_{i+1}) \in E(T')} |Leaves(v_{i+1})| \cdot w(e),$$

where $w(e)$ stands for the weight of the edge e in T'. Observe that for $v_{i+1} \in V_{i+1}$, $|Leaves(v_{i+1})| = 2^{\log N - (i+1)}$. It follows that

$$\sum_{e=(v_i, v_{i+1}) \in E(T')} |Leaves(v_{i+1})| \cdot w(e) \;=\; \sum_{i=0}^{\log N - 1} \sum_{e \in E_i} |Leaves(v_{i+1})| \cdot w(e) \;=\; \sum_{i=0}^{\log N - 1} 2^{\log N - (i+1)} \cdot W_i.$$

This completes the proof. ∎

We are now ready to prove our refined bound on the weight of the SPT T, which was constructed in Section 2.1.

Lemma 2.12 *Suppose that* $\sum_{v\in V\setminus\{rt\}} dist(rt,v) \le \alpha\cdot\beta$, *for some pair* $\alpha\ge 1, \beta > 0$ *of numbers. Then the weight* $w(T)$ *of* $T = T(H')$ *satisfies* $w(T) \le \beta + \lceil\log\alpha\rceil\cdot w(H') + W^*(M,rt)$.

Remark: Clearly, $\sum_{v\in V\setminus\{rt\}} dist(rt,v) \le (n-1)\cdot w(MST(M))$. By substituting $\alpha = n-1, \beta = w(MST(M))$, we get an upper bound of $w(MST(M)) + \lceil\log(n-1)\rceil\cdot w(H') + W^*(M,rt)$ on $w(T)$, which is slightly larger than the upper bound given in Corollary 2.8. On the other hand, we get significantly better bounds on $w(T)$ whenever $\sum_{v\in V\setminus\{rt\}} dist(rt,v) \ll (n-1)\cdot w(MST(M))$.
Proof: Since $w(T) = w(T') + \rho(rt') \le w(T') + W^*(M,rt)$, it suffices to show that $w(T') \le \beta + \lceil\log\alpha\rceil\cdot w(H')$.
Suppose first that $\lceil\log\alpha\rceil \ge \log N$. By Lemma 2.7, we have $w(T') \le \log N\cdot w(H')$, and so $w(T') \le \log N\cdot w(H') \le \beta + \lceil\log\alpha\rceil\cdot w(H')$. We henceforth assume that $\lceil\log\alpha\rceil \le \log N - 1$.
By construction, we have $w(T') = \sum_{i=0}^{\log N-1} W_i$ and $\sum_{v\in V\setminus\{rt\}} d_T(rt,v) \ge \sum_{v\in V\setminus\{rt\}} d_{T'}(rt',v)$. Since T is an SPT for M, we conclude that

$$\alpha\cdot\beta \;\ge\; \sum_{v\in V\setminus\{rt\}} dist(rt,v) \;=\; \sum_{v\in V\setminus\{rt\}} d_T(rt,v) \;\ge\; \sum_{v\in V\setminus\{rt\}} d_{T'}(rt',v).$$

Hence, by Claim 2.11,

$$\begin{aligned}\alpha\cdot\beta \;\ge\; \sum_{v\in V\setminus\{rt\}} d_{T'}(rt',v) \;&=\; \sum_{i=0}^{\log N-1} 2^{\log N-(i+1)}\cdot W_i \;\ge\; \sum_{i=0}^{\log N-(\lceil\log\alpha\rceil+1)} 2^{\log N-(i+1)}\cdot W_i \\ &\ge\; 2^{\lceil\log\alpha\rceil}\cdot\left(\sum_{i=0}^{\log N-(\lceil\log\alpha\rceil+1)} W_i\right) \;\ge\; \alpha\cdot\left(\sum_{i=0}^{\log N-(\lceil\log\alpha\rceil+1)} W_i\right),\end{aligned}$$

and so

$$\sum_{i=0}^{\log N-(\lceil\log\alpha\rceil+1)} W_i \;\le\; \beta.$$

Also, Lemma 2.7 implies that

$$\sum_{\log N-\lceil\log\alpha\rceil}^{\log N-1} W_i \;\le\; \lceil\log\alpha\rceil\cdot w(H').$$

Altogether, we have

$$w(T') \;=\; \sum_{i=0}^{\log N-1} W_i \;=\; \left(\sum_{i=0}^{\log N-(\lceil\log\alpha\rceil+1)} W_i\right) + \left(\sum_{\log N-\lceil\log\alpha\rceil}^{\log N-1} W_i\right) \;\le\; \beta + \lceil\log\alpha\rceil\cdot w(H'). \qquad\blacksquare$$

Now we proceed to extending our construction of shortest-paths trees from Section 2.1 to a construction of shallow-light trees. As was discussed above, this construction can be seen as a hybrid of our construction from Section 2.1 with the construction of [7] of spanning shallow-light trees. Its analysis is closely related to that of [7], except that it relies on Lemma 2.12 for the analysis of weight.

Consider an n-point metric $M = (V, dist)$, let $T = T(M)$ be an MST of M rooted at an arbitrary designated point $rt\in V$, and let $L = L(M)$ be a Hamiltonian path for M of weight at most $2\cdot w(MST(M))$. By Claim 2.9, such a Hamiltonian path can be computed in $O(n^2)$ time.
Fix a parameter $\theta \le 2$. The value of θ will determine the values of the root-stretch and lightness of the constructed tree.

We start with identifying a set of "break-points" $\mathcal{B} = \{B_1, B_2, \ldots, B_k\}$, $\mathcal{B} \subseteq V$, $k \geq 2$. The break-point B_1 is the vertex $v_1 = rt$. The break-point B_{i+1}, $i \in [k-1]$, is the first vertex in L after B_i such that

$$d_T(B_i, B_{i+1}) > \theta \cdot dist(rt, B_{i+1}).$$

Let $M_{\mathcal{B}}$ be the sub-metric of M induced by the point set of $\mathcal{B}$. Also, let $L' = (B_2, \ldots, B_k\}$ be the sub-path of L that contains the breakpoints of $B \backslash \{rt\}$. By the triangle inequality, $w(L') \leq w(L) \leq 2 \cdot w(MST(M))$. By Corollary 2.8, we can build a Steiner SPT $T_{\mathcal{B}} = T_{\mathcal{B}}(L')$ of $M_{\mathcal{B}}$ rooted at rt with small weight. Denote the set of Steiner points in $T_{\mathcal{B}}$ by $S_{\mathcal{B}}$.
Let $\tilde{G} = (V \cup S_{\mathcal{B}}, E(T) \cup E(T_{\mathcal{B}}))$ be the graph obtained from the union of the two trees T and $T_{\mathcal{B}}$. Finally, we define $S(T)$ to be an SPT over $\tilde{G}$ rooted at rt.

The following claim implies that the sum of distances in T, taken over all pairs of consecutive break-points, is not too large. It follows from the observation that D visits each edge twice.

Claim 2.13 $\sum_{i=1}^{k-1} d_T(B_i, B_{i+1}) \leq 2 \cdot w(T) = 2 \cdot w(MST(M))$.

The next lemma bounds the root-stretch of the constructed tree $S(T)$.

Lemma 2.14 *For any vertex $v \in V \setminus \{rt\}$, it holds that $d_{S(T)}(rt, v) \leq (1 + 2\theta) \cdot dist(rt, v)$.*

Proof: Consider an arbitrary vertex $v \in V$. First, recall that $S(T)$ is an SPT over $\tilde{G}$ rooted at rt, and so $d_{S(T)}(rt, v) = d_{\tilde{G}}(rt, v)$. Clearly, the lemma holds if v is a breakpoint, as in this case we have

$$d_{S(T)}(rt, v) = d_{\tilde{G}}(rt, v) \leq d_{T_{\mathcal{B}}}(rt, v) = dist(rt, v).$$

We henceforth assume that v is not a breakpoint. Let i be the index in $[k-1]$ such that v is located between B_i and B_{i+1} in L. Since B_i is a break-point, it holds that $d_{\tilde{G}}(rt, B_i) \leq d_{T_{\mathcal{B}}}(rt, B_i) = dist(rt, B_i)$. Clearly, $d_{\tilde{G}}(B_i, v) \leq d_T(B_i, v)$. By the triangle inequality, $d_{\tilde{G}}(rt, v) \leq d_{\tilde{G}}(rt, B_i) + d_{\tilde{G}}(B_i, v)$. Altogether,

$$d_{S(T)}(rt, v) = d_{\tilde{G}}(rt, v) \leq d_{\tilde{G}}(rt, B_i) + d_{\tilde{G}}(B_i, v) \leq dist(rt, B_i) + d_T(B_i, v).$$

Since v was not identified as a break-point, necessarily

$$d_T(B_i, v) \leq \theta \cdot dist(rt, v), \tag{3}$$

and so

$$d_{S(T)}(rt, v) \leq dist(rt, B_i) + \theta \cdot dist(rt, v). \tag{4}$$

By the triangle inequality and Equation (3),

$$\begin{aligned} dist(rt, B_i) &\leq dist(rt, v) + dist(B_i, v) \leq dist(rt, v) + d_T(B_i, v) \\ &\leq dist(rt, v) + \theta \cdot dist(rt, v) = (1 + \theta) \cdot dist(rt, v). \end{aligned} \tag{5}$$

Plugging Equation (5) in Equation (4), we obtain

$$d_{S(T)}(rt, v) \leq (1 + \theta) \cdot dist(rt, v) + \theta \cdot dist(rt, v) = (1 + 2\theta) \cdot dist(rt, v). \quad \blacksquare$$

Next, we bound the weight of the constructed tree $S(T)$.

Lemma 2.15 $w(S(T)) \leq 2 \cdot \left(\left\lceil \log\left(\frac{2}{\theta}\right) \right\rceil + 3/2\right) \cdot w(MST(M))$.

Proof: By the choice of break-points, for each index $i \in [k-1]$, $dist(rt, B_{i+1}) < \frac{1}{\theta} \cdot d_T(B_i, B_{i+1})$. By Claim 2.13, $\sum_{i=1}^{k-1} d_T(B_i, B_{i+1}) \leq 2 \cdot w(MST(M))$. Therefore,

$$\sum_{i=1}^{k-1} dist(rt, B_{i+1}) < \frac{1}{\theta} \cdot \sum_{i=1}^{k-1} d_T(B_i, B_{i+1}) \leq \frac{2}{\theta} \cdot w(MST(M)).$$

Consider the metric $M_{\mathcal{B}} = (\mathcal{B}, dist)$, and set $\alpha = \frac{2}{\theta}$, $\beta = w(MST(M))$. Notice that

$$\sum_{B \in \mathcal{B} \setminus \{rt\}} dist(rt, B) = \sum_{i=1}^{k-1} dist(rt, B_{i+1}) \leq \alpha \cdot \beta.$$

Since $\theta \leq 2$, we have $\alpha \geq 1$. Clearly, $\beta = w(MST(M)) > 0$. Hence, by Lemma 2.12, the weight $w(T_{\mathcal{B}})$ of $T_{\mathcal{B}} = T_{\mathcal{B}}(L')$ satisfies

$$\begin{aligned} w(T_{\mathcal{B}}) &\leq \beta + \lceil \log \alpha \rceil \cdot w(L') + W^*(M_{\mathcal{B}}, rt) \\ &= w(MST(M)) + \left\lceil \log \left(\frac{2}{\theta} \right) \right\rceil \cdot w(L') + W^*(M_{\mathcal{B}}, rt) \\ &\leq 2 \cdot w(MST(M)) + \left\lceil \log \left(\frac{2}{\theta} \right) \right\rceil \cdot 2 \cdot w(MST(M)) \end{aligned}$$

By construction, $w(S(T)) \leq w(\tilde{G}) = w(T) + w(T_{\mathcal{B}})$, and so

$$\begin{aligned} w(S(T)) &\leq w(MST(M)) + 2 \cdot w(MST(M)) + \left\lceil \log \left(\frac{2}{\theta} \right) \right\rceil \cdot 2 \cdot w(MST(M)) \\ &= 2 \cdot \left(\left\lceil \log \left(\frac{2}{\theta} \right) \right\rceil + 3/2 \right) \cdot w(MST(M)). \quad \blacksquare \end{aligned}$$

Note also that for any metric M, the weight of the Minimum Steiner Tree for M (denoted $SMT(M)$) is greater or equal to half the weight of the Minimum Spanning Tree for M, i.e., $w(SMT(M)) \geq \frac{1}{2} \cdot w(MST(M))$. It follows that $w(S(T)) \leq 4 \cdot \left(\left\lceil \log \left(\frac{2}{\theta} \right) \right\rceil + 3/2 \right) \cdot w(SMT(M))$.

Finally, we analyze the running time of the construction.

Lemma 2.16 *The tree $S(T)$ can be constructed in $O(n^2)$ time.*

Proof: First, note that the MST T for M can be computed within $O(n^2)$ time. Also, as was mentioned above, Claim 2.9 implies that the Hamiltonian path L can be computed within $O(n^2)$ time as well. It is easy to see that $O(n)$ time suffices to identify the set $\mathcal{B} = \{B_1, B_2, ..., B_k\}$ of break-points. An additional amount of $O(n)$ time requires to construct the sub-path L' of L. By Corollary 2.8, the Steiner SPT $T_{\mathcal{B}} = T_{\mathcal{B}}(L')$ consists of $O(k) = O(n)$ vertices, and given the Hamiltonian path L' for $M_{\mathcal{B}}$, it can be computed within $O(k^2) = O(n^2)$ time. Consequently, the graph $\tilde{G}$ can be constructed in $O(n^2)$ time, and it consists of $O(n)$ vertices and edges. The final step of the algorithm is the construction of an SPT over $\tilde{G}$, which can be carried out in $O(n \cdot \log n)$ time. The lemma follows. $\blacksquare$

Set $\epsilon = 2\theta$. Lemmas 2.14, 2.15 and 2.16 imply the following corollary.

Corollary 2.17 *For any n-point metric $M = (V, dist)$, a designated point $rt \in V$ and a number $0 < \epsilon < \frac{1}{2}$, there exists a Steiner tree of M rooted at rt, having root-stretch at most $(1+\epsilon)$, lightness $O(\log \frac{1}{\epsilon})$ and $O(n)$ vertices. Moreover, this construction can be implemented within $O(n^2)$ time.*

3 Lower Bounds for Steiner SLTs

In this section we show that there exist n-point metrics for which any Steiner SPT has lightness $\Omega(\log n)$. We then employ this result and show that for any $\epsilon > 0$ there exist metrics for which any Steiner tree that approximates all distances from a designated root vertex by a factor of at most $(1+\epsilon)$ has lightness $\Omega(\log \frac{1}{\epsilon})$. In view of our upper bounds from Sections 2, these lower bounds are tight up to constant factors.

