

Narrow-Shallow-Low-Light Trees with and without Steiner Points *

Michael Elkin^{†§}

Shay Solomon^{†‡}

Abstract

We show that for every set \mathcal{S} of n points in the plane and a designated point $rt \in \mathcal{S}$, there exists a tree T that has small maximum degree, depth and weight. Moreover, for every point $v \in \mathcal{S}$, the distance between rt and v in T is within a factor of $(1+\epsilon)$ close to their Euclidean distance $\|rt, v\|$. We call these trees *narrow-shallow-low-light* (NSLLTs). We demonstrate that our construction achieves optimal (up to constant factors) tradeoffs between *all* parameters of NSLLTs. Our construction extends to point sets in \mathbb{R}^d , for an arbitrarily large constant d . The running time of our construction is $O(n \cdot \log n)$.

We also study this problem in *general metric spaces*, and show that NSLLTs with small maximum degree, depth and weight can always be constructed if one is willing to compromise the root-distortion. On the other hand, we show that the increased root-distortion is inevitable, even if the point set \mathcal{S} resides in a Euclidean space of dimension $\Theta(\log n)$.

In addition, we show that if one is allowed to use Steiner points then it is possible to achieve root-distortion of $(1 + \epsilon)$ together with small maximum degree, depth and weight for *general metric spaces*.

Finally, we establish some lower bounds on the power of Steiner points in the context of Euclidean spanning trees and spanners.

*A preliminary version of this paper appeared in ESA'09.

[†]Department of Computer Science, Ben-Gurion University of the Negev, POB 653, Beer-Sheva 84105, Israel. E-mail: {elkinm, shayso}@cs.bgu.ac.il

Partially supported by the Lynn and William Frankel Center for Computer Sciences.

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

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

1 Introduction

1.1 Euclidean Spaces

Given a set \mathcal{S} of n points in \mathbb{R}^d and a designated root vertex rt , we want to construct a spanning tree T for \mathcal{S} rooted at rt that enjoys a number of useful properties. First, we want T to be *light*, that is, to be not much heavier than the minimum spanning tree of \mathcal{S} (denoted $MST(\mathcal{S})$). Second, we want it to be *low*, i.e., to have a small depth¹. Third, we want T to be *shallow*, meaning that for every vertex v in T , the distance $dist_T(rt, v)$ between rt and v in T should not be much greater than the Euclidean distance $\|rt, v\|$. (The maximum ratio $\max \left\{ \frac{dist_T(rt, v)}{\|rt, v\|} : v \in \mathcal{S} \right\}$ will be called the *root-stretch* or *root-distortion* of T .) Fourth, the tree T should be *narrow*, that is, to have a small maximum degree.

Each of these requirements has a natural network-design analogue. The weight of T corresponds to the total cost of building and maintaining the network. The depth and the root-stretch of the tree correspond to communication delays experienced by network end-users. The maximum degree of T corresponds to the load experienced by the relay stations or network routers. Finally, the tree structure of the designed network may be necessary for some applications. In other applications which can be executed in a network that contains cycles, having a cycle-free network may still be very advantageous. Consequently, the problem of designing trees that enjoy all these properties is a basic problem in the area of *geometric network design*. Similar problems arise in the context of VLSI design [1, 11, 12, 13], telecommunications and distributed computing [5, 6], road network design and medical imaging [15].

Obviously, some of these requirements come at the expense of others, and there are inherent tradeoffs between the different parameters. In a seminal STOC'95 paper on Euclidean spanners, Arya et al. [3] have shown that for every set \mathcal{S} of n points in \mathbb{R}^d there exists a rooted spanning tree (T, rt) with depth $O(\log n)$, maximum degree $O(\epsilon^{-d+1})$, and root-stretch at most $(1 + \epsilon)$, for arbitrarily small $\epsilon > 0$. However, their trees (called *single-sink spanners*) may have weight² that is greater by a factor of n than $w(MST(\mathcal{S}))$. Recently, Dinitz et al. [14] devised a construction of narrow-low-light trees (henceforth, NLLTs) that enjoys a small lightness (i.e., $O(\log n)$), small depth (i.e., $O(\log n)$) and constant maximum degree. However, the resulting trees may have arbitrarily large root-stretch. (The NLLTs construction of [14] applies, in fact, to general metric spaces.) In this paper we fill in the gap and devise a single construction that combines all the useful properties of the single-sink spanners construction of [3] and the NLLTs construction of [14]. Specifically, we show that for every set \mathcal{S} of n points in \mathbb{R}^d , a designated point $rt \in \mathcal{S}$, an integer $\ell = 1, 2, \dots, O(\log n)$ and a number $\epsilon > 0$, there exists a rooted spanning tree (\widehat{T}_1, rt) with lightness $O(\ell \cdot (\epsilon^{-1}))$ (“light”), depth $O(\ell \cdot n^{1/\ell})$ (“low”), maximum degree $O(\epsilon^{-d+1})$ (“narrow”) and root-stretch at most $(1 + \epsilon)$ (“shallow”). There also exists a rooted spanning tree (\widehat{T}_2, rt) with lightness $O(\ell \cdot n^{1/\ell} \cdot (\epsilon^{-1}))$, depth $O(\ell)$, maximum degree $O(n^{1/\ell} + (\epsilon^{-d+1}))$ and root-stretch at most $(1 + \epsilon)$. Moreover, the former (respectively, latter) tree can be constructed in time $O(n \cdot \log n)$ (resp., $O(n \cdot \ell) = O(n \cdot \log n)$) for $d = 2$, and in time $O(n \cdot \log n \cdot (\log \epsilon^{-1}) + (\epsilon^{-d+1}))$ (resp., $O(n \cdot \ell \cdot (\log \epsilon^{-1}) + (\epsilon^{-d+1}))$) for $d \geq 3$. We remark that the running time bounds presented in this paper hold in the real-RAM model, assuming that the floor function is available and that the word size is at least $\log n$.

Our results generalize and improve the previous constructions of [3, 14]. Specifically, substituting $\ell = O(\log n)$ in our results we obtain a construction of trees that enjoy all properties of the construction of Arya et al. [3], and, *in addition*, have small lightness (specifically, $O(\log n \cdot (\epsilon^{-1}))$). Also, our trees combine the same optimal combination between the lightness and depth with the optimal maximum

¹The *depth* of a rooted tree (T, rt) , denoted $h(T)$, is the maximum number of hops in a path connecting the root vertex rt with a leaf z of T .

²For convenience, we will henceforth refer to the normalized notion of weight, called *lightness*, which is defined as the ratio between the weight and $w(MST(\mathcal{S}))$.

degree as the NLLTs of Dinitz et al. [14] do ³, and, *in addition*, achieve root-stretch at most $(1 + \epsilon)$. (See Table 1 for a concise comparison of our and previous constructions.)

1.2 General Metric Spaces

We also study the problem of constructing trees that satisfy all the aforementioned four properties (henceforth, *narrow-shallow-low-light* trees, or shortly, NSLLTs) in *general* metric spaces. We strengthen the results of Dinitz et al. [14], and demonstrate that one can trade maximum degree for root-stretch.

We start with building *shallow-low-light* trees (henceforth, SLLTs), i.e., trees that achieve small depth and lightness together with root-stretch at most $(1 + \epsilon)$, but may have arbitrarily large maximum degree. Specifically, we show that for every n -point metric space M , a designated point $rt \in M$, an integer $\ell = 1, 2, \dots, O(\log n)$ and a number $\epsilon > 0$, there exists a rooted spanning tree (T'_1, rt) with lightness $O(\ell \cdot (\epsilon^{-1}))$, depth $O(\ell \cdot n^{1/\ell})$ and root-stretch at most $(1 + \epsilon)$. There also exists a rooted spanning tree (T'_2, rt) with lightness $O(\ell \cdot n^{1/\ell} \cdot (\epsilon^{-1}))$, depth $O(\ell)$ and root-stretch at most $(1 + \epsilon)$.

Our upper bound argument then proceeds by showing that one can achieve small depth and lightness together with the optimal maximum degree, at the expense of increasing the root-stretch. Specifically, we demonstrate that for every n -point metric space M , a designated point $rt \in M$ and an integer $\ell = 1, 2, \dots, O(\log n)$, there exists a rooted spanning tree (\hat{T}_1, rt) with lightness $O(\ell)$, depth $O(\ell \cdot n^{1/\ell})$, constant maximum degree and root-stretch $O(\log n)$. There also exists a rooted spanning tree (\hat{T}_2, rt) with lightness $O(\ell \cdot n^{1/\ell})$, depth $O(\ell)$, maximum degree $O(n^{1/\ell})$ and root-stretch $O(\ell)$. In other words, these trees achieve the optimal tradeoff between the lightness and depth, *together with the optimal maximum degree*, at the expense of increasing the root-stretch from $(1 + \epsilon)$ to $O(\log n)$ and $O(\ell)$, respectively. In addition, we show that this increase in root-stretch is inevitable as long as one considers general (rather than low-dimensional Euclidean) metric spaces. Specifically, we show that our tradeoff between the maximum degree D and root-stretch $O(\frac{\log n}{\log D})$ cannot be improved even if M is a set of n points in Euclidean space of dimension $d = \Omega(\log n)$. We also extend this lower bound and show that in any dimension $d = O(\log n)$, the root-stretch is at least $\Omega(\frac{d}{\log D})$.

On the bright side, we demonstrate that this inherent tradeoff between the maximum degree and root-stretch is only valid when considering *spanning trees*. The situation changes drastically if one is allowed to add *Steiner points*, that is, points that do not belong to the original point set of M . In this case the root-stretch can be improved all the way down to $(1 + \epsilon)$, without increasing any of the other three parameters by more than a factor of $O(\epsilon^{-1})$. (Specifically, the lightness may grow by a factor of $O(\epsilon^{-1})$, while the bounds on the degree and the diameter stay intact.) We also show that the lower bound of [14] on the tradeoff between the lightness and depth of spanning trees also applies to trees that may include Steiner points (henceforth, *Steiner trees*). Consequently, similarly to the case of spanning trees, our tradeoffs between the four involved parameters are optimal with respect to Steiner trees as well. (See also Section 1.3.)

All our constructions for general metric spaces can be implemented in time $O(n^2)$, which is linear in the size of the input. If an *MST*, or a constant approximation of an *MST*, is given as a part of the input, then our constructions can be implemented in time $O(\text{SORT}(n)) = O(n \cdot \log n)$, where $\text{SORT}(n)$ is the time required to sort n distances. Moreover, if our metric space M is the induced metric of graph G with m edges, and G is given as a part of the input, then our constructions can be implemented in $O(m + n \cdot \log n)$ time.

We remark that our bounds (see, e.g., Table 1) are stated using the O -notation, which hides small constants. These constants depend on constants that are involved in previous constructions of spanning

³The optimality of our tradeoff between the lightness and depth in the entire range of parameters follows from lower bounds of [14]. The optimality of our tradeoff between depth and maximum degree, again in the entire range of parameters, is obvious.

			Ψ	h	Δ	ϱ
[26, 6]	general graphs (SLTs)		$O(\epsilon^{-1})$	∞	∞	$(1 + \epsilon)$
[3]	d -dimensional Euclidean spaces (single-sink spanners)		∞	$O(\log n)$	$O(\epsilon^{-d+1})$	$(1 + \epsilon)$
[14]	general metrics (NLLTs)	I	$O(\ell)$	$O(\ell \cdot n^{1/\ell})$	$O(1)$	∞
		II	$O(\ell \cdot n^{1/\ell})$	$O(\ell)$	$O(n^{1/\ell})$	∞
New	d -dimensional Euclidean spaces (NSLLTs)	I	$O(\ell \cdot (\epsilon^{-1}))$	$O(\ell \cdot n^{1/\ell})$	$O(\epsilon^{-d+1})$	$(1 + \epsilon)$
		II	$O(\ell \cdot n^{1/\ell} \cdot (\epsilon^{-1}))$	$O(\ell)$	$O(n^{1/\ell} + (\epsilon^{-d+1}))$	$(1 + \epsilon)$
New	general metrics (SLLTs)	I	$O(\ell \cdot (\epsilon^{-1}))$	$O(\ell \cdot n^{1/\ell})$	∞	$(1 + \epsilon)$
		II	$O(\ell \cdot n^{1/\ell} \cdot (\epsilon^{-1}))$	$O(\ell)$	∞	$(1 + \epsilon)$
New	general metrics (NSLLTs)	I	$O(\ell)$	$O(\ell \cdot n^{1/\ell})$	$O(1)$	$O(\log n)$
		II	$O(\ell \cdot n^{1/\ell})$	$O(\ell)$	$O(n^{1/\ell})$	$O(\ell)$
New	general metrics (Steiner NSLLTs)	I	$O(\ell \cdot (\epsilon^{-1}))$	$O(\ell \cdot n^{1/\ell})$	$O(1)$	$(1 + \epsilon)$
		II	$O(\ell \cdot n^{1/\ell} \cdot (\epsilon^{-1}))$	$O(\ell)$	$O(n^{1/\ell})$	$(1 + \epsilon)$

Table 1: A concise summary of previous and new constructions of NSLLTs and related trees. The column marked with Ψ (respectively, Δ ; resp., Δ ; resp., ϱ) provides upper bounds on lightness (resp., depth; resp., maximum degree; resp., root-stretch). The three upper rows of the table summarize known constructions, while the four bottom rows describe our results. In each of the five bottom rows there are two subrows, marked by I and II. The subrows marked by I (respectively, II) describe results for the range $h = \Omega(\log n)$ (resp., $h = O(\log n)$) of depth.

trees with related properties, specifically in the constructions of [3, 26, 14]. We made no attempt to optimize these constants. However, all constants in Table 1, excluding the coefficients of ϵ^{-d+1} , are not large, i.e., at most 50. On the other hand, the coefficients of ϵ^{-d+1} in the second (results of Arya et al. [3]) and fourth (our results) rows of the table (in the degree column) grow super-exponentially with the Euclidean dimension d .

1.3 Lower Bounds for Euclidean Spanners

We have proved two lower bounds on the tradeoffs between different parameters of NSLLTs. These lower bounds were mentioned above. Both these lower bounds have implications for Euclidean spanners. Next, we discuss these implications.

Our lower bound on the tradeoff between the lightness and depth parameters of Steiner trees implies directly a lower bound on the tradeoff between these parameters for Euclidean Steiner spanners. Specifically, Dinitz et al. [14] considered the 1-dimensional Euclidean space ϑ_n with n points $1, 2, \dots, n$ on the x -axis, and showed that any spanning tree of ϑ_n with depth $o(\log n)$ has weight $\omega(n \cdot \log n)$, and vice versa, i.e., any spanning tree of ϑ_n with weight $o(n \cdot \log n)$ has depth $\omega(\log n)$. This result implies that no construction of Euclidean spanners may guarantee hop-diameter⁴ $O(\log n)$ and lightness $o(\log n)$, or vice versa. Consequently, the construction of Arya et al. [3] of Euclidean spanners with both hop-diameter and lightness $O(\log n)$ is optimal. However, the lower bound of [14] does not preclude the existence of *Steiner* spanners with hop-diameter $O(\log n)$ and lightness $o(\log n)$, or vice versa. In the current paper we show that Steiner points do not help in this context, and thus the construction of Arya et al. [3] cannot be improved even if one allows the spanner to use (arbitrarily many) Steiner points.

