

STEINER SHALLOW-LIGHT TREES ARE EXPONENTIALLY LIGHTER THAN SPANNING ONES*

MICHAEL ELKIN[†] AND SHAY SOLOMON[‡]

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 (resp., 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, Baratz, and Peleg [*Proceedings of the 9th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 1990, pp. 177–187, *Efficient Broadcast and Light-Weight Spanners*, manuscript, 1991] and Khuller, Raghavachari, and Young [*Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1993, pp. 243–250]. 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. In the regime $\epsilon = 0$ 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.

Key words. minimum spanning tree, shortest-path tree, shallow-light tree, Steiner point, Steiner tree

AMS subject classification. 05C85

DOI. 10.1137/13094791X

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., [21, Chaps. 24, 25 and the references therein]).

*Received by the editors December 6, 2013; accepted for publication (in revised form) June 24, 2015; published electronically August 4, 2015. A preliminary version of this paper appeared in *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2011, pp. 373–382.

<http://www.siam.org/journals/sicomp/44-4/94791.html>

[†]Department of Computer Science, Ben-Gurion University of the Negev, Beer-Sheva 84105, Israel (elkinm@cs.bgu.ac.il). The work of this author was supported by BSF grant 2008430, ISF grant 87209011, and the Lynn and William Frankel Center for Computer Sciences.

[‡]Department of Computer Science and Applied Mathematics, The Weizmann Institute of Science, Rehovot 76100, Israel (shay.solomon@weizmann.ac.il). The work of this author was supported by the Koshland Center for Basic Research. Part of this work was done while this author was a graduate student in the Department of Computer Science, Ben-Gurion University of the Negev, under the support of Clore Fellowship grant 81265410, BSF grant 2008430, and ISF grant 87209011.

The desirability of having a single tree that combines the properties of the MST and the SPT was realized already in the mid-1980s [13, 33]. The specific notion of *shallow-light tree* (shortly, SLT) was introduced in [6, 7, 34].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) $w(T) \leq \beta \cdot w(MST(G))$. A tree T that satisfies only the first requirement will be referred to as an α -shortest-path tree (shortly, α -SPT) of G . The parameter α (resp., β) will be called the *root-distortion* (resp., *lightness*) of the tree T . Awerbuch, Baratz, and Peleg [6, 7] and Khuller, Raghavachari, and Young [34] 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 [34]. 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 [15, 38, 47], in wireless and sensor networks [10, 22, 45, 48], and in the message-passing model of distributed computing [6, 7]. Other applications of SLTs include network and VLSI-circuit design [17–19, 44] and routing [50]. SLTs were also used for constructing spanners [7, 40]. Closely related tree structures, such as *light approximate routing trees* and *shallow-low-light trees*, were investigated in [23, 25, 50]. Moreover, some of the constructions in [23, 25, 50] are based on constructions 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 , 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, Raghavachari, and Young [34] 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.

A noteworthy 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 SPTs) 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 SPT 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. [3] 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 $e = (u, v) \in E$, $\mathbb{E}_{T \in \mathcal{D}}[\frac{d_T(u, v)}{w(e)}] = 2^{O(\sqrt{\log n \log \log n})}$. Such a distribution is

¹Khuller, Raghavachari, and Young [34] called the same notion *light approximate spanning tree* (shortly, *LAST*).

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.

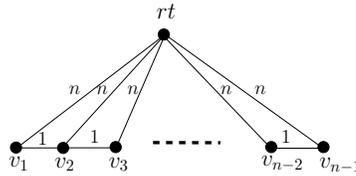


FIG. 1. All edges incident to rt have weight n , while other edges have unit weight. The SPT with respect to rt is the star rooted at rt ; its weight is $n(n-1)$. The MST contains a single edge (rt, v_i) of weight n and the entire path $(v_1, v_2, \dots, v_{n-1})$ of unit-weight edges. Its weight is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

called *probabilistic tree embedding* [8], and the value $\max_{e=(u,v) \in E} \mathbb{E}[\frac{d_T(u,v)}{w(e)}]$ is called the *distortion* of the embedding. Bartal [8,9] and Fakcharoenphol, Rao, and Talwar [26] showed that using Steiner points one can drastically improve the bound of [3], and they devised probabilistic tree embeddings with distortion $O(\log n)$. The bound $O(\log n)$ is also known to be optimal up to constant factors [3,8]. However, Konjevod, Ravi, and Salman [36] and Gupta [29] demonstrated that the same bounds (up to constant factors) as those of Bartal [8,9] and Fakcharoenphol, Rao, and Talwar [26] can be obtained without Steiner points, i.e., by using only *spanning trees* of the metric M_G induced by G . (Such spanning trees are, however, allowed to use *Steiner edges* rather than Steiner points.) Moreover, more recent studies [1,2,24] 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 occurs in the context of *low-light trees*, which combine small lightness with small depth [23,25]. 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 same (up to constant factors) lightness and depth as those of T [25].

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. [4], 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* [14].

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 [4, 14, 25, 29, 36]. 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 sharp contrast to all these examples, shallow-light trees can be *exponentially* improved by using Steiner points.

Interestingly, we also demonstrate that while Steiner *points* are very helpful in the context of SLTs, this is not the case for Steiner *edges*. 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-distortion are bounded by the weight and root-distortion of T , respectively. (Hence there exist metrics (and not just graphs) for which any $(1 + \epsilon)$ -SPT has lightness $1 + \Omega(\frac{1}{\epsilon})$.)

1.3. Related work. Shallow-low-light trees were studied in [23,25]. In addition to small root-distortion and lightness, these trees have small depth. However, similarly to the SLTs of Awerbuch, Baratz, and Peleg [6,7] and Khuller, Raghavachari, and Young [34], the shallow-low-light trees of [23,25] 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 [30,37,39]. Some of this research considered Steiner SLTs [16,30,37]. However, the word ‘‘Steiner’’ is used in these papers with 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 [16,30,37] 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 [28]. Heuristics for finding SLTs were developed in [35].

There is also a vast literature on the Euclidean minimum Steiner tree [11,12,27,31]. 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 SPT (with respect to an arbitrary vertex $rt \in C_n$) that uses only Euclidean Steiner points has lightness $\Omega(n)$. More generally, 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 $\Omega(\sqrt{\frac{1}{\epsilon}})$. (See Appendix B for details.) 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 SPTs and Steiner shallow-light trees, respectively. We remark that lower bounds on the power of Euclidean Steiner points were shown in [5,43].

We note that the Steiner trees, of Bartal [8,9] and of Fakcharoenphol, Rao, and Talwar [26] have a particularly useful structure; specifically, they are hierarchically well-separated trees (HSTs). (See [8] for the definition.) The Steiner trees that we devise in the current paper do not have this structure; i.e., they are not HSTs.

1.4. 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. Matching lower bounds for Steiner SPTs and SLTs are provided in section 4. The argument which shows that Steiner edges do not help appears in section 5. In Appendix A we prove a technical lemma, used in the lower bounds of section 4. In Appendix B we show that any construction of SLTs for Euclidean metrics that employs only Euclidean Steiner points and has root-distortion at most $1 + \epsilon$ has lightness $\Omega(\sqrt{\frac{1}{\epsilon}})$.

1.5. 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 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 SPTs 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 SPT 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 $V \setminus \{rt\}$.