Let $\mathcal{P}_n$ be the family of all n-point 1-dimensional Euclidean metrics, such that the distance between any two consecutive points is at least 1. We denote the diameter of a metric $M \in \mathcal{P}_n$ by $diam(M)$, and observe that $diam(M) \geq n-1$. Given a metric $M \in \mathcal{P}_n$, a number $r \geq \frac{1}{2} \cdot diam(M)$ and a number $\epsilon > 0$, we say that a rooted Steiner tree (T, rt) of M is an *r-tree* (respectively, *(r,ϵ)-tree*) for M, if $d_T(rt, v) = r$ (resp., $r \leq d_T(rt, v) < (1+\epsilon) \cdot r$), for every point v in M. Note that if T is either an r-tree or an *(r,ϵ)-tree* for M, for an arbitrary number $\epsilon > 0$, then the root vertex rt of T must be a Steiner point, i.e., $rt \not\in M$.

Lemma 3.1 *Let M be an arbitrary metric in $\mathcal{P}_n$, and let $r \geq \frac{1}{2} \cdot diam(M) \geq \frac{1}{2} \cdot (n-1)$. Then any r-tree T for M has weight $w(T)$ at least $g(n,r) = r - \frac{1}{2} \cdot (n-1) + \frac{1}{2} \cdot n \cdot \log n$.*

Proof: The proof is by induction on $n, n \geq 1$, for all values of $r \geq \frac{1}{2} \cdot diam(M)$. The basis $n = 1$ is trivial, as $g(1, r) = r$.
Induction Step: We assume that the statement holds for all smaller values of $n, n \geq 2$, and prove it for n. Let (T, rt) be an r-tree for M with a minimum number of Steiner points, taken over all r-trees for M of minimum weight. Next, we show that $w(T) \geq r - \frac{1}{2} \cdot (n-1) + \frac{1}{2} \cdot n \cdot \log n$.
Denote the children of rt in T by $c_1, c_2, \ldots, c_k$, with $k \geq 1$. Fix an arbitrary index $i \in [k]$. We denote the subtree of T rooted at c_i by T_i and the set of required points in T_i by V_i. Also, we write n_i as a shortcut for $|V_i|$ and w_i as a shortcut for $d_T(rt, c_i) = w(rt, c_i)$. (Clearly, $\sum_{i=1}^{k} n_i = n$.) We argue that $n_i \geq 1$; indeed, otherwise the tree obtained from T by removing from it the subtree T_i is an r-tree for M, having weight no greater than that of T and less Steiner points, yielding a contradiction. Consider the metric M_i induced by the point set of V_i. Observe that $M_i \in \mathcal{P}_{n_i}$. Also, note that the subtree T_i of T dominates M_i, and so it must hold that $r - w_i \geq \frac{1}{2} \cdot diam(M_i)$. It follows that T_i is an $(r - w_i)$-tree for M_i. Hence, by the induction hypothesis, we have $w(T_i) \geq g(n_i, r - w_i) = (r - w_i) - \frac{1}{2} \cdot (n_i - 1) + \frac{1}{2} \cdot n_i \cdot \log n_i$.

By construction, $w(T) = \sum_{i=1}^{k} (w_i + w(T_i))$. Consequently,

$$
\begin{aligned}
w(T) &\geq \sum_{i=1}^{k} \left(w_i + \left((r - w_i) - \frac{1}{2} \cdot (n_i - 1) + \frac{1}{2} \cdot n_i \cdot \log n_i \right) \right) \\
&= \sum_{i=1}^{k} \left(r - \frac{1}{2} \cdot (n_i - 1) + \frac{1}{2} \cdot n_i \cdot \log n_i \right) \\
&= \sum_{i=1}^{k} \left(r - \frac{1}{2} \cdot (n - 1) + \frac{1}{2} \cdot (n - n_i) + \frac{1}{2} \cdot n_i \cdot \log n_i \right) \\
&\geq r - \frac{1}{2} \cdot (n-1) + \sum_{i=1}^{k} \left(\frac{1}{2} \cdot (n - n_i) + \frac{1}{2} \cdot n_i \cdot \log n_i \right) \\
&= r - \frac{1}{2} \cdot (n-1) + \frac{1}{2} \cdot \left(n \cdot (k-1) + \sum_{i=1}^{k} (n_i \cdot \log n_i) \right) \\
&\geq r - \frac{1}{2} \cdot (n-1) + \frac{1}{2} \cdot \left(\sum_{i=1}^{k} \left(n_i \cdot \log \left(\frac{n}{n_i} \right) \right) + \sum_{i=1}^{k} (n_i \cdot \log n_i) \right)
\end{aligned}
$$

$$= \; r - \frac{1}{2} \cdot (n-1) + \frac{1}{2} \cdot \sum_{i=1}^{k} (n_i \cdot \log n) \; = \; r - \frac{1}{2} \cdot (n-1) + \frac{1}{2} \cdot n \cdot \log n.$$

(The last inequality follows from Lemma A.1 that appears in Appendix A.) ▮

Consider the 1-dimensional Euclidean metric $\zeta_n \in \mathcal{P}_n$ that consists of n vertices $v_1, v_2, \ldots, v_n$ that lie on the x-axis with coordinates $1, 2, \ldots, n$, respectively. Now extend the metric ζ_n to include an additional vertex rt, such that the distance between rt and v_i is equal to $\frac{n-1}{2}$, for each $i \in [n]$. Denote the resulting $(n+1)$-point metric by $\tilde{\zeta}_{n+1}$. Lemma 3.1 implies that any Steiner SPT for $\tilde{\zeta}_{n+1}$ rooted at rt has lightness $\Omega(\log n)$. We remark that $\tilde{\zeta}_{n+1}$ is not a Euclidean metric. However, the same bound of $\Omega(\log n)$ on the lightness of Steiner SPTs can be obtained for simple Euclidean 2-dimensional point sets. For example, with a slight abuse of notation, let C_n denote a set of n points that are uniformly spaced on the boundary of a circle with radius $\frac{n}{2\pi}$ (rather than unit radius as in Section 1.3), centered at the origin $(0,0)$, and define $\tilde{C}_{n+1} = C_n \cup \{(0,0)\}$. It is not hard to see that the proof of Lemma 3.1 carries through (with minor adjustments) also if T is an r-tree for C_n. In other words, any SPT for the point set $\tilde{C}_{n+1}$ rooted at $rt = (0,0)$ has lightness $\Omega(\log n)$.

Theorem 3.2 *For any sufficiently large integer n, there exists a Euclidean 2-dimensional n-point metric M and a designated point $rt \in M$, such that every Steiner SPT rooted at rt has lightness $\Omega(\log n)$.*

We use the following lemma to prove Lemma 3.4, which, in turn, enables us to generalize Theorem 3.2 for Steiner SLTs.

Lemma 3.3 *Let M be an arbitrary metric in $\mathcal{P}_n$ and let $r \geq \frac{1}{2} \cdot diam(M) \geq \frac{1}{2} \cdot (n-1)$. Also, let (T, rt) be an $(r, \frac{1}{r})$-tree for M with a minimum number of Steiner points, taken over all $(r, \frac{1}{r})$-trees for M of minimum weight. Then: (1) All leaves of T are required points. (2) All inner vertices of T are Steiner points. (3) There are at most $2n-1$ edges in T.*

Proof: The first assertion of the lemma is obvious.

To prove the second assertion of the lemma, suppose for contradiction that there is an inner vertex v in T that belongs to M, and let l be some leaf in the subtree T_v of T rooted at v. The first assertion of this lemma implies that l belongs to M. Since T is an $(r, \frac{1}{r})$-tree for M and both v and l belong to M, we have $d_T(rt, v) \geq r$ and $d_T(rt, l) < (1 + \frac{1}{r}) \cdot r = r + 1$. However, since T dominates M, it must hold that $d_T(v, l) \geq d_M(v, l) \geq 1$, and so

$$d_T(rt, l) \; = \; d_T(rt, v) + d_T(v, l) \; \geq \; r + 1,$$

yielding a contradiction.

To prove the third assertion, it suffices to show that every inner vertex in T, except for maybe the root vertex rt, has at least two children. Suppose for contradiction that there is an inner vertex $v \neq rt$ with only one child u, and let $\pi(v)$ be the parent of v in T. Denote by $w(e)$ the weight of an edge e in T. The second assertion of this lemma implies that both v and its parent $\pi(v)$ are Steiner points. Thus, we can remove v from T by replacing the two edges $(\pi(v), v)$ and (v, u) that are incident to v in T with a single edge $(\pi(v), u)$ of the same weight $w(\pi(v), v) + w(v, u)$. The resulting tree is an $(r, \frac{1}{r})$-tree for M, having less Steiner points than T and the same weight, a contradiction. Lemma 3.3 follows. ▮

Lemma 3.4 *Let M be an arbitrary metric in $\mathcal{P}_n$ and let $r \geq \frac{1}{2} \cdot diam(M) \geq \frac{1}{2} \cdot (n-1)$. Then any $(r, \frac{1}{r})$-tree T for M has weight $w(T)$ at least $g(n, r) - 2n = \left(r - \frac{1}{2} \cdot (n-1) + \frac{1}{2} \cdot n \cdot \log n\right) - 2n$.*

Proof: Let (T, rt) be an $(r, \frac{1}{r})$-tree for M with a minimum number of Steiner points, taken over all $(r, \frac{1}{r})$-trees for M of minimum weight.

Let w be the weight function of T, and denote by $w(e)$ the weight of an edge e in T. Also, denote the vertex-set and edge-set of T by V and E, respectively. For a vertex v in V, let $\pi(v)$ denote the parent of v in T, let $Ch(v)$ denote the set of children of v in T, and denote by $\delta(v)$ the maximum distance between v and a leaf in the subtree T_v of T rooted at v. Consider an arbitrary edge $e = (\pi(v), v) \in E$. Observe that

$$\delta(\pi(v)) \;=\; \max\{\delta(u) + w(\pi(v), u) \mid u \in Ch(\pi(v))\} \;\geq\; \delta(v) + w(\pi(v), v). \tag{6}$$

Next, we define a new weight function w' over the edge-set E of T. Specifically, for each edge $e = (\pi(v), v) \in E$, set $w'(e) = \delta(\pi(v)) - \delta(v)$. For convenience, we denote by T' the tree induced by the edge-set E of T and the new weight function w'.

Claim 3.5 *The tree (T', rt) is a $\delta(rt)$-tree for M. Moreover, its weight $w'(T')$ is greater by an additive term of at most $2n$ than the weight $w(T)$ of the original tree T, i.e., $w'(T') \leq w(T) + 2n$.*

Proof: First, Equation (6) implies that for every edge $e \in E$, $w'(e) \geq w(e)$. Since T dominates M, it follows that T' dominates M as well. Also, note that $\delta(rt) \geq r \geq \frac{1}{2} \cdot diam(M)$.

Next, we prove that for every vertex $v \in V$ and any leaf l in the subtree T'_v of T' rooted at v, $d_{T'}(v, l) = \delta(v)$. The proof is by induction on the depth $h = h(v)$ of v. The basis $h = 0$ is trivial.
Induction Step: We assume that the statement holds for all children of v, and prove it for v. Let u be the child of v, such that the leaf l belongs to the subtree T'_u of T' rooted at u. By the induction hypothesis, $d_{T'}(u, l) = \delta(u)$. Also, by construction,

$$d_{T'}(v, l) \;=\; w'(v, u) + d_{T'}(u, l) \;=\; \delta(v) - \delta(u) + \delta(u) = \delta(v),$$

which proves the induction step. It follows that $d_{T'}(rt, l) = \delta(rt)$, for every leaf l in T'. By the first two assertions of Lemma 3.3, all points of M are leaves in T', implying that the distance in T' between rt and every point of M is $\delta(rt)$. Hence, T' is a $\delta(rt)$-tree for M.

It remains to bound the weight $w'(T')$ of T'. Since T is an $(r, \frac{1}{r})$-tree for M and all leaves of T belong to M, it follows that $r \leq d_T(rt, l) < (1 + \frac{1}{r}) \cdot r = r + 1$, for every leaf l in T. Consequently, for every two leaves l_1 and l_2 in T, $|d_T(rt, l_1) - d_T(rt, l_2)| < 1$. More generally, if l_1 and l_2 are descendants of some vertex x in T, then we have

$$|d_T(x, l_1) - d_T(x, l_2)| \;=\; |d_T(rt, l_1) - d_T(rt, l_2)| \;<\; 1. \tag{7}$$

Consider now an arbitrary edge $e = (\pi(v), v) \in E$, and let u be a child of $\pi(v)$, such that $\delta(\pi(v)) = \delta(u) + w(\pi(v), u)$. Also, let l_u (respectively, l_v) be a leaf in the subtree T_u (resp., T_v), such that $\delta(u) = d_T(u, l_u)$ (resp., $\delta(v) = d_T(v, l_v)$). Notice that both l_u and l_v are descendants of $\pi(v)$ in T and $d_T(\pi(v), l_u) \geq d_T(\pi(v), l_v)$. Thus, Equation (7) yields $d_T(\pi(v), l_u) - d_T(\pi(v), l_v) < 1$. Also, we have $d_T(\pi(v), l_u) = \delta(u) + w(\pi(v), u)$ and $d_T(\pi(v), l_v) = \delta(v) + w(\pi(v), v)$. Altogether,

$$\begin{aligned} w'(e) \;&=\; w'(\pi(v), v) \;=\; \delta(\pi(v)) - \delta(v) \;=\; (\delta(u) + w(\pi(v), u)) - \delta(v) \\ &=\; d_T(\pi(v), l_u) - (d_T(\pi(v), l_v) - w(\pi(v), v)) \;<\; w(\pi(v), v) + 1. \end{aligned}$$

We have proved that for any edge $e \in E$, $w'(e) < w(e) + 1$. (See Figure 5 for an illustration.) By the third assertion of Lemma 3.3, there are at most $2n - 1$ edges in E, and so

$$w'(T') \;=\; \sum_{e \in E} w'(e) \;<\; \sum_{e \in E} (w(e) + 1) \;=\; \sum_{e \in E} w(e) + |E| \;\leq\; w(T) + 2n - 1,$$

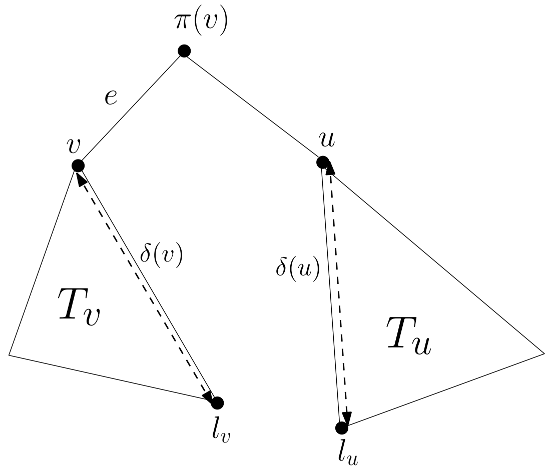

Figure 5: The vertex $\pi(v)$ and its two subtrees T_v and T_u.

which completes the proof of Claim 3.5. ▮

Claim 3.5 implies that T' is a $\delta(rt)$-tree for M, with $\delta(rt) \geq r \geq \frac{1}{2} \cdot diam(M)$ and $w'(T') \leq w(T) + 2n$. By Lemma 3.1,

$$w'(T') \;\geq\; \delta(rt) - \frac{1}{2} \cdot (n-1) + \frac{1}{2} \cdot n \cdot \log n \;\geq\; r - \frac{1}{2} \cdot (n-1) + \frac{1}{2} \cdot n \cdot \log n,$$

and so

$$w(T) \;\geq\; w'(T') - 2n \;\geq\; r - \frac{1}{2} \cdot (n-1) + \frac{1}{2} \cdot n \cdot \log n - 2n.$$

Lemma 3.4 follows. ▮

Lemma 3.4 implies that any Steiner tree for $\tilde{\zeta}_{n+1}$ rooted at rt with root-stretch less than $1 + \frac{2}{n-1}$ has lightness $\Omega(\log n)$. More generally, consider the metric $\tilde{\zeta}_{k+1}$ that consists of $k+1$ points, for some parameter $k \leq n$. Extend this metric by adding to it $n-k$ points, that are "located" (arbitrarily) at tiny distances from the already existing points of $\tilde{\zeta}_{k+1} \setminus \{rt\}$. Clearly, any Steiner tree for the resulting $(n+1)$-point metric rooted at rt with root-stretch less than $1 + \frac{2}{k-1}$ has lightness $\Omega(\log k)$. Also, similarly to above, the same bound of $\Omega(\log k)$ on the lightness of Steiner $(1 + \frac{2}{k-1})$-SPTs can be obtained for simple Euclidean 2-dimensional metrics. Setting $\epsilon = \frac{2}{k-1}$, we obtain the following result.

