

Steiner Shallow-Light Trees are Exponentially Lighter than Spanning Ones

Michael Elkin

Department of Computer Science
Ben-Gurion University of the Negev
Beer-Sheva 84105, Israel
Email: elkinm@cs.bgu.ac.il

Shay Solomon

Department of Computer Science
Ben-Gurion University of the Negev
Beer-Sheva 84105, Israel
Email: shayso@cs.bgu.ac.il

Abstract— For a pair of parameters $\alpha, \beta \geq 1$, a spanning tree T of a weighted undirected n -vertex graph $G = (V, E, w)$ is called an (α, β) -*shallow-light tree* (shortly, (α, β) -SLT) of G with respect to a designated vertex $rt \in V$ if (1) it approximates all distances from rt to the other vertices up to a factor of α , and (2) its weight is at most β times the weight of the minimum spanning tree $MST(G)$ of G . The parameter α (respectively, β) is called the *root-distortion* (resp., *lightness*) of the tree T . Shallow-light trees (SLTs) constitute a fundamental graph structure, with numerous theoretical and practical applications. In particular, they were used for constructing spanners, in network design, for VLSI-circuit design, for various data gathering and dissemination tasks in wireless and sensor networks, in overlay networks, and in the message-passing model of distributed computing.

Tight tradeoffs between the parameters of SLTs were established by Awerbuch et al. [5], [6] and Khuller et al. [33]. They showed that for any $\epsilon > 0$ there always exist $(1 + \epsilon, O(\frac{1}{\epsilon}))$ -SLTs, and that the upper bound $\beta = O(\frac{1}{\epsilon})$ on the lightness of SLTs cannot be improved. In this paper we show that using Steiner points one can build SLTs with *logarithmic lightness*, i.e., $\beta = O(\log \frac{1}{\epsilon})$. This establishes an *exponential separation* between spanning SLTs and Steiner ones.

One particularly remarkable point on our tradeoff curve is $\epsilon = 0$. In this regime our construction provides a *shortest-path tree* with weight at most $O(\log n) \cdot w(MST(G))$. Moreover, we prove matching lower bounds that show that all our results are tight up to constant factors.

Finally, on our way to these results we settle (up to constant factors) a number of open questions that were raised by Khuller et al. [33] in SODA'93.

Keywords-minimum spanning tree; shortest-path tree; Steiner points; Steiner trees;

1. INTRODUCTION

1.1. Main Results

A *minimum spanning tree* (henceforth, MST) of a weighted undirected graph $G = (V, E, w)$, $w : E \rightarrow \mathbb{R}^+$, is a spanning tree $T = (V, H, w)$ of G with minimum weight $w(T) = \sum_{e \in H} w(e)$. A *shortest-path tree* (henceforth, SPT) of G with respect to a designated vertex $rt \in V$ is a spanning

tree $T = (V, H, w)$ that satisfies that for every vertex $v \in V$, the distance $d_T(rt, v)$ between rt and v in T is equal to the distance $d_G(rt, v)$ between them in G . Both the MST and the SPT are among the most fundamental and well-studied graph constructs (see, e.g., [20], ch. 24, 25, and the references therein).

The desirability of having a single tree that combines the properties of the MST and the SPT was realized already in the mid-eighties [12], [32]. The specific notion of *shallow-light tree* (shortly, SLT) was introduced in [5], [6], [33].¹ For a pair of parameters $\alpha, \beta \geq 1$, a spanning tree T of G with respect to a designated vertex $rt \in V$ is called an (α, β) -SLT if (1) for every vertex $v \in V$, $d_T(rt, v) \leq \alpha \cdot d_G(rt, v)$, and (2) $w(T) \leq \beta \cdot w(MST(G))$. The parameter α (respectively, β) will be called the *root-distortion* (resp., *lightness*) of the tree T . Awerbuch et al. [5], [6] and Khuller et al. [33] demonstrated that for every $\epsilon > 0$, a $(1 + \epsilon, O(\frac{1}{\epsilon}))$ -SLT exists for every graph G , and that this tradeoff is tight [33]. Since then SLTs were shown to have numerous applications. In particular, they were found useful for various data gathering and dissemination tasks in overlay networks [14], [46], [37], in wireless and sensor networks [47], [21], [9], [44], and in the message-passing model of distributed computing [5], [6]. Other applications of SLTs include network and VLSI-circuit design [16], [17], [18], [43], and routing [48]. SLTs were also used for constructing spanners [6], [40]. Closely related tree structures, such as *light approximate routing trees* and *shallow-low-light trees* were investigated in [48], [22], [24]. Moreover, some of the constructions in [48], [22], [24] are based on constructions of SLTs. See also the introduction of [39] for additional applications of SLTs.

This variety of both theoretical and practical applications testifies that SLTs constitute a fundamental graph structure of independent interest. In this paper we explore the impact of *Steiner points* on SLTs. A *Steiner tree* for a graph $G = (V, E, w)$ is a tree $T = (V', H, w')$ with $V' \supseteq V$ and $w' : H \rightarrow \mathbb{R}^+$, that *dominates* the metric M_G induced by G ,

Both authors are partially supported by the Lynn and William Frankel Center for Computer Sciences and by the BSF grant No. 2008430. The research of Shay Solomon has been supported by the Clore Fellowship grant No. 81265410.

¹Khuller et al. [33] called the same notion *light approximate spanning tree*, or shortly, *LAST*.




Figure 1. All edges incident to rt have weight n , while other edges have unit weight. The SPT with respect to rt is the star rooted at rt ; its weight is $n(n-1)$. The MST contains a single edge (rt, v_i) of weight n , and the entire path $(v_1, v_2, \dots, v_{n-1})$ of unit-weight edges. Its weight is $2n-2$. Hence the lightness of the SPT with respect to rt in this example is $\frac{n}{2}$.

i.e., for every pair of original vertices $u, v \in V$, $d_T(u, v) \geq d_G(u, v)$. For a pair of parameters $\alpha, \beta \geq 1$, we say that the Steiner tree T is a *Steiner (α, β) -SLT* for G with respect to a designated vertex $rt \in V$ if (1) for every vertex $v \in V$, $d_T(rt, v) \leq \alpha \cdot d_G(rt, v)$, and (2) $w(T) \leq \beta \cdot w(MST(G))$.²

We demonstrate that using Steiner points one can drastically improve the lightness of SLTs. Specifically, we show that for every $\epsilon > 0$, every graph G and every designated vertex rt in G , there exists a $(1 + \epsilon, O(\log \frac{1}{\epsilon}))$ -SLT with respect to rt , and that this result is tight up to a constant factor (hidden by the O -notation). As was mentioned above, a lower bound of Khuller et al. [33] shows that the lightness of spanning SLTs with the same root-distortion is $\Omega(\frac{1}{\epsilon})$, i.e., we establish an *exponential separation* between the lightness of Steiner and spanning SLTs.

One particularly remarkable point on our tradeoff curve is $\epsilon = 0$, i.e., when we do not allow any distortion whatsoever. The lightness of our Steiner trees that preserve distances from a designated vertex (i.e., Steiner shortest-path trees) is $O(\log n)$, where n is the number of vertices; this is again tight up to a constant factor. Note also that there are graphs for which any spanning shortest-path tree has lightness $\Omega(n)$. (See Figure 1 for an illustration.)