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 ,

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

Given (1), we have $\rho(x) \leq w(MST(M))$ by the triangle inequality for any vertex x of T' . In particular, it holds that $\rho(rt') \leq w(MST(M))$. We show (Proposition 2.8) 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 SPT of M with respect to rt of weight $w(T) = O(\log n) \cdot w(H') + \rho(rt')$. (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 will be

$$\begin{aligned} \Psi(T) &= \frac{w(T)}{w(MST(M))} = \frac{O(\log n) \cdot w(H') + \rho(rt')}{w(MST(M))} \\ &\leq \frac{O(\log n) \cdot w(H') + w(MST(M))}{w(MST(M))} = O(\log n). \end{aligned}$$

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 , $L(x) =$

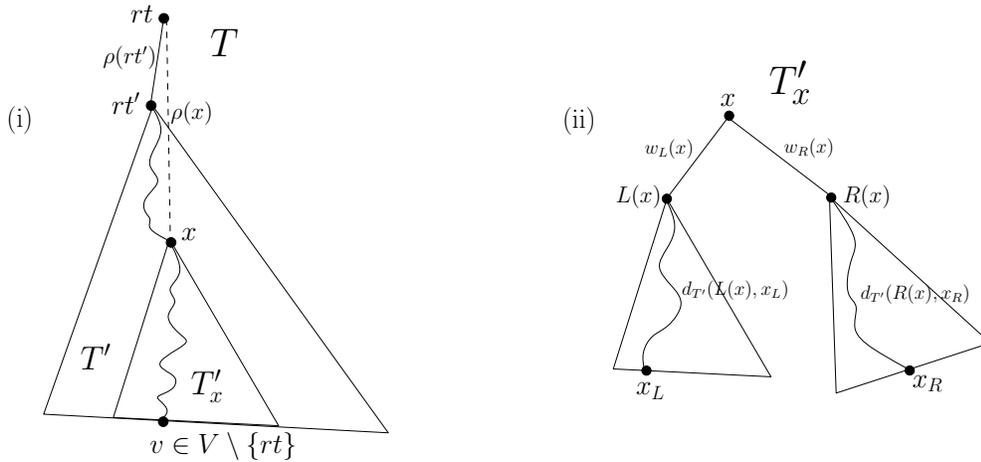


FIG. 2. (i) The tree (T, rt) is obtained from T' by adding to it an edge (rt, rt') of weight $\rho(rt')$. By (1), $d_{T'}(x, v) = \text{dist}(rt, v) - \rho(x)$. Hence $d_T(rt, x) = d_T(rt, v) - d_{T'}(x, v) = (d_T(rt, v) - \text{dist}(rt, v)) + \rho(x) = \rho(x)$. The last equation holds because, as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

$R(x) = \text{NULL}$ and $\text{Leaves}(x) = \{x\}$, whereas for an inner vertex x , $\text{Leaves}(x) = \text{Leaves}(L(x)) \cup \text{Leaves}(R(x))$. Observe that $\text{Leaves}(rt') = \{p_1, p_2, \dots, p_{\bar{n}}\} = V \setminus \{rt\}$. 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 $\rho(x) = \text{dist}(rt, x)$, and we 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 $x_L \in \text{Leaves}(L(x)), x_R \in \text{Leaves}(R(x))$. Let $\Delta(x_L, x_R) = \text{dist}(x_L, x_R) - (d_{T'}(L(x), x_L) + d_{T'}(R(x), x_R))$. 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 by Δ_x , is defined as the maximum distance surplus over all pairs (x_L, x_R) with $x_L \in \text{Leaves}(L(x)), x_R \in \text{Leaves}(R(x))$, i.e.,

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

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 $\rho(x) = \text{dist}(rt, v) - d_{T'}(x, v)$. (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. $w_L(x) + w_R(x) = \max\{\Delta_x, |\delta_x|\}$, and
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 SPTs 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, $\rho(L(x))$ (resp., $\rho(R(x))$) indicates how close the root rt is 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 $Leaves(L(x))$ (resp., $Leaves(R(x))$); it is convenient to think of this value as a “distance” between rt and the set $Leaves(L(x))$ (resp., $Leaves(R(x))$). 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 $Leaves(L(x))$ are located “farther” from rt than the vertices of $Leaves(R(x))$, 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 $Leaves(x)$. 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 $Leaves(x)$.

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 dominance 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 nonnegative.

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 $v \in Leaves(x)$ can be selected for setting $\rho(x)$.

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 we prove it for x . Consider an arbitrary pair v, z of vertices in $Leaves(x)$.

Next, we show that $dist(rt, v) - d_{T'}(x, v) = dist(rt, z) - d_{T'}(x, z)$, 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 z 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 $v, z \in T'_{L(x)}$, then $dist(rt, v) - d_{T'}(L(x), v) = dist(rt, z) - d_{T'}(L(x), z)$. However, $dist(rt, v) - d_{T'}(x, v) = (dist(rt, v) - d_{T'}(L(x), v)) - w_L(x)$, and similarly, $dist(rt, z) - d_{T'}(x, z) = (dist(rt, z) - d_{T'}(L(x), z)) - w_L(x)$. Hence $dist(rt, v) - d_{T'}(x, v) = dist(rt, z) - d_{T'}(x, z)$, as required. (See Figure 3 for an illustration.)

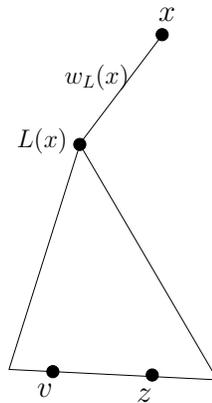


FIG. 3. *The case $v, z \in Leaves(L(x))$.*

The case when $v, z \in T'_{R(x)}$ is analogous.

We may henceforth suppose without loss of generality that $v \in Leaves(L(x))$ and $z \in Leaves(R(x))$. By construction, $d_{T'}(x, v) = d_{T'}(L(x), v) + w_L(x)$ and $d_{T'}(x, z) =$

$d_{T'}(R(x), z) + w_R(x)$. By the third statement of Observation 2.1 and the induction hypothesis,

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

Lemma 2.2 shows that for an inner vertex x and a leaf $v \in \text{Leaves}(x)$, $\text{dist}(rt, v) = d_{T'}(x, v) + \rho(x)$. The next lemma shows that $\rho(x) \geq 0$. Hence $\text{dist}(rt, v) \geq d_{T'}(x, v)$. Eventually we will need to guarantee $\text{dist}(rt, v) = d_T(rt, v)$, where $T = T' \cup \{(rt, rt')\}$ is our ultimate SPT. 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 $\rho(x) = \rho(rt')$. In general, the distance between rt and x in T would need to be equal to $\rho(x)$; indeed, to guarantee $\text{dist}(rt, v) = d_T(rt, v)$, it must hold that $d_T(rt, x) = d_T(rt, v) - d_{T'}(x, v) = (d_T(rt, v) - \text{dist}(rt, v)) + \rho(x) = \rho(x)$. (See Figure 2(i) for an illustration.)

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 we 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 $x_L \in \text{Leaves}(L(x))$ and $x_R \in \text{Leaves}(R(x))$ be two vertices for which $\Delta_x = \Delta(x_L, x_R) = \text{dist}(x_L, x_R) - d_{T'}(L(x), x_L) - d_{T'}(R(x), x_R)$. By Lemma 2.2, $\rho(L(x)) = \text{dist}(rt, x_L) - d_{T'}(L(x), x_L)$ and $\rho(R(x)) = \text{dist}(rt, x_R) - d_{T'}(R(x), x_R)$. Altogether,

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

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

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

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

$$\rho(x) = \text{dist}(rt, x_L) - d_{T'}(x, x_L) = \text{dist}(rt, x_L) - d_{T'}(L(x), x_L) - w_L(x) \geq 0. \quad \square$$

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, z of vertices in $\text{Leaves}(x)$, $d_{T'}(v, z) \geq \text{dist}(v, z)$.*

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 we prove it for x .

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

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

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 $Leaves(x) = \{p_{f(x)}, p_{f(x)+1}, \dots, p_{l(x)}\}$. For a pair i, j of indices in $[\tilde{n}]$, $i \leq j$, let $\mathcal{W}(i, j) = \sum_{k=i}^{j-1} dist(p_k, p_{k+1})$ denote the sum of all edge weights along the subpath $(p_i, p_{i+1}, \dots, p_j)$ of H' . We will soon show (in Lemma 2.6) that the sum of weights of the two edges descending from a vertex x is dominated by the weight of the subpath $(p_{f(x)}, p_{f(x)+1}, \dots, p_{l(x)})$ of the Hamiltonian path H' traversing all the leaves of the subtree T'_x rooted at x . To this end we first argue that there exists a leaf p_i of T'_x whose distance from x in T'_x is dominated by the weight of the subpath of H' between p_i and $p_{l(x)}$, and we also argue that there exists another leaf p_j , $i \leq j$, whose distance from x in T'_x is dominated by the weight of the subpath of H' between $p_{f(x)}$ and p_j . We will later use this technical statement to prove Lemma 2.6.