Theorem 3.6 *For any sufficiently large integer n and any parameter $\epsilon = \Omega(\frac{1}{n})$, $\epsilon \leq \frac{1}{2}$, there exists a Euclidean 2-dimensional n-point metric M and a designated point $rt \in M$, such that every Steiner tree rooted at rt with root-stretch less than $1 + \epsilon$ has lightness $\Omega(\log \frac{1}{\epsilon})$.*

Next, we strengthen Theorem 3.6 in two ways. Specifically, we present a metric $\vartheta = \vartheta_{n,k}$ for which the same tradeoff of $1+\epsilon$ versus $\Omega(\log \frac{1}{\epsilon})$ between the root-stretch and lightness holds for *any* root vertex. Moreover, we demonstrate that this tradeoff cannot be improved even if we consider *average* root-stretch rather than (worst-case) root-stretch. For the analysis of this metric, we will use Lemmas 3.1 and 3.4.

Let n, k be an arbitrary pair of integers, such that $n \geq 2k \geq 4$, and let α be some tiny number, with $0 < \alpha \ll \frac{1}{n}$. In what follows we assume for simplicity that n is even and k divides $n/2$, but the general case can be handled similarly. Let $V = \bigcup_{\ell=1}^{k} V_\ell$ and $U = \bigcup_{\ell=1}^{k} U_\ell$, where for each index $\ell \in [k]$, $V_\ell = \{v_\ell^{(1)}, v_\ell^{(2)}, \ldots, v_\ell^{(\frac{n}{2k})}\}$ and $U_\ell = \{u_\ell^{(1)}, u_\ell^{(2)}, \ldots, u_\ell^{(\frac{n}{2k})}\}$. Define $\vartheta = \vartheta_{n,k} = (\mathcal{V}, dist)$ to be the n-point metric, where $\mathcal{V} = V \cup U$, and the distance function $dist$ is set as follows. (See Figure 6 for an illustration.)

1. For any index $\ell \in [k]$ and any pair of distinct indices $i, j \in [\frac{n}{2k}]$, $dist(v_\ell^{(i)}, v_\ell^{(j)}) = dist(u_\ell^{(i)}, u_\ell^{(j)}) = \alpha$.
2. For any pair of distinct indices $\ell, q \in [k]$ and any pair of indices $i, j \in [\frac{n}{2k}]$, $dist(v_\ell^{(i)}, v_q^{(j)}) = dist(u_\ell^{(i)}, u_q^{(j)}) = |\ell - q|$.
3. For any pair of indices $\ell, q \in [k]$ and any pair of indices $i, j \in [\frac{n}{2k}]$, $dist(v_\ell^{(i)}, u_q^{(j)}) = \frac{k-1}{2}$.

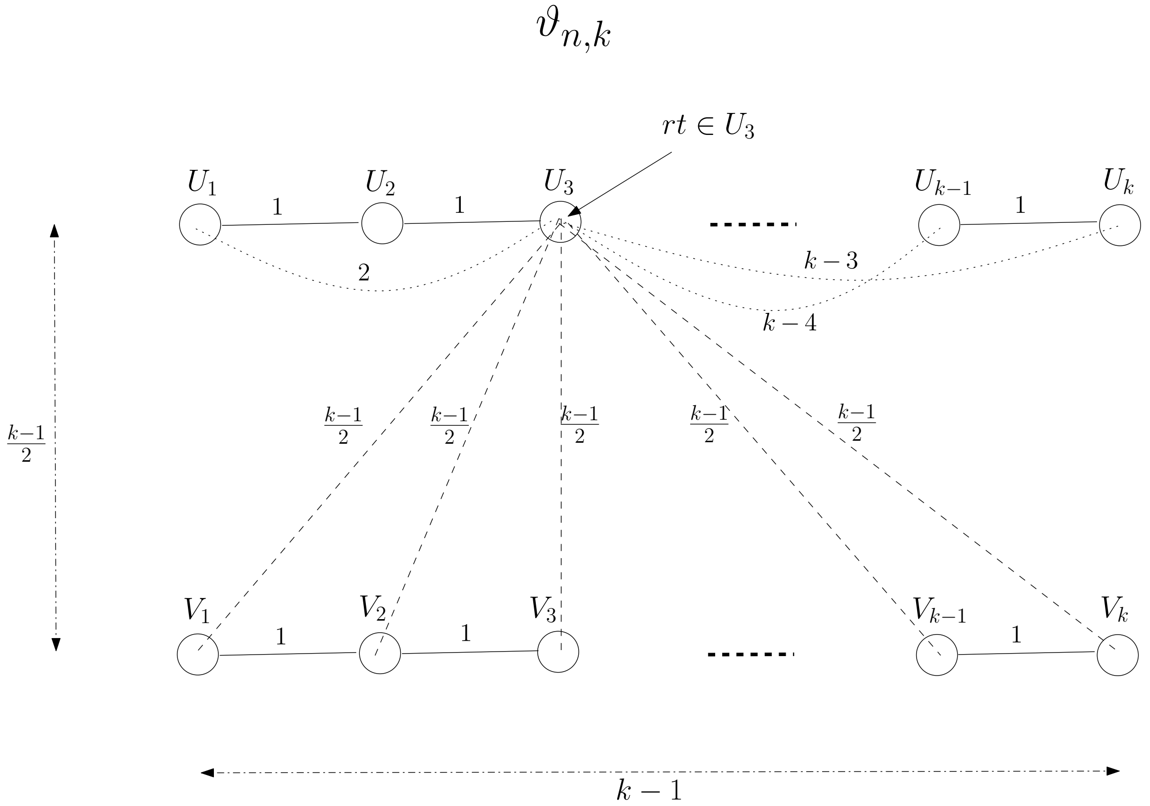

Figure 6: An illustration of the metric $\vartheta_{n,k}$. For each index ℓ, the circles around V_ℓ and U_ℓ designate the $(\frac{n}{2k})$-point sets corresponding to them. The distance between all pairs of points that belong to the same set V_ℓ or U_ℓ, $\ell \in [k]$, is equal to some tiny number $0 < \alpha \ll \frac{1}{n}$. All solid lines in the figure have length 1. These lines designate the distance between a point in V_ℓ (respectively, U_ℓ) and a point in $V_{\ell+1}$ (resp., $U_{\ell+1}$), $\ell \in [k-1]$. All distances between an arbitrary designated point $rt \in U_3$ and some point in U_ℓ, $\ell \in [k], |3-\ell| \geq 2$, are depicted in the figure by dotted lines. Finally, the distance between a point in V and a point in U is equal to $\frac{k-1}{2}$. All distances between $rt \in U_3$ and some point in V_ℓ, $\ell \in [k]$, are depicted in the figure by dashed lines.

A point set $W \subseteq V$ (respectively, $W \subseteq U$) is called V-elementary (resp., U-elementary), if $|W \cap V_\ell| \leq 1$ (resp., $|W \cap U_\ell| \leq 1$), for each index $\ell \in [k]$. We say that W is *elementary* if it is either V-elementary or U-elementary.

Observation 3.7 *For any elementary point set W, the sub-metric $\vartheta(W)$ of $\vartheta_{n,k}$ induced by W belongs to $\mathcal{P}_{|W|}$.*

Lemma 3.8 *Consider the metric $\vartheta_{n,k}$, for an arbitrary pair n, k of integers, such that $n \geq 2k \geq 4$, and let T be a Steiner tree of $\vartheta_{n,k}$ rooted at an arbitrary point $rt \in \vartheta_{n,k}$. Then: (1) If the root-stretch of (T, rt) is less than $1 + \frac{2}{k-1}$ then its weight $w(T)$ is at least $\frac{1}{2} \cdot k \cdot \log k - 2k$. (2) If the average root-stretch of (T, rt) is at most $1 + \frac{1}{2 \cdot (k-1)}$ then its weight $w(T)$ is at least $\frac{1}{4} \cdot k \cdot \log k - k$.*

Proof: Suppose without loss of generality that $rt \in U$.

To prove the first assertion, assume that the root-stretch of (T, rt) is less than $1 + \frac{2}{k-1}$, and consider an arbitrary V-elementary point set W of size k, for example, take $W = \{v_1^{(1)}, v_2^{(1)}, \ldots, v_k^{(1)}\}$. Observation 3.7 implies that $\vartheta(W)$ belongs to $\mathcal{P}_k$. By definition, the distance in $\vartheta_{n,k}$ between rt and any point in W is equal to $\frac{k-1}{2}$. Also, note that $diam(\vartheta(W)) \leq diam(\vartheta_{n,k}) = k-1$, and so $\frac{k-1}{2} \geq \frac{1}{2} \cdot diam(\vartheta(W))$. Since the root-stretch of (T, rt) is less than $1 + \frac{2}{k-1}$, it follows that (T, rt) is a $(\frac{k-1}{2}, \frac{2}{k-1})$-tree for $\vartheta(W)$. By Lemma 3.4, the weight $w(T)$ of T is at least $g(k, \frac{k-1}{2}) - 2k = (\frac{k-1}{2} - \frac{1}{2} \cdot (k-1) + \frac{1}{2} \cdot k \cdot \log k) - 2k = \frac{1}{2} \cdot k \cdot \log k - 2k$.

Next, we prove the second assertion of the lemma. Assume that the average root-stretch of (T, rt) is at most $1 + \frac{1}{2 \cdot (k-1)}$. Denote by V' the set of all points v in V, such that $Str_T(rt, v) < 1 + \frac{2}{k-1}$, and define $V'' = V \setminus V'$. Also, write $\mathcal{V}^* = \mathcal{V} \setminus \{rt\}$. Clearly, for each point v in $\mathcal{V}^* \setminus V''$, $Str_T(rt, v) \geq 1$, and for each point v in V'', $Str_T(rt, v) \geq 1 + \frac{2}{k-1}$. Observe that

$$\begin{aligned} 1 + \frac{1}{2 \cdot (k-1)} &\geq AvgStr(T, rt) = \frac{\sum_{v \in \mathcal{V}^*} Str_T(rt, v)}{n-1} \\ &= \frac{\sum_{v \in \mathcal{V}^* \setminus V''} Str_T(rt, v)}{n-1} + \frac{\sum_{v \in V''} Str_T(rt, v)}{n-1} \\ &\geq \frac{|\mathcal{V}^* \setminus V''|}{n-1} + \frac{|V''| \cdot \left(1 + \frac{2}{k-1}\right)}{n-1} = \frac{|\mathcal{V}^*|}{n-1} + \frac{|V''| \cdot \frac{2}{k-1}}{n-1} = 1 + \frac{|V''| \cdot \frac{2}{k-1}}{n-1}. \end{aligned}$$

implying that $|V''| \leq n/4$. Denote by I the set of all indices ℓ in $[k]$, such that $V' \cap V_\ell \neq \emptyset$. Observe that $V'' \supseteq \bigcup_{\ell \in [k] \setminus I} V_\ell$, and so

$$n/4 \geq |V''| \geq \sum_{\ell \in [k] \setminus I} |V_\ell| = \sum_{\ell \in [k] \setminus I} \frac{n}{2k} = (k - |I|) \cdot \frac{n}{2k}.$$

It follows that $|I| \geq k/2$. For each index $\ell \in I$, let $v(\ell)$ be an arbitrary point in $V' \cap V_\ell$. Consider the V-elementary point set $W = \{v(\ell) \mid \ell \in I\}$, and remove points from it arbitrarily until $|W| = k/2$. Note that $Str_T(rt, v) < 1 + \frac{2}{k-1}$, for each point $v \in W$, implying that (T, rt) is a Steiner tree for $\vartheta(W)$ with root-stretch less than $1 + \frac{2}{k-1}$. By Observation 3.7, $\vartheta(W)$ belongs to $\mathcal{P}_{k/2}$. Also, by definition, the distance in $\vartheta_{n,k}$ between rt and any point in W is equal to $\frac{k-1}{2}$. Finally, note that $diam(\vartheta(W)) \leq diam(\vartheta_{n,k}) = k-1$, and so $\frac{k-1}{2} \geq \frac{1}{2} \cdot diam(\vartheta(W))$. Consequently, (T, rt) is a $(\frac{k-1}{2}, \frac{2}{k-1})$-tree for $\vartheta(W)$. By Lemma 3.4, the weight $w(T)$ of T is at least

$$g\left(\frac{k}{2}, \frac{k-1}{2}\right) - 2 \cdot \frac{k}{2} = \left(\frac{k-1}{2} - \frac{1}{2} \cdot \left(\frac{k}{2} - 1\right) + \frac{1}{2} \cdot \frac{k}{2} \cdot \log\left(\frac{k}{2}\right)\right) - k = \frac{1}{4} \cdot k \cdot \log k - k. \quad \blacksquare$$

Observe that $\frac{5}{2} \cdot (k-1) \leq w(MST(\vartheta_{n,k})) \leq \frac{5}{2} \cdot (k-1) + \alpha \cdot (n-1) \leq \frac{5}{2} \cdot (k-1) + 1$. Substituting k with $\Theta(\frac{1}{\epsilon})$ in Lemma 3.8, we obtain the main result of this section.

Corollary 3.9 *For any integer n and any parameter $\frac{2}{n} \leq \epsilon \leq \frac{1}{2}$, every Steiner tree T of $\vartheta_{n,k}$ rooted at an arbitrary point $rt \in \vartheta_{n,k}$ that has average root-stretch at most $1 + \epsilon$ must have lightness at least $\Omega(\log \frac{1}{\epsilon})$, where $k = \lfloor \frac{1}{2\epsilon} \rfloor + 1$.*

Remarks: (1) Observe that $1 + \epsilon \leq 1 + \frac{1}{2 \cdot (k-1)}$. To apply Lemma 3.8, we need to have $n \geq 2k \geq 4$. Indeed, since $\epsilon \leq \frac{1}{2}$, it holds that $\frac{1}{2\epsilon} \geq 1$, and so $k = \lfloor \frac{1}{2\epsilon} \rfloor + 1 \geq \lfloor 1 \rfloor + 1 = 2$. Also, since $\frac{2}{n} \leq \epsilon \leq \frac{1}{2}$, we have $n \geq \frac{2}{\epsilon} = \frac{1}{\epsilon} + \frac{1}{\epsilon} \geq 2 \cdot \lfloor \frac{1}{2\epsilon} \rfloor + 2 = 2k \geq 4$. (2) The metric $\vartheta_{n,k}$ is not Euclidean. However, the same (up to constant factors) lower bound as the one established in this statement can be obtained also for Euclidean 2-dimensional metrics.