⁴The *hop-diameter* or *unweighted diameter* of a possibly weighted graph $G = (V, E, w)$ is the maximum unweighted distance between a pair of vertices in G .

Interestingly, there are (at least) two possible notions of Steiner points in the context of Euclidean spanners. One notion, to which we refer to as *in-metric Steiner point*, is a point in the same Euclidean space in which the original n points lie. If we build a Euclidean spanner for a point set \mathcal{S} in \mathbb{R}^d , then an *in-metric (Euclidean) Steiner spanner* may use points in \mathbb{R}^d that do not belong to the set \mathcal{S} . This notion is more common in the geometric literature. Another notion, to which we refer to as *off-metric Steiner point*, is a point that may not belong to the original Euclidean space. For example, for a point set $\mathcal{S} \subseteq \mathbb{R}^2$, an off-metric Steiner point may belong to \mathbb{R}^3 . An *off-metric (Euclidean) Steiner spanner* for a point set \mathcal{S} is a weighted undirected graph $G = (V, E, w)$, with $V \supseteq \mathcal{S}$, such that for every pair of original points $u, v \in \mathcal{S}$, $\|u, v\| \leq \text{dist}_G(u, v)$. Surprisingly, our lower bounds apply not only to in-metric Steiner spanners but rather to this very general notion of off-metric Steiner spanners. Henceforth, in the sequel we restrict the attention to off-metric Steiner points and spanners, and write “Steiner points” and “Steiner spanners” as shortcuts for “off-metric Steiner points” and “off-metric Steiner spanners”, respectively.

Our lower bound on the tradeoff between the maximum degree D and root-stretch $\Omega(\frac{d}{\log D})$ of spanning trees for point sets in \mathbb{R}^d , $d = O(\log n)$, implies that if $d = \omega(1)$ is super-constant, then either the maximum degree or the root-stretch is super-constant as well. Consequently, no construction of Euclidean spanners for a super-constant dimension can possibly achieve simultaneously constant maximum degree and stretch. On the other hand, for any constant dimension d , Arya et al. [3] have built spanners with stretch at most $(1 + \epsilon)$ for arbitrarily small $\epsilon > 0$ and constant maximum degree D . Hence our lower bound implies that this result of Arya et al. [3] cannot be extended to super-constant dimension.

1.4 Proof Overview

Both our Euclidean and general constructions of NSLLTs are based on the following insight. To construct a tree that enjoys the optimal combination of all four parameters, one can construct two different trees each of which is good with respect to only three out of the four parameters, and combine them into a single tree. Specifically, we start with constructing a tree that achieves small maximum degree, root-stretch and depth, henceforth *narrow-shallow-low tree* (NSLoT), from scratch. Then we proceed to building an SLLT. To this end we adapt the shallow-light trees (henceforth, SLTs) construction of Awerbuch et al. [6] to our needs, and then plug it on top of the construction of NLLTs due to Dinitz et al. [14]. While the root vertex in the resulting SLLT T' is prone to suffer from an arbitrarily large degree, we demonstrate that all other vertices necessarily achieve the optimal (up to a constant factor) degree. Next, to reduce the degree of the root vertex rt in T' , we manipulate with the star subtree Z rooted at rt . The vertex set $V(Z)$ of the subtree Z contains the root vertex rt of T' and all its children c_1, c_2, \dots, c_q , and the edge set of Z is the set $\{(rt, c_1), (rt, c_2), \dots, (rt, c_q)\}$. Then we construct an NSLoT \tilde{T} for the point set $V(Z)$. Finally, we remove the star Z from the SLLT T' , and replace it with the NSLoT \tilde{T} . We show that the resulting tree \hat{T} is an NSLLT, i.e., enjoys *all the four desired properties*. (See Figure 1 for an illustration.)

We remark that reducing the root-degree of an SLLT directly appears to be difficult. To illustrate the problem, consider a clique on n vertices, with a designated root vertex rt . All edges of the clique that are adjacent to rt are assigned unit weight, and all other edges are assigned weight 2. Let S denote the star that contains all edges adjacent to rt . Obviously, S is an optimal SLLT for this clique, as it enjoys optimal root-stretch, lightness and depth. However, its maximum degree is $n - 1$, and it cannot be substantially reduced without increasing significantly both the depth and the root-stretch.

Our Euclidean construction of NSLoTs is based on the construction of Arya et al. [3], which was, in turn, inspired by the work of Ruppert and Seidel [36]. However, our construction provides a general tradeoff between the maximum degree D and the depth $O(\log_D n)$, while in the construction of [3] the maximum degree is $O(1)$ and the depth is $O(\log n)$. (In addition, whenever $D = \omega(1)$, the running time of our generalized construction is $O(n \cdot \log_D n)$, which is better than the running time $O(n \cdot \log n)$ in [3].)




Figure 1: The root vertex rt of the SLLT T' may have a large degree. The star subtree Z of T' is replaced by the NSLoT \hat{T} to obtain the NSLLT \hat{T} .

This extension is a natural one, and we provide it for completeness.

Our construction of NSLoTs for general metric spaces is based on different principles. We show that it achieves the optimal tradeoff between the three involved parameters. We anticipate that this construction will be useful for other applications.

Finally, in Section 6 we show that the tradeoff of [14] between the lightness and hop-diameter of Euclidean spanners cannot be improved by using Steiner points. To this end we demonstrate that any Steiner tree can be “cleaned” from Steiner points, while increasing the depth and lightness by only a small factor. This result is reminiscent of the work by Gupta [22] that shows that as far as *maximum* stretch and lightness are concerned, one can do without Steiner points. However, our argument is substantially different from that of [22], since, in particular, the hop-diameter parameter exhibits a different behavior than the maximum stretch.

1.5 Related Work

Euclidean spanners are being a subject of ongoing intensive research since the mid-eighties. See the recent book by Narasimhan and Smid [34] for an excellent survey on this subject. Euclidean single-sink spanners were studied by Arya et al. [3]; see also [34], Chapter 4.2. Lukovszki [29, 30] devised fault-tolerant constructions of single-sink spanners. Single-sink spanners were also used in maintenance algorithms for wireless networks [20, 31]. Farshi and Gudmundsson [16] conducted an experimental study of single-sink spanners.

Arya et al. [3] have also devised a construction of $(1+\epsilon)$ -spanners with hop-diameter $O(\log n)$, lightness $O(\log^2 n)$ and maximum degree $O(\epsilon^{-2d+2})$. This result is incomparable to our result. On the one hand, the spanner of [3] is not a tree⁵. Furthermore, its lightness is $O(\log^2 n)$, and not $O(\log n)$ as in our construction. On the other hand, it provides a $(1+\epsilon)$ -approximation to *all* distances, while our tree provides a $(1+\epsilon)$ -approximation to distances between a designated root vertex rt and all other vertices. Also, we remark that, as far as we know, there is no simple way to “extract” a good NSLLT from the

⁵Moreover, it is well-known (see, e.g., [15]) that *no* construction of Euclidean spanners that guarantees stretch $o(n)$ for *all* pairs of vertices may possibly be cycle-free.

$(1 + \epsilon)$ -spanner of [3]. In particular, suppose that we start with constructing a spanner H with the above properties, and then compute a shortest-paths tree (henceforth, *SPT*) T rooted at rt over H . This tree has root-stretch at most $(1 + \epsilon)$, lightness $O(\log^2 n)$ and maximum degree $O(\epsilon^{-2d+2})$, but its depth may be much larger than $O(\log n)$. On the other hand, suppose that we compute T by taking the union of spanner paths between rt and v , for all $v \in V$, $v \neq rt$. The resulting subgraph T satisfies all the four requirements (though its lightness $O(\log^2 n)$ is still significantly larger than that of our tree), but it is *not necessarily a tree*. To convert it into a tree one needs to remove some edges that close cycles. However, this may result in either an arbitrarily large depth or an arbitrarily large root-stretch.

Trees that have small lightness and guarantee root-stretch at most $(1 + \epsilon)$, but do not necessarily have small depth or small maximum degree, are called *shallow-light trees* (henceforth SLTs). SLTs were studied by a number of authors, including Awerbuch et al. [5, 6], Khuller et al. [26], Alpert et al. [1] and Cong et al. [11, 12, 13]. Salowe et al. [37] studied trees that combine small lightness with small “bottleneck” size; see [37] for further details. Papadimitriou and Vazirani [35], Monma and Suri [33], Fekete et al. [17] and Chan [8] devised constructions of light trees with small maximum degree for low-dimensional Euclidean point sets. See also the survey of Eppstein [15] for other references to works that study geometric spanning trees. Approximation algorithms for closely related problems were devised in [2, 10, 25, 28, 32].

In [23] Jia et al. studied *universal trees*. Consider a spanning tree T of an n -point metric space M , rooted at a designated point $rt \in M$. For a subset $Q \subseteq M$ of *terminal* points, let T_Q denote the subtree of T restricted to the vertex set Q . In other words, the subtree T_Q is obtained from T by pruning all vertices of T whose subtrees contain no terminal points. For a parameter $\alpha \geq 1$, the tree T is said to be α -*approximate universal tree* with respect to M and rt if for every subset $Q \subseteq M$ of terminal points the tree T_Q has weight which is greater by at most a factor of α than the weight of the minimum spanning tree $MST(Q)$ of the set Q . Jia et al. [23] devised a construction of $O(\frac{\log^4 n}{\log \log n})$ -approximate (respectively, $O(\log n)$ -approximate) universal trees for general (resp., doubling) metrics. Since one can substitute $Q = V$ these trees have relatively small lightness. However, there is no guarantee whatsoever on their maximum degree, depth or root-stretch.

Another line of research that focuses on spanning trees that are good for all subsets of the original metric space is the study of *single-sink buy-at-bulk* problem. In this context together with the metric space M and the root point rt one is given a concave cost function f . For a subset $Q \subseteq M$ of terminal points, a tree T that spans Q and an edge e of T , let $j(e)$ denote the number of terminal points $v \in Q$ that satisfy that the path in T that connects rt and v passes through the edge e . The *cost* of e with respect to the function f and the tree T is defined by $C(e) = f(j(e))$. The cost of the tree T is $C(T) = \sum_{e \in T} C(e)$. The *single-sink buy-at-bulk* (henceforth, *SSBAB*) problem [4, 21] asks for constructing a tree T such that for *every* subset $Q \subseteq M$ of points, the cost $C(T_Q)$ of the tree T restricted to the subset Q is close to the optimal one. Goel and Estrin [18] introduced the *simultaneous SSBAB* problem in which one wishes to build a tree T which is good (in the above sense) for all concave non-decreasing functions f that satisfy $f(0) = 0$ *simultaneously*. See also [19] and the references therein for most recent results on this problem. We remark that this thread of research is similar in spirit to our research, as we also aim to construct a spanning tree which is good with respect to a number of distinct parameters simultaneously. On the other hand, the results of [18, 19] are incomparable to ours, because the optimized parameters are different.

1.6 Structure of the Paper

In Section 2 we define the basic notions and present the notation that is used throughout the paper. In Section 3 we describe our constructions of NSLoTs. Therein we start with presenting our Euclidean construction of NSLoTs. Then we proceed with describing two constructions for general metric spaces. The first (respectively, second) of them is a construction of spanning (resp., Steiner) NSLoTs for general

metric spaces.

In Section 4 we present our construction of SLLTs. In Section 5 we employ the construction of SLLTs from Section 4 in conjunction with the constructions of NSLoTs from Section 3 to devise our ultimate constructions of NSLLTs. Similarly to Section 3, there are three constructions in Section 5. The first one is the Euclidean one. The second and the third are for general metric spaces, where the second builds spanning NSLLTs and the third builds Steiner NSLLTs. Finally, Section 6 is devoted to our lower bounds for Euclidean Steiner spanners.

2 Preliminaries

An n -point metric space $M = (V, dist)$ can be viewed as the complete graph $G(M) = (V, \binom{V}{2}, dist)$ in which for every pair of points $x, y \in V$, the weight of the edge $e = (x, y)$ between x and y in $G(M)$ is defined by $w(x, y) = dist(x, y)$.

For a rooted tree (T, rt) and a vertex v in T , the *level* of v in T is the unweighted distance between the root vertex rt of T and v in T . Denote by $deg(T, v)$ the degree of a vertex v in T , and define $\Delta(T) = \max\{deg(T, v) : v \in V\}$ to be the maximum degree of a vertex in T . Also, define $\lambda(T) = \max\{deg(T, v) : v \in V, v \neq rt\}$ to be the maximum degree of a non-root vertex in T . For any two vertices $u, v \in V(T)$, their weighted distance in T is denoted by $dist_T(u, v)$. For a positive integer D , a rooted tree in which every vertex has at most D children is called a D -ary tree.

A tree T is called a *Steiner tree* of a metric space $M = (V, dist)$ if it spans a superset of V and if for any pair of points $u, v \in V$, $dist_T(u, v) \geq dist(u, v)$. Let T be either a spanning or a Steiner tree of M rooted at an arbitrary designated vertex rt . We define the *stretch* between two vertices u and v in V to be $\zeta_T(u, v) = \frac{dist_T(u, v)}{dist(u, v)}$, and the *root-stretch* of (T, rt) to be $\varrho(T, rt) = \max\{\zeta_T(rt, v) : v \in V\}$.

For a positive integer n , we denote the set $\{1, 2, \dots, n\}$ by $[n]$.

3 Narrow-Shallow-Low Trees (NSLoTs)

In this section we devise our constructions of *narrow-shallow-low* trees (NSLoTs), i.e., trees that achieve small maximum degree, depth and root-stretch simultaneously. Even though our NSLoTs may be quite heavy, we do require their weight to be bounded by $O(\sum_{v \in M} dist(rt, v))$, where rt is the designated root. We call the quantity $\sum_{v \in M} dist(rt, v)$ the *star-weight* of the point set M with respect to rt , and denote it $W^*(M, rt)$.

The following statement summarizes the properties of our construction of NSLoTs for point sets in the plane.