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

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 $[\tilde{n}]$ such that $x = p_k$. Thus $f(x) = l(x) = k$, and for $i = j = k$ we have $d_{T'}(x, p_i) = d_{T'}(x, p_j) = \mathcal{W}(i, l(x)) = \mathcal{W}(f(x), j) = 0$.

Induction step. We assume that the statement holds for the two children $L(x)$ and $R(x)$ of x , and we prove it for x .

Suppose first that $|\delta_x| \leq \Delta_x$. By the first two assertions 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 $[f(x), l(x)]$. It follows that

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

(The penultimate inequality holds by the triangle inequality.) Similarly, we get that

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

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 $[f(L(x)), l(L(x))]$, with $d_{T'}(L(x), p_i) \leq \mathcal{W}(i, l(L(x)))$ and $d_{T'}(L(x), p_j) \leq$

$\mathcal{W}(f(L(x)), j)$. Since $w_L(x) = 0$, we have $d_{T'}(x, p_i) = d_{T'}(L(x), p_i)$ 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 $d_{T'}(x, p_i) \leq \mathcal{W}(i, l(L(x))) \leq \mathcal{W}(i, l(x))$ and $d_{T'}(x, p_j) \leq \mathcal{W}(f(L(x)), j) \leq \mathcal{W}(f(x), j)$. \square

The next lemma is the key to our weight analysis. It 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 subpath of the Hamiltonian path H that traverses all vertices from $Leaves(x)$.

LEMMA 2.6. *For any vertex x in T' , $w_L(x) + w_R(x) \leq \mathcal{W}(f(x), l(x))$.*

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. We assume that the statement holds for the two children $L(x)$ and $R(x)$ of x , and we 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

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

(The penultimate 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 $d_{T'}(L(x), p_a) \leq \mathcal{W}(f(L(x)), a)$. 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

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

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

Notice that the depth of T' is $\log \tilde{n}$. 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 $i \in [0, \log \tilde{n}]$. 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 $i \in [0, \log \tilde{n} - 1]$, and denote by W_i the sum of all edge weights in E_i . Observe that the weight $w(T')$ of T' satisfies $w(T') = \sum_{i=0}^{\log \tilde{n}-1} W_i$. The following lemma implies that $w(T') \leq \log \tilde{n} \cdot w(H')$.

LEMMA 2.7. *For each index $i \in [0, \log \tilde{n}]$, $W_i \leq w(H')$.*

Proof. Fix an arbitrary index $i \in [0, \log \tilde{n}]$. 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 $f(x) \leq l(x) < f(x') \leq l(x')$ or $f(x') \leq l(x') < f(x) \leq l(x)$

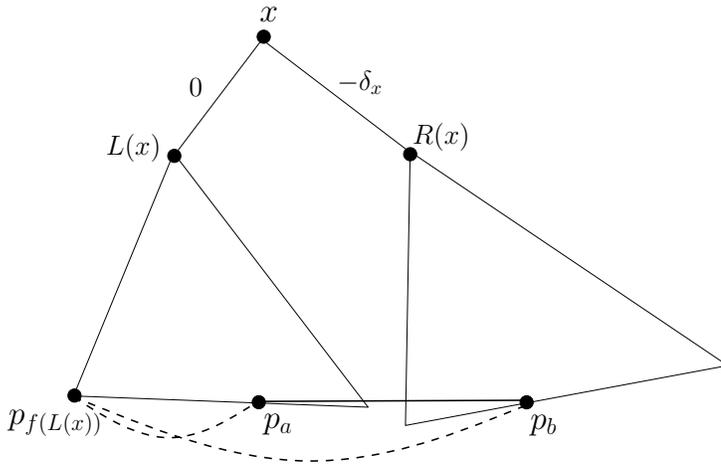


FIG. 4. An illustration for the inequality $\text{dist}(p_a, p_b) + \mathcal{W}(f(L(x)), a) \leq \mathcal{W}(f(L(x)), b)$.

must hold. It follows that $\sum_{x \in V_i} \mathcal{W}(f(x), l(x)) \leq \mathcal{W}(1, \tilde{n}) = w(H')$. Lemma 2.6 implies that $w_L(x) + w_R(x) \leq \mathcal{W}(f(x), l(x))$ for any $x \in V_i$. Altogether,

$$W_i = \sum_{x \in V_i} w_L(x) + w_R(x) \leq \sum_{x \in V_i} \mathcal{W}(f(x), l(x)) \leq w(H'). \quad \square$$

Note that the tree T' consists of $2\tilde{n} - 1 = 2n - 3 = O(n)$ vertices. Also, it is easy to verify that given the metric M and the Hamiltonian path H' , the tree T' can be constructed in $O(n^2)$ time. (Indeed, the bottleneck in the construction time of T' is the computation of the Δ_x values for all vertices x of T' . All these values can be computed in the obvious way by considering all $O(n^2)$ distance pairs.)

Next, we consider the tree T that is obtained from T' by adding to it an edge (rt, rt') of weight $\rho(rt')$. Lemmas 2.2, 2.3, 2.4, and 2.7 imply the following result.

PROPOSITION 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 binary Steiner SPT for M (with nonnegative edge weights), having weight at most $\log \tilde{n} \cdot w(H') + w(\text{MST}(M))$ and $O(n)$ vertices. Moreover, given the Hamiltonian path H' , the tree T can be constructed in $O(n^2)$ time.*

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

By Lemma 2.4, for any two points v and z in $V \setminus \{rt\}$, $d_T(v, z) = d_{T'}(v, z) \geq \text{dist}(v, z)$. Also, by Lemma 2.2 in the particular case $x = rt'$, for any point v in $V \setminus \{rt\}$ we have $d_T(rt, v) = d_{T'}(rt', v) + \rho(rt') = \text{dist}(rt, v)$. It follows that (T, rt) is a Steiner SPT for M .

To bound the weight of the tree, first note that for any point $v \in V \setminus \{rt\}$, $\rho(rt') = \text{dist}(rt, v) - d_{T'}(rt', v)$, and so $\rho(rt') \leq \text{dist}(rt, v) \leq w(\text{MST}(M))$. By Lemma 2.7, we have $w(T') \leq \log \tilde{n} \cdot w(H')$. It follows that $w(T) = w(T') + \rho(rt') \leq \log \tilde{n} \cdot w(H') + w(\text{MST}(M))$.

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' . \square

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 2.8, we take H' to be $L(M')$. Then the weight bound is reduced to $O(\log n) \cdot w(MST(M))$ and the overall construction time is reduced to $O(n^2)$.

THEOREM 2.9. *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, the maximum degree of the tree may increase as a result of this operation. See also remarks 3 and 4 after Corollary 3.7.

3. Steiner shallow-light trees. Theorem 2.9 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 2.9 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 [7], Awerbuch, Baratz, and Peleg 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 2.9, 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 nonroot 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 2.8 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.

We start with the following simple observation.

CLAIM 3.1. $\sum_{v \in V \setminus \{rt\}} d_{T'}(rt', v) = \sum_{i=0}^{\log \tilde{n}-1} 2^{\log \tilde{n}-(i+1)} \cdot W_i$.

Proof. Fix an arbitrary index $i \in [0, \log \tilde{n} - 1]$, 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 $v_{i+1} \in V_{i+1}$, $|Leaves(v_{i+1})| = 2^{\log \tilde{n}-(i+1)}$. It follows that

$$\begin{aligned} \sum_{e=(v_i, v_{i+1}) \in E(T')} |Leaves(v_{i+1})| \cdot w(e) &= \sum_{i=0}^{\log \tilde{n}-1} \sum_{e \in E_i} |Leaves(v_{i+1})| \cdot w(e) \\ &= \sum_{i=0}^{\log \tilde{n}-1} 2^{\log \tilde{n}-(i+1)} \cdot W_i. \end{aligned}$$

This completes the proof. \square

We are now ready to prove our refined bound on the weight of the SPT T constructed in section 2.