1.2. Steiner Points in Other Metric Structures

The impact of Steiner points on metric trees was extensively studied in a few related settings. Most notably, it is a subject of intensive investigation in the context of *probabilistic tree embeddings* and *average-distortion trees*. These two settings are essentially equivalent, and so we will only discuss the former one. Alon et al. [2] showed that for every n -vertex 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* [7], and the value $\max_{e=(u,v) \in E} \mathbb{E} \left[\frac{d_T(u, v)}{w(e)} \right]$ is called the

²Alternatively, one can require here the weight of T to be no greater than β times the weight of the *minimum Steiner tree* of G , $SMT(G)$. However, since $\frac{1}{2} \cdot w(MST(G)) \leq w(SMT(G)) \leq w(MST(G))$, these two definitions are identical up to the constant factor of 2. We will henceforth ignore this subtlety.

distortion of the embedding. Bartal and Fakcharoenphol et al. [7], [8], [25] showed that using Steiner points one can drastically improve the bound of [2], and devised probabilistic tree embeddings with distortion $O(\log n)$. The bound $O(\log n)$ is also known to be optimal up to constant factors [2], [7]. However, Konjevod et al. [35] and Gupta [28] demonstrated that the same bounds (up to constant factors) as those of Bartal and Fakcharoenphol et al. [7], [8], [25] can be obtained without Steiner points, i.e., by using only *spanning trees* of the metric M_G induced by G . Moreover, more recent studies [23], [1] showed that nearly the same bounds can be obtained by using spanning trees of the original graph G . Therefore, it turns out that Steiner points do not really help to improve probabilistic tree embeddings.

A similar situation occurs in the context of *low-light trees*, which combine small lightness with small depth [22], [24]. It is known that Steiner points do not help in this context either, i.e., that any Steiner tree T can be converted into a spanning tree with the the same (up to constant factors) lightness and depth as those of T [24].

Steiner points were also studied in the context of *graph spanners*. Given a graph $G = (V, E, w)$ and a number $\alpha \geq 1$, an α -*spanner* $G' = (V, H, w)$, $H \subseteq E$, is a subgraph that satisfies $d_{G'}(u, v) \leq \alpha \cdot d_G(u, v)$, for every pair of vertices $u, v \in V$. A *Steiner α -spanner* $G' = (V', H', w')$, $V' \supseteq V$, is a graph that satisfies $d_G(u, v) \leq d_{G'}(u, v) \leq \alpha \cdot d_G(u, v)$, for every pair of original vertices $u, v \in V$. Althöfer et al. [3], extending previous bounds due to Peleg and Schäffer [41], showed that for every n -vertex graph G there exists an α -spanner with $n^{1+O(\frac{1}{\alpha})}$ edges, and that there exist n -vertex graphs for which every *Steiner α -spanner* requires $n^{1+\Omega(\frac{1}{\alpha})}$ edges. In other words, Steiner points cannot help to significantly improve the bounds on the number of edges required for graph spanners in general. The situation is similar in the context of another variety of spanners called *distance preservers* [13].

To summarize, Steiner points were studied in many different settings in the context of probabilistic tree embeddings, low-light trees, graph spanners and distance preservers [3], [35], [28], [13], [24]. In all these settings they are known not to help much in improving inherent tradeoffs between the involved parameters. In this paper we show that in a sharp contrast to all these examples, shallow-light trees can be *exponentially* improved by using Steiner points!

1.3. Spanning Shallow-Light Trees

We also proved a number of results concerning *spanning* SLTs (rather than Steiner SLTs). In the Conclusions section of their paper [33] Khuller et al. asked four questions. In this paper we settle three of them up to constant factors.

First, Khuller et al. [33] showed that there are graphs for which any spanning tree with root-distortion at most $1 + \epsilon$ has lightness $\Omega(\frac{1}{\epsilon})$, and vice versa. They asked whether the

same lower bound applies to Euclidean graphs. We answer this question in the affirmative, and show that there exist configurations of points in the Euclidean plane for which any spanning tree with root-distortion at most $1 + \epsilon$ rooted at a designated root-vertex rt has lightness at least $\Omega(\frac{1}{\epsilon})$, for any $\epsilon > 0$. We also show that the opposite direction is true too, i.e., there are configurations of points in the plane for which any spanning tree with lightness at most $1 + \epsilon$ has root-distortion at least $\Omega(\frac{1}{\epsilon})$.

Second, Khuller et al. [33] suggested to relax the notion of root-distortion, and replace it with the notion of *average root-distortion*. A spanning tree T for a graph G is said to have average root-distortion 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-distortion of a tree is bounded above by its root-distortion.) They asked whether their lower bound applies when replacing root-distortion by average root-distortion. We show that there exist configurations of points in the Euclidean plane for which any spanning tree with average root-distortion (respectively, lightness) at most $1 + \epsilon$ with respect to a certain designated root vertex rt has lightness (resp., average root-distortion) at least $\Omega(\frac{1}{\epsilon})$.

Third, Khuller et al. [33] asked whether their lower bound applies if one is allowed to select the root vertex at will. We show that there are configurations of points in the Euclidean plane for which any spanning tree with average root-distortion (respectively, lightness) at most $1 + \epsilon$ with respect to *any root vertex* rt has lightness (resp., average root-distortion) at least $\Omega(\frac{1}{\epsilon})$.

1.4. Related Work

Shallow-low-light trees were studied in [22], [24]. In addition to small root-distortion and lightness, these trees have small depth. However, similarly to the SLTs of Awerbuch et al. [5], [6] and Khuller et al. [33], the shallow-low-light trees of [22], [24] exhibit inverse-linear tradeoff between root-distortion $1 + \epsilon$ and lightness $\Omega(\frac{1}{\epsilon})$.

SLTs were also studied from the viewpoint of approximation algorithms [38], [36], [39], [29]. Some of this research considered Steiner SLTs [36], [15], [29]. However, the word "Steiner" is used in these papers in a different meaning. Specifically, as a part of the input one is given a graph $G = (V, E, w)$ and a subset $U \subseteq V$ of terminals. A Steiner tree in the context of [36], [15], [29] is a tree that spans U but is allowed to use vertices from $V \setminus U$ as well. Online approximation algorithms for SLTs were devised in [27]. Heuristics for finding SLTs were developed in [34].

There is also a vast literature on the Euclidean minimum Steiner tree [26], [30], [11], [10]. In this context the input graph is a Euclidean one, and Steiner points are required to belong to the Euclidean plane as well. We remark that

although some of our lower bound examples in this paper are Euclidean, we focus on the general metric scenario. In particular, even if the input graph is Euclidean, our constructions may use Steiner points that do not belong to the plane. This is for a good reason, though; it is easy to see that for a set C_n of n equally spaced points on a circle in the Euclidean plane, any shortest-path tree (with respect to an arbitrary vertex $rt \in C_n$) that uses only Euclidean Steiner points has lightness $\Omega(n)$. More generally, we show that any shallow-light tree for C_n with root-distortion at most $1 + \epsilon$ (with respect to an arbitrary vertex $rt \in C_n$) that uses only Euclidean Steiner points has lightness $\Omega(\sqrt{\frac{1}{\epsilon}})$. These lower bounds are exponentially larger than our logarithmic upper bounds $O(\log n)$ and $O(\log \frac{1}{\epsilon})$ on the lightness of Steiner shortest-path trees and Steiner shallow-light trees, respectively. We remark that lower bounds on the power of Euclidean Steiner points were shown in [42], [4].