Proposition 3.1 *Let $k \geq 8$ and $\theta = 2\pi/k$. For any set V of n points in the plane, an arbitrary designated point rt and a positive integer $2 \leq D \leq n-1$, there exists a $(2D+k)$ -ary rooted spanning tree (T_θ, rt) for V with weight at most $\frac{1}{\cos\theta - \sin\theta} \cdot W^*(V, rt)$, depth at most $\log_D n$ and root-stretch at most $\frac{1}{\cos\theta - \sin\theta}$. Moreover, T_θ can be constructed in $O(n \cdot \log_D n)$ time.*

Remarks: (1) For large k , $\frac{1}{\cos\theta - \sin\theta} = 1 + O(\theta)$. Hence we get a tree with maximum degree $O(D + (\theta^{-1}))$, depth at most $\log_D n$, root-stretch $1 + O(\theta)$ and weight $O(W^*(V, rt))$. (2) The requirement to have $k \geq 8$ is necessary to guarantee a root-stretch of at most $\frac{1}{\cos\theta - \sin\theta}$. In particular, Claim 3.3 below does not hold for smaller values of k .

Proof: For any $D \geq \frac{n-10}{2}$, the star graph rooted at rt satisfies the conditions of the proposition. We henceforth assume that $D < \frac{n-10}{2}$.

If we rotate the positive x -axis by angles $i \cdot \theta$, $0 \leq i < k$, then we get k rays. Each pair of successive rays defines a *cone* that spans an angle of θ and whose apex is at the origin. Denote by $\mathcal{C} = \mathcal{C}_\theta = \{C_1, C_2, \dots, C_k\}$ the collection of the resulting k cones. (Note that these k cones partition the plane.) For a cone C_i of \mathcal{C} and a point p in the plane, we define $C_i(p) = C_i + p = \{x+p : x \in C_i\}$ to be the cone obtained from C_i by *translating* it such that its apex is at p . We define $\mathcal{C}(p) = \mathcal{C}_\theta(p) = \{C_1(p), C_2(p), \dots, C_k(p)\}$ to be the collection obtained from \mathcal{C} by translating each cone in it such that its apex is at p .

We denote by $V_i(p) = V \cap C_i(p)$ the subset of V contained in a cone $C_i(p)$ of $\mathcal{C}(p)$. Observe that the collection $\{V_1(p), V_2(p), \dots, V_k(p)\}$ is a partition of $V \setminus \{p\}$. For each $i \in [k]$, we define $n_i = |V_i(p)|$. Let $\mathcal{P}(p)$ be the collection obtained from $\{V_1(p), V_2(p), \dots, V_k(p)\}$ by partitioning each set $V_i(p)$ in it (arbitrarily) into $\lceil \frac{n_i}{\lfloor n/D \rfloor} \rceil$ subsets of size at most $\lfloor n/D \rfloor$ each.

Claim 3.2 *For any point set V and any point p in the plane, $|\mathcal{P}(p)| \leq 2D + k$.*

Proof: Observe that

$$|\mathcal{P}(p)| = \sum_{i=1}^k \left\lceil \frac{n_i}{\lfloor n/D \rfloor} \right\rceil \leq \sum_{i=1}^k \left(\frac{n_i}{\lfloor n/D \rfloor} + 1 \right) = \frac{1}{\lfloor n/D \rfloor} \left(\sum_{i=1}^k n_i \right) + k.$$

Since the collection $\{V_1(p), V_2(p), \dots, V_k(p)\}$ is a partition of $V \setminus \{p\}$, we have $\sum_{i=1}^k n_i \leq n$. Hence, $|\mathcal{P}(p)| \leq \frac{n}{\lfloor n/D \rfloor} + k \leq 2D + k$. (The last inequality holds for all $D \leq n/2$.) ■

Observe that if D divides n then

$$|\mathcal{P}(p)| = \sum_{i=1}^k \left\lceil \frac{n_i}{n/D} \right\rceil \leq \sum_{i=1}^k \left(\frac{n_i}{n/D} + 1 \right) = D + k.$$

The tree $T = T_\theta$ is constructed recursively as follows. First, we compute a partition $\mathcal{P}(rt) = \{P_1(rt), P_2(rt), \dots, P_m(rt)\}$ of $V \setminus \{rt\}$ as described above, where $m = |\mathcal{P}(rt)| \leq 2D + k$. For each $i \in [m]$, let $rt(i)$ be the point in $P_i(rt)$ whose orthogonal projection onto the bisector of the cone in $\mathcal{C}(rt)$ that contains it is closest to rt . For each $i \in [m]$, we set $rt(i)$ to be a child of the root vertex rt , and construct a rooted tree $(T_i, rt(i))$ for the subset $P_i(rt)$ recursively. The recursion stops if a subset has size one.

Next, we analyze the properties of the constructed tree T . First, it is easy to see that T is a $(2D + k)$ -ary spanning tree of V rooted at rt , and that its depth is at most $\log_D n$.

To bound the root-stretch of T we use the following claim from [7] (see also [34], Lemma 4.1.4, page 66).

Claim 3.3 *Let $k \geq 8$ and $\theta = 2\pi/k$. Let p and q be two distinct points in the plane, and let C be the cone of $\mathcal{C}_\theta(p)$ such that $q \in C$. Let r be a point in C such that the orthogonal projection of r onto the bisector of C is at least as close to p as the orthogonal projection of q onto the bisector of C . Then*

$$\|p, r\| + \frac{1}{\cos \theta - \sin \theta} \cdot \|r, q\| \leq \frac{1}{\cos \theta - \sin \theta} \cdot \|p, q\|.$$

The next lemma shows that the root-stretch of T is small.

Lemma 3.4 *For each point v in V , $\text{dist}_T(rt, v) \leq \frac{1}{\cos \theta - \sin \theta} \cdot \|rt, v\|$.*

Proof: The proof is by induction on $n = |V|$. The basis $n = 1$ is trivial.

Induction Step: We assume that the statement holds for all smaller values of n and prove it for n .

If $v = rt$, then we have $dist_T(rt, v) = 0 \leq \frac{1}{\cos\theta - \sin\theta} \cdot \|rt, v\|$.

We henceforth assume that $v \neq rt$. Let $i \in [m]$ be the index such that $v \in P_i(rt)$. By construction,

$$dist_T(rt, v) = dist_T(rt, rt(i)) + dist_T(rt(i), v) = \|rt, rt(i)\| + dist_{T_i}(rt(i), v).$$

By the induction hypothesis, $dist_{T_i}(rt(i), v) \leq \frac{1}{\cos\theta - \sin\theta} \cdot \|rt(i), v\|$. It follows that

$$dist_T(rt, v) \leq \|rt, rt(i)\| + \frac{1}{\cos\theta - \sin\theta} \cdot \|rt(i), v\|. \quad (1)$$

Observe that $rt(i)$ and v belong to the same cone of $\mathcal{C}(rt)$. That cone spans an angle of θ and has its apex at rt . Moreover, the orthogonal projection of $rt(i)$ onto the bisector of that cone is at least as close to rt as that of v . Hence, by Claim 3.3, the right-hand side of Equation (1) satisfies

$$\|rt, rt(i)\| + \frac{1}{\cos\theta - \sin\theta} \cdot \|rt(i), v\| \leq \frac{1}{\cos\theta - \sin\theta} \cdot \|rt, v\|,$$

and we are done. \blacksquare

The next lemma establishes an upper bound on the weight of T . It follows directly from Lemma 3.4.

Lemma 3.5 $w(T) \leq \frac{1}{\cos\theta - \sin\theta} \cdot W^*(V, rt)$.

Proof: For $v \neq rt$, denote by $\pi(v)$ the parent of v in T . Observe that

$$w(T) = \sum_{v \in V, v \neq rt} dist_T(\pi(v), v) \leq \sum_{v \in V, v \neq rt} dist_T(rt, v).$$

By Lemma 3.4, we get that

$$w(T) \leq \sum_{v \in V, v \neq rt} dist_T(rt, v) \leq \sum_{v \in V, v \neq rt} \frac{1}{\cos\theta - \sin\theta} \cdot \|rt, v\| = \frac{1}{\cos\theta - \sin\theta} \cdot W^*(V, rt). \quad \blacksquare$$

Finally, we show that this construction can be implemented efficiently.

Lemma 3.6 *The tree T can be constructed in $O(n \cdot \log_D n)$ time in the real-RAM model.*

Proof: First, we note that the partition $\mathcal{P}(rt) = \{P_1(rt), P_2(rt), \dots, P_m(rt)\}$ can be computed in $c \cdot n$ time, for a sufficiently large constant c .

Next, we prove by induction on $n = |V|$ that the tree T can be constructed in $c'' \cdot n \cdot \log_D n$ time, where $c'' = c + c'$, and c' is a sufficiently large constant. The basis $n = 1$ is trivial.

Induction Step: We assume that the statement holds for all smaller values of n and prove it for n . The time required to connect the root vertex rt to its children $rt(1), rt(2), \dots, rt(m)$ is bounded from above by $c' \cdot n$. By the induction hypothesis, for each $i \in [m]$, the tree T_i can be constructed in $c'' \cdot |P_i(rt)| \cdot \log_D |P_i(rt)|$ time. Note that $\max\{|P_i(rt)| : i \in [m]\} \leq \lfloor n/D \rfloor \leq n/D$. It follows that the overall running time required to construct T is no greater than

$$c \cdot n + c' \cdot n + \sum_{i=1}^m c'' \cdot |P_i(rt)| \cdot \log_D |P_i(rt)| \leq c'' \cdot n + c'' \cdot \log_D(n/D) \cdot \sum_{i=1}^m |P_i(rt)| \leq c'' \cdot n \cdot \log_D n. \quad \blacksquare$$

Lemmas 3.4, 3.5 and 3.6 imply the validity of Proposition 3.1. \blacksquare

The following statement generalizes Proposition 3.1 for point sets in \mathbb{R}^d , for arbitrarily large constant $d \geq 2$. Its proof is a natural generalization of the corresponding proof for Proposition 3.1, and thus most details are omitted.

Proposition 3.7 *Let θ be a real number such that $0 < \theta < \pi/4$, and let $d \geq 2$ be an integer constant. Define $k = 2d! \cdot \left[\sqrt{2(d-1)/(1-\cos\theta)} \right]^{d-1} = O(\theta^{-d+1})$. For any set V of n points in \mathbb{R}^d , an arbitrary designated point rt and an integer $2 \leq D \leq n-1$, there exists a $(2D+k)$ -ary rooted spanning tree (T_θ, rt) of M with weight at most $\frac{1}{\cos\theta - \sin\theta} \cdot W^*(V, rt)$, depth at most $\log_D n$ and root-stretch at most $\frac{1}{\cos\theta - \sin\theta} \cdot$. Moreover, T_θ can be constructed in $O(n \cdot \log_D n \cdot (\log \theta^{-1}) + (\theta^{-d+1}))$ time.*

Remark: For large k , $\frac{1}{\cos\theta - \sin\theta} = 1 + O(\theta)$. Hence we get a tree with maximum degree $O(D + (\theta^{-d+1}))$, depth at most $\log_D n$, root-stretch $1 + O(\theta)$ and weight $O(W^*(V, rt))$.

Proof: For any $D \geq \frac{n-k-1}{2}$, the star graph rooted at rt satisfies the conditions of the proposition. We henceforth assume that $D < \frac{n-k-1}{2}$.

To generalize the collection \mathcal{C} of two-dimensional cones presented in the proof of Proposition 3.1 for any constant dimension d , we use the θ -frame construction from [30] (cf. [34], Theorem 5.2.8, page 97). Refer also to [38] and [36] for earlier constructions of θ -frame having a larger number of cones. The construction of [30] provides a collection $\mathcal{C} = \mathcal{C}^d$ of $k = O(\theta^{-d+1})$ simplicial d -dimensional cones with disjoint interiors that cover \mathbb{R}^d , each cone of which has *angular diameter* at most θ and has its apex at the origin. (Refer to [34], page 93, for the definitions of simplicial cone and angular diameter.) Moreover, the collection \mathcal{C} can be constructed in $O(k) = O(\theta^{-d+1})$ time. The translated cone $C_i(p)$ and the collection $\mathcal{C}(p)$ of translated cones are defined in the same way as in the two-dimensional case.

The d -dimensional analogues for the collection $\mathcal{P}(p)$ and the tree $T = T_\theta$ are constructed similarly to the two-dimensional case. Claim 3.2 and Lemmas 3.4 and 3.5 hold true also in the d -dimensional case. The following claim is the d -dimensional analogue of Lemma 3.6.

Lemma 3.8 *The tree T can be constructed in $O(n \cdot \log_D n \cdot (\log \theta^{-1}) + (\theta^{-d+1}))$.*

Proof: First, we construct the collection \mathcal{C} in $O(\theta^{-d+1})$ time. Next, we preprocess \mathcal{C} to answer *point location* queries of the following form: Given any two points p, q in \mathbb{R}^d , compute in $O(\log \theta^{-1})$ time a cone in $\mathcal{C}(p)$ that contains q . The preprocessing phase can be performed in $O(\theta^{-d+1})$ time (cf. [34], Theorem 5.3.2, page 98). Since both phases of constructing and preprocessing \mathcal{C} require together at most $O(\theta^{-d+1})$ time, we disregard them in the sequel.

We claim that the partition $\mathcal{P}(rt) = \{P_1(rt), P_2(rt), \dots, P_m(rt)\}$ can be computed in $c \cdot n \cdot (\log \theta^{-1})$ time, for a sufficiently large constant c . To see this, first note that the above preprocessing phase enables us to determine in $O(\log \theta^{-1})$ time to which cone in $\mathcal{C}(rt)$ does an arbitrary point in \mathbb{R}^d belong. Hence the partition $\{V_1(rt), V_2(rt), \dots, V_k(rt)\}$ of $V \setminus rt$ can be computed in $O(n \cdot (\log \theta^{-1}))$ time. For each $i \in [m]$, the time required to partition the set $V_i(rt)$ *arbitrarily* into $\left\lceil \frac{n_i}{\lfloor n/D \rfloor} \right\rceil$ subsets of size at most $\lfloor n/D \rfloor$ each is linear in its size n_i , implying the claim.

We prove by induction on $n = |V|$ that T can be constructed in $c'' \cdot n \cdot \log_D n \cdot (\log \theta^{-1})$ time, where $c'' = c + \frac{c'}{\log \theta^{-1}}$, and c' is a sufficiently large constant. Note that for $\theta \leq \pi/4$, $\log \theta^{-1} \geq 1/3$, and so c'' is also a constant. The basis $n = 1$ is trivial.