4 Lower Bounds for Euclidean Spanning SLTs

In this section we address three out of the four open questions posed by Khuller et al. [23]. More specifically, up to constant factors we settle these questions.

As before, let C_n denote a set of n points that are uniformly spaced around the boundary C of the unit circle, centered at the origin $(0,0)$, and define $\tilde{C}_{n+1} = C_n \cup \{(0,0)\}$.

4.1 (Average) root-stretch $1+\epsilon$ implies lightness $\Omega(\frac{1}{\epsilon})$

In this section we establish a tradeoff of $(1+\epsilon)$ versus $\Omega(\frac{1}{\epsilon})$ between the root-stretch and lightness of spanning SLTs. We show that this tradeoff holds for any choice of root-vertex, even if we consider average root-stretch rather than (worst-case) root-stretch.

The *root-degree* of a rooted tree (T, rt), denoted $\gamma(T, rt)$, is the degree of rt in T.

Lemma 4.1 *Let n be a sufficiently large integer, let T be a spanning tree for $\tilde{C}_{n+1}$ rooted at $rt = (0,0)$, and write $\gamma = \gamma(T, rt)$. Then: (1) $rtStr(T, rt) \geq 1 + (\frac{2}{\gamma} - \frac{2}{n})$. (2) $AvgStr(T, rt) \geq 1 + \left(\frac{1}{\gamma} - \frac{\gamma}{n^2}\right)$.*

Proof: Denote the γ neighbors of rt in clockwise order along C by $v_1, v_2, \ldots, v_\gamma$, and define $v_{\gamma+1} = v_1$. Let $S_i = (w_0^{(i)} = v_i, w_1^{(i)}, \ldots, w_{n_i}^{(i)} = v_{i+1})$ be the sequence of all points that we traverse when going from v_i to v_{i+1} along C in clockwise order, for each index $i \in [\gamma]$. We have $\sum_{i=1}^{\gamma} n_i = n$. Fix an arbitrary pair i, j of indices, $i \in [\gamma], j \in [0, n_i]$, and consider the point $w_j^{(i)}$. Since $n_i \leq n$, we have $\min\{j, n_i - j\} \leq \frac{n}{2}$, and so the circular distance between $w_j^{(i)}$ and at least one of the points v_i or v_{i+1} is at most π. Consequently,

$$\min\left\{\|v_i - w_j^{(i)}\|, \|v_{i+1} - w_j^{(i)}\|\right\} \;\geq\; \frac{2}{\pi} \cdot \min\left\{j \cdot \frac{2\pi}{n}, (n_i - j) \cdot \frac{2\pi}{n}\right\}.$$

Let $P = P\left(rt, w_j^{(i)}\right)$ be the path from rt to $w_j^{(i)}$ in T. It is easy to verify that

$$w(P) \;\geq\; 1 + \min\left\{\|v_i - w_j^{(i)}\|, \|v_{i+1} - w_j^{(i)}\|\right\} \;\geq\; 1 + \frac{2}{\pi} \cdot \min\left\{j \cdot \frac{2\pi}{n}, (n_i - j) \cdot \frac{2\pi}{n}\right\}.$$

Also, we have $\|rt - w_j^{(i)}\| = 1$. It follows that

$$Str_T\left(rt, w_j^{(i)}\right) \;=\; \frac{d_T(rt, w_j^{(i)})}{\|rt - w_j^{(i)}\|} \;=\; w(P) \;\geq\; 1 + \frac{2}{\pi} \cdot \min\left\{j \cdot \frac{2\pi}{n}, (n_i - j) \cdot \frac{2\pi}{n}\right\}. \tag{8}$$

Since $\sum_{i=1}^{\gamma} n_i = n$, there is an index $k \in [\gamma]$, such that $n_k \geq \frac{n}{\gamma}$. By Equation (8),

$$rtStr(T, rt) \;\geq\; Str_T\left(rt, w^{(k)}_{\lfloor \frac{n_k}{2} \rfloor}\right) \;\geq\; 1 + \frac{2}{\pi} \cdot \left\lfloor \frac{n_k}{2} \right\rfloor \cdot \frac{2\pi}{n} \;\geq\; 1 + 2 \cdot \left(\frac{n_k - 1}{n}\right) \geq 1 + \left(\frac{2}{\gamma} - \frac{2}{n}\right),$$

which proves the first assertion of the lemma.

Next, we bound the average root-stretch of T.

We argue that

$$\sum_{j=0}^{n_i - 1} Str_T\left(rt, w_j^{(i)}\right) \;\geq\; n_i + \frac{n_i^2 - 1}{n}. \tag{9}$$

We restrict the attention to the case when n_i is odd. (The case when n_i is even can be handled similarly.) By Equation (8),

$$\begin{aligned}\sum_{j=0}^{n_i-1} Str_T\left(rt, w_j^{(i)}\right) &\geq \sum_{j=0}^{n_i-1}\left(1+\frac{2}{\pi}\cdot\min\left\{j\cdot\frac{2\pi}{n},(n_i-j)\cdot\frac{2\pi}{n}\right\}\right)\\ &= n_i+\frac{4}{\pi}\cdot\left(\sum_{j=1}^{(n_i-1)/2} j\cdot\frac{2\pi}{n}\right) \;=\; n_i+\frac{8}{n}\cdot\left(\sum_{j=1}^{(n_i-1)/2} j\right)\\ &= n_i+\frac{8}{n}\cdot\left(\frac{n_i^2-1}{8}\right) \;=\; n_i+\frac{n_i^2-1}{n},\end{aligned}$$

which proves Equation (9). Recall that $\sum_{i=1}^{\gamma} n_i = n$. Hence, by the Cauchy-Schwarz inequality, $\sum_{i=1}^{\gamma} n_i^2 \geq \frac{n^2}{\gamma}$. By Equation (9),

$$\begin{aligned}\sum_{v\in\tilde{C}_{n+1}\setminus\{rt\}} Str_T(rt,v) &= \sum_{i=1}^{\gamma}\sum_{j=0}^{n_i-1} Str_T\left(rt,w_j^{(i)}\right) \;\geq\; \sum_{i=1}^{\gamma}\left(n_i+\frac{n_i^2-1}{n}\right)\\ &= n+\frac{1}{n}\cdot\sum_{i=1}^{\gamma}(n_i^2-1) \;=\; n+\frac{1}{n}\cdot\sum_{i=1}^{\gamma} n_i^2-\frac{\gamma}{n}\\ &\geq n+\frac{1}{n}\cdot\frac{n^2}{\gamma}-\frac{\gamma}{n} \;=\; n+\frac{n}{\gamma}-\frac{\gamma}{n}.\end{aligned}$$

It follows that

$$AvgStr(T,rt) \;=\; \frac{\sum_{v\in\tilde{C}_{n+1}\setminus\{rt\}} Str_T(rt,v)}{n} \;\geq\; \frac{n+\frac{n}{\gamma}-\frac{\gamma}{n}}{n} \;=\; 1+\left(\frac{1}{\gamma}-\frac{\gamma}{n^2}\right),$$

which proves the second assertion of the lemma. Lemma 4.1 follows. ▮

Lemma 4.1 implies the following corollary.

Corollary 4.2 *Let n be a sufficiently large integer, and let T be an arbitrary spanning tree for $\tilde{C}_{n+1}$ rooted at $rt=(0,0)$. If the lightness $\Psi(T)$ of T it at most α, for an arbitrary parameter $\alpha\geq 1$, then (1) $rtStr(T,rt)\geq 1+\left(\frac{2}{(2\pi+1)\cdot\alpha}-\frac{2}{n}\right)$, and (2) $AvgStr(T,rt)\geq 1+\left(\frac{1}{(2\pi+1)\cdot\alpha}-\frac{(2\pi+1)\cdot\alpha}{n^2}\right)$.*

Remark: For $\alpha\leq\frac{n}{2(2\pi+1)}$, we obtain (1) $rtStr(T,rt)\geq 1+\frac{1}{(2\pi+1)\cdot\alpha}=1+\Omega\left(\frac{1}{\alpha}\right)$, and (2) $AvgStr(T,rt)\geq 1+\frac{3}{4\cdot(2\pi+1)\cdot\alpha}=1+\Omega\left(\frac{1}{\alpha}\right)$.

Proof: Note that the weight of every edge in T that is incident to rt is equal to 1. Hence, $w(T)\geq\gamma=\gamma(T,rt)$. Also, the weight $w(MST(\tilde{C}_{n+1}))$ of the MST for $\tilde{C}_{n+1}$ is smaller than $2\pi+1$. Consequently,

$$\alpha \;\geq\; \Psi(T) \;=\; \frac{w(T)}{w(MST(\tilde{C}_{n+1}))} \;\geq\; \frac{\gamma}{2\pi+1},$$

and so $\gamma\leq(2\pi+1)\cdot\alpha$. By Lemma 4.1,

$$rtStr(T,rt) \;\geq\; 1+\left(\frac{2}{\gamma}-\frac{2}{n}\right) \;\geq\; 1+\left(\frac{2}{(2\pi+1)\cdot\alpha}-\frac{2}{n}\right)$$

and

$$AvgStr(T,rt) \;\geq\; 1+\left(\frac{1}{\gamma}-\frac{\gamma}{n^2}\right) \;\geq\; 1+\left(\frac{1}{(2\pi+1)\cdot\alpha}-\frac{(2\pi+1)\cdot\alpha}{n^2}\right). \quad ▮$$

Corollary 4.2 provides a near-optimal tradeoff between the lightness and (average) root-stretch of spanning SLTs. However, this tradeoff holds for a specific root vertex. Next, we demonstrate that the same tradeoff also holds when the root vertex is selected at will.

Lemma 4.3 *Let n be a sufficiently large integer, let rt be an arbitrary point in C_n, and let T be a spanning tree for C_n rooted at rt. If the lightness of T is at most α, for an arbitrary parameter $1 \leq \alpha \leq \frac{1}{8\pi} \cdot n$, then the average root-stretch of (T, rt) is at least $1 + \Omega\left(\frac{1}{\alpha}\right)$.*

Proof: We suppose without loss of generality that rt is the bottom-most point of the circle C, i.e., $rt = (0, -1)$. Let U (respectively, D) denote the set of all points in C with y-coordinate at least $\frac{1}{2}$ (resp., at most $-\frac{1}{2}$). Also, let $V_U = U \cap C_n$ (respectively, $V_D = D \cap C_n$) be the set of all points in C_n with y-coordinate at least $\frac{1}{2}$ (resp., at most $-\frac{1}{2}$). An edge e of T is called *semi-vertical* if it connects a point in V_U with a point in V_D. Denote by SV the set of all semi-vertical edges of T. Clearly, for any semi-vertical edge $e = (u, v)$ in T, $w(e) = \|u, v\| \geq 1$, and so the weight $w(T)$ of T satisfies $w(T) \geq \sum_{e \in SV} w(e) \geq |SV|$. Since the weight $w(MST(C_n))$ of the MST for C_n is smaller than 2π, it follows that

$$\alpha \;\geq\; \Psi(T) \;=\; \frac{w(T)}{w(MST(C_n))} \;\geq\; \frac{|SV|}{2\pi},$$

and so $|SV| \leq 2\pi \cdot \alpha$. Denote the endpoints of the edges in SV that belong to V_U in clockwise order by $u_1, u_2, \ldots, u_s$, and observe that $s \leq |SV|$. Let u_0 and u_{s+1} be the points of V_U that are closest to the left-most point $(-\frac{\sqrt{3}}{2}, \frac{1}{2})$ and the right-most point $(\frac{\sqrt{3}}{2}, \frac{1}{2})$ of U, respectively. Since the length of the arc U is equal to $\frac{2\pi}{3}$, it follows that $|V_U| \approx \frac{n}{3}$. In what follows we assume for simplicity that $|V_U| = \frac{n}{3}$, $u_0 \neq u_1$ and $u_s \neq u_{s+1}$; the general case is handled similarly. (See Figure 7 for an illustration.)

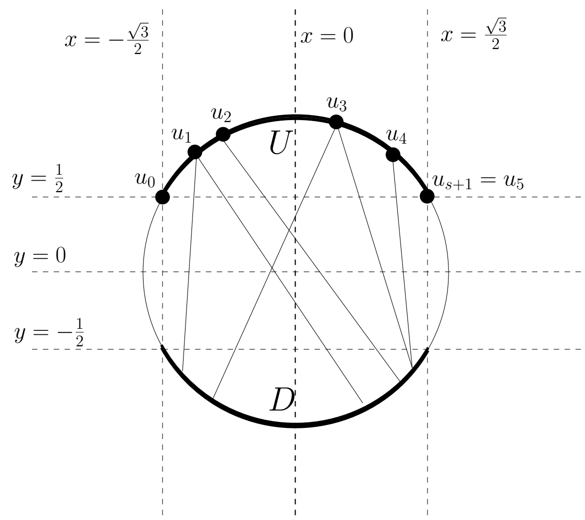

Figure 7: An illustration of the circle C and the semi-vertical edges of SV. The two arcs U and D of C are marked by bold lines. The six semi-vertical edges of SV in the figure are depicted by solid lines, whereas the four endpoints u_1, u_2, u_3 and u_4 of these edges that belong to V_U, as well as the two boundary points u_0 and $u_{s+1} = u_5$ of V_U, are depicted by black dots.