1.5. Structure of the Paper

In Section 2 we present our construction of Steiner SPTs with logarithmic lightness. The construction of Steiner SLTs with root-distortion at most $1 + \epsilon$ and lightness $O(\log \frac{1}{\epsilon})$ is described in Section 3. In Section 4 we provide our lower bounds for Euclidean spanning SLTs. Some proofs are omitted from this extended abstract.

1.6. Preliminaries

Let $T = (T, rt)$ be either a spanning or a Steiner tree of a graph $G = (V, E, w)$ rooted at some designated point rt . The *distortion* between a pair u, v of vertices in T is defined as $\varphi_T(u, v) = \frac{d_T(u, v)}{d_G(u, v)}$. The *root-distortion* and *average root-distortion* of (T, rt) are defined as

$$\chi(T, rt) = \max \{ \varphi_T(rt, v) \mid v \in V \setminus \{rt\} \}$$

and

$$\lambda(T, rt) = \frac{\sum_{v \in V \setminus \{rt\}} \varphi_T(rt, v)}{|V| - 1},$$

respectively. We remark that this definition of average root-distortion is somewhat different than the one used by Khuller et al. [33]. (See Section 1.3.) Nevertheless, all bounds on average root-distortion presented in this paper apply with respect to both definitions. We denote by $\Psi(T) = \frac{w(T)}{w(MST(G))}$ the *lightness* of a spanning or a Steiner tree T of G . The *depth* of a tree T rooted at a vertex rt is the maximum unweighted distance between rt and a vertex v in T . For a pair of points u, v in the plane, denote by $\|u - v\|$ their Euclidean distance. Finally, for a pair k, n of integers, $0 \leq k \leq n$, denote the sets $\{k, k + 1, \dots, n\}$ and $\{1, 2, \dots, n\}$ by $[k, n]$ and $[n]$, respectively.

2. STEINER SHORTEST-PATH TREES WITH LOGARITHMIC LIGHTNESS

In this section we show that for every n -vertex graph $G = (V, E, w)$ and for every vertex $rt \in V$, there exists a

Steiner shortest-path tree T with respect to rt with logarithmic lightness, i.e., $\Psi(T) = O(\log n)$. Note that without loss of generality it is enough to prove this for metrics. Indeed, a Steiner tree of a graph G is also a Steiner tree of the metric M_G induced by G , and vice versa.

In Section 3 we harness this construction 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 an arbitrary 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 $w(T') = O(\log n) \cdot w(H')$. 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 ,

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

We show (Proposition 2.7) that the rooted tree (T, rt) , $T = T(H')$, obtained by adding to T' an edge (rt, rt') of weight $w(rt, rt') = \rho(rt')$, is a Steiner shortest-path tree of M with respect to rt of weight $w(T) = O(\log n) \cdot w(H')$. (See Figure 2(i) for an illustration.) In particular, if we take H' to be a Hamiltonian path for M' of weight $O(w(MST(M'))) = O(w(MST(M)))$ (e.g., if H' is a Traveling Salesman Path for M'), then the lightness of T is

$$\Psi(T) = \frac{w(T)}{w(MST(M))} = \frac{O(\log n) \cdot w(H')}{w(MST(M))} = O(\log n).$$

Write $\tilde{n} = n - 1$ and $H' = (p_1, p_2, \dots, p_{\tilde{n}})$, and suppose for simplicity that \tilde{n} is an integer power of 2. To construct T' , we start by building a skeleton of a full balanced binary tree rooted at rt' with \tilde{n} leaves, denoted from left to right by $\ell_1, \ell_2, \dots, \ell_{\tilde{n}}$. For each index $i \in [\tilde{n}]$, the leaf ℓ_i corresponds to the point p_i . All $\tilde{n} - 1$ inner vertices of T' are Steiner points.

Before describing 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 , $L(x) = R(x) = NULL$ and $Leaves(x) = \{x\}$, whereas for an inner vertex x , $Leaves(x) = Leaves(L(x)) \cup Leaves(R(x))$. Observe that $Leaves(rt') = \{p_1, p_2, \dots, p_{\tilde{n}}\} = V \setminus \{rt\}$. 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 associate with 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

$$\delta_x = \rho(L(x)) - \rho(R(x)).$$

The variable δ_x will be referred to as the *disbalance* of the vertex x . Note that the disbalance may be negative. Consider a pair of leaves $x_L \in Leaves(L(x))$, $x_R \in Leaves(R(x))$. Let $\Delta(x_L, x_R) = dist(x_L, x_R) - (d_{T'}(L(x), x_L) + d_{T'}(R(x), x_R))$. To guarantee that the tree T' will dominate the metric M' , we need to make sure that the weights $w_L(x)$ and $w_R(x)$ will satisfy $w_L(x) + w_R(x) \geq \Delta(x_L, x_R)$. (See Figure 2(ii) for an illustration.) We call $\Delta(x_L, x_R)$ the *distance surplus* of the pair (x_L, x_R) . The *distance surplus* of the vertex x , denoted Δ_x , is defined as the maximum distance surplus over all pairs (x_L, x_R) with $x_L \in Leaves(L(x))$, $x_R \in Leaves(R(x))$, i.e.,

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

It follows that the choice of the weights $w_L(x)$ and $w_R(x)$ for the edges $(x, L(x))$ and $(x, R(x))$, respectively, needs to satisfy $w_L(x) + w_R(x) \geq \Delta_x$.

Given the values of disbalance δ_x and surplus Δ_x of the vertex x determined as above, we set the weights $w_L(x)$ and $w_R(x)$ of its descending edges as follows. (Some of the edges in the tree that we construct may have zero weight. This will be easily corrected later.)

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* leaf v in $Leaves(x)$, and set $\rho(x) = dist(rt, v) - d_{T'}(x, v)$. (Lemma 2.2 below implies that any choice of v leads to the same value of $\rho(x)$.)

Observation 2.1: For any vertex x in T' ,

- 1) $w_L(x), w_R(x) \geq 0$,
- 2) $w_L(x) + w_R(x) = \max\{\Delta_x, |\delta_x|\}$,
- 3) $w_L(x) - w_R(x) = \delta_x$, or equivalently, $\rho(L(x)) - w_L(x) = \rho(R(x)) - w_R(x)$.

Next we provide some intuition for the construction. When the algorithm assigns weights $w_L(x)$ and $w_R(x)$, the two subtrees $T'_{L(x)}$ and $T'_{R(x)}$ are already constructed. Intuitively, these trees can be viewed as Steiner shortest-path trees rooted at $L(x)$ and $R(x)$, respectively. Observe that at this stage we are also given two parameters, $\rho(L(x))$ and $\rho(R(x))$. Intuitively, $\rho(L(x))$ (respectively, $\rho(R(x))$) indicates how close is the root rt to the set of original vertices that belong to $Leaves(L(x))$ (resp., $Leaves(R(x))$); this is not the actual distance between rt and this set of vertices, but rather the value that needs to be assigned to the edge $(rt, L(x))$ (resp., $(rt, R(x))$) to obtain a Steiner SPT for




Figure 2. (i) The tree (T, rt) is obtained from T' by adding to it an edge (rt, rt') of weight $\rho(rt')$. By (1), $d_{T'}(x, v) = \text{dist}(rt, v) - \rho(x)$. Hence $d_T(rt, x) = d_T(rt, v) - d_{T'}(x, v) = (d_T(rt, v) - \text{dist}(rt, v)) + \rho(x) = \rho(x)$. The last equation holds because, as we will show in the sequel, T is an SPT with respect to rt . (ii) The path that connects x_L with x_R in T'_x has weight $d_{T'}(L(x), x_L) + w_L(x) + w_R(x) + d_{T'}(R(x), x_R)$.