Induction Step: We assume that the statement holds for all smaller values of n and prove it for n . The time required to connect rt to its children $rt(1), rt(2), \dots, rt(m)$ is bounded from above by $c' \cdot n$. By the induction hypothesis, for each $i \in [m]$, the tree T_i can be constructed in $c'' \cdot |P_i(rt)| \cdot \log_D |P_i(rt)| \cdot (\log \theta^{-1})$ time. Note that $\max\{|P_i(rt)| : i \in [m]\} \leq \lfloor n/D \rfloor \leq n/D$. We get that the overall running time required to construct T is at most

$$c \cdot n \cdot (\log \theta^{-1}) + c' \cdot n + \sum_{i=1}^m c'' \cdot |P_i(rt)| \cdot \log_D |P_i(rt)| \cdot (\log \theta^{-1}).$$

Observe that $c \cdot n \cdot (\log \theta^{-1}) + c' \cdot n = c'' \cdot n \cdot (\log \theta^{-1})$. Also, we have

$$\begin{aligned} \sum_{i=1}^m c'' \cdot |P_i(rt)| \cdot \log_D |P_i(rt)| \cdot (\log \theta^{-1}) &\leq c'' \cdot \log_D(n/D) \cdot (\log \theta^{-1}) \cdot \sum_{i=1}^m |P_i(rt)| \\ &\leq c'' \cdot n \cdot \log_D(n/D) \cdot (\log \theta^{-1}). \end{aligned}$$

It follows that the time required to construct T is at most

$$c'' \cdot n \cdot (\log \theta^{-1}) + c'' \cdot n \cdot \log_D(n/D) \cdot (\log \theta^{-1}) = c'' \cdot n \cdot \log_D n \cdot (\log \theta^{-1}). \quad \blacksquare$$

This completes the proof of Proposition 3.7. \blacksquare

Next, we devise a construction of NSLoTs for general metric spaces.

Proposition 3.9 *For any n -point metric space $M = (V, \text{dist})$, an arbitrary designated point rt and a positive integer $2 \leq D \leq n-1$, there exists a D -ary rooted spanning tree (T, rt) of M with weight at most $2 \cdot W^*(M, rt)$, depth at most $\lceil \log_D n \rceil$ and root-stretch at most $2 \cdot \lceil \log_D n \rceil$.*

Proof: Let $V = (rt = v_0, v_1, \dots, v_{n-1})$. Suppose without loss of generality that the n points $rt = v_0, v_1, \dots, v_{n-1}$ are ordered according to their distance from rt , i.e., $0 = \text{dist}(rt, v_0) \leq \text{dist}(rt, v_1) \leq \dots \leq \text{dist}(rt, v_{n-1})$.

In what follows we construct a rooted tree (T, rt) that satisfies the conditions of the proposition. The D points v_1, v_2, \dots, v_D become the children of $rt = v_0$ in T , the next D points $v_{D+1}, v_{D+2}, \dots, v_{2D}$ become the children of v_1 , the next D points $v_{2D+1}, v_{2D+2}, \dots, v_{3D}$ become the children of v_2 , and so on. Generally, the point v_i becomes the child of point $v_{\lceil \frac{i}{D} \rceil - 1}$ in T , for each index $i \in [n-1]$. (See Figure 2.a for an illustration.)




Figure 2: a) A 3-ary NSLoT for a 12-point metric space. b) A 3-ary Steiner NSLoT for a 9-point metric space $M = (V, \text{dist})$. The Steiner points are d_1, d_2 and d_3 , and all other points $rt = v_0, v_1, \dots, v_8$ belong to the original point set V . Edges of weight ϵ are depicted by thin lines. Edges of greater weight (specifically, edges $(v_i, \pi(v_i))$ of weight $\text{dist}(rt, v_i)$) are depicted by bold lines.

Observation 3.10 (1) For any pair of vertices v and w , such that v is an ancestor of w in T , $\text{dist}(rt, v) \leq \text{dist}(rt, w)$, (2) (T, rt) is a D -ary rooted spanning tree of M , (3) The depth $h(T)$ of the rooted tree (T, rt) is no greater than $\lceil \log_D n \rceil$.

The next lemma provides the required bound on the weight of the constructed tree T .

Lemma 3.11 $w(T) \leq 2 \cdot W^*(M, rt)$.

Proof: Denote by $\pi(v)$ the parent of v in T , $v \neq rt$. By the triangle inequality and the first assertion of Observation 3.10,

$$\text{dist}_T(\pi(v), v) = \text{dist}(\pi(v), v) \leq \text{dist}(rt, \pi(v)) + \text{dist}(rt, v) \leq 2 \cdot \text{dist}(rt, v).$$

Consequently,

$$w(T) = \sum_{v \in V, v \neq rt} \text{dist}_T(\pi(v), v) \leq \sum_{v \in V, v \neq rt} 2 \cdot \text{dist}(rt, v) = 2 \cdot W^*(M, rt). \quad \blacksquare$$

Next, we analyze the root-stretch of T .

Lemma 3.12 *For a vertex v of level i in T , $\text{dist}_T(rt, v) \leq (2 \cdot i - 1) \cdot \text{dist}(rt, v)$.*

Proof: The proof is by induction on the level of v in T . The basis $i = 0$ is trivial.

Induction Step: We assume that the statement holds for all smaller values of $i \geq 1$ and prove it for i . Since $i \geq 1$, v has a parent in T , denoted $\pi(v)$. By construction,

$$\text{dist}_T(rt, v) = \text{dist}_T(rt, \pi(v)) + \text{dist}_T(\pi(v), v) = \text{dist}_T(rt, \pi(v)) + \text{dist}(\pi(v), v).$$

Observe that the level of $\pi(v)$ in T is $i - 1$, and so by the induction hypothesis, $\text{dist}_T(rt, \pi(v)) \leq (2 \cdot i - 3) \cdot \text{dist}(rt, \pi(v))$. By the triangle inequality, $\text{dist}(\pi(v), v) \leq \text{dist}(rt, \pi(v)) + \text{dist}(rt, v)$. Hence,

$$\text{dist}_T(rt, v) = \text{dist}_T(rt, \pi(v)) + \text{dist}(\pi(v), v) \leq (2 \cdot i - 2) \cdot \text{dist}(rt, \pi(v)) + \text{dist}(rt, v).$$

By the first assertion of Observation 3.10, $\text{dist}(rt, \pi(v)) \leq \text{dist}(rt, v)$. This completes the proof. \blacksquare

The third assertion of Observation 3.10 and Lemma 3.12 imply that the root-stretch of (T, rt) is at most $2 \cdot \lceil \log_D n \rceil$, which concludes the proof of Proposition 3.9. \blacksquare

The next statement implies that the upper bound given in Proposition 3.9 is tight up to constant factors. In particular, it shows that the tradeoff D versus $O(\frac{\log n}{\log D})$ between the maximum degree and root-stretch established there cannot be improved even for Euclidean spaces of dimension $\Theta(\log n)$.

Proposition 3.13 *There exists a set V of n points in $\mathbb{R}^{O(\log n)}$, such that for any positive integer $2 \leq D \leq n - 1$ and any point $v \in V$, every D -ary spanning tree of V rooted at $rt = v$ has weight $\Omega(W^*(V, rt))$, depth at least $\lfloor \log_D n \rfloor$ and root-stretch $\Omega(\log_D n)$.*

Proof: Let $\mathcal{E}_n = \{e_1, e_2, \dots, e_n\}$, $e_1 = (1, 0, \dots, 0)$, $e_2 = (0, 1, 0, \dots, 0)$, \dots , $e_n = (0, \dots, 0, 1)$, be the point set of the n vectors of the standard basis in \mathbb{R}^n . By the Flattening Lemma due to Johnson and Lindenstrauss [24], for any constant $\epsilon > 0$, there exists n points e'_1, e'_2, \dots, e'_n in $\mathbb{R}^{O(\log n)}$ that satisfy that for every pair of distinct indices $i, \ell \in [n]$,

$$\|e_i, e_\ell\| \leq \|e'_i, e'_\ell\| \leq (1 + \epsilon) \cdot \|e_i, e_\ell\|. \quad (2)$$

Let (T', e'_j) be some D -ary spanning tree for $\mathcal{E}'_n = \{e'_1, e'_2, \dots, e'_n\}$ rooted at an arbitrary point $e'_j \in \mathcal{E}'_n$. Since T' is D -ary, its depth is at least $\lfloor \log_D n \rfloor$. To bound the weight and root-stretch of T' , consider the corresponding spanning tree (T, e_j) of \mathcal{E}_n rooted at e_j . (The edge set $E(T)$ of T is given by $E(T) = \{(e_i, e_\ell) : (e'_i, e'_\ell) \in E(T')\}$.) Observe that the weight of any edge in T is $\sqrt{2}$. Hence, by Equation (2), the total weight $w(T')$ of T' satisfies

$$w(T') \geq w(T) = \sqrt{2} \cdot (n - 1) = \sum_{i \in [n], i \neq j} \|e_j, e_i\| \geq \sum_{i \in [n], i \neq j} \frac{\|e'_j, e'_i\|}{1 + \epsilon} = \frac{W^*(\mathcal{E}'_n, e'_j)}{1 + \epsilon}.$$

Observe that T is D -ary, and so its depth is at least $\lfloor \log_D n \rfloor$. It follows that the distance in T between e_j and some vertex farthest away from it, denoted e_h , is at least $\sqrt{2} \cdot \lfloor \log_D n \rfloor$, implying that

$$\frac{\text{dist}_T(e_j, e_h)}{\|e_j, e_h\|} \geq \lfloor \log_D n \rfloor. \quad (3)$$

Note also that if $P' = (e'_j = e'^{(0)}, e'^{(1)}, \dots, e'_h = e'^{(k)})$ is the unique path between e'_j and e'_h in T' , then $P = (e_j = e^{(0)}, e^{(1)}, \dots, e_h = e^{(k)})$ is the corresponding path between e_j and e_h in T . Equation (2) yields $\|e'^{(i)}, e'^{(i+1)}\| \geq \|e^{(i)}, e^{(i+1)}\|$, for every index $i \in [0, k-1]$. Hence,

$$\text{dist}_{T'}(e'_j, e'_h) = \sum_{i=0}^{k-1} \|e'^{(i)}, e'^{(i+1)}\| \geq \sum_{i=0}^{k-1} \|e^{(i)}, e^{(i+1)}\| = \text{dist}_T(e_j, e_h). \quad (4)$$

Altogether, Equations (2), (3) and (4) imply that

$$\zeta_{T'}(e'_j, e'_h) = \frac{\text{dist}_{T'}(e'_j, e'_h)}{\|e'_j, e'_h\|} \geq \frac{\text{dist}_T(e_j, e_h)}{(1+\epsilon) \cdot \|e_j, e_h\|} \geq \frac{\lfloor \log_D n \rfloor}{1+\epsilon}.$$

Hence $\varrho(T', e'_j) \geq \zeta_{T'}(e'_j, e'_h) \geq \frac{\lfloor \log_D n \rfloor}{1+\epsilon}$, and we are done. \blacksquare

Proposition 3.13 should be compared with Propositions 3.1 and 3.7. Specifically, as long as the dimension d is constant, one can obtain NSLoTs with root-stretch at most $(1+\epsilon)$, while for $d = \Omega(\log n)$ it is no longer possible. The next statement extends the lower bound on the tradeoff between the maximum degree D and root-stretch $\Omega(\frac{\log n}{\log D})$ established in Proposition 3.13 to any dimension $d = O(\log n)$. In particular, it shows that whenever $d = \omega(1)$ is super-constant, it is no longer possible to achieve simultaneously constant maximum degree and root-stretch.

Proposition 3.14 *For any parameter $d = 1, 2, \dots, \log n$, there exists a set \tilde{V} of n points in $\mathbb{R}^{O(d)}$, such that for any positive integer $2 \leq D \leq n-1$ and any point $v \in \tilde{V}$, every D -ary spanning tree \tilde{T} of \tilde{V} rooted at $rt = v$ has root-stretch $\Omega(\frac{d}{\log D})$.*

Proof: Set $k = 2^d$. We assume for simplicity that k divides n . Put the k vectors e_1, e_2, \dots, e_k of \mathcal{E}_k in one group G_1 and replicate this group n/k times to get $G_1, \dots, G_{n/k}$. Relocate these n/k groups (by moving all points in a group together) sufficiently far one from another. Let \tilde{T} be a D -ary spanning tree for the resulting k -dimensional point set rooted at an arbitrary point rt , and suppose without loss of generality that the root vertex rt of \tilde{T} belongs to G_1 . Since \tilde{T} is D -ary, the distance in \tilde{T} between rt and some vertex in G_1 furthest away from it, denoted v_{max} , is at least $\sqrt{2} \cdot \lfloor \log_D k \rfloor$, implying that

$$\varrho(\tilde{T}, rt) \geq \zeta_{\tilde{T}}(rt, v_{max}) = \frac{\text{dist}_{\tilde{T}}(rt, v_{max})}{\|rt, v_{max}\|} \geq \lfloor \log_D k \rfloor = \left\lfloor \frac{d}{\log D} \right\rfloor.$$

To reduce the dimension of this point set from 2^d to $O(d)$, we employ the Flattening Lemma due to Johnson and Lindenstrauss [24]. \blacksquare

Proposition 3.14 implies that the constructions of single-sink spanners and constant degree $(1+\epsilon)$ -spanners of Arya et al. [3] cannot be extended to super-constant dimensions.

Interestingly, the lower bound of $\Omega(\log_D n)$ on the root-stretch of NSLoTs for general metric spaces that is guaranteed by Proposition 3.13 ceases to hold if one is allowed to use additional Steiner points. Next, we demonstrate that Steiner points can be employed to improve the root-stretch from $\Theta(\log_D n)$ (see Propositions 3.9 and 3.13) all the way down to $(1+\epsilon')$, for arbitrarily small $\epsilon' > 0$. Moreover, this improved bound of $(1+\epsilon')$ on the root-stretch comes together with maximum degree $O(D)$. The constructions of d -dimensional Euclidean NSLoTs (see Propositions 3.1 and 3.7) can also guarantee a root-stretch of at most $(1+\epsilon')$, but the maximum degree of these NSLoTs will be $O(D + (\epsilon')^{-d+1})$ rather than $O(D)$. Finally, we remark that if one allows zero distances between vertices in the Steiner NSLoT, the root-stretch can be further improved to 1.

Proposition 3.15 *For any n -point metric space $M = (V, dist)$, an arbitrary designated point rt , a positive integer $2 \leq D \leq n - 1$ and a number $\epsilon' > 0$, there exists a D -ary Steiner rooted tree (T, rt) of M with $O(n/D)$ Steiner points, weight at most $(1 + \epsilon') \cdot W^*(M, rt)$, depth at most $\lceil \log_D n \rceil$ and root-stretch at most $(1 + \epsilon')$. Moreover, if one allows zero distances between vertices in the tree, the weight and root-stretch of T can be improved to $W^*(M, rt)$ and 1, respectively.*

Proof: Consider an n -point metric space M , and let $rt \in M$ be the designated root point. Let $D \geq 2$ be an integer, serving as the degree parameter. Suppose first that $n - 1$ is an integer power of D . We form the full D -ary tree T rooted at rt , whose $n - 1$ leaves are the $n - 1$ points of $V \setminus \{rt\}$. The remaining $\frac{n-2}{D-1} - 1$ vertices of T (excluding rt) are Steiner points. The weight assignment for edges of T is set in the following way. For each point $v \in V \setminus \{rt\}$, the weight of the edge $(v, \pi(v))$ that connects it to its parent in T is set as $dist(rt, v)$. All other edges are assigned weight 0. If one prefers to avoid using weights 0, one can simply use an arbitrarily small number $\epsilon \leq \frac{\epsilon'}{n} \cdot w_{min}$, where w_{min} is the minimum distance between a pair of points in M . It is easy to verify that the resulting tree is a D -ary Steiner NSLoT with $O(n/D)$ Steiner points, depth $\log_D(n - 1)$, root-stretch at most $(1 + \epsilon')$, and weight at most $(1 + \epsilon') \cdot W^*(M, rt)$.