LEMMA 3.2. *Suppose that $\sum_{v \in V \setminus \{rt\}} \text{dist}(rt, v) \leq \xi \cdot \eta$ 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 $\sum_{v \in V \setminus \{rt\}} \text{dist}(rt, v) \leq (n - 1) \cdot w(MST(M))$, i.e., the assumption of the lemma holds for $\xi = n - 1, \eta = w(MST(M))$. In this case we get an upper bound of $2 \cdot w(MST(M)) + \lceil \log(n - 1) \rceil \cdot w(H')$ on $w(T)$, which is slightly larger than the upper bound given in Proposition 2.8. On the other hand, we get significantly better bounds on $w(T)$ whenever $\sum_{v \in V \setminus \{rt\}} \text{dist}(rt, v) \ll (n - 1) \cdot w(MST(M))$.

Proof. Since $w(T) = w(T') + \rho(rt') \leq w(T') + w(MST(M))$, it suffices to show that $w(T') \leq \eta + \lceil \log \xi \rceil \cdot w(H')$. Recall that $\tilde{n} = n - 1$.

Suppose first that $\lceil \log \xi \rceil \geq \log \tilde{n}$. By Lemma 2.7, 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

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

(The last inequality follows from Lemma 2.3.) Since T is an SPT for M , we get that

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

By Claim 3.1,

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

Therefore,

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

(The second inequality holds since $\xi \geq 1$.) Hence

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

Also, Lemma 2.7 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,

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

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 $T^* = T^*(M)$ 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

$$d_{T^*}(B_i, B_{i+1}) > \theta \cdot \text{dist}(rt, B_{i+1}).$$

(It is possible to have just a single break-point $B_1 = rt$.) Let $M_{\mathcal{B}}$ be the submetric of M induced by the point set \mathcal{B} . Also, let $L' = (B_2, \dots, B_k)$ be the subpath 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 2.8 and Lemma 3.2, 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 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 3.3. $\sum_{i=1}^{k-1} d_{T^*}(B_i, B_{i+1}) \leq 2 \cdot w(T^*) = 2 \cdot w(\text{MST}(M))$.

The next lemma bounds the root-distortion of the constructed tree S .

LEMMA 3.4. *For any vertex $v \in V \setminus \{rt\}$, it holds that $d_S(rt, v) \leq (1 + 2\theta) \cdot \text{dist}(rt, v)$.*

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

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

We henceforth assume that v is not a break-point. 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 $d_{\tilde{G}}(rt, B_i) \leq d_{T_{\mathcal{B}}}(rt, B_i) = \text{dist}(rt, B_i)$. 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,

$$d_S(rt, v) = d_{\tilde{G}}(rt, v) \leq d_{\tilde{G}}(rt, B_i) + d_{\tilde{G}}(B_i, v) \leq \text{dist}(rt, B_i) + d_{T^*}(B_i, v).$$

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

$$(2) \quad d_{T^*}(B_i, v) \leq \theta \cdot \text{dist}(rt, v),$$

and so

$$(3) \quad d_S(rt, v) \leq \text{dist}(rt, B_i) + \theta \cdot \text{dist}(rt, v).$$

By the triangle inequality and (2),

$$(4) \quad \begin{aligned} \text{dist}(rt, B_i) &\leq \text{dist}(rt, v) + \text{dist}(B_i, v) \leq \text{dist}(rt, v) + d_{T^*}(B_i, v) \\ &\leq \text{dist}(rt, v) + \theta \cdot \text{dist}(rt, v) = (1 + \theta) \cdot \text{dist}(rt, v). \end{aligned}$$

Plugging (4) in (3), we obtain

$$d_S(rt, v) \leq (1 + \theta) \cdot \text{dist}(rt, v) + \theta \cdot \text{dist}(rt, v) = (1 + 2\theta) \cdot \text{dist}(rt, v). \quad \square$$

Next, we bound the weight of the constructed tree S .

LEMMA 3.5. $w(S) \leq 2 \cdot (\lceil \log(\frac{2}{\theta}) \rceil + 2) \cdot w(MST(M))$.

Proof. By the choice of break-points, for each index $i \in [k-1]$, $\text{dist}(rt, B_{i+1}) < \frac{1}{\theta} \cdot d_{T^*}(B_i, B_{i+1})$. By Claim 3.3, $\sum_{i=1}^{k-1} d_{T^*}(B_i, B_{i+1}) \leq 2 \cdot w(MST(M))$. Therefore,

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

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

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

Since $\theta \leq 2$, we have $\xi \geq 1$. Clearly, $\eta = w(MST(M)) > 0$. Hence, by Lemma 3.2, the weight $w(T_{\mathcal{B}})$ of $T_{\mathcal{B}} = T_{\mathcal{B}}(L')$ satisfies

$$\begin{aligned} w(T_{\mathcal{B}}) &\leq \eta + \lceil \log \xi \rceil \cdot w(L') + w(MST(M_{\mathcal{B}})) \\ &\leq w(MST(M)) + \left\lceil \log \left(\frac{2}{\theta} \right) \right\rceil \cdot w(L') + 2 \cdot w(MST(M)) \\ &\leq 3 \cdot w(MST(M)) + \left\lceil \log \left(\frac{2}{\theta} \right) \right\rceil \cdot 2 \cdot w(MST(M)). \end{aligned}$$

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

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

Note also that for any metric M the weight of the minimum Steiner tree for M (denoted by $SMT(M)$) is greater than or equal to half the weight of the MST for M , i.e., $w(SMT(M)) \geq \frac{1}{2} \cdot w(MST(M))$. It follows that $w(S) \leq 4 \cdot (\lceil \log(\frac{2}{\theta}) \rceil + 2) \cdot w(SMT(M))$.

Finally, we analyze the running time of the construction.

LEMMA 3.6. *The tree S 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 mentioned above, the Hamiltonian path L can be computed within

$O(n^2)$ time as well. It is easy to see that $O(n^2)$ time suffices to identify the set $\mathcal{B} = \{B_1, B_2, \dots, B_k\}$ of break-points. An additional amount of $O(n)$ time requires one to construct the subpath L' of L . By Proposition 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. \square

Set $\epsilon = 2\theta$. Lemmas 3.4, 3.5, and 3.6 imply the following corollary.

COROLLARY 3.7. *For any n -point metric M , a designated point $rt \in V$, and a number $0 < \epsilon < \frac{1}{2}$, there exists 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 Corollary 3.7.
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 Theorem 2.9 and Corollary 3.7 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 $\omega \approx 2.3727$ is the matrix multiplication exponent [20, 49] and in time $O(n^3 \sqrt{\frac{\log \log n}{\log n}})$ if G is an arbitrary weighted graph [46].
4. One cannot significantly improve the running time of our constructions from Theorem 2.9 and Corollary 3.7 even in the *cell-probe model* of computation, which is perhaps the strongest model of computation for data structures, subsuming the common word-RAM model; it is assumed that the memory is divided into fixed-size cells (words), and the cost of an operation is just the number of cells it reads or writes. (See [42] for a more detailed description of the cell-probe model.) 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 [32] that any deterministic (resp., 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 [32] extends to Steiner trees as well.

4. 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 and 3, 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 $\text{diam}(M)$ and observe that $\text{diam}(M) \geq n - 1$. Given a metric

$M \in \mathcal{P}_n$, a number $r \geq \frac{1}{2} \cdot \text{diam}(M)$, and a number $\epsilon > 0$, we say that a rooted Steiner tree (T, rt) of M is an r -tree (resp., (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 . Observe that if T is an r -tree and $r > 0$, then the root vertex rt of T must be a Steiner point, i.e., $rt \notin M$.