$\text{Leaves}(L(x))$ (resp., $\text{Leaves}(R(x))$); it is convenient to think of this value as a “distance” between rt and the set $\text{Leaves}(L(x))$ (resp., $\text{Leaves}(R(x))$). Our objective at this point is to merge the two subtrees $T'_{L(x)}$ and $T'_{R(x)}$ into a single Steiner SPT rooted at x . As a part of this merging operation we need to balance these two subtrees. This is done using the disbalance parameter $\delta_x = \rho(L(x)) - \rho(R(x))$. If $\delta_x > 0$ then $w_L(x)$ needs to be greater than $w_R(x)$. The reason for that is that in this case the vertices of $\text{Leaves}(L(x))$ are located “farther” from rt than the vertices of $\text{Leaves}(R(x))$, and so by setting $w_L(x)$ to be greater than $w_R(x)$ we compensate for this. The case $\delta_x < 0$ is symmetric. If $\delta_x = 0$ then we set $w_L(x)$ to be equal to $w_R(x)$. The third statement of Observation 2.1 demonstrates the intuitive meaning of the disbalance variable δ_x ; it is the difference between $w_L(x)$ and $w_R(x)$.

Our additional concern during this merging step is to ensure that the resulting tree T'_x will dominate the metric distances between the vertices of $\text{Leaves}(x)$. To this end we employ the distance surplus parameter Δ_x . Specifically, our choice of values for $w_L(x)$ and $w_R(x)$ guarantees that $w_L(x) + w_R(x) \geq \Delta_x$. As was discussed above, this, in turn, guarantees that T'_x dominates the metric distances between the vertices of $\text{Leaves}(x)$.

There are two cases that we encounter in the merging process. The first one is $|\delta_x| > \Delta_x$, and the second one is the complementary case. In the first case we can assign weights $w_L(x)$ and $w_R(x)$ regardless of the distance surplus parameter Δ_x . (Indeed, in this case $|w_L(x) - w_R(x)| = |\delta_x| > \Delta_x$; since $w_L(x), w_R(x) \geq 0$, we have $w_L(x) + w_R(x) \geq |w_L(x) - w_R(x)| > \Delta_x$.) In other words, if $|\delta_x| > \Delta_x$, then the balancing operation implicitly takes care of the domination condition. On the other hand, if $|\delta_x| \leq \Delta_x$, then this is no longer the case. In this case we explicitly make sure that $w_L(x) + w_R(x) \geq \Delta_x$. In addition, we also need

$|w_L(x) - w_R(x)| = |\delta_x|$ to hold, in order to balance the two subtrees. These two conditions are achieved simultaneously by assigning one of the edges weight $\frac{\Delta_x + \delta_x}{2}$, and the other one weight $\frac{\Delta_x - \delta_x}{2}$. Since $\Delta_x \geq |\delta_x|$, both these weights are non-negative.

Finally, we also need to guarantee that the resulting tree has small weight. Interestingly, it turns out that the greedy approach works. Specifically, we set the edge weights $w_L(x)$ and $w_R(x)$ to the minimum values that are possible under the other limitations; the other limitations are imposed by the requirements that the resulting tree needs to be a dominating one, and that it needs to be an SPT.

Next, we turn to the formal analysis of our construction.

The following lemma shows that the numbers $\rho(x)$ satisfy (1), i.e., that indeed any leaf $v \in \text{Leaves}(x)$ can be selected for setting $\rho(x)$. Moreover, it implies that for an inner vertex x and a leaf $v \in \text{Leaves}(x)$, $\text{dist}(rt, v) \geq d_{T'}(x, v)$. Eventually we will need to guarantee $\text{dist}(rt, v) = d_T(rt, v)$, 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 2(i) for an illustration.)

Lemma 2.2: For any vertex x in T' and any vertex v in $\text{Leaves}(x)$, $\text{dist}(rt, v) - d_{T'}(x, v) = \rho(x) \geq 0$.

The proof of this lemma is omitted from this extended abstract

Consider an arbitrary pair v, z of points in $V \setminus \{rt\}$, let x be their least common ancestor in T' , and suppose without loss of generality that $v \in \text{Leaves}(L(x))$, $z \in \text{Leaves}(R(x))$. By the second statement of Observation 2.1,

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

Hence,

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

Corollary 2.3: The tree T' dominates the metric M' .

Next, we analyze the weight of the tree T' . For a vertex x in T' , let $f(x)$ and $l(x)$, standing for *first* and *last* of x , respectively, $f(x) \leq l(x)$, be the indices in $[\tilde{n}]$ for which $\text{Leaves}(x) = \{p_{f(x)}, p_{f(x)+1}, \dots, p_{l(x)}\}$. For a pair i, j of indices in $[\tilde{n}]$, $i \leq j$, let $\mathcal{W}(i, j) = \sum_{k=i}^{j-1} \text{dist}(p_k, p_{k+1})$ denote the sum of all edge weights along the subpath $(p_i, p_{i+1}, \dots, p_j)$ of H' . We start the weight analysis with the following claim.

Claim 2.4: For any vertex x in T' , there exist indices a and b in $[f(x), l(x)]$, $a \leq b$, such that $d_{T'}(x, p_a) \leq \mathcal{W}(a, l(x))$ and $d_{T'}(x, p_b) \leq \mathcal{W}(f(x), b)$.

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: 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, $w_L(x), w_R(x) \leq \max\{\Delta_x, |\delta_x|\} = \Delta_x$. Let $p_a \in \text{Leaves}(L(x))$ and $p_b \in \text{Leaves}(R(x))$ be two points for which $\Delta_x = \Delta(p_a, p_b) = \text{dist}(p_a, p_b) - (d_{T'}(L(x), p_a) + d_{T'}(R(x), p_b))$. Clearly both a and b are indices in $[f(x), l(x)]$. It follows that

$$\begin{aligned} d_{T'}(x, p_a) &\leq \Delta_x + d_{T'}(L(x), p_a) \\ &= \text{dist}(p_a, p_b) - d_{T'}(R(x), p_b) \\ &\leq \text{dist}(p_a, p_b) \\ &\leq \mathcal{W}(a, b) \leq \mathcal{W}(a, l(x)). \end{aligned}$$

Similarly, we get that $d_{T'}(x, p_b) \leq \mathcal{W}(f(x), b)$.

Otherwise, $|\delta_x| > \Delta_x$. Suppose without loss of generality that $w_L(x) \leq w_R(x)$. In this case $w_L(x) = 0$. By induction hypothesis, there exist indices a and b in $[f(L(x)), l(L(x))]$, with $d_{T'}(L(x), p_a) \leq \mathcal{W}(a, l(L(x)))$ and $d_{T'}(L(x), p_b) \leq \mathcal{W}(f(L(x)), b)$. Since $w_L(x) = 0$, we have $d_{T'}(x, p_a) = d_{T'}(L(x), p_a)$ and $d_{T'}(x, p_b) = d_{T'}(L(x), p_b)$. Also, $[f(L(x)), l(L(x))] \subset [f(x), l(x)]$. Thus, a and b serve as two indices in $[f(x), l(x)]$ for which $d_{T'}(x, p_a) \leq \mathcal{W}(a, l(L(x))) \leq \mathcal{W}(a, l(x))$ and $d_{T'}(x, p_b) \leq \mathcal{W}(f(L(x)), b) \leq \mathcal{W}(f(x), b)$. ■

The next lemma is the key to our weight analysis.