This construction generalizes in the obvious way to the case where $n - 1$ is not an integer power of D , with the depth of the tree becoming $\lceil \log_D(n - 1) \rceil$. (See Figure 2.b for an illustration.) ■

4 Shallow-Low-Light-Trees

In this section we demonstrate that the NLLTs of [14] can be transformed into SLLTs. As a result of this transformation, the root-stretch becomes as small as $(1 + \epsilon)$. However, the root-degree becomes arbitrarily large. Our argument is closely related to that of Awerbuch et al. [6].

Consider an n -point metric space $M = (V, dist)$, and a spanning tree T of M rooted at an arbitrary designated root vertex rt , having weight $w(T)$ and depth $h(T)$. We construct a spanning tree $S(T)$ of M that has the same weight and depth, up to constant factors, and that also achieves root-stretch at most $(1 + \epsilon)$. (Thus, if the original tree T is low and light, the constructed tree $S(T)$ is shallow, low and light.) Moreover, the degree of any non-root vertex in $S(T)$ is greater by at most one unit than the corresponding degree in T . On the negative side, the root-degree of $S(T)$ may be much larger than that of T .

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 $L = (v_1 = rt, v_2, \dots, v_n)$ the resulting Hamiltonian path of M . Fix a parameter θ to be a positive real number. The value of θ determines the values of other parameters 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 the vertex $v_1 = rt$. The break-point B_{i+1} , $i \in [k - 1]$, is the first vertex in L after B_i such that

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

Let \mathcal{S} be the set of edges in M connecting rt with all other break-points, i.e., $\mathcal{S} = \{(rt, B_{i+1}) : i \in [k - 1]\}$. Let $\tilde{G} = (V, E(T) \cup \mathcal{S})$ be the graph obtained from T by adding to it all edges in \mathcal{S} .

Finally, we define $S(T)$ to be a shortest-paths tree (henceforth, SPT) over \tilde{G} rooted at rt . (See Figure 3 for an illustration.)

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 4.1 $\sum_{i=1}^{k-1} dist_T(B_i, B_{i+1}) \leq 2 \cdot w(T)$.




Figure 3: A spanning tree T of an 18-point metric space is depicted on the left. The vertices of T are ordered according to a pre-order traversal of T . The six breakpoints $B_1 = v_1$, $B_2 = v_3$, $B_3 = v_8$, $B_4 = v_{13}$, $B_5 = v_{15}$ and $B_6 = v_{18}$ are enclosed within dashed circles. The edges in \mathcal{S} are depicted by dashed arrows. The graph \tilde{G} is obtained from T by adding to it all edges in \mathcal{S} . The tree $S(T)$, which is an SPT over \tilde{G} rooted at $v_1 = rt$, is depicted on the right. It contains only three out of the five edges in \mathcal{S} , and its root-degree is greater by two than the root-degree of T .

The following two lemmas imply that the weight and root-stretch of the constructed tree $S(T)$ are not much greater than those of the original tree T . In fact, if we set θ to be a small constant, these two parameters do not increase by more than a small constant factor.

Lemma 4.2 $w(S(T)) \leq (1 + 2/\theta) \cdot w(T)$.

Proof: Observe that

$$w(S(T)) \leq w(\tilde{G}) = w(T) + w(\mathcal{S}) = w(T) + \sum_{i=1}^{k-1} \text{dist}(rt, B_{i+1}). \quad (5)$$

By the choice of break-points, for each index $i \in [k-1]$, $\text{dist}(rt, B_{i+1}) < \frac{1}{\theta} \cdot \text{dist}_T(B_i, B_{i+1})$. By Claim 4.1, $\sum_{i=1}^{k-1} \text{dist}_T(B_i, B_{i+1}) \leq 2 \cdot w(T)$. Therefore,

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

implying that the right-hand side of Equation (5) is at most $(1 + 2/\theta) \cdot w(T)$, and we are done. \blacksquare

Lemma 4.3 For any vertex $v \in V$, it holds that $\text{dist}_{S(T)}(rt, v) \leq (1 + 2\theta) \cdot \text{dist}(rt, v)$.

Proof: Consider an arbitrary vertex $v \in V$. First, recall that $S(T)$ is an SPT over \tilde{G} rooted at rt , and so $\text{dist}_{S(T)}(rt, v) = \text{dist}_{\tilde{G}}(rt, v)$. Clearly, the lemma holds if v is a breakpoint, as in this case we have $\text{dist}_{S(T)}(rt, v) = \text{dist}_{\tilde{G}}(rt, v) = \text{dist}(rt, v)$. We henceforth assume that v is not a breakpoint. Let i be the index in $[k-1]$ such that v is located between B_i and B_{i+1} in L . Since B_i is a break-point, it holds that $\text{dist}_{\tilde{G}}(rt, B_i) = \text{dist}(rt, B_i)$. Clearly, $\text{dist}_{\tilde{G}}(B_i, v) \leq \text{dist}_T(B_i, v)$. By the triangle inequality, $\text{dist}_{\tilde{G}}(rt, v) \leq \text{dist}_{\tilde{G}}(rt, B_i) + \text{dist}_{\tilde{G}}(B_i, v)$. Altogether,

$$\text{dist}_{S(T)}(rt, v) = \text{dist}_{\tilde{G}}(rt, v) \leq \text{dist}_{\tilde{G}}(rt, B_i) + \text{dist}_{\tilde{G}}(B_i, v) \leq \text{dist}(rt, B_i) + \text{dist}_T(B_i, v).$$

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

$$\text{dist}_T(B_i, v) \leq \theta \cdot \text{dist}(rt, v), \quad (6)$$

and so

$$\text{dist}_{S(T)}(rt, v) \leq \text{dist}(rt, B_i) + \theta \cdot \text{dist}(rt, v). \quad (7)$$

By the triangle inequality and Equation (6),

$$\begin{aligned} \text{dist}(rt, B_i) &\leq \text{dist}(rt, v) + \text{dist}(B_i, v) \leq \text{dist}(rt, v) + \text{dist}_T(B_i, v) \\ &\leq \text{dist}(rt, v) + \theta \cdot \text{dist}(rt, v) = (1 + \theta) \cdot \text{dist}(rt, v). \end{aligned} \quad (8)$$

Plugging Equation (8) in Equation (7), we obtain

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

The next lemma shows that the depth of the constructed tree is not much greater than that of the original tree.

Lemma 4.4 $h(S(T)) \leq 2 \cdot h(T) - 1$.

Proof: Observe that a path P from the root vertex rt to a vertex v in $S(T)$ may use an edge e of \mathcal{S} only as its first edge. If it does so, then the sub-path $P \setminus \{e\}$ of P forms a path in the original tree T not traversing the root vertex rt , implying that the number of edges in P is bounded from above by $1 + 2(h(T) - 1)$. Otherwise, the entire path P forms a path in the original tree T starting at rt , and so it consists of at most $h(T)$ edges. \blacksquare

Observe that all edges in $\mathcal{S} = \{(rt, B_{i+1}) : i \in [k - 1]\}$ connecting rt with the $k - 1$ breakpoints B_2, B_3, \dots, B_k may be present in $S(T)$, and so it is possible to have $\text{deg}(S(T), rt) = \Omega(k)$. As the number k of breakpoints may in general be arbitrarily large, the root-degree in $S(T)$ may be unbounded. On the other hand, consider a non-root vertex v in $S(T)$. The set of its neighbors in \tilde{G} is comprised of all its neighbors in T , and possibly the root vertex rt . It follows that $\text{deg}(S(T), v) \leq \text{deg}(\tilde{G}, v) \leq \text{deg}(T, v) + 1$, implying that

$$\lambda(S(T)) \leq \lambda(T) + 1. \quad (9)$$

Finally, we argue that the tree $S(T)$ can be constructed efficiently.

Lemma 4.5 *Given the original tree T , the tree $S(T)$ can be constructed in linear time.*

Proof: Clearly, the Hamiltonian path L can be constructed in $O(n)$ time. It is easy to see that within $O(n)$ time one can identify the set $\mathcal{B} = \{B_1, B_2, \dots, B_k\}$ of break-points, implying that the graph \tilde{G} can be constructed in $O(n)$ time. The final step of the algorithm is the construction of an SPT over \tilde{G} . It is easy to see that \tilde{G} is planar. Klein et al. [27] showed that an SPT can be computed within linear time in planar graphs, and so the final step of the algorithm can be carried out in linear time. The lemma follows. \blacksquare

Set $\epsilon = 2\theta$. Lemmas 4.2, 4.3, 4.4 and 4.5 imply the following corollary.

Corollary 4.6 *For any n -point metric space M , a designated point $rt \in V$, a number $\epsilon > 0$ and a rooted spanning tree (T, rt) , there exists a rooted spanning tree $(S(T), rt)$ of M with weight $O(w(T) \cdot (\epsilon^{-1}))$, depth at most $2h(T) - 1$ and root-stretch at most $(1 + \epsilon)$. Moreover, the maximum degree $\lambda(S(T))$ of a non-root vertex in $S(T)$ is at most $\lambda(T) + 1$. In addition, if T is given as a part of the input, then $S(T)$ can be constructed in linear time.*

Consider the NLLTs construction of Dinitz et al. [14] for general metric spaces. For a metric space M , a designated point $rt \in M$ and an integer $\ell = O(\log n)$, it provides a rooted NLLT (T_1, rt) with weight $w(T_1) = O(\ell) \cdot w(MST(M))$, depth $h(T_1) = O(\ell \cdot n^{1/\ell})$ and constant maximum degree. In addition, in the complementary range of depth $h = O(\log n)$ this construction provides a rooted NLLT (T_2, rt) with weight $w(T_2) = O(\ell \cdot n^{1/\ell}) \cdot w(MST(M))$, depth $h(T_2) = O(\ell)$ and maximum degree $O(n^{1/\ell})$. The NLLTs construction of [14] can be implemented within $O(n^2)$ time in general metric spaces, and within linear time in Euclidean low-dimensional ones. (The linear time implementation of the NLLTs in [14] is based on a recent work of Chan [9].)

By employing Corollary 4.6 in conjunction with the NLLTs construction of [14] in the range $h = \Omega(\log n)$, we obtain a rooted SLLT $(S(T_1), rt)$ for M with weight $w(S(T_1)) = O(\ell \cdot (\epsilon^{-1})) \cdot w(MST(M))$, depth $h(S(T_1)) = O(\ell \cdot n^{1/\ell})$ and root-stretch at most $(1 + \epsilon)$. Also, all vertices of $S(T_1)$ except its root have optimal degree $O(1)$, and the root vertex rt may have arbitrarily large degree. Similarly, by employing Corollary 4.6 in conjunction with the NLLTs construction of [14] in the complementary range of depth $h = O(\log n)$, we obtain a rooted SLLT $(S(T_2), rt)$ for M with weight $w(S(T_2)) = O(\ell \cdot n^{1/\ell} \cdot (\epsilon^{-1})) \cdot w(MST(M))$, depth $h(S(T_2)) = O(\ell)$ and root-stretch at most $(1 + \epsilon)$. Also, all vertices of $S(T_2)$ except its root have optimal degree $O(n^{1/\ell})$, whereas the root vertex rt may have arbitrarily large degree. In addition, given the trees T_1 and T_2 , both trees $S(T_1)$ and $S(T_2)$ can be constructed in linear time. Together with the time needed to construct the NLLTs T_1 and T_2 , we get that the SLLTs $S(T_1)$ and $S(T_2)$ can be constructed within $O(n^2)$ time in general metric spaces, and within linear time in Euclidean low-dimensional ones. Summarizing, we have proved the following result.

Theorem 4.7 (SLLTs for general metric spaces) *For any n -point metric space M , a designated point $rt \in M$, an integer ℓ and a number $\epsilon > 0$, there exists a spanning rooted tree (T'_1, rt) for the metric space M with lightness $O(\ell \cdot (\epsilon^{-1}))$, depth $O(\ell \cdot n^{1/\ell})$ and root-stretch at most $(1 + \epsilon)$. Moreover, the maximum degree of a non-root vertex in T'_1 is $O(1)$. There also exists a rooted tree (T'_2, rt) for M with lightness $O(\ell \cdot n^{1/\ell} \cdot (\epsilon^{-1}))$, depth $O(\ell)$ and root-stretch at most $(1 + \epsilon)$. Moreover, the maximum degree of a non-root vertex in T'_2 is $O(n^{1/\ell})$. Both trees can be constructed within $O(n^2)$ time in general metric spaces, and within linear time in Euclidean low-dimensional ones.*

5 Narrow-Shallow-Low-Light Trees

In this section we present a general technique for constructing NSLLTs out of SLLTs and NSLoTs. Then we employ this technique in conjunction with the constructions of SLLTs and NSLoTs from Sections 4 and 3, respectively, to obtain our constructions of NSLLTs. The latter constructions exhibit optimal tradeoffs between all the four involved parameters.

Consider an n -point metric space M , and let T' be an SLLT for M rooted at some designated point $rt \in M$. Next, we argue that by using NSLoTs one can significantly reduce the degree of rt , while only slightly increasing other parameters of T' . Let Z be the star subtree of T' rooted at rt . In other words, the vertex set of Z is $V(Z) = \{rt, c_1, c_2, \dots, c_q\}$, where c_1, c_2, \dots, c_q are the children of rt in T' . Also, the weights of edges (rt, c_i) agree in T' and Z , for all indices $i \in [q]$. Let \tilde{T} be some NSLoT rooted at rt for the metric space M_Z induced by the points in $V(Z)$. (Observe that the weight $w(Z)$ of Z is equal to the star-weight $W^*(M_Z, rt)$ of M_Z with respect to rt , i.e., $w(Z) = W^*(M_Z, rt) \leq w(T')$.) Finally, let \hat{T} be the spanning tree of M obtained from the SLLT T' by replacing the star Z with the NSLoT \tilde{T} . (See Figure 1 for an illustration.)

The properties of the resulting tree are summarized in the following statement.

Proposition 5.1 (1) $w(\hat{T}) = w(T') - w(Z) + w(\tilde{T})$, (2) $h(\hat{T}) \leq h(T') - 1 + h(\tilde{T})$, (3) $\lambda(\hat{T}) \leq \lambda(T') - 1 + \lambda(\tilde{T})$, (4) $deg(\hat{T}, rt) = deg(\tilde{T}, rt)$, (5) $\varrho(\hat{T}, rt) \leq \varrho(\tilde{T}, rt) \cdot \varrho(T', rt)$.