LEMMA 4.1. *Let M be an arbitrary metric in \mathcal{P}_n , and let $r \geq \frac{1}{2} \cdot \text{diam}(M) \geq \frac{1}{2} \cdot (n - 1)$. 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 $r \geq \frac{1}{2} \cdot \text{diam}(M)$. 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 we 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, \dots, 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 non-Steiner points in T_i (i.e., those belonging to \mathcal{P}_n) by V_i . Also, we write n_i as a shorthand for $|V_i|$ and w_i as a shorthand 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 fewer Steiner points, yielding a contradiction. Consider the metric M_i induced by the point set V_i . Observe that $M_i \in \mathcal{P}_{n_i}$, and for every $v \in V_i$, $d_{T_i}(c_i, v) = r - w_i$. Moreover, $\text{diam}(M_i) \geq n_i - 1$. Also, note that the subtree T_i of T dominates M_i , and so it must hold that $r - w_i \geq \frac{1}{2} \cdot \text{diam}(M_i)$. 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,

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

(The last inequality follows from Lemma A.1, which appears in Appendix A.) \square

Consider the 1-dimensional Euclidean metric $\zeta_n \in \mathcal{P}_n$ that consists of n vertices v_1, v_2, \dots, v_n that lie on the x -axis with coordinates $1, 2, \dots, n$, 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 4.1 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 $\frac{n}{2}$ (rather than unit radius as in section 1.3), centered at the origin $(0,0)$, and define $\tilde{C}_{n+1} = C_n \cup \{(0,0)\}$. Note that the circular distance between a pair of consecutive points in C_n is equal to $\pi \cdot \frac{n}{n-1} \approx \pi$. It is not hard to see that the proof of Lemma 4.1 carries through (with minor adjustments) also if T is an r -tree for C_n . Specifically, the properties of M that are needed for carrying out the proof of Lemma 4.1 are the following: (1) The distance in M between any pair of points is at least 1. (2) For any subset M' of M with $n' \leq n$ points, $\text{diam}(M') \geq n' - 1$. In particular, $\text{diam}(M) \geq n - 1$. (3) Any r -tree of M or any submetric M' of M , with $r > 0$, has a Steiner root vertex. It is easy to verify that C_n (defined above) satisfies these three properties. To summarize, any SPT for the point set \tilde{C}_{n+1} rooted at $rt = (0,0)$ has lightness $\Omega(\log n)$.

THEOREM 4.2. *For any sufficiently large integer n , there exist 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 $\Omega(\log n)$. (The statement holds even for trees that can use non-Euclidean Steiner points.)*

Our next objective is to generalize Theorem 4.2 for Steiner SLTs. We first prove a lemma which provides some structural information about $(r, \frac{1}{r})$ -trees for metrics M from \mathcal{P}_n .

LEMMA 4.3. *Let M be an arbitrary metric in \mathcal{P}_n , and let $r \geq \frac{1}{2} \cdot \text{diam}(M) \geq \frac{1}{2} \cdot (n-1)$. 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 the following hold:*

- (1) *All leaves of T are non-Steiner points; i.e., they all belong to \mathcal{P}_n .*
- (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 fewer Steiner points than T and the same weight, yielding a contradiction. Lemma 4.3 follows. \square

Now we are ready to extend Lemma 4.1 to Steiner SLTs.

LEMMA 4.4. *Let M be an arbitrary metric in \mathcal{P}_n , and let $r \geq \frac{1}{2} \cdot \text{diam}(M) \geq \frac{1}{2} \cdot (n - 1)$. Then any $(r, \frac{1}{r})$ -tree T for M has weight $w(T)$ at least $g(n, r) - 2n = (r - \frac{1}{2} \cdot (n - 1) + \frac{1}{2} \cdot n \cdot \log n) - 2n$.*

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

$$(5) \quad \delta(\pi(v)) = \max\{\delta(u) + w(\pi(v), u) \mid u \in Ch(\pi(v))\} \geq \delta(v) + w(\pi(v), v).$$

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 4.5. *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, (5) 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 $\delta(rt) \geq r \geq \frac{1}{2} \cdot \text{diam}(M)$.

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 we 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 4.3, 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

$$(6) \quad |d_T(x, l_1) - d_T(x, l_2)| = |d_T(rt, l_1) - d_T(rt, l_2)| < 1.$$

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 (resp., 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, (6) 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.)

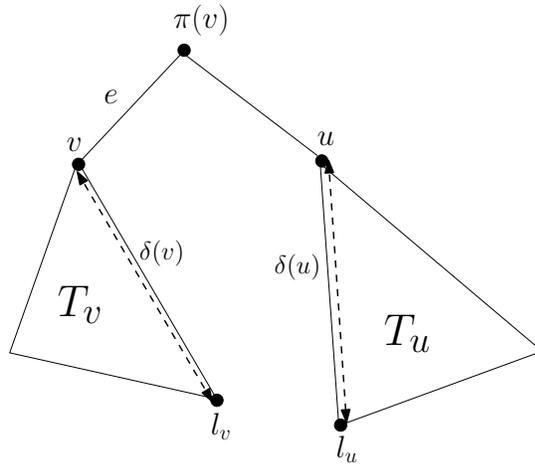


FIG. 5. The vertex $\pi(v)$ and its two subtrees T_v and T_u .

By the third assertion of Lemma 4.3, there are at most $2n - 1$ edges in E , and so

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

which completes the proof of Claim 4.5. \square

Claim 4.5 implies that T' is a $\delta(rt)$ -tree for M , with $\delta(rt) \geq r \geq \frac{1}{2} \cdot \text{diam}(M)$ and $w'(T') \leq w(T) + 2n$. By Lemma 4.1,

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

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 4.4 follows. \square

Lemma 4.4 implies that any Steiner tree for $\tilde{\zeta}_{n+1}$ rooted at rt with root-distortion 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-distortion less than $1 + \frac{2}{k-1}$ has lightness $\Omega(\log k)$. Also, similarly to above, by locating many points of tiny distances from points of C_k we obtain the same bound of $\Omega(\log k)$ on the lightness of Steiner $(1 + \frac{2}{k-1})$ -SPTs for

simple Euclidean 2-dimensional metrics. Setting $\epsilon = \frac{2}{k-1}$, we obtain the following result.

THEOREM 4.6. *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-distortion at most $1 + \epsilon$ has lightness $\Omega(\log \frac{1}{\epsilon})$.*

Next, we strengthen Theorem 4.6 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-distortion and lightness holds for *any* root vertex. Moreover, we demonstrate that this tradeoff cannot be improved even if we consider *average* root-distortion rather than (worst-case) root-distortion. For the analysis of this metric, we will use Lemma 4.4.

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 $\ell \in [k]$, $V_\ell = \{v_\ell^{(1)}, v_\ell^{(2)}, \dots, v_\ell^{(\frac{n}{2k})}\}$ and $U_\ell = \{u_\ell^{(1)}, u_\ell^{(2)}, \dots, u_\ell^{(\frac{n}{2k})}\}$. 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 $i, j \in [\frac{n}{2k}]$, $dist(v_\ell^{(i)}, u_q^{(j)}) = \frac{k-1}{2}$.

A point set $W \subseteq V$ (resp., $W \subseteq U$) 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 4.7. *For any elementary point set W , the submetric $\vartheta(W)$ of $\vartheta_{n,k}$ induced by W belongs to $\mathcal{P}_{|W|}$.*

LEMMA 4.8. *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 the following hold:*

(1) *If the root-distortion 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-distortion 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-distortion 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 $W = \{v_1^{(1)}, v_2^{(1)}, \dots, v_k^{(1)}\}$. Observation 4.7 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-distortion 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 4.4, the weight $w(T)$ of T is at least $g(k, \frac{k-1}{2}) - 2k = (\frac{k-1}{2} - \frac{1}{2} \cdot (k-1) + \frac{1}{2} \cdot k \cdot \log k) - 2k = \frac{1}{2} \cdot k \cdot \log k - 2k$.