Let $S_i = (w_0^{(i)} = u_i, w_1^{(i)}, \ldots, w_{n_i}^{(i)} = u_{i+1})$ be the sequence of all points that we traverse when going from u_i to u_{i+1} along U in clockwise order, for each index $i \in [0, s]$. Note also that $\sum_{i=0}^{s} n_i = |V_U| = \frac{n}{3}$. Fix an arbitrary pair i, j of indices, $i \in [0, s], j \in [0, n_i]$, and consider the point $w_j^{(i)}$. Since the length of

U is $\frac{2}{3} \cdot \pi$, it follows that the circular distance between $w_j^{(i)}$ and at least one of the points u_i or u_{i+1} is at most $\frac{\pi}{3}$. Consequently, we get

$$\min\left\{\|u_i - w_j^{(i)}\|, \|u_{i+1} - w_j^{(i)}\|\right\} \;\geq\; \frac{3}{\pi} \cdot \min\left\{j \cdot \frac{2\pi}{n}, (n_i - j) \cdot \frac{2\pi}{n}\right\}.$$

Let $P = P\left(rt, w_j^{(i)}\right)$ be the path from rt to $w_j^{(i)}$ in T. Next, we show that the weight $w(P)$ of P satisfies

$$w(P) \;\geq\; \|rt - w_j^{(i)}\| + c \cdot \min\left\{j \cdot \frac{2\pi}{n}, (n_i - j) \cdot \frac{2\pi}{n}\right\}, \tag{10}$$

for some constant $0 < c < 1$ to be determined later. It is easy to verify that Equation (10) holds true if P contains a point in $C_n \setminus (V_U \cup V_D)$. We restrict the attention to the complementary case where all points in P belong to $V_U \cup V_D$, which is more difficult. In this case P must contain at least one semi-vertical edge e. Denote the endpoint of e that belongs to V_U by u, and note that $u = u_\ell$, for some index $\ell \in [s]$. By the triangle inequality, $w(P) \geq \|rt - u\| + \|u - w_j^{(i)}\|$. Simple geometric considerations imply that

$$\|rt - u\| + \|u - w_j^{(i)}\| \;\geq\; \|rt - w_j^{(i)}\| + \tilde{c} \cdot \|u - w_j^{(i)}\|,$$

for an appropriate constant $0 < \tilde{c} < 1$. Also, observe that

$$\|u - w_j^{(i)}\| \;\geq\; \min\left\{\|u_i - w_j^{(i)}\|, \|u_{i+1} - w_j^{(i)}\|\right\}.$$

It follows that

$$\begin{aligned} w(P) \;&\geq\; \|rt - u\| + \|u - w_j^{(i)}\| \;\geq\; \|rt - w_j^{(i)}\| + \tilde{c} \cdot \|u - w_j^{(i)}\| \\ &\geq\; \|rt - w_j^{(i)}\| + \tilde{c} \cdot \min\left\{\|u_i - w_j^{(i)}\|, \|u_{i+1} - w_j^{(i)}\|\right\} \\ &\geq\; \|rt - w_j^{(i)}\| + \left(\tilde{c} \cdot \frac{3}{\pi}\right) \cdot \min\left\{j \cdot \frac{2\pi}{n}, (n_i - j) \cdot \frac{2\pi}{n}\right\}, \end{aligned}$$

for some constant $0 < c = \left(\tilde{c} \cdot \frac{3}{\pi}\right) < 1$, which proves Equation (10). Observe that $\|rt - w_j^{(i)}\| \leq 2$ and define $c' = \frac{c}{2}$. It follows that

$$\begin{aligned} Str_T\left(rt, w_j^{(i)}\right) \;&=\; \frac{d_T\left(rt, w_j^{(i)}\right)}{\|rt - w_j^{(i)}\|} \;=\; \frac{w\left(P\right)}{\|rt - w_j^{(i)}\|} \;\geq\; \frac{\|rt - w_j^{(i)}\| + c \cdot \min\left\{j \cdot \frac{2\pi}{n}, (n_i - j) \cdot \frac{2\pi}{n}\right\}}{\|rt - w_j^{(i)}\|} \\ &\geq\; 1 + c' \cdot \min\left\{j \cdot \frac{2\pi}{n}, (n_i - j) \cdot \frac{2\pi}{n}\right\}. \end{aligned} \tag{11}$$

Next, we argue that

$$\sum_{j=0}^{n_i - 1} Str_T\left(rt, w_j^{(i)}\right) \;\geq\; n_i + \frac{\pi \cdot c' \cdot (n_i^2 - 1)}{2n}. \tag{12}$$

We restrict the attention to the case when n_i is odd. (The case when n_i is even can be handled similarly.) By Equation (11),

$$\sum_{j=0}^{n_i - 1} Str_T\left(rt, w_j^{(i)}\right) \;\geq\; \sum_{j=0}^{n_i - 1} \left(1 + c' \cdot \min\left\{j \cdot \frac{2\pi}{n}, (n_i - j) \cdot \frac{2\pi}{n}\right\}\right)$$

$$= \quad n_i + 2 \cdot c' \cdot \left(\sum_{j=1}^{(n_i-1)/2} j \cdot \frac{2\pi}{n} \right) \quad = \quad n_i + \frac{4\pi \cdot c'}{n} \cdot \left(\sum_{j=1}^{(n_i-1)/2} j \right)$$

$$= \quad n_i + \frac{4\pi \cdot c'}{n} \cdot \left(\frac{n_i^2 - 1}{8} \right) \quad = \quad n_i + \frac{\pi \cdot c' \cdot (n_i^2 - 1)}{2n},$$

which proves Equation (12).

Recall that $\sum_{i=0}^{s} n_i = \frac{n}{3}$ and $s \leq |SV| \leq 2\pi \cdot \alpha$. Hence, by the Cauchy-Schwarz inequality,

$$\sum_{i=0}^{s} n_i^2 \quad \geq \quad \frac{\left(\sum_{i=0}^{s} n_i\right)^2}{s+1} \quad = \quad \frac{n^2}{9 \cdot (s+1)} \quad \geq \quad \frac{n^2}{9 \cdot (2\pi \cdot \alpha + 1)}. \tag{13}$$

By Equations (12) and (13),

$$\begin{aligned}
\sum_{v \in V_U} Str_T(rt, v) \quad &= \quad \sum_{i=0}^{s} \sum_{j=0}^{n_i - 1} Str_T\left(rt, w_j^{(i)}\right) \quad \geq \quad \sum_{i=0}^{s} \left(n_i + \frac{\pi \cdot c' \cdot (n_i^2 - 1)}{2n} \right) \\
&= \quad \frac{n}{3} + \frac{\pi \cdot c'}{2n} \cdot \sum_{i=0}^{s} (n_i^2 - 1) \quad = \quad \frac{n}{3} + \frac{\pi \cdot c'}{2n} \cdot \sum_{i=0}^{s} n_i^2 - \frac{\pi \cdot c' \cdot (s+1)}{2n} \\
&\geq \quad \frac{n}{3} + \frac{\pi \cdot c'}{2n} \cdot \frac{n^2}{9 \cdot (2\pi \cdot \alpha + 1)} - \frac{\pi \cdot c' \cdot (2\pi \cdot \alpha + 1)}{2n} \quad \geq \quad \frac{n}{3} + \Omega\left(\frac{n}{\alpha}\right).
\end{aligned}$$

(The last inequality holds for all $\alpha \leq \frac{1}{8\pi} \cdot n$.) Notice that the point set $W = (C_n \setminus \{rt\}) \setminus V_U$ consists of $\frac{2n}{3} - 1$ points, and for every point $v \in W$, we have $Str_T(rt, v) \geq 1$. Hence, $\sum_{v \in W} Str_T(rt, v) \geq \frac{2n}{3} - 1$. It follows that

$$\begin{aligned}
AvgStr(T, rt) \quad &= \quad \frac{\sum_{v \in C_n \setminus \{rt\}} Str_T(rt, v)}{n-1} \quad = \quad \frac{\sum_{v \in V_U} Str_T(rt, v)}{n-1} + \frac{\sum_{v \in W} Str_T(rt, v)}{n-1} \\
&\geq \quad \frac{\frac{n}{3} + \Omega\left(\frac{n}{\alpha}\right)}{n-1} + \frac{\frac{2n}{3} - 1}{n-1} \quad = \quad 1 + \Omega\left(\frac{1}{\alpha}\right). \qquad \blacksquare
\end{aligned}$$

4.2 Lightness $1 + \epsilon$ implies (average) root-stretch $\Omega(\frac{1}{\epsilon})$

In this section we establish a tradeoff of $(1 + \epsilon)$ versus $\Omega(\frac{1}{\epsilon})$ between the lightness and root-stretch of spanning SLTs. We show that this tradeoff holds for any choice of root-vertex, even if we consider average root-stretch rather than (worst-case) root-stretch.

Let $k \geq 1$ be an arbitrary integer and let $\epsilon < 1$ be a sufficiently small number. For each index $i \in [k+1]$, define h_i to be the horizontal segment connecting point $(0, (i-1) \cdot \epsilon)$ with point $(1, (i-1) \cdot \epsilon)$. Also, for each index $i \in [\lceil k/2 \rceil]$, define r_i to be the vertical segment connecting point $(1, (2i-2) \cdot \epsilon)$ with point $(1, (2i-1) \cdot \epsilon)$, and for each index $i \in [\lfloor k/2 \rfloor]$, define l_i to be the vertical segment connecting point $(0, (2i-1) \cdot \epsilon)$ with point $(0, 2i \cdot \epsilon)$. Let $C_{k,\epsilon}$ be the 2-dimensional (ladder) curve obtained by concatenating $h_1 \circ r_1 \circ h_2 \circ l_1 \circ \ldots \circ r_{k/2} \circ h_k \circ l_{k/2} \circ h_{k+1}$, if k is even, and $h_1 \circ r_1 \circ h_2 \circ l_1 \circ \ldots \circ r_{(k-1)/2} \circ h_{k-1} \circ l_{(k-1)/2} \circ h_k \circ r_{(k+1)/2} \circ h_{k+1}$, if k is odd. (See Figure 8 for an illustration.)

Lemma 4.4 *Let n be a sufficiently large integer and ϵ be a sufficiently small number, such that $\frac{4}{n-4} < \epsilon < 1$, and write $\delta_1 = \delta_1(n, \epsilon) = \frac{2+\epsilon}{\lceil n/2 \rceil - 1}$. There exists a set $S_1' = S_{1,\epsilon}'(n)$ of n points that lie on the curve $C_{1,\epsilon}$ and a designated point $rt \in S_1'$, such that every spanning tree T of S_1' rooted at rt with lightness less than $1 + \frac{\epsilon - \delta_1}{2+\epsilon} \approx 1 + \frac{\epsilon}{2}$ must have root-stretch at least $\frac{2}{\epsilon}$ and average root-stretch at least $\frac{1}{\epsilon}$.*

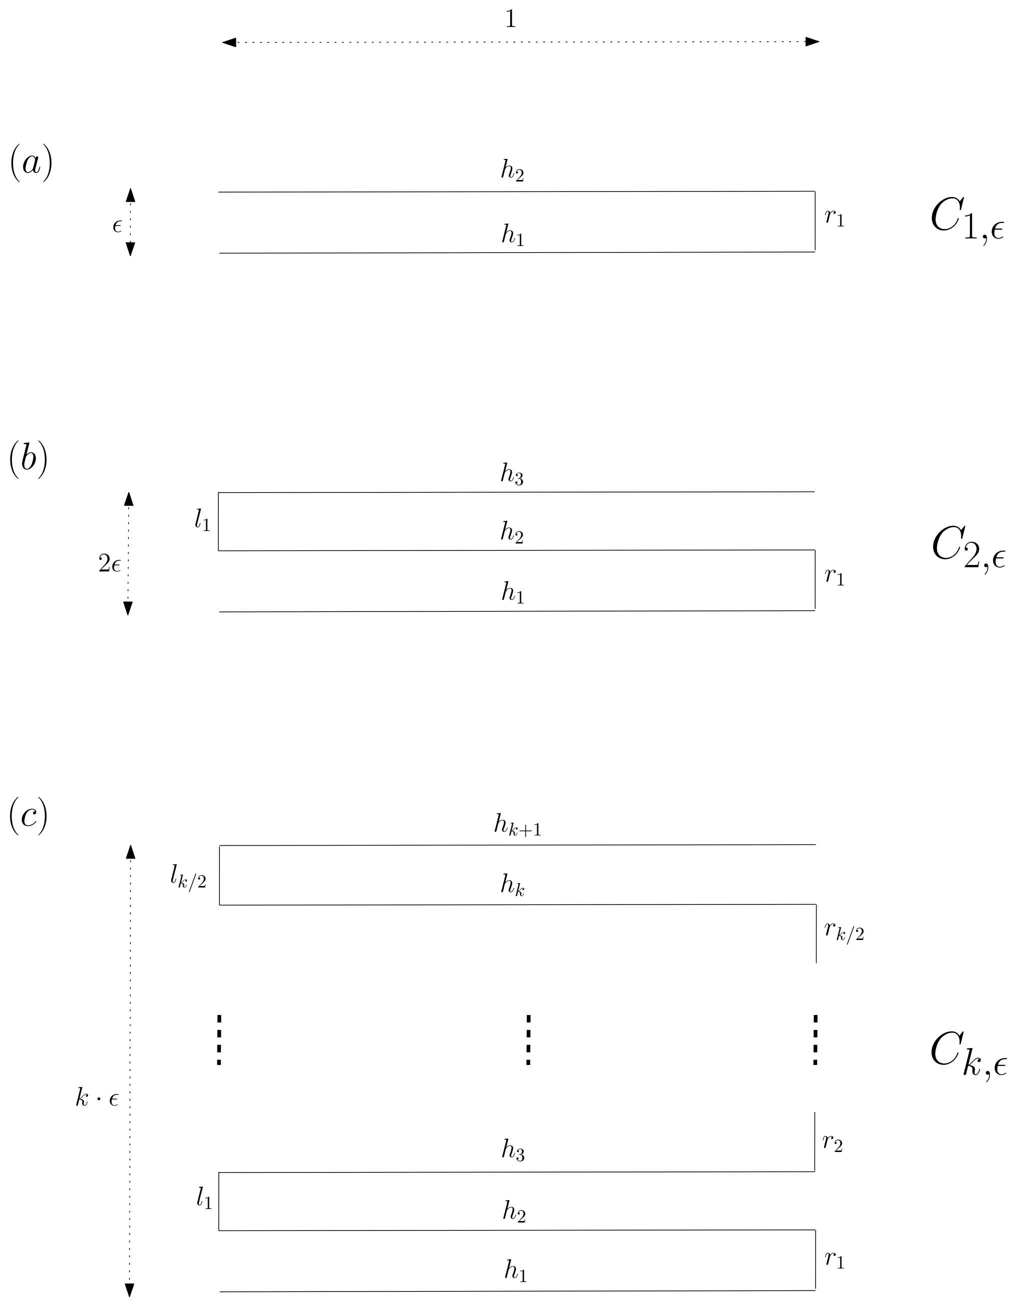

Figure 8: An illustration of the curve $C_{k,\epsilon}$, when (a) $k = 1$, (b) $k = 2$, and (c) k is an arbitrary even integer.

Proof: Denote the endpoints $(0,0)$ and $(0,\epsilon)$ of the curve $C_{1,\epsilon}$ by A and B, respectively. Let $S_1 = S_{1,\epsilon}(\lceil n/2 \rceil)$ be a set of $\lceil n/2 \rceil$ points that are distributed uniformly on the curve $C_{1,\epsilon}$, so that (1) both points A and B belong to S_1, and (2) the distance between any pair of consecutive points of S_1 along $C_{1,\epsilon}$ is equal to δ_1. Since $\frac{4}{n-4} < \epsilon < 1$, it follows that $\delta_1 < \epsilon$, and so the distance between any pair of consecutive points of S_1 along $C_{1,\epsilon}$ is strictly smaller than the length ϵ of the vertical segment r_1. Consequently, the minimum distance between a pair of points in S_1 is equal to δ_1, and the weight of the MST for S_1 is equal to the perimeter $2+\epsilon$ of $C_{1,\epsilon}$. Let $S' = S'_\epsilon(\lfloor n/2 \rfloor)$ be a set of $\lfloor n/2 \rfloor$ points that reside arbitrarily close to point $B = (0,\epsilon)$. For the sake of simplicity, we view all points of S' as copies of point B. Let $S'_1 = S'_{1,\epsilon}(n)$ be the n-point set obtained from the union of S_1 and S', i.e., $S'_1 = S_1 \cup S'$.