Remark: This statement remains valid if \tilde{T} is a Steiner NSLoT for M_Z , with \hat{T} becoming a Steiner tree rather than a spanning tree for M .

Proof: All assertions of the proposition but the last are straightforward. We henceforth prove only the last assertion. Define $s' = \varrho(T', rt)$, and $\tilde{s} = \varrho(\tilde{T}, rt)$. Let v be a point in V satisfying $\zeta_{\hat{T}}(rt, v) = \varrho(\hat{T}, rt)$. Next, we show that $\varrho(\hat{T}, rt) = \zeta_{\hat{T}}(rt, v) \leq \tilde{s} \cdot s'$.

For each $k \in [q]$, we denote by T'_k (respectively, \hat{T}_k) the subtree of T' (resp., \hat{T}) rooted at c_k . The assertion is trivial if $v = rt$. We henceforth assume that $v \neq rt$. Let j be the index in $[q]$ such that v is in T'_j . It follows that

$$\begin{aligned} \text{dist}_{\hat{T}}(rt, v) &= \text{dist}_{\hat{T}}(rt, c_j) + \text{dist}_{\hat{T}}(c_j, v) \leq \tilde{s} \cdot \text{dist}_{T'}(rt, c_j) + \text{dist}_{T'}(c_j, v) \\ &\leq \tilde{s} \cdot (\text{dist}_{T'}(rt, c_j) + \text{dist}_{T'}(c_j, v)) = \tilde{s} \cdot \text{dist}_{T'}(rt, v) \leq \tilde{s} \cdot s' \cdot \text{dist}(rt, v). \end{aligned}$$

Hence,

$$\varrho(\hat{T}, rt) = \zeta_{\hat{T}}(rt, v) = \frac{\text{dist}_{\hat{T}}(rt, v)}{\text{dist}(rt, v)} \leq \tilde{s} \cdot s'. \quad \blacksquare$$

Consider the SLLTs construction of Theorem 4.7 for general metric spaces. For a metric space M , a designated point $rt \in M$, an integer $\ell = O(\log n)$ and a number $\epsilon > 0$, it provides a rooted SLLT $(T'_1, rt) = (T'_1(\ell), rt)$ with weight $w(T'_1) = O(\ell \cdot (\epsilon^{-1}) \cdot w(MST(M)))$, depth $h(T'_1) = O(\ell \cdot n^{1/\ell})$ and root-stretch $\varrho(T'_1, rt) \leq (1 + \epsilon)$. Moreover, all vertices of T'_1 except its root vertex rt have optimal degree $O(1)$. The degree of the root may, however, be arbitrarily large. In addition, in the complementary range of depth $h = O(\log n)$ this construction provides a rooted SLLT $(T'_2, rt) = (T'_2(\ell), rt)$ with weight $w(T'_2) = O(\ell \cdot n^{1/\ell} \cdot (\epsilon^{-1}) \cdot w(MST(M)))$, depth $h(T'_2) = O(\ell)$ and root-stretch $\varrho(T'_2, rt) \leq (1 + \epsilon)$. Similarly to T'_1 , all vertices of T'_2 but the root vertex rt have optimal degree $O(n^{1/\ell})$, and the root may have arbitrarily large degree.

Next, we reduce the root-degree in both trees T'_1 and T'_2 . For $i = 1, 2$, let Z_i be the star subtree of T'_i rooted at rt , and let \tilde{T}_i be some NSLoT rooted at rt for the $(q_i + 1)$ -point metric space M_{Z_i} , where q_i stands for the number of children of rt in T'_i . To construct an NSLLT \hat{T}_i out of T'_i and \tilde{T}_i , we replace the star Z_i with \tilde{T}_i .

Specifically, if M is a set of n points in the plane, then our construction of NSLoTs (Proposition 3.1) in the particular case $D = O(1)$ provides a rooted NSLoT (\tilde{T}_1, rt) for M_{Z_1} with weight $w(\tilde{T}_1) = O(W^*(M_{Z_1}, rt))$, depth $h(\tilde{T}_1) = O(\log(q_1 + 1)) = O(\log n)$, maximum degree $\Delta(\tilde{T}_1) = O(\epsilon^{-1})$ and root-stretch $\varrho(\tilde{T}_1, rt) \leq (1 + \epsilon)$. By Proposition 5.1, replacing the star Z_1 of $T'_1 = T'_1(\ell)$ with \tilde{T}_1 produces a rooted NSLLT (\hat{T}_1, rt) for M with weight that satisfies

$$\begin{aligned} w(\hat{T}_1) &= w(T'_1) - w(Z_1) + w(\tilde{T}_1) = w(T'_1) + O(W^*(M_{Z_1}, rt)) \\ &= O(w(T'_1)) = O(\ell \cdot (\epsilon^{-1}) \cdot w(MST(M))), \end{aligned}$$

depth that satisfies

$$h(\hat{T}_1) \leq h(T'_1) - 1 + h(\tilde{T}_1) = O(\ell \cdot n^{1/\ell}) + O(\log n) = O(\ell \cdot n^{1/\ell}),$$

maximum degree that satisfies

$$\Delta(\hat{T}_1) = \max \left\{ \lambda(\hat{T}_1), \text{deg}(\hat{T}_1, rt) \right\} \leq \max \left\{ \lambda(T'_1) - 1 + \lambda(\tilde{T}_1), \text{deg}(\tilde{T}_1, rt) \right\} = O(\epsilon^{-1}),$$

and root-stretch $\varrho(\hat{T}_1, rt) \leq (1 + \epsilon)^2 = 1 + O(\epsilon)$. In addition, Proposition 3.1 in the particular case $D = (q_2 + 1)^{1/\ell}$ provides a rooted NSLoT (\tilde{T}_2, rt) for M_{Z_2} with weight $w(\tilde{T}_2) = O(W^*(M_{Z_2}, rt))$, depth $h(\tilde{T}_2) = O(\ell)$, maximum degree $\Delta(\tilde{T}_2) = O((q_2 + 1)^{1/\ell} + (\epsilon^{-1})) = O(n^{1/\ell} + (\epsilon^{-1}))$ and root-stretch

$\varrho(\tilde{T}_2, rt) \leq (1 + \epsilon)$. By Proposition 5.1, replacing the star Z_2 of $T'_2 = T'_2(\ell)$ with \tilde{T}_2 produces a rooted NSLLT (\hat{T}_2, rt) for M with weight that satisfies

$$\begin{aligned} w(\hat{T}_2) &= w(T'_2) - w(Z_2) + w(\tilde{T}_2) = w(T'_2) + O(W^*(M_{Z_2}, rt)) \\ &= O(w(T'_2)) = O(\ell \cdot n^{1/\ell} \cdot (\epsilon^{-1})) \cdot w(MST(M)), \end{aligned}$$

depth that satisfies

$$h(\hat{T}_2) \leq h(T'_2) - 1 + h(\tilde{T}_2) = O(\ell),$$

maximum degree that satisfies

$$\Delta(\hat{T}_2) = \max \left\{ \lambda(\hat{T}_2), \deg(\hat{T}_2, rt) \right\} \leq \max \left\{ \lambda(T'_2) - 1 + \lambda(\tilde{T}_2), \deg(\tilde{T}_2, rt) \right\} = O(n^{1/\ell} + (\epsilon^{-1})),$$

and root-stretch $\varrho(\hat{T}_2, rt) \leq (1 + \epsilon)^2 = 1 + O(\epsilon)$. The SLLTs T'_1 and T'_2 provided by Theorem 4.7 can be constructed within $O(n)$ time in Euclidean spaces. By Proposition 3.1, the NSLoT \tilde{T}_1 (respectively, \tilde{T}_2) can be constructed in $O(n \cdot \log n)$ (resp., $O(n \cdot \ell)$) time. Hence the overall time required to construct the NSLLT \hat{T}_1 (respectively, \hat{T}_2) is $O(n \cdot \log n)$ (resp., $O(n \cdot \ell)$). This argument generalizes in a straightforward way to point sets in \mathbb{R}^d , for any constant $d \geq 2$. Summarizing, we have proved the following result.

Theorem 5.2 (NSLLTs for constant dimensional Euclidean spaces) *Let $d \geq 2$ be an integer constant. For a set M of n points in \mathbb{R}^d , a designated point $rt \in M$, an integer $\ell = O(\log n)$ and a number $\epsilon > 0$, there exists a rooted spanning tree (\hat{T}_1, rt) with lightness $O(\ell \cdot (\epsilon^{-1}))$, depth $O(\ell \cdot n^{1/\ell})$, maximum degree $O(\epsilon^{-d+1})$ and root-stretch at most $(1 + \epsilon)$. In addition, there exists a rooted spanning tree (\hat{T}_2, rt) with lightness $O(\ell \cdot n^{1/\ell} \cdot (\epsilon^{-1}))$, depth $O(\ell)$, maximum degree $O(n^{1/\ell} + (\epsilon^{-d+1}))$ and root-stretch at most $(1 + \epsilon)$. For $d = 2$, the tree \hat{T}_1 (respectively, \hat{T}_2) can be constructed in time $O(n \cdot \log n)$ (resp., $O(n \cdot \ell)$). For $d \geq 3$, the tree \hat{T}_1 (respectively, \hat{T}_2) can be constructed in time $O(n \cdot \log n \cdot (\log \epsilon^{-1}) + (\epsilon^{-d+1}))$ (resp., $O(n \cdot \ell \cdot (\log \epsilon^{-1}) + (\epsilon^{-d+1}))$).*

Similarly, by employing our construction of NSLoTs for general metric spaces (Proposition 3.9) in the particular cases of $D = O(1)$ and $D = (q_2 + 1)^{1/\ell}$ in conjunction with the SLLTs construction of Theorem 4.7 (substituting $\epsilon = \frac{1}{2}$, or some other fixed constant), we get two rooted NSLLTs for general metric spaces (\hat{T}_1, rt) and (\hat{T}_2, rt) , respectively, whose properties are summarized in the following theorem.

Theorem 5.3 (NSLLTs for general metric spaces) *For a general n -point metric space M , a designated point $rt \in M$ and an integer $\ell = O(\log n)$, there exists a rooted spanning tree (\hat{T}_1, rt) with lightness $O(\ell)$, depth $O(\ell \cdot n^{1/\ell})$, constant maximum degree and root-stretch $O(\log n)$. In addition, there exists a rooted spanning tree (\hat{T}_2, rt) with lightness $O(\ell \cdot n^{1/\ell})$, depth $O(\ell)$, maximum degree $O(n^{1/\ell})$ and root-stretch $O(\ell)$. Both trees can be constructed in $O(n^2)$ time.*

Finally, by employing our construction of Steiner NSLoTs for general metric spaces (Proposition 3.15) in the particular cases of $D = O(1)$ and $D = (q_2 + 1)^{1/\ell}$ in conjunction with the SLLTs construction of Theorem 4.7, we get two Steiner rooted NSLLTs for general metric spaces (\hat{T}_1, rt) and (\hat{T}_2, rt) , respectively, whose properties are summarized in the following theorem.

Theorem 5.4 (Steiner NSLLTs for general metric spaces) *For a general n -point metric space M , a designated point $rt \in M$, an integer $\ell = O(\log n)$ and a number $\epsilon > 0$, there exists a Steiner rooted tree (\hat{T}_1, rt) with lightness $O(\ell \cdot (\epsilon^{-1}))$, depth $O(\ell \cdot n^{1/\ell})$, constant maximum degree and root-stretch at most $(1 + \epsilon)$. In addition, there exists a Steiner rooted tree (\hat{T}_2, rt) with lightness $O(\ell \cdot n^{1/\ell} \cdot (\epsilon^{-1}))$, depth $O(\ell)$, maximum degree $O(n^{1/\ell})$ and root-stretch at most $(1 + \epsilon)$. Both trees can be constructed in $O(n^2)$ time.*

Remark: Assuming that ϵ is constant, all constructions of NSLLTs presented in this section (Theorems 5.2, 5.3 and 5.4) exhibit optimal (up to constant factors) tradeoffs between all the four involved parameters. In other words, any improvement in one of the parameters must come at the expense of the others. Indeed, all constructions achieve the same tradeoffs between the depth and maximum degree and between the lightness and depth. The optimality of the former tradeoff is obvious. The optimality of the latter tradeoff is proved in [14] for spanning trees, and in Section 6 we extend it to Steiner trees. The tradeoff between the maximum degree and root-stretch given in Theorem 5.3 is optimal by Proposition 3.13. The bound of $(1 + \epsilon)$ on the root-stretch given in Theorems 5.2 and 5.4, for arbitrarily small constant $\epsilon > 0$, is optimal as well. The general case when ϵ is non-constant is more involved. That is, if one is interested in smaller root-stretch $\varrho \leq 1 + \epsilon, \epsilon = o(1)$, then it is not clear to the authors whether the dependences of the lightness and maximum degree on ϵ can be improved. As for the dependence of the lightness on ϵ , there is a lower bound of Khuller et al. [26] that shows that there are n -vertex graphs for which any spanning tree T with root-stretch $\varrho(T) \leq 1 + \epsilon$ has lightness at least $\Psi(T) = \Omega(\epsilon^{-1})$. However, this lower bound is not known to apply neither to Steiner trees nor to metric spaces. Moreover, for an SLLT T with root-stretch at most $1 + \epsilon$ and depth $O(\ell)$ (respectively, $O(\ell \cdot n^{1/\ell})$) one can hope to achieve lightness $\Psi(T) = O(\ell \cdot n^{1/\ell}) + f(\epsilon^{-1})$ (resp., $\Psi(T) = O(\ell) + f(\epsilon^{-1})$), for an appropriate function f . Observe that in our construction, the lightness is the product $O(\ell \cdot n^{1/\ell}) \cdot (\epsilon^{-1})$ (or $O(\ell) \cdot (\epsilon^{-1})$), rather than the sum. On the other hand, the dependence of the maximum degree on ϵ (for d -dimensional Euclidean spaces), i.e., $\Delta(T) = O(\epsilon^{-d+1})$, appears to be the correct one. Though a proof for this claim is currently out of our reach, we note that the same dependence is present in the previous construction of single-sink spanners of Arya et al. [3] (see the second row of Table 1).

6 Lower Bounds for Euclidean Steiner Spanners

In this section we extend the lower bounds of Dinitz et al. [14] to Steiner trees. Dinitz et al. [14] considered the 1-dimensional Euclidean space ϑ_n with n points $1, 2, \dots, n$ on the x -axis, and showed that any spanning tree of ϑ_n with depth $o(\log n)$ has weight $\omega(n \cdot \log n)$, and vice versa, i.e., any spanning tree of ϑ_n with weight $o(n \cdot \log n)$ has depth $\omega(\log n)$. This result implies that no construction of Euclidean spanners may guarantee hop-diameter $O(\log n)$ and lightness $o(\log n)$, or vice versa. However, the lower bound of [14] does not preclude the existence of *Steiner* spanners with hop-diameter $O(\log n)$ and lightness $o(\log n)$, or vice versa. In this section we prove that Steiner points do not help in this context, and thus the construction of Arya et al. [3] cannot be improved even if one allows the spanner to use (arbitrarily many) Steiner points.