Lemma 2.5: For any vertex x in T' , $w_L(x) + w_R(x) \leq \mathcal{W}(f(x), l(x))$.

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: 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 $p_i \in \text{Leaves}(L(x))$ and $p_j \in \text{Leaves}(R(x))$ be two points for which $\Delta_x = \Delta(p_i, p_j) =$

$\text{dist}(p_i, p_j) - (d_{T'}(L(x), p_i) + d_{T'}(R(x), p_j))$. Hence,

$$\begin{aligned} w_L(x) + w_R(x) &= \Delta_x \leq \text{dist}(p_i, p_j) \\ &\leq \mathcal{W}(i, j) \leq \mathcal{W}(f(x), l(x)). \end{aligned}$$

Otherwise, $|\delta_x| > \Delta_x$. Suppose without loss of generality that $w_L(x) \leq w_R(x)$. In this case $|\delta_x| = -\delta_x$, and so $w_L(x) + w_R(x) = -\delta_x = \rho(R(x)) - \rho(L(x))$. By Claim 2.4, there exists an index i in $[f(L(x)), l(L(x))]$, such that $d_{T'}(L(x), p_i) \leq \mathcal{W}(f(L(x)), i)$. By Lemma 2.2, $\rho(L(x)) = \text{dist}(rt, p_i) - d_{T'}(L(x), p_i)$ and $\rho(R(x)) = \text{dist}(rt, p_j) - d_{T'}(R(x), p_j)$, for an arbitrary index j in $[f(R(x)), l(R(x))]$. It follows that

$$\begin{aligned} w_L(x) + w_R(x) &= \rho(R(x)) - \rho(L(x)) \\ &= \text{dist}(rt, p_j) - d_{T'}(R(x), p_j) \\ &\quad - \text{dist}(rt, p_i) + d_{T'}(L(x), p_i) \\ &\leq \text{dist}(rt, p_j) - \text{dist}(rt, p_i) \\ &\quad + d_{T'}(L(x), p_i) \\ &\leq \text{dist}(p_i, p_j) + \mathcal{W}(f(L(x)), i) \\ &\leq \mathcal{W}(f(L(x)), j) \leq \mathcal{W}(f(x), l(x)). \end{aligned}$$

Note that the depth of T' is $\log \tilde{n}$. The *level* of a vertex in T' is defined as its unweighted distance from rt . Denote by V_i the set of all vertices in T' of level i , for each $i \in [0, \log \tilde{n}]$. 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 $i \in [0, \log \tilde{n} - 1]$, and denote by W_i the sum of all edge weights in E_i . By Lemma 2.5, for each index $i \in [0, \log \tilde{n} - 1]$,

$$\begin{aligned} W_i &= \sum_{e \in E_i} w(e) = \sum_{x \in V_i} (w_L(x) + w_R(x)) \\ &\leq \sum_{x \in V_i} \mathcal{W}(f(x), l(x)) \leq w(H'). \end{aligned} \quad (2)$$

Corollary 2.6: The weight $w(T')$ of T' satisfies

$$w(T') = \sum_{i=0}^{\log \tilde{n} - 1} W_i \leq \log \tilde{n} \cdot w(H').$$

The tree T' consists of $2\tilde{n} - 1 = 2n - 3 = O(n)$ vertices. It is easy to verify that given the metric M , the tree T' can be constructed in $O(n^2)$ time, disregarding the time needed to compute the Hamiltonian path H' .

Next, we consider the tree T that is obtained from T' by adding to it an edge (rt, rt') of weight $\rho(rt')$. Note that $\rho(rt') \leq w(MST(M))$. Lemma 2.2, Corollary 2.3 and Corollary 2.6 imply the following result.

Proposition 2.7: 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 binary Steiner SPT for M (with non-negative edge weights), weight at most $\log \tilde{n} \cdot w(H') + w(MST(M))$ and $O(n)$ vertices.

Given the metric $M = (V, dist)$ one can construct a Hamiltonian path $L(M')$ for $M' = (V \setminus \{rt\}, dist)$ with weight at most $2 \cdot w(MST(M')) = O(w(MST(M)))$ within $O(n^2)$ time. To optimize the bounds on the weight and construction time of the tree T in Proposition 2.7, we take H' to be $L(M')$. Then the weight bound becomes $w(T) = O(\log n) \cdot w(MST(M))$, and the overall construction time is reduced to $O(n^2)$.

Theorem 2.8: Given an n -point metric M , the rooted tree (T, rt) returned by our construction is a binary Steiner SPT with lightness $O(\log n)$ and $O(n)$ vertices. The running time of this construction is $O(n^2)$.

Remark: The Steiner tree T that we constructed may contain edges of zero weight. All these edges may be contracted without affecting the lightness and distance properties of the resulting tree. On the other hand, its maximum degree may increase as a result of this operation. See also remarks (3) and (4) after Theorem 3.2.

The next lower bound states that our construction of Steiner SPTs of Theorem 2.8 is optimal. The proof of this lower bound is omitted from this extended abstract.

Theorem 2.9: For any sufficiently large integer n , there exists an n -point metric M and a designated point $rt \in M$, such that every Steiner SPT rooted at rt has lightness $\Omega(\log n)$.

3. STEINER SHALLOW-LIGHT TREES

Theorem 2.8 shows that for any n -point metric M and a designated point rt there exists a Steiner tree T that preserves the distances between rt and all other points of M and has lightness at most $O(\log n)$. In this section we generalize Theorem 2.8 and show that for any $\epsilon > 0$ there is 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 [6] Awerbuch et al. devised a construction of *spanning* SLTs with lightness $O(\frac{1}{\epsilon})$. 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. It is pretty obvious that the resulting tree satisfies the desired distance properties. Also, by Theorem 2.8, its lightness is $O(\log |\mathcal{B}|)$. If we could show that $|\mathcal{B}| = O(\frac{1}{\epsilon})$, 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 break-point, and we will thus get $|\mathcal{B}| = n - 1$. To overcome this obstacle we refine the bound on $w(T)$ from Proposition 2.7, 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 of points in M . We then use this refined bound to analyze the weight of our shallow-light trees.

The following lemma establishes our refined bound on the weight of the SPT T from Section 2.

Lemma 3.1: Suppose that $\sum_{v \in V \setminus \{rt\}} dist(rt, v) \leq \xi \cdot \eta$, for some pair $\xi \geq 1, \eta > 0$ of numbers. Then the weight $w(T)$ of $T = T(H')$ satisfies

$$w(T) \leq \eta + \lceil \log \xi \rceil \cdot w(H') + w(MST(M)).$$

Remark: Observe that $\sum_{v \in V \setminus \{rt\}} dist(rt, v) \leq (n - 1) \cdot w(MST(M))$, i.e., the assumption of the lemma holds for $\xi = n - 1, \eta = w(MST(M))$. Hence, this lemma generalizes the weight bound from Section 2.