Let T be a spanning tree of S'_1 rooted at $rt = A = (0,0) \in S'_1$, with $\Psi(T) < 1 + \frac{\epsilon - \delta_1}{2+\epsilon}$. First, we argue that there are no edges in T of weight at least ϵ. Indeed, otherwise we have $w(T) \geq 2 + \epsilon + (\epsilon - \delta_1)$, and so

$$\Psi(T) \;=\; \frac{w(T)}{2+\epsilon} \;\geq\; \frac{2+\epsilon+(\epsilon-\delta_1)}{2+\epsilon} \;=\; 1 + \frac{\epsilon-\delta_1}{2+\epsilon},$$

a contradiction.

Observe that $B \in S'_1$. Moreover, $B \in S_1$, and by our assumption above, S' contains $\lfloor n/2 \rfloor$ additional copies of this point, denoted $B'_1, B'_2, \ldots, B'_{\lfloor n/2 \rfloor}$. Also, notice that $\|rt - B\| = \epsilon$. Consider the path $P(rt, B) = (v_1 = rt, v_2, \ldots, v_m = B)$ from rt to B in T. Let v_i be the last vertex in $P(rt, B)$ with y-coordinate smaller than that of h_2. Since there are no edges in T of weight at least ϵ, it follows that v_i belongs to r_1. Hence, by the triangle inequality, the weights of the sub-paths $P(rt, v_i)$ and $P(v_i, B)$ of $P(rt, B)$ from rt to v_i and from v_i to B are at least $\|rt - v_i\| \geq 1$ and $\|v_i - B\| \geq 1$, respectively. Consequently,

$$d_T(rt, B) \;=\; w(P(rt,B)) \;=\; w(P(rt,v_i)) + w(P(v_i,B)) \;\geq\; \|rt - v_i\| + \|v_i - B\| \;\geq\; 2,$$

implying that

$$rtStr(T, rt) \;\geq\; Str_T(rt, B) \;=\; \frac{d_T(rt,B)}{\|rt - B\|} \;\geq\; \frac{2}{\epsilon}.$$

Also, for each point $B_i \in S'$, $Str_T(rt, B_i) = Str_T(rt, B) \geq \frac{2}{\epsilon}$. Since $|S'| = \lfloor n/2 \rfloor$, it follows that

$$\begin{aligned} AvgStr(T, rt) \;&=\; \frac{\sum_{v \in S'_1 \setminus \{rt\}} Str_T(rt, v)}{n-1} \;\geq\; \frac{Str_T(rt,B)}{n-1} + \frac{\sum_{B_i \in S'} Str_T(rt, B_i)}{n-1} \\ &\geq\; \frac{\frac{2}{\epsilon}}{n-1} + \frac{\lfloor n/2 \rfloor \cdot \frac{2}{\epsilon}}{n-1} \;\geq\; \frac{1}{\epsilon}. \quad \blacksquare \end{aligned}$$

The next lemma strengthens Lemma 4.4 in the following sense. Lemma 4.4 provides an example in which there exists a designated root vertex rt such that any spanning tree with lightness at most $1+\epsilon$ has root-stretch $\Omega(\frac{1}{\epsilon})$ with respect to the vertex rt. In the following lemma we provide an example in which *for any* choice of the root vertex rt, the above statement holds. On the other hand, while Lemma 4.4 achieves the above tradeoff also for *average* root-stretch, the following lemma only achieves it for (worst-case) root-stretch.

Lemma 4.5 *Let n be a sufficiently large integer and ϵ be a sufficiently small number, such that $\frac{3}{n-3} < \epsilon < 1$, and write $\delta_2 = \delta_2(n,\epsilon) = \frac{3+2\epsilon}{n-1}$. Also, let $S_2 = S_{2,\epsilon}(n)$ be a set of n points that are distributed uniformly on the curve $C_{2,\epsilon}$, so that the distance between any pair of consecutive points of S_2 along $C_{2,\epsilon}$ is equal to δ_2. Then for any designated point $rt \in S_2$, every rooted spanning tree (T, rt) of S_2 with lightness less than $1 + \frac{\epsilon - \delta_2}{3+2\epsilon} \approx 1 + \frac{\epsilon}{3}$ must have root-stretch at least $\frac{2-\delta_2}{2\epsilon+\delta_2} \approx \frac{1}{\epsilon}$.*

Proof: Since $\frac{3}{n-3} < \epsilon < 1$, it follows that $\delta_2 < \epsilon$, and so the distance between any pair of consecutive points of S_2 along $C_{2,\epsilon}$ is strictly smaller than the length ϵ of the vertical segments r_1 and l_1. Consequently, the minimum distance between a pair of points in S_2 is equal to δ_2, and the weight of the MST for S_2 is equal to the perimeter $3 + 2\epsilon$ of $C_{2,\epsilon}$.

Let T be a spanning tree of S_2 rooted at an arbitrary designated point $rt \in S_2$, with $\Psi(T) < 1 + \frac{\epsilon - \delta_2}{3+2\epsilon}$. First, we argue that there are no edges in T of weight at least ϵ. Indeed, otherwise we have $w(T) \geq 3 + 2\epsilon + (\epsilon - \delta_2)$, and so

$$\Psi(T) \;=\; \frac{w(T)}{3+2\epsilon} \;\geq\; \frac{3+2\epsilon+(\epsilon-\delta_2)}{3+2\epsilon} \;=\; 1 + \frac{\epsilon-\delta_2}{3+2\epsilon},$$

a contradiction.

For each index $i \in [3]$, decompose h_i into two sub-segments h_i^L and h_i^R, where h_i^L connects point $(0, (i-1)\cdot\epsilon)$ with point $(\frac{1}{2}, (i-1)\cdot\epsilon)$ and h_i^R connects point $(\frac{1}{2}, (i-1)\cdot\epsilon)$ with point $(1, (i-1)\cdot\epsilon)$. Next, we define a mapping $f : S_2 \to S_2$. Consider an arbitrary point $v \in S_2$. If v belongs to h_1^L (respectively, h_1^R), then $f(v)$ is defined to be the point in $h_3^L \cap S_2$ (resp., $h_3^R \cap S_2$) closest to v, and vice versa, if v belongs to h_3^L (resp., h_3^R), then $f(v)$ is defined to be the point in $h_1^L \cap S_2$ (resp., $h_1^R \cap S_2$) closest to v. Also, if v belongs to h_2^L (respectively, h_2^R), then $f(v)$ is defined to be the point in $h_1^L \cap S_2$ (resp., $h_3^R \cap S_2$) closest to v. Finally, if v belongs to r_1 (respectively, l_1), then $f(v)$ is defined to be the point in $h_3^R \cap S_2$ (resp., $h_1^L \cap S_2$) closest to v. (See Figure 9 for an illustration.)

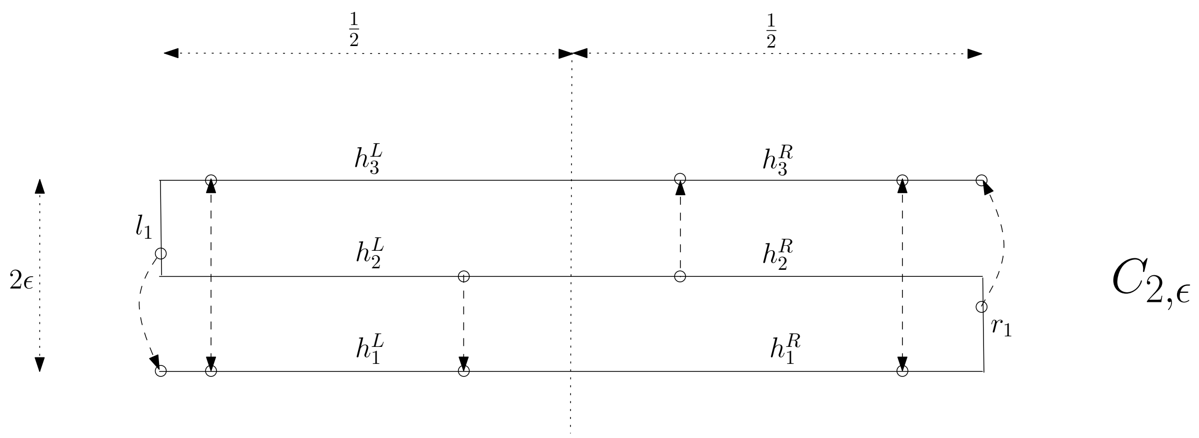

Figure 9: An illustration of the mapping f.

Next, we show that for any point $v \in S_2$, $Str_T(v, f(v)) \geq \frac{2-\delta_2}{2\epsilon+\delta_2}$. In particular, by substituting v with rt we will get $rtStr(T, rt) \geq Str_T(rt, f(rt)) \geq \frac{2-\delta_2}{2\epsilon+\delta_2}$, which concludes the proof. We restrict the attention to the case when v belongs to h_1^L. The other cases can be handled similarly. By definition, $f(v)$ is the point in $h_3^L \cap S_2$ that is closest to v. Hence, we have $\|v - f(v)\| \leq 2\epsilon + \delta_2$. Consider the path $P(v, f(v)) = (v_1 = v, v_2, \ldots, v_m = f(v))$ from v to $f(v)$ in T. Let v_i be the last vertex in $P(v, f(v))$ with y-coordinate smaller than that of h_2, and let v_j be the first vertex that comes after v_i in $P(v, f(v))$ and has y-coordinate greater than that of h_2. Since there are no edges in T of weight at least ϵ, it follows that v_i belongs to r_1 and v_j belongs to l_1. Hence, by the triangle inequality, the weight of the sub-path $P(v_i, v_j)$ of $P(v, f(v))$ from v_i to v_j is at least $\|v_i - v_j\| \geq 1$. Also, the weights of the sub-paths $P(v, v_i)$ and $P(v_j, f(v))$ of $P(v, f(v))$ from v to v_i and from v_j to $f(v)$ are at least $\|v - v_i\|$ and $\|v_j - f(v)\|$, respectively. Notice that $\|v - v_i\| + \|v_j - f(v)\| \geq 1 - \delta_2$. It follows that

$$\begin{aligned} d_T(v, f(v)) \;&=\; w(P(v, f(v))) \;=\; w(P(v, v_i)) + w(P(v_i, v_j)) + w(P(v_j, f(v))) \\ &\geq\; \|v - v_i\| + \|v_i - v_j\| + \|v_j - f(v)\| \;\geq\; 2 - \delta_2, \end{aligned}$$

and so

$$Str_T(v, f(v)) \;=\; \frac{d_T(v, f(v))}{\|v - f(v)\|} \;\geq\; \frac{2-\delta_2}{2\epsilon+\delta_2}. \qquad \blacksquare$$

The following lemma provides an example that achieves the same (up to constant factors) tradeoff between the lightness and average root-stretch as Lemma 4.4, for *every* choice of root-vertex. Therefore it is stronger than both Lemma 4.4 and Lemma 4.5.

Lemma 4.6 *Let n be a sufficiently large integer, let $\epsilon = \Omega(\frac{1}{\sqrt{n}})$ be a sufficiently small number, define $k = \lceil \frac{1}{\epsilon} \rceil$, and suppose that $k \cdot (k+2) \leq \frac{n-1}{14}$. Also, let $S = S_{k,\epsilon}(n)$ be a set of n points that are distributed uniformly on the curve $C = C_{k,\epsilon}$. Then for every designated point $rt \in S$, any rooted spanning tree (T, rt) for S with lightness at most $1 + \frac{\epsilon}{14}$ has average root-stretch $\Omega(\frac{1}{\epsilon})$.*

Proof: In what follows we assume for simplicity that $\frac{1}{\epsilon}$ is an even number, and so $k = \frac{1}{\epsilon}$ is even as well. The general case is handled similarly. Observe that the perimeter of C is equal to $k + 1 + k \cdot \epsilon = k + 2$. Since the points of S are distributed uniformly on C, the distance between any pair of consecutive points of S along C is $\frac{k+2}{n-1}$. Also, since $k \cdot (k+2) \leq \frac{n-1}{14}$, it follows that $\frac{k+2}{n-1} \leq \frac{1}{14} \cdot \epsilon < \epsilon$, and so the distance between any pair of consecutive points of S along C is strictly smaller than the length ϵ of each of the vertical segments l_i and r_i, $i \in [k/2]$. Consequently, the minimum distance between a pair of points in S is equal to $\frac{k+2}{n-1}$, and the weight of the MST for S is equal to the perimeter $k+2$ of C.

Let T be a spanning tree of S rooted at an arbitrary point $rt \in S$, with $\Psi(T) \leq 1 + \frac{\epsilon}{14}$. An edge e of T is called *heavy* if its weight $w(e)$ is at least ϵ. Let H be the set of all heavy edges in T. For an edge $e \in E(T)$, define $f(e) = w(e) - \frac{k+2}{n-1}$. Since the minimum distance between a pair of points in S is equal to $\frac{k+2}{n-1}$, it follows that $f(e) \geq 0$, for each edge $e \in E(T)$. Also, since $\frac{k+2}{n-1} \leq \frac{1}{14} \cdot \epsilon$, it follows that $f(e) \geq \frac{13}{14} \cdot w(e)$, for each edge $e \in H$. Consequently,

$$\begin{aligned} w(T) \;&=\; \sum_{e \in E(T)} w(e) \;=\; \sum_{e \in E(T)} \left(f(e) + \frac{k+2}{n-1} \right) \;=\; \left(\sum_{e \in E(T)} f(e) \right) + (k+2) \\ &\geq\; \left(\sum_{e \in H} f(e) \right) + (k+2) \;\geq\; \frac{13}{14} \cdot \left(\sum_{e \in H} w(e) \right) + (k+2). \end{aligned} \tag{14}$$

Claim 4.7 *(1) $\sum_{e \in H} w(e) \leq \frac{1}{13} + \frac{2\epsilon}{13}$. (2) $|H| \leq \frac{1}{13} \cdot (k+2)$.*

Proof: Equation (14) implies that

$$1 + \frac{\epsilon}{14} \;\geq\; \Psi(T) \;=\; \frac{w(T)}{k+2} \;\geq\; 1 + \frac{13}{14 \cdot (k+2)} \cdot \left(\sum_{e \in H} w(e) \right). \tag{15}$$

It follows that

$$\sum_{e \in H} w(e) \;\leq\; \frac{\epsilon \cdot (k+2)}{13} \;=\; \frac{1}{13} + \frac{2\epsilon}{13},$$

which proves the first assertion of this claim.

Also, by Equation (15),

$$\frac{\epsilon}{14} \;\geq\; \frac{13}{14 \cdot (k+2)} \cdot \left(\sum_{e \in H} w(e) \right) \;\geq\; \frac{13}{14 \cdot (k+2)} \cdot |H| \cdot \epsilon,$$

and so $|H| \leq \frac{1}{13} \cdot (k+2)$. Claim 4.7 follows. ▮

We say that an horizontal segment h_i is *dirty*, $i \in [k+1]$, if it intersects at least one edge in H. Otherwise it is *clean*

Claim 4.8 *(1) The number of dirty segments is bounded above by* $\frac{2}{13} \cdot (k+2)$. *(2) If the straight segment* $s(rt, v)$ *between* rt *and some point* v *in* S *intersects at least* M *clean segments, disregarding the segments* s_{rt} *and* s_v *to which* rt *and* v *belong, respectively, then* $d_T(rt, v) \geq M$.