The following proposition is central in our proof.

Proposition 6.1 *For any weighted rooted tree (T', rt') and a subset $R' = R(T')$ of $V' = V(T')$, there exists a weighted rooted tree (T, rt) over R' with depth no greater than that of T' (i.e., $h(T) \leq h(T')$), weight at most twice the weight of T' (i.e., $w(T) \leq 2 \cdot w(T')$), and which also dominates T' in the following sense: for every pair of vertices u, v in R' , $dist_T(u, v) \geq dist_{T'}(u, v)$. Moreover, T can be constructed in $O(n)$ time.*

Proof: We refer to the vertices of $R' = R(T')$ as the *required* vertices of T' . The root of a tree τ is denoted by $rt(\tau)$. (For an empty tree τ , we write $rt(\tau) = NULL$.) For an inner vertex v in T' , we denote its children by $c_1(v), c_2(v), \dots, c_{ch(v)}(v)$, where $ch(v)$ denotes the number of its children in T' . Denote by T'_v the subtree rooted at v , and denote by $R'_v = R(T'_v)$ the set of required vertices in T'_v . For technical convenience we assume that all edge weights in T' are distinct, as any ties can be broken using lexicographic rules. Also, we suppose without loss of generality that for each vertex v in T' , such that $R'_v \setminus \{v\} \neq \emptyset$, the vertex closest to v in T' (i.e., the one whose weighted distance from v in T' is

minimal) among all vertices in $R'_v \setminus \{v\}$ belongs to the subtree $T'_{c_1(v)}$ rooted at its first child $c_1(v)$. (This requirement can be guaranteed by a straightforward procedure that runs in linear time.) For convenience, we write T'_i and R'_i as shortcuts for $T'_{c_i(rt')}$ and $R'_{c_i(rt')}$, respectively, for all $i \in [ch(rt')]$. Also, we refer to the tree T' (respectively, T) as the *original* (resp., *resulting*) tree.

Denote by $I' = I(T')$ the set of all indices i , such that $i \in [ch(rt')]$ and the set R'_i is non-empty. Note that if $|I'| \geq 1$, then $1 \in I'$.

Next, we present a linear time procedure *Prune* that accepts as input a non-empty weighted rooted tree (T', rt') over $V' = V(T')$, and transforms T' into a weighted rooted tree (T, rt) over the set R' of required vertices which satisfies the conditions of the proposition.

If T' consists of just the single vertex rt' , then the procedure either leaves T' intact if $rt' \in R'$, or it transforms T' into an empty tree if $rt' \notin R'$. Otherwise, $|I'| \geq 2$. In this case for each $i \in [ch(rt')]$, the tree T'_i is recursively transformed into a tree T_i which spans the set R'_i of required vertices in T'_i . Also, all edges connecting the root vertex rt' of T' with its children in T' are removed. The execution of the procedure then splits into two cases.

Case 1: $rt' \in R'$. The root vertex rt' of T' remains the root vertex of T , and for each index i in I' , an edge of weight $2 \cdot dist_{T'}(rt', rt(T_i))$ connecting the root vertex rt' of T with the vertex $rt(T_i)$ is added.

Case 2: $rt' \notin R'$. Recall that if T'_1 contains no required vertices, then T' contains no required vertices as well, and so we set $T = NULL$. Otherwise, T_1 is non-empty. In this case we set the vertex $rt(T_1)$ to be the root of T , and for each index i in $I' \setminus \{1\}$, an edge of weight $2 \cdot dist_{T'}(rt', rt(T_i))$ connecting the root vertex $rt(T_1)$ of T with the vertex $rt(T_i)$ is added. Also, the original root vertex rt' of T' is removed. (See Figure 4 for an illustration.)

Remark: For each index i in I' , the tree T_i is non-empty, and its root vertex $rt(T_i)$ is a (required) vertex in T' . Hence the distance $dist_{T'}(rt', rt(T_i))$ for such indices is well-defined.

First, it is easy to verify that the Procedure *Prune* can be implemented in linear time.

Second, observe that the resulting tree T spans the vertex set R' . Moreover, for each index i in I' , the tree T_i spans the vertex set R'_i , and this property holds recursively.

The following lemma shows that the depth $h(T)$ of the resulting tree T is no greater than the depth $h(T')$ of the original tree T' . (We define the depth of an empty tree to be zero.)

Lemma 6.2 $h(T) \leq h(T')$.

Proof: By induction on depth. ■

Next, we show that the weight of T is at most twice greater than that of T' .

Lemma 6.3 *If T is non-empty, then $w(T) \leq 2 \cdot (w(T') - dist_{T'}(rt', rt(T)))$.*

Proof: The proof is by induction on the depth $h = h(T')$ of the original tree T' . The basis $h = 0$ is trivial.

Induction Step: We assume the correctness of the claim for all smaller values of h and prove it for h . Observe that for every index i in I' , the vertex $rt(T_i)$ belongs to the subtree T'_i of T' . Hence

$$dist_{T'}(rt', rt(T_i)) = dist_{T'}(rt', rt(T'_i)) + dist_{T'}(rt(T'_i), rt(T_i)), \quad (10)$$

and by the induction hypothesis,

$$w(T_i) \leq 2 \cdot \left(w(T'_i) - dist_{T'_i}(rt(T'_i), rt(T_i)) \right) = 2 \cdot \left(w(T'_i) - dist_{T'}(rt(T'_i), rt(T_i)) \right). \quad (11)$$

(See Figure 5 for an illustration.) The analysis splits into two cases.




Figure 4: The procedure Prune transforms a rooted spanning tree (T', rt') over V' into a rooted spanning tree (T, rt) over the set R' of required vertices of T' . The tree T depicted at the bottom (respectively, right) part of the figure corresponds to the case $rt' \in R'$ (resp., $rt' \notin R'$). In both cases, for each index $i \in [3]$, the rooted tree $(T'_i, rt(T'_i))$ is recursively transformed into a rooted spanning tree $(T_i, rt(T_i))$ over R'_i , and the edge $(rt', rt(T'_i))$ is removed. Also, in the former (respectively, latter) case the root vertex $rt(T)$ of T is set as rt' (resp., $rt(T_1)$), and for each $i \in [3]$ (resp., $i \in [2, 3]$), the edge $(rt(T), rt(T_i))$ of weight $2 \cdot dist_{T'}(rt', rt(T_i))$ is added.




Figure 5: For every index $i \in [3]$, the vertex $rt(T_i)$ belongs to the subtree T'_i of T' . Thus the path connecting rt' with $rt(T_i)$ in T' passes through the vertex $rt(T'_i)$. Moreover, the required vertex closest to rt' in T' (excluding rt' itself if $rt' \in R'$) is $rt(T_1)$.

Case 1: $rt' \in R'$. In this case $rt(T) = rt'$, and so $dist_{T'}(rt', rt(T)) = 0$. By construction,

$$w(T) = \sum_{i \in I'} 2 \cdot dist_{T'}(rt', rt(T_i)) + \sum_{i \in I'} w(T_i).$$

By Equations (10) and (11),

$$\begin{aligned} w(T) &= \sum_{i \in I'} 2 \cdot dist_{T'}(rt', rt(T'_i)) + \sum_{i \in I'} 2 \cdot dist_{T'}(rt(T'_i), rt(T_i)) + \sum_{i \in I'} w(T_i) \\ &\leq 2 \cdot \sum_{i \in I'} dist_{T'}(rt', rt(T'_i)) + 2 \cdot \sum_{i \in I'} w(T'_i) \\ &\leq 2 \cdot \sum_{i=1}^{ch(rt')} dist_{T'}(rt', rt(T'_i)) + 2 \cdot \sum_{i=1}^{ch(rt')} w(T'_i) \\ &= 2 \cdot w(T') = 2 \cdot (w(T') - dist_{T'}(rt', rt(T))). \end{aligned}$$

Case 2: $rt' \notin R'$. In this case $rt(T) = rt(T_1)$. Since T is non-empty, $rt(T) = rt(T_1) \neq NULL$, implying that T_1 is non-empty as well, and so $1 \in I'$. By construction,

$$w(T) = \sum_{i \in I' \setminus \{1\}} 2 \cdot dist_{T'}(rt', rt(T_i)) + \sum_{i \in I'} w(T_i).$$

By Equations (10) and (11),

$$\begin{aligned}
w(T) &= \sum_{i \in I' \setminus \{1\}} 2 \cdot \text{dist}_{T'}(rt', rt(T'_i)) + \sum_{i \in I' \setminus \{1\}} 2 \cdot \text{dist}_{T'}(rt(T'_i), rt(T_i)) + \sum_{i \in I'} w(T_i) \\
&= 2 \cdot \left(\sum_{i \in I'} \text{dist}_{T'}(rt', rt(T'_i)) - \text{dist}_{T'}(rt', rt(T_1)) \right) \\
&\quad + 2 \cdot \left(\sum_{i \in I'} \text{dist}_{T'}(rt(T'_i), rt(T_i)) - \text{dist}_{T'}(rt(T_1), rt(T_1)) \right) + \sum_{i \in I'} w(T_i) \\
&\leq 2 \cdot \sum_{i \in I'} \text{dist}_{T'}(rt', rt(T'_i)) + 2 \cdot \sum_{i \in I'} w(T'_i) - 2 \cdot \text{dist}_{T'}(rt', rt(T_1)) \\
&\leq 2 \cdot \sum_{i=1}^{ch(rt')} \text{dist}_{T'}(rt', rt(T'_i)) + 2 \cdot \sum_{i=1}^{ch(rt')} w(T'_i) - 2 \cdot \text{dist}_{T'}(rt', rt(T_1)) \\
&= 2 \cdot (w(T') - \text{dist}_{T'}(rt', rt(T_1))) = 2 \cdot (w(T') - \text{dist}_{T'}(rt', rt(T))). \quad \blacksquare
\end{aligned}$$

The factor 2 in the inequality $w(T) \leq 2 \cdot w(T')$ cannot be improved. Indeed, let T' be the n -vertex star with all edges having unit weight. Also, let the set $R' = R(T')$ of the required vertices contain all the $n - 1$ leaves. Observe that $w(T') = n - 1$, and for any two vertices u and v in R' , $\text{dist}_{T'}(u, v) = 2$. On the other hand, for any spanning tree T of R' that dominates T' , the weight of any edge is at least 2. Hence, $w(T) \geq 2(n - 2) \geq (2 - o(1)) \cdot w(T')$.

The next claim shows that the required vertex closest to rt' in the original tree T' (i.e., the one whose weighted distance from rt' in T' is minimal) is the root vertex $rt(T)$ of the resulting tree T . We use this property to prove Lemma 6.5.

Claim 6.4 *For a vertex v in T' such that $|R'_v| \geq 1$, denote by $\delta(v)$ the vertex closest to v in T' among all vertices in R'_v . Then $\delta(rt') = rt(T)$.*

Proof: The proof is by induction on the depth $h = h(T')$ of the original tree T' .

Basis: $h = 0$. In this case T' consists of just the single vertex rt' , and $rt' \in R'$. By construction, $rt(T) = rt'$, and so $\delta(rt') = rt' = rt(T)$.

Induction Step: We assume the correctness of the statement for all smaller values of h and prove it for h . If $rt' \in R'$, then $rt(T) = rt'$, and clearly, $\delta(rt') = rt' = rt(T)$.

Otherwise $rt(T) = rt(T_1)$. Recall that the vertex $\delta(rt')$ closest to rt' in T' among all required vertices belongs to T'_1 , and so $\delta(rt') = \delta(rt(T'_1))$. (See Figure 5 for an illustration.) By the induction hypothesis, $\delta(rt(T'_1)) = rt(T_1)$. Altogether, we have

$$\delta(rt') = \delta(rt(T'_1)) = rt(T_1) = rt(T). \quad \blacksquare$$

The following lemma shows that the resulting tree T dominates the original tree T' .

Lemma 6.5 *For every pair of required vertices u, v in R' , $\text{dist}_T(u, v) \geq \text{dist}_{T'}(u, v)$.*

Proof: The proof is by induction on the depth $h = h(T')$ of the original tree T' . The basis $h = 0$ holds vacuously.

Induction Step: We assume the correctness of the statement for all smaller values of h and prove it for h . If $|R'| \leq 1$, then the lemma holds vacuously. Otherwise, consider an arbitrary pair u, v of vertices in R' . If u and v belong to the same subtree T'_i of the original tree T' , $i \in I'$, then the lemma follows from

the induction hypothesis. Otherwise, the analysis splits into two cases.

Case 1: $rt' \in R'$. Then $rt(T) = rt'$. Let w be a required vertex in T'_k , for some index $k \in I'$. Observe that w ends up in the subtree T_k of the resulting tree T , and $rt(T_k)$ belongs to the subtree T'_k of the original tree T' . Consequently,

$$dist_T(rt', w) = dist_T(rt', rt(T_k)) + dist_T(rt(T_k), w),$$

and by the induction hypothesis,

$$dist_T(rt(T_k), w) = dist_{T_k}(rt(T_k), w) \geq dist_{T'_k}(rt(T_k), w) = dist_{T'}(rt(T_k), w).$$

By the triangle inequality,

$$dist_{T'}(rt', w) \leq dist_{T'}(rt', rt(T_k)) + dist_{T'}(rt(T_k), w).$$

By construction, $dist_T(rt', rt(T_k)) = 2 \cdot dist_{T'}(rt', rt(T_k))$. Altogether,

$$\begin{aligned} dist_T(rt', w) &= dist_T(rt', rt(T_k)) + dist_T(rt(T_k), w) \\ &\geq 2 \cdot dist_{T'}(rt', rt(T_k)) + dist_{T'}(rt(T_k), w) \geq dist_{T'}(rt', w). \end{aligned} \quad (12)$$

If one among u and v is the root vertex rt' , then the lemma follows from Equation (12). Otherwise, there exist indices i and j in I' , such that u is in T'_i and v is in T'_j . Observe that

$$dist_{T'}(u, v) = dist_{T'}(u, rt') + dist_{T'}(rt', v). \quad (13)$$

Also, u and v end up in different subtrees of the resulting tree T , specifically, $u \in T_i$ and $v \in T_j$. Thus,

$$dist_T(u, v) = dist_T(u, rt') + dist_T(rt', v). \quad (14)$$

Equations (12), (13) and (14) yield

$$dist_T(u, v) = dist_T(u, rt') + dist_T(rt', v) \geq dist_{T'}(u, rt') + dist_{T'}(rt', v) = dist_{T'}(u, v).$$

Case 2: $rt' \notin R'$. Then $rt(T) = rt(T_1)$, and there exist indices i and j in I' , such that u is in T'_i and v is in T'_j . Observe that u (respectively, v) belongs to the subtree T_i (resp., T_j) of the resulting tree T , and $rt(T_i)$ (resp., $rt(T_j)$) belongs to the subtree T'_i (resp., T'_j) of the original tree T' . Hence, by the induction hypothesis,

$$dist_T(u, rt(T_i)) = dist_{T_i}(u, rt(T_i)) \geq dist_{T'_i}(u, rt(T_i)) = dist_{T'}(u, rt(T_i)) \quad (15)$$

and

$$dist_T(rt(T_j), v) = dist_{T_j}(rt(T_j), v) \geq dist_{T'_j}(rt(T_j), v) = dist_{T'}(rt(T_j), v). \quad (16)$$

Next, we argue that

$$dist_T(rt(T_i), rt(T_1)) + dist_T(rt(T_1), rt(T_j)) \geq dist_{T'}(rt(T_i), rt') + dist_{T'}(rt', rt(T_j)). \quad (17)$$

Suppose without loss of generality that $dist_{T'}(rt(T_i), rt') \leq dist_{T'}(rt', rt(T_j))$. By Claim 6.4, the required vertex $\delta(rt')$ closest to rt' in the original tree T' is $rt(T) = rt(T_1)$, implying that $j \geq 2$. Hence, by construction,

$$dist_T(rt(T_1), rt(T_j)) = 2 \cdot dist_{T'}(rt', rt(T_j)).$$

It follows that

$$\begin{aligned} \text{dist}_T(\text{rt}(T_i), \text{rt}(T_1)) + \text{dist}_T(\text{rt}(T_1), \text{rt}(T_j)) &\geq \text{dist}_T(\text{rt}(T_1), \text{rt}(T_j)) = 2 \cdot \text{dist}_{T'}(\text{rt}', \text{rt}(T_j)) \\ &\geq \text{dist}_{T'}(\text{rt}(T_i), \text{rt}') + \text{dist}_{T'}(\text{rt}', \text{rt}(T_j)), \end{aligned}$$

and we are done.