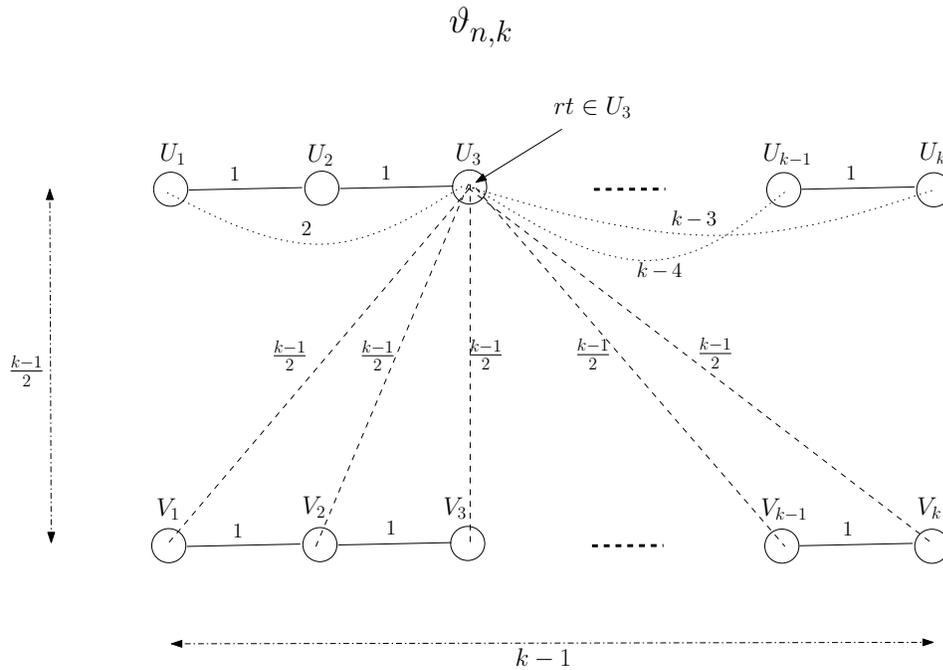


FIG. 6. An illustration of the metric $\vartheta_{n,k}$. For each index ℓ , the circles around V_ℓ and U_ℓ designate the $(\frac{2}{k})$ -point sets corresponding to them. The distance between all pairs of points that belong to the same set V_ℓ or U_ℓ , $\ell \in [k]$, is equal to some tiny number $0 < \alpha \ll \frac{1}{n}$. All solid lines in the figure have length 1. These lines designate the distance between a point in V_ℓ (resp., 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.

Next, we prove the second assertion of the lemma. Assume that the average root-distortion 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 $\varphi_T(rt, v) < 1 + \frac{2}{k-1}$, and define $V'' = V \setminus V'$. Also, write $\mathcal{V}^* = \mathcal{V} \setminus \{rt\}$. Clearly, for each point v in $\mathcal{V}^* \setminus V''$, $\varphi_T(rt, v) \geq 1$, and for each point v in V'' , $\varphi_T(rt, v) \geq 1 + \frac{2}{k-1}$. Observe that

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

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 $\varphi_T(rt, v) < 1 + \frac{2}{k-1}$ for each point $v \in W$, implying that (T, rt) is a Steiner tree for $\vartheta(W)$ with root-distortion less than $1 + \frac{2}{k-1}$. By Observation 4.7, $\vartheta(W)$ 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 $\text{diam}(\vartheta(W)) \leq \text{diam}(\vartheta_{n,k}) = k - 1$, and so $\frac{k-1}{2} \geq \frac{1}{2} \cdot \text{diam}(\vartheta(W))$. Consequently, (T, rt) is a $(\frac{k-1}{2}, \frac{2}{k-1})$ -tree for $\vartheta(W)$. By Lemma 4.4, the weight $w(T)$ of T is at least

$$\begin{aligned} g\left(\frac{k}{2}, \frac{k-1}{2}\right) - 2 \cdot \frac{k}{2} &= \left(\frac{k-1}{2} - \frac{1}{2} \cdot \left(\frac{k}{2} - 1\right) + \frac{1}{2} \cdot \frac{k}{2} \cdot \log\left(\frac{k}{2}\right)\right) - k \\ &= \frac{1}{4} \cdot k \cdot \log k - k. \quad \square \end{aligned}$$

Observe that $\frac{5}{2} \cdot (k-1) \leq w(\text{MST}(\vartheta_{n,k})) \leq \frac{5}{2} \cdot (k-1) + \alpha \cdot (n-1) \leq \frac{5}{2} \cdot (k-1) + 1$. Substituting k with $\Theta(\frac{1}{\epsilon})$ in Lemma 4.8, we obtain the main result of this section.

COROLLARY 4.9. *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-distortion 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 4.8, 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.

5. 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-distortion (with respect to some designated root vertex rt). Finally, a *metric-spanning SLT* of G is a metric-spanning tree of G with the same properties (small lightness and root-distortion).

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

LEMMA 5.1. 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-distortion and the average root-distortion of (T', rt) are no greater than the root-distortion and the average root-distortion 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 $P_e = (v_1 = x, v_2 = z, v_3, \dots, v_k = y), k \geq 3$. Since P_e is a shortest path between x and y in G , it follows that

$$(7) \quad d_G(x, z) + d_G(z, y) = w(P_e) = d_G(x, y).$$

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 Figure 7(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 Figure 7(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

$$(8) \quad d_T(p, z) + d_G(z, y) < d_T(p, x) + d_G(x, y).$$

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 Figure 7(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 (8), $d_{T'}(rt, y) < d_T(rt, y)$. 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 that for all other vertices v , $d_{T'}(rt, v) = d_T(rt, v)$.

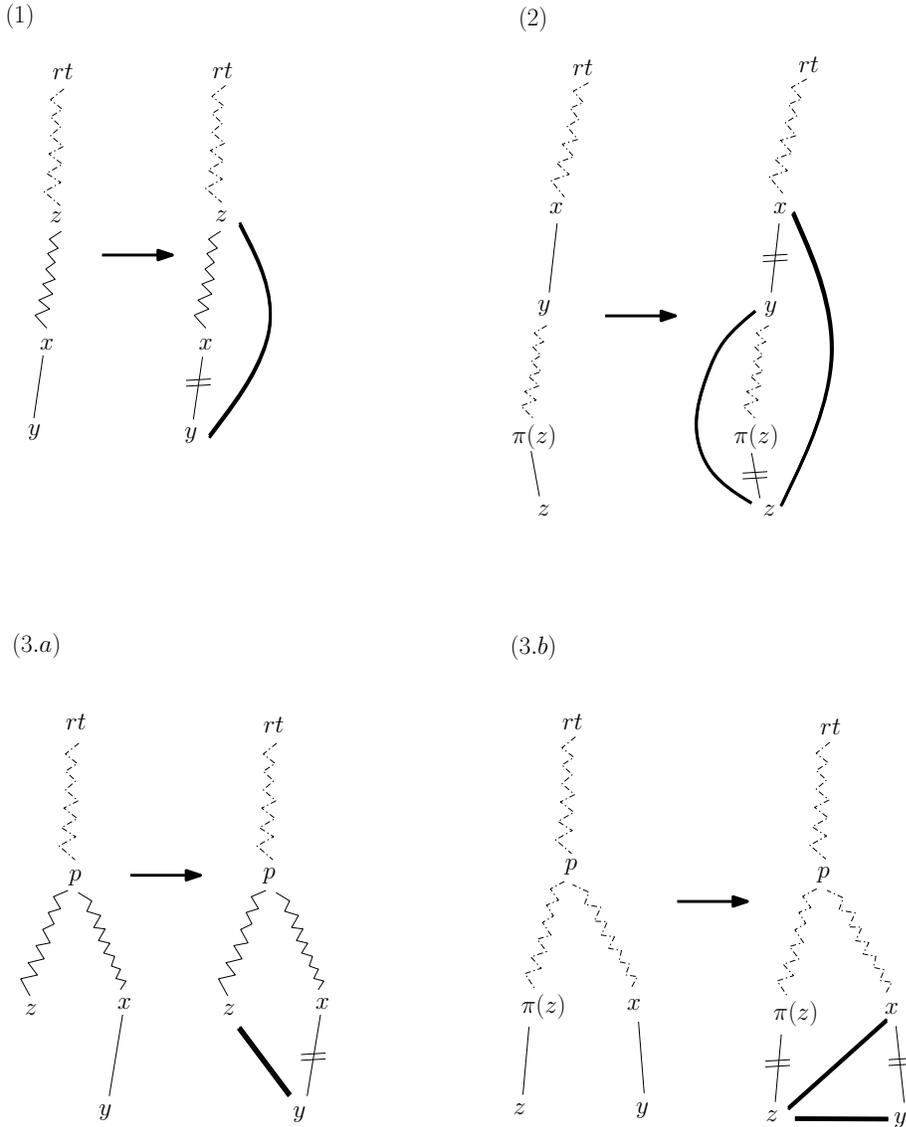


FIG. 7. An illustration of the four cases that are considered in the proof of Lemma 5.1. In each case it is shown how to transform the tree (T, rt) into another tree (T', rt) that satisfies the conditions of the lemma. Paths that may contain zero or more edges are depicted by zigzag lines. A path that 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.