Proof: As $w(T) = w(T') + \rho(rt') \leq w(T') + w(MST(M))$, it suffices to show that $w(T') \leq \eta + \lceil \log \xi \rceil \cdot w(H')$. Write $\tilde{n} = n - 1$. Suppose first that $\lceil \log \xi \rceil \geq \log \tilde{n}$. By Corollary 2.6, we have $w(T') \leq \log \tilde{n} \cdot w(H')$, and so

$$w(T') \leq \log \tilde{n} \cdot w(H') \leq \eta + \lceil \log \xi \rceil \cdot w(H').$$

We henceforth assume that $\lceil \log \xi \rceil \leq \log \tilde{n} - 1$. By construction, we have $w(T') = \sum_{i=0}^{\log \tilde{n} - 1} W_i$ and

$$\sum_{v \in V \setminus \{rt\}} d_T(rt, v) \geq \sum_{v \in V \setminus \{rt\}} d_{T'}(rt', v).$$

Since T is an SPT for M , we get that

$$\begin{aligned} \xi \cdot \eta &\geq \sum_{v \in V \setminus \{rt\}} dist(rt, v) = \sum_{v \in V \setminus \{rt\}} d_T(rt, v) \\ &\geq \sum_{v \in V \setminus \{rt\}} d_{T'}(rt', v). \end{aligned}$$

It is easy to show by double counting that

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

Therefore,

$$\begin{aligned} \xi \cdot \eta &\geq \sum_{v \in V \setminus \{rt\}} d_{T'}(rt', v) \\ &= \sum_{i=0}^{\log \tilde{n} - 1} 2^{\log \tilde{n} - (i+1)} \cdot W_i \\ &\geq \sum_{i=0}^{\log \tilde{n} - (\lceil \log \xi \rceil + 1)} 2^{\log \tilde{n} - (i+1)} \cdot W_i \\ &\geq \xi \cdot \left(\sum_{i=0}^{\log \tilde{n} - (\lceil \log \xi \rceil + 1)} W_i \right), \end{aligned}$$

and so

$$\sum_{i=0}^{\log \tilde{n} - (\lceil \log \xi \rceil + 1)} W_i \leq \eta.$$

Also, (2) implies that

$$\sum_{\log \tilde{n} - \lceil \log \xi \rceil}^{\log \tilde{n} - 1} W_i \leq \lceil \log \xi \rceil \cdot w(H').$$

Altogether,

$$\begin{aligned}
w(T') &= \sum_{i=0}^{\log \tilde{n}-1} W_i \\
&= \left(\sum_{i=0}^{\log \tilde{n} - (\lceil \log \xi \rceil + 1)} W_i \right) + \left(\sum_{i=\log \tilde{n} - \lceil \log \xi \rceil}^{\log \tilde{n}-1} W_i \right) \\
&\leq \eta + \lceil \log \xi \rceil \cdot w(H').
\end{aligned}$$

■

Now we proceed to extending our construction of Steiner SPTs from Section 2 to a construction of Steiner SLTs. Consider an n -point metric $M = (V, \text{dist})$, let T be an MST of M rooted at an arbitrary point $rt \in V$, and let D be an Euler tour of T starting at rt . For every vertex $v \in V$, remove from D all occurrences of v except for the first one, and denote by $L = L(M)$ the resulting Hamiltonian path of M . It is easy to verify that L can be constructed in $O(n^2)$ time, and $w(L) \leq 2 \cdot w(T) = 2 \cdot w(\text{MST}(M))$. Fix a parameter $\theta < \frac{1}{4}$. The value of θ will determine the values of the root-distortion and lightness of the constructed tree. We start with identifying a set of “break-points” $\mathcal{B} = \{B_1, B_2, \dots, B_k\}$, $\mathcal{B} \subseteq V$. The break-point B_1 is rt . The break-point B_{i+1} , $i \in [k-1]$, is the first vertex in L after B_i such that

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

Let $M_{\mathcal{B}}$ be the sub-metric of M induced by the point set of \mathcal{B} . Also, let $L' = (B_2, \dots, B_k)$ be the sub-path of L that contains the break-points of $\mathcal{B} \setminus \{rt\}$. By the triangle inequality, $w(L') \leq w(L) \leq 2 \cdot w(\text{MST}(M))$. By Proposition 2.7 and Lemma 3.1, 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 to be an SPT over \tilde{G} rooted at rt . The properties of the constructed Steiner tree S are summarized in the following theorem.

Theorem 3.2: For any n -point metric M , a designated point $rt \in M$ and a number $0 < \epsilon < \frac{1}{2}$, $\epsilon = 2\theta$, there is a Steiner tree S of M rooted at rt , having root-distortion at most $1 + \epsilon$, lightness $O(\log \frac{1}{\epsilon})$ and $O(n)$ vertices. The running time of this construction is $O(n^2)$.

Remarks:

- 1) The Steiner tree S constructed above may contain edges of zero weight. These edges, however, may be given arbitrarily small positive weights, without violating any of the bounds of Theorem 3.2.
- 2) Even though the maximum degree of this construction may be large, it can be easily decreased to $O(1)$ without affecting any of the other parameters.
- 3) If the input is an n -point metric M , the running time $O(n^2)$ of our constructions from Theorems 2.8 and 3.2 is linear in the input size. For the general case

where the input is a graph G rather than a metric M induced by G , one needs to compute the metric first. In this case the running time of our constructions is dominated by the time needed to compute the metric M . This task can be carried out in time $O(n^\omega)$ if G is an unweighted graph, where $\omega \approx 2.376$ is the matrix multiplication exponent [19], and in time $O(n^3 \sqrt{\frac{\log \log n}{\log n}})$ if G is an arbitrary weighted graph [45].

- 4) One cannot significantly improve the running time of our constructions from Theorems 2.8 and 3.2 in the cell-probe model of computation. In this model each distance can be fetched in $O(1)$ time, and one does not need to spend $\Omega(n^2)$ time just to read the entire input. Yet, it is known [31] that any deterministic (respectively, randomized) algorithm that computes a spanning tree with lightness β in this model requires $\Omega(n^2) - O(n \cdot \beta) = \Omega(n^2)$ (resp., $\Omega(\frac{n^2}{\beta})$) time. It is easy to see that the lower bound of [31] extends to Steiner trees as well.

The next lower bound states that our construction of Steiner SLTs of Theorem 3.2 is optimal. The proof of this lower bound is omitted from this extended abstract.

Theorem 3.3: For any sufficiently large integer n and any parameter $\epsilon = \Omega(\frac{1}{n})$, there exists an n -point metric M and a designated point $rt \in M$, such that every Steiner tree rooted at rt with root-distortion at most $1 + \epsilon$ has lightness $\Omega(\log \frac{1}{\epsilon})$.

4. LOWER BOUNDS FOR EUCLIDEAN SPANNING SLTs

In this section we establish a tradeoff of $1 + \epsilon$ versus $\Omega(\frac{1}{\epsilon})$ between the average root-distortion and lightness of Euclidean *spanning* trees. We demonstrate that this tradeoff holds for *any* choice of root-vertex.

Let C_n be a set of n points that are uniformly spaced around the boundary C of the unit circle, centered at the origin $(0, 0)$ in the plane. Also, let $\tilde{C}_{n+1} = C_n \cup \{(0, 0)\}$. Observe that for both these point configurations C_n and \tilde{C}_{n+1} , the weight of the MST is $O(1)$.