Finally, observe that

$$\text{dist}_T(u, v) = \text{dist}_T(u, \text{rt}(T_i)) + \text{dist}_T(\text{rt}(T_i), \text{rt}(T_1)) + \text{dist}_T(\text{rt}(T_1), \text{rt}(T_j)) + \text{dist}_T(\text{rt}(T_j), v). \quad (18)$$

Plugging Equations (15), (16) and (17) in Equation (18) yields

$$\begin{aligned} \text{dist}_T(u, v) &= \text{dist}_T(u, \text{rt}(T_i)) + (\text{dist}_T(\text{rt}(T_i), \text{rt}(T_1)) + \text{dist}_T(\text{rt}(T_1), \text{rt}(T_j))) + \text{dist}_T(\text{rt}(T_j), v) \\ &\geq \text{dist}_{T'}(u, \text{rt}(T_i)) + (\text{dist}_{T'}(\text{rt}(T_i), \text{rt}') + \text{dist}_{T'}(\text{rt}', \text{rt}(T_j))) + \text{dist}_{T'}(\text{rt}(T_j), v) \\ &\geq \text{dist}_{T'}(u, v). \end{aligned}$$

(The last inequality holds by the triangle inequality.) \blacksquare

Lemmas 6.2, 6.3 and 6.5 imply the validity of Proposition 6.1. \blacksquare

Dinitz et al. [14] analyzed the 1-dimensional metric space ϑ_n with n points v_1, v_2, \dots, v_n lying on the x -axis with coordinates $1, 2, \dots, n$, respectively, and have shown that for any parameter $\ell = O(\log n)$, any spanning tree τ for ϑ_n with depth $h(\tau) = o(\ell)$ has weight $w(\tau) = \omega(\ell \cdot n^{1+1/\ell}) = \omega(\ell \cdot n^{1/\ell}) \cdot w(\text{MST}(\vartheta_n))$, and vice versa, i.e., if $w(\tau) = o(\ell \cdot n) = o(\ell) \cdot w(\text{MST}(\vartheta_n))$, then $h(\tau) = \omega(\ell \cdot n^{1/\ell})$. Proposition 6.1 enables us to extend this lower bound to Steiner trees.

Corollary 6.6 *For a positive integer $\ell = O(\log n)$, any Steiner tree for ϑ_n with depth $o(\ell)$ has lightness $\omega(\ell \cdot n^{1/\ell})$, and vice versa, i.e., any Steiner tree for ϑ_n with lightness $o(\ell)$ has depth $\omega(\ell \cdot n^{1/\ell})$.*

Proof: Define $R' = \{v_1, v_2, \dots, v_n\}$, and let (T', rt') be an arbitrary Steiner rooted tree for ϑ_n spanning some superset V' of R' . By definition, for every pair v_i, v_j of points in R' , $\text{dist}_{T'}(v_i, v_j) \geq |i - j|$. By Proposition 6.1, there exists a weighted rooted tree (T, rt) over R' , having depth $h(T) \leq h(T')$ and weight $w(T) \leq 2 \cdot w(T')$, and which also dominates T' in the following sense: for every pair v_i, v_j of points in R' , $\text{dist}_T(v_i, v_j) \geq \text{dist}_{T'}(v_i, v_j)$. It follows that

$$\text{dist}_T(v_i, v_j) \geq \text{dist}_{T'}(v_i, v_j) \geq |i - j|. \quad (19)$$

We transform T into a spanning tree τ of ϑ_n in the obvious way, i.e., each edge in τ is assigned a weight that is equal to the Euclidean distance between the endpoints of that edge. Equation (19) implies that $w(\tau) \leq w(T)$, and so $w(\tau) \leq w(T) \leq 2 \cdot w(T')$. Clearly, $h(\tau) = h(T) \leq h(T')$. Consequently, if $h(T') = o(\ell)$, then $h(\tau) \leq h(T') = o(\ell)$, and so the aforementioned lower bound of [14] yields $w(T') \geq w(\tau)/2 = \omega(\ell \cdot n^{1/\ell}) \cdot w(\text{MST}(\vartheta_n))$. Symmetrically, if the weight $w(T')$ of T' is bounded from above by $o(\ell) \cdot w(\text{MST}(\vartheta_n))$, then so is the weight $w(\tau)$ of τ , which, by the lower bound of [14], yields $h(T') \geq h(\tau) = \omega(\ell \cdot n^{1/\ell})$. \blacksquare

In particular, Corollary 6.6 shows that any Steiner tree for ϑ_n has either depth or lightness at least $\Omega(\log n)$. On the other hand, Arya et al. [3] devised a construction of Euclidean $(1 + \epsilon)$ -spanners with both hop-diameter and lightness at most $O(\log n)$. Corollary 6.6 implies that the result of [3] cannot be improved even if one allows the spanner to use Steiner points.

Theorem 6.7 (Tight lower bounds for Euclidean Steiner spanners) *Any Euclidean (possibly Steiner) spanner for ϑ_n with hop-diameter $o(\log n)$ has lightness $\omega(\log n)$, and vice versa, i.e., any Euclidean (possibly Steiner) spanner for ϑ_n with lightness $o(\log n)$ has hop-diameter $\omega(\log n)$.*

Proof: Suppose for contradiction that for any set M of n points in the plane there exists a Steiner spanner with hop-diameter $o(\log n)$ and lightness $O(\log n)$ or vice versa. Let H be such a spanner for ϑ_n , and let T be a breadth-first-search tree over H rooted at an arbitrary point $v_i \in \vartheta_n$. Let $\Lambda(H)$ denote the hop-diameter of H . Obviously, $w(T) \leq w(H)$, and $h(T) \leq \Lambda(H)$. Hence, if $(w(H) = o(\log n) \cdot w(MST(\vartheta_n))$ and $\Lambda(H) = O(\log n)$), then $(w(T) = o(\log n) \cdot w(MST(\vartheta_n))$ and $h(T) = O(\log n)$), contradicting Corollary 6.6. The complementary case is that $(w(H) = O(\log n) \cdot w(MST(\vartheta_n))$ and $\Lambda(H) = o(\log n)$). However, similarly, it follows that $(w(T) = O(\log n) \cdot w(MST(\vartheta_n))$ and $h(T) = o(\log n)$), again contradicting Corollary 6.6. ■

Acknowledgements

We are grateful to the anonymous reviewers for their thoughtful and helpful comments.

References

- [1] C. J. Alpert, T. C. Hu, J. H. Huang, A. B. Kahng, and D. Karger. Prim-Dijkstra tradeoffs for improved performance-driven routing tree design. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 14(7):890–896, 1995.
- [2] E. Althaus, S. Funke, S. Har-Peled, J. Könemann, E. A. Ramos, and M. Skutella. Approximating k -hop minimum-spanning trees. *Oper. Res. Lett.*, 33(2):115–120, 2005.
- [3] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. H. M. Smid. Euclidean spanners: short, thin, and lanky. In *Proc. of 27th STOC*, pages 489–498, 1995.
- [4] B. Awerbuch and Y. Azar. Buy-at-bulk network design. In *Proc. of 38th FOCS*, pages 542–547, 1997.
- [5] B. Awerbuch, A. Baratz, and D. Peleg. Cost-sensitive analysis of communication protocols. In *Proc. of 9th PODC*, pages 177–187, 1990.
- [6] B. Awerbuch, A. Baratz, and D. Peleg. Efficient broadcast and light-weight spanners. *Manuscript*, 1991.
- [7] P. Bose, J. Gudmundsson, and P. Morin. Ordered theta graphs. *Comput. Geom.*, 28(1):11–18, 2004.
- [8] T. M. Chan. Euclidean bounded-degree spanning tree ratios. *Discrete & Computational Geometry*, 32(2):177–194, 2004.
- [9] T. M. Chan. Well-separated pair decomposition in linear time? *Inf. Process. Lett.*, 107(5):138–141, 2008.
- [10] M. Charikar, C. Chekuri, T. Cheung, Z. Dai, A. Goel, S. Guha, and M. Li. Approximation algorithms for directed Steiner problems. *J. Algorithms*, 33(1):73–91, 1999.
- [11] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong. Performance-driven global routing for cell based ics. In *Proc. of 9th ICCD*, pages 170–173, 1991.
- [12] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong. Provably good algorithms for performance-driven global routing. In *Proc. of 5th ISCAS*, pages 2240–2243, 1992.
- [13] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong. Provably good performance-driven global routing. *IEEE Trans. on CAD of Integrated Circuits and Sys.*, 11(6):739–752, 1992.
- [14] Y. Dinitz, M. Elkin, and S. Solomon. Low-light trees, and tight lower bounds for Euclidean spanners. *Discrete & Computational Geometry*, 43(4):736–783, 2010.
- [15] D. Eppstein. Spanning trees and spanners. *Technical report, Department of Information and Computer Science, University of California, Irvine*, (96-16), 1996.
- [16] M. Farshi and J. Gudmundsson. Experimental study of geometric t -spanners: A running time comparison. In *Proc. of 6th WEA*, pages 270–284, 2007.
- [17] S. P. Fekete, S. Khuller, M. Klemmstein, B. Raghavachari, and N. E. Young. A network-flow technique for finding low-weight bounded-degree spanning trees. *J. Algorithms*, 24(2):310–324, 1997.
- [18] A. Goel and D. Estrin. Simultaneous optimization for concave costs: single sink aggregation or single source buy-at-bulk. In *Proc. of 14th SODA*, pages 499–505, 2003.
- [19] A. Goel and I. Post. One tree suffices: A simultaneous $o(1)$ -approximation for single-sink buy-at-bulk. In *To appear in Proc. of 51st FOCS*, 2010.
- [20] M. Grünewald, T. Lukovszki, C. Schindelbauer, and K. Vollbert. Distributed maintenance of resource efficient wireless network topologies. In *Proc. of 8th Euro-Par*, pages 935–946, 2002.

- [21] S. Guha, A. Meyerson, and K. Munagala. A constant factor approximation for the single sink edge installation problems. In *Proc. of 33rd STOC*, pages 383–388, 2001.
- [22] A. Gupta. Steiner points in tree metrics don't (really) help. In *Proc. of 12th SODA*, pages 220–227, 2001.
- [23] L. Jia, G. Lin, G. Noubir, R. Rajaraman, and R. Sundaram. Universal approximations for tsp, steiner tree, and set cover. In *Proc. of 37th STOC*, pages 386–395, 2005.
- [24] W. B. Johnson and J. Lindenstrauss. Extensions of lipschitz mappings into a hilbert space. *Contemp. Math.*, 26:189–206, 1984.
- [25] E. Kantor and D. Peleg. Approximate hierarchical facility location and applications to the shallow Steiner tree and range assignment problems. In *Proc. of 6th CIAC*, pages 211–222, 2006.
- [26] S. Khuller, B. Raghavachari, and N. E. Young. Balancing minimum spanning and shortest path trees. In *Proc. of 4th SODA*, pages 243–250, 1993.
- [27] P. N. Klein, S. Rao, M. Rauch, and S. Subramanian. Faster shortest-path algorithms for planar graphs. In *Proc. of 26th STOC*, pages 27–37, 1994.
- [28] G. Kortsarz and D. Peleg. Approximating the weight of shallow Steiner trees. *Discrete Applied Mathematics*, 93(2-3):265–285, 1999.
- [29] T. Lukovszki. New results of fault tolerant geometric spanners. In *Proc. of 6th WADS*, pages 193–204, 1999.
- [30] T. Lukovszki. New results on geometric spanners and their applications. *Ph.D. thesis, Department of Computer-Science, University of Paderborn, Paderborn, Germany*, 1999.
- [31] T. Lukovszki, C. Schindelhauer, and K. Volbert. Resource efficient maintenance of wireless network topologies. *J. UCS*, 12(9):1292–1311, 2006.
- [32] M. V. Marathe, R. Ravi, R. Sundaram, S. S. Ravi, D. J. Rosenkrantz, and H. B. H. III. Bicriteria network design problems. *J. Algorithms*, 28(1):142–171, 1998.
- [33] C. L. Monma and S. Suri. Transitions in geometric minimum spanning trees. *Discrete & Computational Geometry*, 8:265–293, 1992.
- [34] G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.
- [35] C. H. Papadimitriou and U. V. Vazirani. On two geometric problems related to the traveling salesman problem. *J. Algorithms*, 5(2):231–246, 1984.
- [36] J. Ruppert and R. Seidel. Approximating the d -dimensional complete Euclidean graph. In *Proc. of 3rd CCCG*, pages 207–210, 1991.
- [37] J. S. Salowe, D. S. Richards, and D. E. Wrege. Mixed spanning trees: a technique for performance-driven routing. In *Proc. of 3rd GLS-VLSI*, pages 62–66, 1993.
- [38] A. C. Yao. On constructing minimum spanning trees in k -dimensional spaces and related problems. *SIAM J. Comput.*, 11(4):721–736, 1982.