Case 3.b. In the complementary subcase, i.e., when (8) does not hold, we have

$$(9) \quad d_T(p, z) + d_G(z, y) \geq d_T(p, x) + d_G(x, y).$$

(Observe that it is possible that $p = x$ and/or $p = \pi(z)$.) Plugging (7) in (9), we obtain

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

implying that

$$(10) \quad d_T(p, z) \geq d_T(p, x) + d_G(x, z).$$

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 Figure 7(3.b) for an illustration.) Exactly as in Case 2, it follows that $w(T') < w(T)$. Also, (10) 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. \square

Lemma 5.1 implies the following corollaries.

COROLLARY 5.2. *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 (resp., S_2) the set of all metric-spanning trees of G rooted at rt with root-distortion (resp., average root-distortion) at most α . Suppose that S_1 (resp., S_2) is nonempty, 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 5.1 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. \square

Now we 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 5.3. *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-distortion (resp., average root-distortion) at most α and lightness at most β , then there is also a graph-spanning tree of G rooted at rt with root-distortion (resp., average root-distortion) at most α and lightness at most β .*

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

LEMMA A.1. *For any positive integers $n_1, n_2, \dots, n_k, k \geq 1$, it holds that $\sum_{i=1}^k (n_i \cdot \log(\frac{n}{n_i})) \leq n \cdot (k - 1)$, 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 $f(x) = x \cdot \log(\frac{n}{x}) + (n - x) \cdot \log(\frac{n}{n-x})$. (Note that $f(\cdot)$ is a scaling of the binary entropy function $H(\cdot)$, i.e., $f(x) = n \cdot H(\frac{x}{n})$.) 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 $\sum_{i=1}^k (n_i \cdot \log(\frac{n}{n_i})) = n \cdot \log 1 = 0 = n \cdot (k - 1)$.

Induction step. We assume that the statement holds for all smaller values of $k, k \geq 2$, and we prove it for k . Define $N_k = \sum_{i=1}^k (n_i \cdot \log(\frac{n}{n_i}))$ and $N_{k-1} = \sum_{i=1}^{k-1} (n_i \cdot \log(\frac{n-n_k}{n_i}))$. 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,

$$N_k = (N_k - N_{k-1}) + N_{k-1} \leq n + (n - n_k) \cdot (k - 2) \leq n + n \cdot (k - 2) = n \cdot (k - 1). \quad \square$$

Appendix B. SLTs with Euclidean Steiner points. Let C_n denote a set of n points that are uniformly spaced on the boundary of the unit circle C centered at the origin, and define $\tilde{C}_{n+1} = C_n \cup \{(0, 0)\}$. We will show that any $(1 + \epsilon)$ -SPT (T, rt) for \tilde{C}_{n+1} rooted at $rt = (0, 0)$ which may use only Euclidean Steiner points must have lightness at least $\Omega(\sqrt{\frac{1}{\epsilon}})$.

Partition the circle C into $t = \Theta(\sqrt{\frac{1}{\epsilon}})$ arcs A_1, \dots, A_t of angle $\frac{2\pi}{t} = \Theta(\sqrt{\epsilon})$ each. Consider two arbitrary points p_i and p_j that reside at the middle of two distinct arcs A_i and A_j , respectively. They are at circular distance at least $\Omega(\sqrt{\epsilon})$ from each other. As the root-distortion is at most $1 + \epsilon$, we know that the distortions in T between rt and p_i and between rt and p_j is at most $1 + \epsilon$. The next lemma follows.

LEMMA B.1. *The paths in T from rt to p_i and from rt to p_j cannot intersect within the annulus \mathcal{A} with inner radius $\frac{1}{2}$ and outer radius 1 centered at the origin.*

Proof. Let q_i (resp., q_j) be the point lying in the middle of the segment connecting rt with p_i (resp., p_j); note that q_i and q_j lie on the boundary of the annulus \mathcal{A} . Let q be the point lying in the middle of the arc connecting q_i and q_j on the boundary of \mathcal{A} . Observe that the circular distance between q_i and q_j is $\Omega(\sqrt{\epsilon})$. By symmetry considerations, if the tree paths between rt and p_i and between rt and p_j intersect within the annulus \mathcal{A} , the maximum root-distortion is minimized if the intersection point is q . However, in the latter case, a straightforward calculation shows that the root-distortion is $\|rt, q\| + \|q, p_i\| = \frac{1}{2} + \|q, p_i\| \geq 1 + \Omega(\epsilon)$. By setting the constant hidden by the O -notation in $t = O(\sqrt{\frac{1}{\epsilon}})$ appropriately, we can guarantee here root-distortion greater than $1 + \epsilon$. \square

It follows that each point p_i at the middle of arc A_i contributes $\frac{1}{2}$ fresh units to the weight of T for each $i = [t]$. Since there are $t = \Theta(\sqrt{\frac{1}{\epsilon}})$ such arcs, it follows that the weight of T is at least $\frac{1}{2} \cdot t = \Omega(\sqrt{\frac{1}{\epsilon}})$. Note also that the weight of the MST for \tilde{C}_{n+1} is $O(1)$, which completes the argument.

Remark. The above lower bound for \tilde{C}_{n+1} holds for the specific choice of root vertex $rt = (0, 0)$. However, this argument can be easily strengthened to hold for C_n and any choice of root vertex $rt \in C_n$.

Acknowledgments. 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 Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2008, pp. 781–790.
- [2] I. ABRAHAM AND O. NEIMAN, *Using petal-decompositions to build a low stretch spanning tree*, in Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC), 2012, pp. 395–406.
- [3] N. ALON, R. M. KARP, D. PELEG, AND D. WEST, *A graph-theoretic game and its application to the k -server problem*, SIAM J. Comput., 24 (1995), pp. 78–100.
- [4] I. ALTHÖFER, G. DAS, D. P. DOBKIN, D. JOSEPH, AND J. SOARES, *On sparse spanners of weighted graphs*, Discrete Comput. Geom., 9 (1993), pp. 81–100.
- [5] 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., 40 (2008), pp. 207–219.