Proof: To prove the first assertion, note that each edge $e \in H$ intersects at most $\left\lfloor \frac{w(e)}{\epsilon} \right\rfloor + 1$ horizontal segments. Hence, the number of dirty segments is bounded above by

$$\sum_{e \in H} \left(\left\lfloor \frac{w(e)}{\epsilon} \right\rfloor + 1 \right) \leq \left(\frac{1}{\epsilon} \cdot \sum_{e \in H} w(e) \right) + |H| \leq \frac{1}{\epsilon} \cdot \left(\frac{1}{13} + \frac{2\epsilon}{13} \right) + \frac{1}{13} \cdot (k+2) = \frac{2}{13} \cdot (k+2).$$

(The last inequality follows from Claim 4.7.)

Next, we prove the second assertion of this claim. Denote by $P(rt, v) = (v_1 = rt, v_2, \ldots, v_m = v)$ the path from rt to v in T, and suppose without loss of generality that the y-coordinate of v is greater than that of rt. Let $h_{i_1}, h_{i_2}, \ldots, h_{i_M}$ denote the clean segments that intersect $s(rt, v)$ by increasing order of y coordinate, excluding s_{rt} and s_v, with $i_1 < i_2 < \ldots < i_M$. Notice that the y-coordinate of h_{i_1} (respectively, h_{i_M}) is larger (resp., smaller) than that of rt (resp., v). Fix an arbitrary index $j \in [M]$. Denote by $pred(i_j)$ the last vertex in $P(rt, v)$ with y-coordinate smaller than that of h_{i_j}. Also, denote by $succ(i_j)$ the first vertex that comes after $pred(i_j)$ in $P(rt, v)$ and has y-coordinate greater than that of h_{i_j}. Since h_{i_j} is clean, no edge of weight at least ϵ intersects it. Consequently, $pred(i_j)$ and $succ(i_j)$ belong to the two vertical segments that intersect h_{i_j} from below and above, respectively. That is, if i_j is even, then $pred(i_j)$ belongs to $r_{(i_j/2)}$ and $succ(i_j)$ belongs to $l_{(i_j/2)}$, and if i_j is odd, then $pred(i_j)$ belongs to $l_{\lfloor i_j/2 \rfloor}$ and $succ(i_j)$ belongs to $r_{\lceil i_j/2 \rceil}$. Hence, by the triangle inequality, the weight of the sub-path of $P(rt, v)$ from $pred(i_j)$ to $succ(i_j)$, denoted P_j, is at least $\|pred(i_j) - succ(i_j)\| \geq 1$. Also, note that the y-coordinate of $succ(i_j)$ is smaller than that of $h_{i_{j+1}}$, for each index $j \in [M-1]$. Thus, by definition, either $succ(i_j) = pred(i_{j+1})$ holds, or $succ(i_j)$ comes before $pred(i_{j+1})$ in $P(rt, v)$, $j \in [M-1]$. Denote by $\tilde{P}_j$ the sub-path of $P(rt, v)$ from $succ(i_j)$ to $pred(i_{j+1})$, $j \in [M-1]$. Also, let $\tilde{P}_0$ and $\tilde{P}_M$ be the sub-paths of $P(rt, v)$ from rt to $pred(i_1)$ and from $succ(i_M)$ to v, respectively. Hence, all paths $\{P_j\}_{j \in [M]}$ and $\{\tilde{P}_j\}_{j \in [0,M]}$ are pairwise edge-disjoint, and we have $P(rt, v) = \tilde{P}_0 \circ P_1 \circ \tilde{P}_1 \circ P_2 \circ \tilde{P}_2 \circ \ldots \circ P_M \circ \tilde{P}_M$. It follows that

$$\begin{aligned} d_T(rt, v) &= w(P(rt, v)) = \sum_{j=1}^{M} w(P_j) + \sum_{j=0}^{M} w(\tilde{P}_j) \geq \sum_{j=1}^{M} w(P_j) \\ &\geq \sum_{j=1}^{M} \|pred(i_j) - succ(i_j)\| \geq M. \end{aligned}$$

Claim 4.8 follows. ▮

Denote by S' the set of all points v in S, such that the straight segment $s(rt, v)$ between rt and v intersects at least $k/4 - 1$ horizontal segments, disregarding the segments s_{rt} and s_v to which rt and v belong, respectively. It is easy to verify that $|S'| \geq n/2$. By the first assertion of Claim 4.8, at most $\frac{2}{13} \cdot (k+2)$ segments are dirty, implying that $s(rt, v)$ intersects at least $k/4 - 1 - \frac{2}{13} \cdot (k+2)$ clean segments, disregarding s_{rt} and s_v. The second assertion of Claim 4.8 implies that $d_T(rt, v) \geq k/4 - 1 - \frac{2}{13} \cdot (k+2)$. Also, observe that $diam(S) = \sqrt{2}$, and so $\|rt - v\| \leq diam(S) = \sqrt{2}$. It follows that

$$Str_T(rt, v) = \frac{d_T(rt, v)}{\|rt - v\|} \geq \frac{k/4 - 1 - \frac{2}{13} \cdot (k+2)}{\sqrt{2}} = \Omega(k) = \Omega\left(\frac{1}{\epsilon}\right).$$

Since $|S'| \geq n/2$, it follows that

$$AvgStr(T, rt) \;=\; \frac{\sum_{v \in S \setminus \{rt\}} Str_T(rt, v)}{n-1} \;\geq\; \frac{\sum_{v \in S'} Str_T(rt, v)}{n-1} \;\geq\; \frac{(n/2) \cdot \Omega\left(\frac{1}{\epsilon}\right)}{n-1} \;=\; \Omega\left(\frac{1}{\epsilon}\right).$$

This completes the proof of Lemma 4.6. ▮

5 A Lower Bound for Euclidean Spanners

In this section we provide a tight lower bound of $\Omega(\epsilon^{-d+1})$ on the maximum degree of Euclidean spanners.

As before, let C_n denote a set of n points that are uniformly spaced around the boundary C of the unit circle, centered at the origin $(0,0)$, and define $\tilde{C}_{n+1} = C_n \cup \{(0,0)\}$.

Let $\epsilon = \Omega(\frac{1}{n})$ be an arbitrary parameter, and let T be an arbitrary spanning tree of $\tilde{C}_{n+1}$ rooted at $rt = (0,0)$. Lemma 4.1 and Corollary 4.2 imply that if the average root-stretch of T is at most $1+\epsilon$, then both its root-degree $\gamma(T, rt)$ and its lightness are at least $\Omega(\frac{1}{\epsilon})$. We derive the following corollary.

Corollary 5.1 *For any sufficiently large integer n and any parameter $\epsilon = \Omega(\frac{1}{n})$, there exists a Euclidean 2-dimensional metric M and a designated point $rt \in M$, such that every spanning tree rooted at rt with average root-stretch at most $1+\epsilon$ has both root-degree and lightness at least $\Omega(\frac{1}{\epsilon})$.*

The tradeoff $1+\epsilon$ versus $\Omega(\frac{1}{\epsilon})$ between the (average) root-stretch and maximum degree generalizes in a natural way to any constant dimension $d \geq 2$, with the maximum degree bound becoming $\Omega(\epsilon^{-d+1})$. Specifically, let $C^{(d)}$ be the unit d-dimensional sphere centered at the origin $(0,0)$, let $C_n^{(d)}$ denote a set of n points that are uniformly spaced around $C^{(d)}$, and define $\tilde{C}_{n+1}^{(d)} = C_n^{(d)} \cup \{(0,0)\}$.

Lemma 5.2 *Let n be a sufficiently large integer, let $d \geq 2$ be an integer constant, and let $\epsilon = \Omega\left(\left(\frac{1}{n}\right)^{1/(d-1)}\right)$ be an arbitrary parameter. Any spanning tree T of $\tilde{C}_{n+1}^{(d)}$ rooted at $rt = (0,0)$ with average root-stretch at most $1+\epsilon$ has root-degree $\Omega(\epsilon^{-d+1})$.*

It is well-known that the path-greedy spanner of [3] provides, for any set of points in $\mathbb{R}^d$, a $(1+\epsilon)$-spanner with maximum degree $O(\epsilon^{-d+1})$. (See, e.g., Corollary 15.1.3 in [25].) Lemma 5.2 yields the following corollary, which implies that the maximum degree of the path-greedy spanner of [3] is optimal.

Corollary 5.3 *For any sufficiently large integer n, an integer constant $d \geq 2$, and a parameter $\epsilon = \Omega\left(\left(\frac{1}{n}\right)^{1/(d-1)}\right)$, every $(1+\epsilon)$-spanner for $\tilde{C}_{n+1}^{(d)}$ has maximum degree $\Omega(\epsilon^{-d+1})$.*

Proof: Suppose for contradiction that there is a $(1+\epsilon)$-spanner H for $\tilde{C}_{n+1}^{(d)}$ with maximum degree $o(\epsilon^{-d+1})$. Let T be an SPT over H rooted at $rt = (0,0)$, and denote by $\gamma(H)$ the maximum degree of H. Obviously, (T, rt) is a spanning tree for $\tilde{C}_{n+1}^{(d)}$, with $AvgStr(T, rt) \leq rtStr(T, rt) \leq 1+\epsilon$ and $\gamma(T, rt) \leq \gamma(H) = o(\epsilon^{-d+1})$, contradicting Lemma 5.2. ▮

6 Steiner Edges Do Not Help

In this section we show that Steiner edges do not help in the context of shallow-light trees.

We start with a few definitions. A graph G is called a *metric graph* if the edge weights satisfy the triangle inequality. For a metric graph $G = (V, E, w)$, let $M_G = (V, d_G)$ be the metric induced by G. In what follows we view M_G as the complete weighted graph $(V, \binom{V}{2}, d_G)$ over V, in which for every pair of vertices $u, v \in V$, there is an edge of weight $d_G(u, v)$ between u and v in G. (Notice that an MST for G is also an MST for M_G.) An edge that belongs to M_G but does not belong to G, i.e., an edge in $\binom{V}{2} \setminus E$, is called a *Steiner edge*. A spanning tree for the metric M_G induced by G that may contain edges that do not belong to G will be called a *metric-spanning tree of* G. To distinguish metric-spanning trees from spanning trees of G (that use only edges of G), we will call the latter *graph-spanning trees* of G. A *graph-spanning shallow-light tree* (henceforth, *spanning SLT*) of G is a graph-spanning tree of G that has small lightness and root-stretch (with respect to some designated root vertex rt). (See Section 1.3.) Finally, a *metric-spanning SLT* of G is a metric-spanning tree of G with the same properties (small lightness and root-stretch).

In what follows we show that the same tradeoffs between lightness and root-stretch that apply to graph-spanning SLTs apply to metric-spanning SLTs as well.

Lemma 6.1 *Let $G = (V, E, w)$ be an arbitrary metric graph, and let $rt \in V$ be an arbitrary designated vertex. Also, let (T, rt) be a rooted metric-spanning tree of G that contains at least one Steiner edge. Then T can be transformed into a rooted metric-spanning tree (T', rt) of G (that may still contain some Steiner edges), having the following properties:*

- *The weight of T' is* strictly *smaller than the weight of T, i.e., $w(T') < w(T)$.*
- *For every vertex $v \in V \setminus \{rt\}$, the stretch between rt and v in T' is no greater than the stretch between them in T. In particular, both the root-stretch and the average root-stretch of (T', rt) are no greater than the root-stretch and the average root-stretch of (T, rt), respectively.*

Proof: Let $e = (x, y)$ be some Steiner edge in T, with $x = \pi(y)$. Since e does not belong to G, there is a path P_e of weight $d_G(x, y)$ in G between x and y, with $P_e = (v_1 = x, v_2 = z, v_3, ..., v_k = y), k \geq 3$. Since P_e is a shortest path between x and y in G, it follows that

$$d_G(x, z) + d_G(z, y) \;=\; w(P_e) \;=\; d_G(x, y). \tag{16}$$

Since $d_G(x, z), d_G(z, y) > 0$, we conclude that both $d_G(x, z)$ and $d_G(z, y)$ are strictly smaller than $d_G(x, y)$.

The analysis splits into three cases.

Case 1: z is an ancestor of y in T. (Observe that $z \neq x$, but it is possible that $z = \pi(x)$.) We transform T into a metric-spanning tree T' of G by removing the edge (x, y) and adding the edge (z, y), with y becoming a child of z. (See Fig. 1(1) for an illustration.) Observe that $w(T') = w(T) - d_G(x, y) + d_G(z, y) < w(T)$. Note also that $d_{T'}(rt, x) = d_T(rt, x)$ and $d_{T'}(rt, y) < d_T(rt, y)$. More generally, for any vertex v that belongs to the subtree T_y of T rooted at y, $d_{T'}(rt, v) < d_T(rt, v)$. For other vertices v, $d_{T'}(rt, v) = d_T(rt, v)$.

Case 2: z is a descendant of y in T. (Observe that it is possible that $y = \pi(z)$.) We transform T into a metric-spanning tree T' of G by removing the two edges (x, y) and $(\pi(z), z)$ and adding the two edges (x, z) and (z, y), with z becoming a child of x and y becoming a child of z. (See Fig. 1(2) for an illustration.) Observe that $w(T') = w(T) - d_G(x, y) - d_G(\pi(z), z) + d_G(x, z) + d_G(z, y)$. Since $d_G(x, z) + d_G(z, y) = d_G(x, y)$, it follows that $w(T') = w(T) - d_G(\pi(z), z) < w(T)$. Note also that for any vertex v that belongs to the subtree T_z of T rooted at z, $d_{T'}(rt, v) < d_T(rt, v)$. For other vertices v, $d_{T'}(rt, v) = d_T(rt, v)$.

Case 3: z is neither an ancestor nor a descendant of y. In this case let p denote the least common ancestor of z and y in T, i.e., $p = LCA(z, y)$. Note that $p \notin \{z, y\}$. Our analysis splits further into two subcases.