We start with analyzing spanning trees of \tilde{C}_{n+1} that are rooted at a *specific* root vertex, namely $rt = (0, 0)$. Such trees are particularly convenient for analysis. First, the degree of rt in such a tree (henceforth, *root-degree*) cannot be greater than its lightness by more than a constant factor; hence, it is sufficient to understand the tradeoff between the root-distortion and the root-degree. Second, the Euclidean distance between rt and all other points is the same (and equal to 1); in particular, the distortion between rt and every other point in the tree is equal to their distance in it. Then, having established the lower bound for this simple configuration, we extend our argument to the general case where the root vertex can be chosen *arbitrarily*.

Let T be a spanning tree for \tilde{C}_{n+1} rooted at $rt = (0, 0)$, and denote the root-degree of T by γ . Let $v_1, v_2, \dots, v_\gamma$ be

the neighbors of rt in clockwise order along C , starting with an arbitrary one of them, and define $v_{\gamma+1} = v_1$. For each index $i \in [\gamma]$, denote by n_i the number of points in C_n that lie on the arc A_i of C between v_i and v_{i+1} (excluding v_i and including v_{i+1}). Observe that $\sum_{i=1}^{\gamma} n_i = n$.

Next, we provide a lower bound for the average³ root-distortion $\lambda(T, rt)$ of T with respect to rt . Fix an arbitrary index $i \in [\gamma]$, and denote the points of C_n that belong to A_i in clockwise order by p_1, p_2, \dots, p_{n_i} . (Note that $p_{n_i} = v_{i+1}$.) We argue that for each index $j \in [n_i]$, the distortion $\varphi_T(rt, p_j)$ between rt and p_j in T satisfies

$$\varphi_T(rt, p_j) \geq 1 + \Omega\left(\frac{1}{n}\right) \cdot \min\{j, n_i - j\}. \quad (3)$$

Consider the path P_j in T between rt and p_j , for some index $j \in [n_i]$. Since rt is not incident in T to any point in the interior $A'_i = A_i \setminus \{v_i, v_{i+1}\}$ of A_i , the path P_j must take a detour through some point $q \in C_n \setminus A'_i$, and so

$$\begin{aligned} \|q - p_j\| &\geq \min\{\|v_i - p_j\|, \|v_{i+1} - p_j\|\} \\ &= \Omega\left(\frac{1}{n}\right) \cdot \min\{j, n_i - j\}. \end{aligned}$$

It follows that

$$\begin{aligned} \varphi_T(rt, p_j) &= d_T(rt, p_j) \geq \|rt - q\| + \|q - p_j\| \\ &\geq 1 + \Omega\left(\frac{1}{n}\right) \cdot \min\{j, n_i - j\}, \end{aligned}$$

which proves (3). Consequently, we get that $\sum_{j \in [n_i]} \varphi_T(rt, p_j) \geq n_i + \Omega\left(\frac{n_i^2}{n}\right)$. Summing over all indices $i \in [\gamma]$, we obtain

$$\begin{aligned} \sum_{v \in C_n} \varphi_T(rt, v) &\geq \sum_{i \in [\gamma]} \left(n_i + \Omega\left(\frac{n_i^2}{n}\right) \right) \\ &= n + \Omega\left(\frac{1}{n}\right) \cdot \sum_{i \in [\gamma]} n_i^2 \\ &\geq n + \Omega\left(\frac{1}{n}\right) \cdot \frac{n^2}{\gamma} \\ &= n + \Omega\left(\frac{n}{\gamma}\right). \end{aligned}$$

(The last inequality holds by the Cauchy-Schwarz inequality.) We conclude that

$$\chi(T, rt) \geq \lambda(T, rt) = \frac{\sum_{v \in C_n} \varphi_T(rt, v)}{n} \geq 1 + \Omega\left(\frac{1}{\gamma}\right).$$

Recall that the lightness $\Psi(T)$ of T is at least $\Omega(\gamma)$, and set $\epsilon = \Theta\left(\frac{1}{\gamma}\right)$. It follows that if the (average) root-distortion of T is at most $1 + \epsilon$, then both its root-degree and lightness are at least $\Omega\left(\frac{1}{\epsilon}\right)$.

In the full version we show that the same tradeoff between the (average) root-distortion and lightness applies when one

³Note that the (worst-case) root-distortion $\chi(T, rt)$ is bounded from below by the average root-distortion $\lambda(T, rt)$.

is allowed to select the root vertex at will. This result is summarized in the following theorem.

Theorem 4.1: For any sufficiently large integer n , there exists a 2-dimensional Euclidean n -point metric M , such that for any parameter $\epsilon = \Omega\left(\frac{1}{n}\right)$ and any designated point $rt \in M$, every spanning tree for M rooted at rt with (average) root-distortion at most $1 + \epsilon$ has lightness at least $\Omega\left(\frac{1}{\epsilon}\right)$.

We also obtain the following result, which is symmetric to Theorem 4.1.

Theorem 4.2: For any sufficiently large integer n and any parameter $\epsilon = \Omega\left(\frac{1}{n}\right)$, there exists a 2-dimensional Euclidean n -point metric M , such that for any designated point $rt \in M$, every spanning tree for M rooted at rt with lightness at most $1 + \epsilon$ has (average) root-distortion at least $\Omega\left(\frac{1}{\epsilon}\right)$.

We remark that the lower bound of Theorem 4.2 *does not hold* for the point configuration C_n . Indeed, let p_1, p_2, \dots, p_n be the n points of C_n in clockwise order, starting with the point $p_1 = (0, 1)$. Consider the tree T formed by connecting each point p_i with its two closest points p_{i-1} and p_{i+1} in C_n , and removing one edge incident to the point $p_j \in C_n$ closest to $(0, -1)$. The tree T is an MST for C_n , i.e., its lightness is 1. Also, its root-distortion with respect to p_1 is at most $\frac{\pi}{2}$.

Therefore, the lower bound of Theorem 4.2 is proved on a different point configuration.

Up to constant factors Theorems 4.1 and 4.2 settle all three aforementioned questions of Khuller et al. [33] (see Section 1.3).

ACKNOWLEDGMENT

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

REFERENCES

- [1] I. Abraham, Y. Bartal, and O. Neiman, "Nearly tight low stretch spanning trees," in *Proc. of 49th FOCS*, 2008, pp. 781–790.
- [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.*, vol. 24, no. 1, pp. 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*, vol. 9, pp. 81–100, 1993.
- [4] B. Aronov, M. de Berg, O. Cheong, J. Gudmundsson, H. J. Haverkort, M. H. M. Smid, and A. Vigneron, "Sparse geometric graphs with small dilation," *Comput. Geom.*, vol. 40, no. 3, pp. 207–219, 2008.
- [5] B. Awerbuch, A. Baratz, and D. Peleg, "Cost-sensitive analysis of communication protocols," in *Proc. of 9th PODC*, 1990, pp. 177–187.
- [6] —, "Efficient broadcast and light-weight spanners," *Manuscript*, 1991.
- [7] Y. Bartal, "Probabilistic approximations of metric spaces and its algorithmic applications," in *Proc. of 37th FOCS*, 1996, pp. 184–193.