- [6] B. AWERBUCH, A. BARATZ, AND D. PELEG, *Cost-sensitive analysis of communication protocols*, in Proceedings of the 9th Annual ACM Symposium on Principles of Distributed Computing (PODC), 1990, pp. 177–187.
- [7] B. AWERBUCH, A. BARATZ, AND D. PELEG, *Proceedings of the 9th Annual ACM Symposium on Principle of Distributed Computing (PODC)*, 1990, pp. 177–187, *Efficient Broadcast and Light-Weight Spanners*, manuscript, 1991.
- [8] Y. BARTAL, *Probabilistic approximations of metric spaces and its algorithmic applications*, in Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 1996, pp. 184–193.
- [9] Y. BARTAL, *On approximating arbitrary metrics by tree metrics*, in Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC), 1998, pp. 161–168.
- [10] Y. BEN-SHIMOL, A. DVIR, AND M. SEGAL, *SPLAST: A novel approach for multicasting in mobile wireless ad hoc networks*, in Proceedings of the 15th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), 2004, pp. 1011–1015.
- [11] P. BERMAN AND V. RAMAIYER, *Improved approximations for the Steiner tree problem*, *J. Algorithms*, 17 (1994), pp. 381–408.
- [12] M. W. BERN AND P. E. PLASSMANN, *The Steiner problem with edge lengths 1 and 2*, *Inform. Process. Lett.*, 32 (1989), pp. 171–176.
- [13] K. BHARATH-KUMAR AND J. M. JAFFE, *Routing to multiple destinations in computer networks*, *IEEE Trans. Commun.*, 31 (1983), pp. 343–351.
- [14] B. BOLLOBÁS, D. COPPERSMITH, AND M. ELKIN, *Sparse distance preservers and additive spanners*, *SIAM J. Discrete Math.*, 19 (2006), pp. 1029–1055.
- [15] R. BRAYNARD, D. KOSTIC, A. RODRIGUEZ, J. CHASE, AND A. VAHDAT, *Opus: An overlay peer utility service*, in Proceedings of the 5th IEEE Conference on Open Architectures and Network Programming (OPENARCH), 2002.
- [16] M. CHARIKAR, C. CHEKURI, T. CHEUNG, Z. DAI, A. GOEL, S. GUHA, AND M. LI, *Approximation algorithms for directed Steiner problems*, *J. Algorithms*, 33 (1999), pp. 73–91.
- [17] J. CONG, A. B. KAHNG, G. ROBINS, M. SARRAFZADEH, AND C. K. WONG, *Performance-driven global routing for cell based ICS*, in Proceedings of the 9th IEEE International Conference on Computer Design (ICCD), 1991, pp. 170–173.
- [18] J. CONG, A. B. KAHNG, G. ROBINS, M. SARRAFZADEH, AND C. K. WONG, *Provably good algorithms for performance-driven global routing*, in Proceedings of the 5th IEEE International Symposium on Circuits and Systems (ISCAS), 1992, pp. 2240–2243.
- [19] J. CONG, A. B. KAHNG, G. ROBINS, M. SARRAFZADEH, AND C. K. WONG, *Provably good performance-driven global routing*, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* 11 (1992), pp. 739–752.
- [20] D. COPPERSMITH AND S. WINOGRAD, *Matrix multiplication via arithmetic progressions*, *J. Symbolic Comput.*, 9 (1990), pp. 251–280.
- [21] T. H. CORMAN, C. E. LEISERSON, R. L. RIVEST, AND C. STEIN, *Introduction to Algorithms*, 2nd ed., McGraw-Hill, Boston, MA, 2001.
- [22] R. CRISTESCU, B. BEFERULLI-LOZANO, M. VETTERLI, AND R. WATTENHOFER, *Network correlated data gathering with explicit communication: NP-completeness and algorithms*, *IEEE/ACM Trans. Netw.*, 14 (2006), pp. 41–54.
- [23] Y. DINITZ, M. ELKIN, AND S. SOLOMON, *Low-light trees, and tight lower bounds for Euclidean spanners*, *Discrete Comput. Geom.*, 43 (2010), pp. 736–783.
- [24] M. ELKIN, Y. EMEK, D. A. SPIELMAN, AND S.-H. TENG, *Lower-stretch spanning trees*, *SIAM J. Comput.*, 38 (2008), pp. 608–628.
- [25] M. ELKIN AND S. SOLOMON, *Narrow-shallow-low-light trees with and without Steiner points*, *SIAM J. Discrete Math.*, 25 (2011), pp. 181–210.
- [26] J. FAKCHAROENPHOL, S. RAO, AND K. TALWAR, *A tight bound on approximating arbitrary metrics by tree metrics*, in Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC), 2003, pp. 448–455.
- [27] E. N. GILBERT AND H. O. POLLAK, *Steiner minimal trees*, *SIAM J. Appl. Math.*, 16 (1968), pp. 1–29.
- [28] A. GOEL AND K. MUNAGALA, *Balancing Steiner trees and shortest path trees online*, in Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2000, pp. 562–563.
- [29] A. GUPTA, *Steiner points in tree metrics don't (really) help*, in Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2001, pp. 220–227.
- [30] M. T. HAJIAGHAYI, G. KORTSARZ, AND M. R. SALAVATIPOUR, *Approximating buy-at-bulk and shallow-light k -Steiner trees*, *Algorithmica*, 53 (2009), pp. 89–103.
- [31] F. K. HWANG AND D. Z. DU, *Steiner minimal trees on Chinese checkerboards*, *Math. Mag.*, 64 (1991), pp. 332–339.

- [32] P. INDYK, *Sublinear time algorithms for metric space problems*, in Proceedings of the 31th Annual ACM Symposium on Theory of Computing (STOC), 1999, pp. 428–434.
- [33] J. M. JAFFE, *Distributed multi-destination routing: The constraints of local information*, SIAM J. Comput., 14 (1985), pp. 875–888.
- [34] S. KHULLER, B. RAGHAVACHARI, AND N. E. YOUNG, *Balancing minimum spanning and shortest path trees*, in Proceedings of the Fourth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 1993, pp. 243–250.
- [35] V. P. KOMPPELLA, J. PASQUALE, AND G. C. POLYZOS, *Multicast routing for multimedia communication*, IEEE/ACM Trans. Netw., 1 (1993), pp. 286–292.
- [36] G. KONJEVOD, R. RAVI, AND F. S. SALMAN, *On approximating planar metrics by tree metrics*, Inform Process. Lett., 80 (2001), pp. 213–219.
- [37] G. KORTSARZ AND D. PELEG, *Approximating the weight of shallow Steiner trees*, Discrete Appl. Math., 93 (1999), pp. 265–285.
- [38] D. KOSTIC AND A. VAHDAT, *Latency Versus Cost Optimizations in Hierarchical Overlay Networks*, Technical report, Duke University, Durham, NC, 2002.
- [39] M. V. MARATHE, R. RAVI, R. SUNDARAM, S. S. RAVI, D. J. ROSENKRANTZ, AND H. B. HUNT III, *Bicriteria network design problems*, J. Algorithms, 28 (1998), pp. 142–171.
- [40] D. PELEG, *Distributed Computing: A Locality-Sensitive Approach*, SIAM, Philadelphia, 2000.
- [41] D. PELEG AND A. SCHÄFFER, *Graph spanners*, J. Graph Theory, 13 (1989), pp. 99–116.
- [42] M. PÁTRÁŠČU AND E. D. DEMAINE, *Logarithmic lower bounds in the cell-probe model*, SIAM J. Comput., 35 (2006), pp. 932–963.
- [43] S. RAO AND W. D. SMITH, *Approximating geometrical graphs via “spanners” and “banyans,”* in Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC), 1998, pp. 540–550.
- [44] F. S. SALMAN, J. CHERIYAN, R. RAVI, AND S. SUBRAMANIAN, *Approximating the single-sink link-installation problem in network design*, SIAM J. Optim., 11 (2000), pp. 595–610.
- [45] H. SHPUNGIN AND M. SEGAL, *Near optimal multicriteria spanner constructions in wireless ad-hoc networks*, IEEE/ACM Trans. Netw., 18 (2010), pp. 1963–1976.
- [46] T. TAKAOKA, *A new upper bound on the complexity of the all pairs shortest path problem*, Inform Process. Lett., 43 (1992), pp. 195–199.
- [47] J. VOGEL, J. WIDMER, D. FARIN, M. MAUVE, AND W. EFFELSBERG, *Priority-based distribution trees for application-level multicast*, in Proceedings of the 2nd ACM Workshop on Network and System Support for Games (NETGAMES), 2003, pp. 148–157.
- [48] P. VON RICKENBACH AND R. WATTENHOFER, *Gathering correlated data in sensor networks*, in Proceedings of the ACM Joint Workshop on Foundations of Mobile Computing (DIALM-POMC), 2004, pp. 60–66.
- [49] V. VASSILEVSKA WILLIAMS, *Multiplying matrices faster than Coppersmith-Winograd*, in Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC), 2012, pp. 887–898.
- [50] B. Y. WU, K.-M. CHAO, AND C. Y. TANG, *Light graphs with small routing cost*, Networks, 39 (2002), pp. 130–138.