Case 3.a: In the first subcase we have

$$d_T(p, z) + d_G(z, y) \;<\; d_T(p, x) + d_G(x, y). \tag{17}$$

This condition implies that $p \neq x$. Indeed, otherwise we get $d_T(x,z) + d_G(z,y) < d_G(x,y)$, which is a contradiction to the triangle inequality. As in case 1, we transform T into a metric-spanning tree T' of G by removing the edge (x,y) and adding the edge (z,y), with y becoming a child of z. (See Fig. 1(3.a) for an illustration.) Exactly as in case 1, it follows that $w(T') < w(T)$. Observe also that $d_T(rt,y) = d_T(rt,p) + d_T(p,x) + d_G(x,y)$. On the other hand, $d_{T'}(rt,y) = d_T(rt,p) + d_T(p,z) + d_G(z,y)$. By Equation (17), $d_{T'}(rt,y) < d_T(rt,y)$. It is easy to verify that for all vertices v that belong to the subtree T_y, $d_{T'}(rt,v) < d_T(rt,v)$, and for all other vertices v, $d_{T'}(rt,v) = d_T(rt,v)$.
Case 3.b: In the complementary subcase, i.e., when Equation (17) does not hold, we have

$$d_T(p,z) + d_G(z,y) \;\geq\; d_T(p,x) + d_G(x,y). \tag{18}$$

(Observe that it is possible that $p = x$ and/or $p = \pi(z)$.) Plugging Equation (16) in Equation (18), we obtain

$$d_T(p,z) + d_G(z,y) \quad \geq \quad d_T(p,x) + d_G(x,y) \quad = \quad d_T(p,x) + d_G(x,z) + d_G(z,y),$$

implying that

$$d_T(p,z) \;\geq\; d_T(p,x) + d_G(x,z). \tag{19}$$

As in case 2, we transform T into a spanning tree T' of M_G by removing the two edges (x,y) and $(\pi(z),z)$ and adding the two edges (x,z) and (z,y), with z becoming a child of x and y becoming a child of z. (See Fig. 1(3.b) for an illustration.) Exactly as in case 2, it follows that $w(T') < w(T)$. Also, Equation (19) implies that $d_{T'}(rt,z) \leq d_T(rt,z)$. Hence, for all vertices v that belong to the subtree T_z of T rooted at z, $d_{T'}(rt,v) \leq d_T(rt,v)$. Other distances from the root stay unchanged. ▮

Lemma 6.1 implies the following corollaries.

Corollary 6.2 *Let $G = (V,E,w)$ be an arbitrary metric graph, let $rt \in V$ be an arbitrary designated vertex, and let $\alpha \geq 1$ be an arbitrary number. Denote by S_1 (respectively, S_2) the set of all metric-spanning trees of G rooted at rt with root-stretch (resp., average root-stretch) at most α. Suppose that S_1 (respectively, S_2) is non-empty, and let (T_1^*, rt) (resp., (T_2^*, rt)) be a tree of minimum lightness among all trees in S_1 (resp., S_2). Then both T_1^* and T_2^* are graph-spanning trees of G.*

Proof: Indeed, if T_1^* contains a Steiner edge, then by Lemma 6.1, there exists a metric-spanning tree $T_1' \in S_1$, with $w(T_1') < w(T_1^*)$. This contradicts the minimality of T_1^*. The proof for T_2^* is analogous. ▮

Now are are ready to derive the main result of this section. Informally, it states that Steiner edges do not help in the context of SLTs.

Corollary 6.3 *Let $G = (V,E,w)$ be an arbitrary metric graph, let $rt \in V$ be an arbitrary designated vertex, and let $\alpha \geq 1, \beta \geq 1$ be an arbitrary pair of numbers. If there is a metric-spanning tree of G rooted at rt with root-stretch (respectively, average root-stretch) at most α and lightness at most β, then there is also a graph-spanning tree of G rooted at rt with root-stretch (respectively, average root-stretch) at most α and lightness at most β.*

Acknowledgments

We thank Samir Khuller, Michiel Smid, Warren Smith and Noam Solomon for helpful discussions.

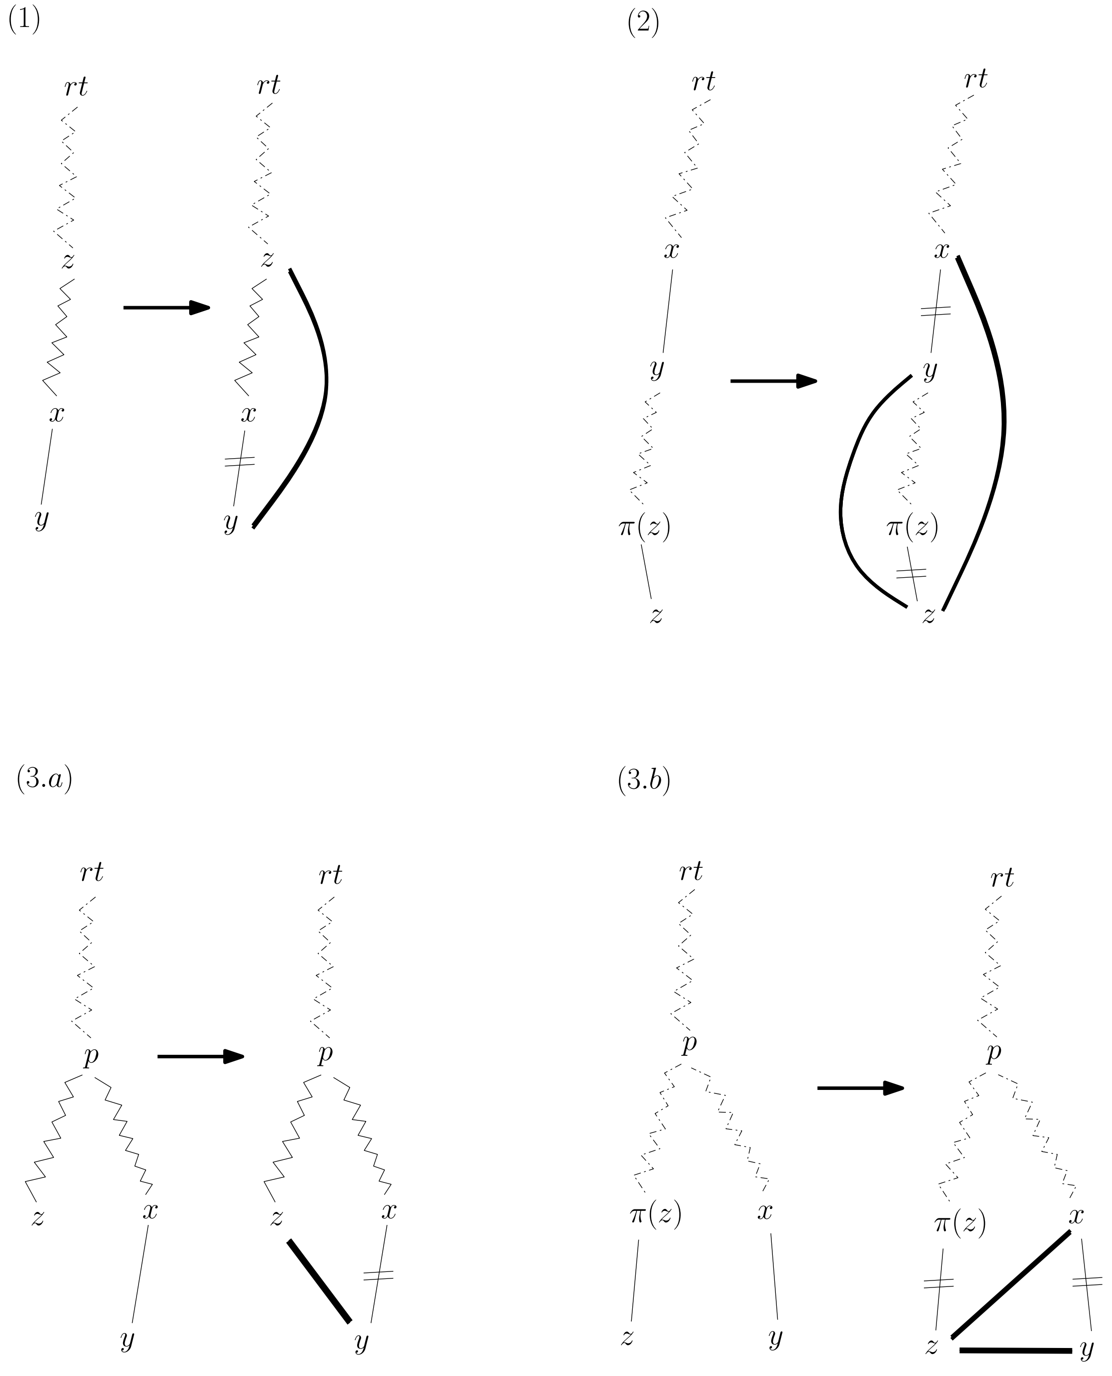

Figure 10: An illustration of the four cases that are considered in the proof of Lemma 6.1. In each case it is shown how to transform the tree (T, rt) into another tree (T', rt) that satisfies the conditions of the lemma. Paths that may contain zero or more edges are depicted by zigzag lines. A path that must contain at least one edge is depicted by a solid zigzag line, whereas a path that might be empty is depicted by a dash-dotted zigzag line. Single edges are depicted by straight lines. Newly added edges are depicted by thick lines, whereas removed edges are depicted by crossed lines.

References

[1] I. Abraham, Y. Bartal, and O. Neiman. Nearly tight low stretch spanning trees. In *Proc. of 49th FOCS*, pages 781–790, 2008.

[2] N. Alon, R. M. Karp, D. Peleg, and D. B. West. A graph-theoretic game and its application to the k-server problem. *SIAM J. Comput.*, 24(1):78–100, 1995.

[3] 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.

[4] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. H. M. Smid. Euclidean spanners: short, thin, and lanky. In *Proc. of 27th STOC*, pages 489–498, 1995.

[5] S. Arya and M. H. M. Smid. Efficient construction of a bounded degree spanner with low weight. *Algorithmica*, 17(1):33–54, 1997.

[6] B. Awerbuch, A. Baratz, and D. Peleg. Cost-sensitive analysis of communication protocols. In *Proc. of 9th PODC*, pages 177–187, 1990.

[7] B. Awerbuch, A. Baratz, and D. Peleg. Efficient broadcast and light-weight spanners. *Manuscript*, 1991.

[8] Y. Bartal. Probabilistic approximations of metric spaces and its algorithmic applications. In *Proc. of 37th FOCS*, pages 184–193, 1996.

[9] Y. Bartal. On approximating arbitrary metrices by tree metrics. In *Proc. of 30th STOC*, pages 161–168, 1998.

[10] B. Bollobás, D. Coppersmith, and M. Elkin. Sparse distance preservers and additive spanners. In *Proc. of 14th SODA*, pages 414–423, 2003.

[11] P. Carmi and M. H. M. Smid. An optimal algorithm for computing angle-constrained spanners. In *Proc. of 21st ISAAC (to appear)*, 2010.

[12] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong. Performance-driven global routing for cell based ics. In *Proc. of 9th ICCD*, pages 170–173, 1991.

[13] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong. Provably good algorithms for performance-driven global routing. In *Proc. of 5th ISCAS*, pages 2240–2243, 1992.

[14] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong. Provably good performance-driven global routing. *IEEE Trans. on CAD of Integrated Circuits and Sys.*, 11(6):739–752, 1992.

[15] T. H. Corman, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, 2nd edition.* McGraw-Hill Book Company, Boston, MA, 2001.

[16] G. Das and G. Narasimhan. A fast algorithm for constructing sparse Euclidean spanners. *Int. J. Comput. Geometry Appl.*, 7(4):297–315, 1997.

[17] Y. Dinitz, M. Elkin, and S. Solomon. Shallow-low-light trees, and tight lower bounds for Euclidean spanners. In *Proc. of 49th FOCS*, pages 519–528, 2008.

[18] M. Elkin, Y. Emek, D. Spielman, and S. Teng. Lower stretch spanning trees. In *Proc. of 37th STOC*, pages 494–503, 2005.

[19] M. Elkin and S. Solomon. Narrow-shallow-low-light trees with and without steiner points. In *Proc. of 17th ESA*, pages 215–226, 2009.

[20] G. Even, J. Naor, S. Rao, and B. Schieber. Divide-and-conquer approximation algorithms via spreading metrics. In *Proc. of 36th FOCS*, pages 62–71, 1995.

[21] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proc. of 35th STOC*, pages 448–455, 2003.

[22] A. Gupta. Steiner points in tree metrics don't (really) help. In *Proc. of 12th SODA*, pages 220–227, 2001.

[23] S. Khuller, B. Raghavachari, and N. E. Young. Balancing minimum spanning and shortest path trees. In *Proc. of 4th SODA*, pages 243–250, 1993.

[24] G. Konjevod, R. Ravi, and F. S. Salman. On approximating planar metrics by tree metrics. *Inf. Process. Lett.*, 80(4):213–219, 2001.

[25] G. Narasimhan and M. Smid. *Geometric Spanner Networks.* Cambridge University Press, 2007.

[26] D. Peleg and A. Schäffer. Graph spanners. *J. Graph Theory*, 13(1):99–116, 1989.

[27] S. Rao and W. D. Smith. Approximating geometrical graphs via "spanners" and "banyans". In *Proc. of 30th STOC*, pages 540–550, 1998.

[28] S. Solomon. An optimal time construction of euclidean sparse spanners with tiny diameter. In *Proc. of 22nd SODA (to appear)*, 2011.

Appendix

A A Technical Lemma

This appendix is devoted to the proof of the following technical lemma.

Lemma A.1 *For any positive integers $n_1, n_2, \ldots, n_k$, $k \geq 1$, it holds that* $\sum_{i=1}^{k} \left(n_i \cdot \log\left(\frac{n}{n_i}\right)\right) \leq n \cdot (k-1)$*, where* $n = \sum_{i=1}^{k} n_i$*.*

Before we prove this lemma, we state the following useful fact.

Fact A.2 *Let n be a fixed positive number. Define* $f(x) = x \cdot \log\left(\frac{n}{x}\right) + (n-x) \cdot \log\left(\frac{n}{n-x}\right)$*. Then for all* $0 < x < n$, $f(x) \leq n$*.*

Next, we turn to the proof of Lemma A.1. The proof is by induction on $k, k \geq 1$.
Basis: $k = 1$. In this case we have $n = n_1$, and so $\sum_{i=1}^{k} \left(n_i \cdot \log\left(\frac{n}{n_i}\right)\right) = n \cdot \log 1 = 0 = n \cdot (k-1)$.
Induction Step: We assume that the statement holds for all smaller values of $k, k \geq 2$, and prove it for k. Define $N_k = \sum_{i=1}^{k} \left(n_i \cdot \log\left(\frac{n}{n_i}\right)\right)$ and $N_{k-1} = \sum_{i=1}^{k-1} \left(n_i \cdot \log\left(\frac{n-n_k}{n_i}\right)\right)$. We need to show that $N_k \leq n \cdot (k-1)$. Observe that $1 \leq n_k \leq n-1$, and so

$$\begin{aligned} N_k - N_{k-1} &= \sum_{i=1}^{k-1} \left(n_i \cdot \log\left(\frac{n}{n-n_k}\right)\right) + n_k \cdot \log\left(\frac{n}{n_k}\right) \\ &= (n-n_k) \cdot \log\left(\frac{n}{n-n_k}\right) + n_k \cdot \log\left(\frac{n}{n_k}\right) \;\leq\; n. \end{aligned}$$

(The last inequality follows from Fact A.2.) By the induction hypothesis for $k-1$, we have $N_{k-1} \leq (n-n_k) \cdot (k-2)$. Consequently,

$$N_k \;=\; (N_k - N_{k-1}) + N_{k-1} \;\leq\; n + (n-n_k) \cdot (k-2) \;\leq\; n + n \cdot (k-2) \;=\; n \cdot (k-1). \qquad \blacksquare$$