- [8] —, “On approximating arbitrary metrics by tree metrics,” in *Proc. of 30th STOC*, 1998, pp. 161–168.
- [9] Y. Ben-Shimol, A. Dvir, and M. Segal, “SPLAST: a novel approach for multicasting in mobile wireless ad hoc networks,” in *Proc. of 15th PIMRC*, 2004, pp. 1011–1015.
- [10] P. Berman and V. Ramaiyer, “Improved approximations for the Steiner tree problem,” *J. Algorithms*, vol. 17, no. 3, pp. 381–408, 1994.
- [11] M. W. Bern and P. E. Plassmann, “The Steiner problem with edge lengths 1 and 2,” *Inf. Process. Lett.*, vol. 32, no. 4, pp. 171–176, 1989.
- [12] K. Bharath-Kumar and J. M. Jaffe, “Routing to multiple destinations in computer networks,” *IEEE Trans. on Commun.*, vol. COM-31, pp. 343–351, 1983.
- [13] B. Bollobás, D. Coppersmith, and M. Elkin, “Sparse distance preservers and additive spanners,” *SIAM J. Discrete Math.*, vol. 19, no. 4, pp. 1029–1055, 2005.
- [14] R. Braynard, D. Kotic, A. Rodriguez, J. Chase, and A. Vahdat, “Opus: an overlay peer utility service,” in *Prof. of 5th OPENARCH*, 2002.
- [15] M. Charikar, C. Chekuri, T. Cheung, Z. Dai, A. Goel, S. Guha, and M. Li, “Approximation algorithms for directed Steiner problems,” *J. Algorithms*, vol. 33, no. 1, pp. 73–91, 1999.
- [16] 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*, 1991, pp. 170–173.
- [17] —, “Provably good algorithms for performance-driven global routing,” in *Proc. of 5th ISCAS*, 1992, pp. 2240–2243.
- [18] —, “Provably good performance-driven global routing,” *IEEE Trans. on CAD of Integrated Circuits and Sys.*, vol. 11, no. 6, pp. 739–752, 1992.
- [19] D. Coppersmith and S. Winograd, “Matrix multiplication via arithmetic progressions,” *J. Symb. Comput.*, vol. 9, no. 3, pp. 251–280, 1990.
- [20] T. H. Corman, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms, 2nd edition*. Boston, MA: McGraw-Hill Book Company, 2001.
- [21] R. Cristescu, B. Beferull-Lozano, M. Vetterli, and R. Wattenhofer, “Network correlated data gathering with explicit communication: NP-completeness and algorithms,” *IEEE/ACM Trans. Netw.*, vol. 14, no. 1, pp. 41–54, 2006.
- [22] Y. Dinitz, M. Elkin, and S. Solomon, “Low-light trees, and tight lower bounds for Euclidean spanners,” *Discrete & Computational Geometry*, vol. 43, no. 4, pp. 736–783, 2010.
- [23] M. Elkin, Y. Emek, D. Spielman, and S.-H. Teng, “Lower stretch spanning trees,” *SIAM J. Comput.*, vol. 38, no. 2, pp. 608–628, 2008.
- [24] M. Elkin and S. Solomon, “Narrow-shallow-low-light trees with and without Steiner points,” *SIAM J. Discrete Math.*, vol. 25, no. 1, pp. 181–210, 2011.
- [25] J. Fakcharoenphol, S. Rao, and K. Talwar, “A tight bound on approximating arbitrary metrics by tree metrics,” in *Proc. of 35th STOC*, 2003, pp. 448–455.
- [26] E. N. Gilbert and H. O. Pollak, “Steiner minimal trees,” *SIAM J. Applied Math.*, vol. 16, no. 1, pp. 1–29, 1968.
- [27] A. Goel and K. Munagala, “Balancing Steiner trees and shortest path trees online,” in *Proc. of 11th SODA*, 2000, pp. 562–563.
- [28] A. Gupta, “Steiner points in tree metrics don’t (really) help,” in *Proc. of 12th SODA*, 2001, pp. 220–227.
- [29] M. T. Hajiaghayi, G. Kortsarz, and M. R. Salavatipour, “Approximating buy-at-bulk and shallow-light k -Steiner trees,” *Algorithmica*, vol. 53, no. 1, pp. 89–103, 2009.
- [30] F. K. Hwang and D. Z. Du, “Steiner minimal trees on chinese checkerboards,” *Mathematics Magazine*, vol. 64, no. 5, pp. 332–339, 1991.
- [31] P. Indyk, “Sublinear time algorithms for metric space problems,” in *Proc. of 31th STOC*, 1999, pp. 428–434.
- [32] J. M. Jaffe, “Distributed multi-destination routing: the constraints of local information,” *SIAM J. Comput.*, vol. 14, pp. 875–888, 1985.
- [33] S. Khuller, B. Raghavachari, and N. E. Young, “Balancing minimum spanning and shortest path trees,” in *Proc. of 4th SODA*, 1993, pp. 243–250.
- [34] V. P. Kompella, J. Pasquale, and G. C. Polyzos, “Multicast routing for multimedia communication,” *IEEE/ACM Trans. Netw.*, vol. 1, no. 3, pp. 286–292, 1993.
- [35] G. Konjevod, R. Ravi, and F. S. Salman, “On approximating planar metrics by tree metrics,” *Inf. Process. Lett.*, vol. 80, no. 4, pp. 213–219, 2001.
- [36] G. Kortsarz and D. Peleg, “Approximating the weight of shallow Steiner trees,” *Discrete Applied Mathematics*, vol. 93, no. 2-3, pp. 265–285, 1999.
- [37] D. Kotic and A. Vahdat, “Latency versus cost optimizations in hierarchical overlay networks,” *Technical report, Duke University*, no. CS-2001-04, 2002.
- [38] M. V. Marathe, R. Ravi, R. Sundaram, S. S. Ravi, D. J. Rosenkrantz, and H. B. Hunt III, “Bicriteria network design problems,” *J. Algorithms*, vol. 28, no. 1, pp. 142–171, 1998.
- [39] J. Naor and B. Schieber, “Improved approximations for shallow-light spanning trees,” in *Proc. of 38th FOCS*, 1997, pp. 536–541.
- [40] D. Peleg, *Distributed Computing: A Locality-Sensitive Approach*. Philadelphia, PA: SIAM, 2000.
- [41] D. Peleg and A. Schäffer, “Graph spanners,” *J. Graph Theory*, vol. 13, no. 1, pp. 99–116, 1989.
- [42] S. Rao and W. D. Smith, “Approximating geometrical graphs via “spanners” and “banyans”,” in *Proc. of 30th STOC*, 1998, pp. 540–550.
- [43] F. S. Salman, J. Cheriyan, R. Ravi, and S. Subramanian, “Approximating the single-sink link-installation problem in network design,” *SIAM Journal on Optimization*, vol. 11, no. 3, pp. 595–610, 2001.
- [44] H. Shpungin and M. Segal, “Near optimal multicriteria spanner constructions in wireless ad-hoc networks,” *IEEE/ACM Transactions on Networking*, vol. 18, no. 6, pp. 1963–1976, 2010.
- [45] T. Takaoka, “A new upper bound on the complexity of the all pairs shortest path problem,” *Inf. Process. Lett.*, vol. 43, no. 4, pp. 195–199, 1992.
- [46] J. Vogel, J. Widmer, D. Farin, M. Mauve, and W. Effelsberg, “Priority-based distribution trees for application-level multicast,” in *Proc. of 2nd NETGAMES*, 2003, pp. 148–157.
- [47] P. von Rickenbach and R. Wattenhofer, “Gathering correlated data in sensor networks,” in *Proc. of DIALM-POMC*, 2004, pp. 60–66.
- [48] B. Y. Wu, K.-M. Chao, and C. Y. Tang, “Light graphs with small routing cost,” *Networks*, vol. 39, no. 3, pp. 130–138, 2002.