

Balancing Degree, Diameter and Weight in Euclidean Spanners

Shay Solomon ^{*} and Michael Elkin ^{**}

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

Abstract. In a seminal STOC'95 paper, Arya et al. [4] devised a construction that for any set S of n points in \mathbb{R}^d and any $\epsilon > 0$, provides a $(1 + \epsilon)$ -spanner with diameter $O(\log n)$, weight $O(\log^2 n)w(MST(S))$, and constant maximum degree. Another construction of [4] provides a $(1 + \epsilon)$ -spanner with $O(n)$ edges and diameter $\alpha(n)$, where α stands for the inverse-Ackermann function. Das and Narasimhan [18] devised a construction with constant maximum degree and weight $O(w(MST(S)))$, but whose diameter may be arbitrarily large. In another construction by Arya et al. [4] there is diameter $O(\log n)$ and weight $O(\log n)w(MST(S))$, but it may have arbitrarily large maximum degree. These constructions fail to address situations in which we are prepared to compromise on one of the parameters, but cannot afford it to be arbitrarily large.

In this paper we devise a novel *unified* construction that trades between maximum degree, diameter and weight gracefully. For a positive integer k , our construction provides a $(1 + \epsilon)$ -spanner with maximum degree $O(k)$, diameter $O(\log_k n + \alpha(k))$, weight $O(k \log_k n \log n)w(MST(S))$, and $O(n)$ edges. For $k = O(1)$ this gives rise to maximum degree $O(1)$, diameter $O(\log n)$ and weight $O(\log^2 n)w(MST(S))$, which is one of the aforementioned results of [4]. For $k = n^{1/\alpha(n)}$ this gives rise to diameter $O(\alpha(n))$, weight $O(n^{1/\alpha(n)}(\log n)\alpha(n)w(MST(S)))$ and maximum degree $O(n^{1/\alpha(n)})$. In the corresponding result from [4] the spanner has the same number of edges and diameter, but its weight and degree may be arbitrarily large. Our construction also provides a similar tradeoff in the complementary range of parameters, i.e., when the weight should be smaller than $\log^2 n$, but the diameter is allowed to grow beyond $\log n$.

1 Introduction

Euclidean Spanners. Consider the weighted complete graph $\mathcal{S} = (S, \binom{S}{2})$ induced by a set S of n points in \mathbb{R}^d , $d \geq 2$. The weight of an edge $(x, y) \in \binom{S}{2}$, for a pair of distinct points $x, y \in S$, is defined to be the Euclidean distance $\|x - y\|$ between x and y . Let $G = (S, E)$ be a spanning subgraph of \mathcal{S} , with $E \subseteq \binom{S}{2}$, and assume that exactly as in \mathcal{S} , for any edge $e = (x, y) \in E$, its weight $w(e)$ in G is defined to be $\|x - y\|$. For a parameter $\epsilon > 0$, the spanning subgraph G is

^{*} This research has been supported by the Clore Fellowship grant No. 81265410.

^{**} This research has been supported by the BSF grant No. 2008430.

called a $(1+\epsilon)$ -*spanner* for the point set S if for every pair $x, y \in S$ of points, the distance $\text{dist}_G(x, y)$ between x and y in G is at most $(1+\epsilon)\|x-y\|$. Euclidean spanners were introduced¹ more than twenty years ago by Chew [17]. Since then they evolved into an important subarea of Computational Geometry [28, 3, 18, 4, 19, 6, 34, 1, 12, 20]. (See also the recent book by Narasimhan and Smid on Euclidean spanners [31], and the references therein.) Also, Euclidean spanners have numerous applications in geometric approximation algorithms [34, 24, 25], geometric distance oracles [24, 26, 25], Network Design [27, 30] and in other areas.

In many of these applications one is required to construct a $(1+\epsilon)$ -spanner $G = (S, E)$ that satisfies a number of useful properties. First, the spanner should contain $O(n)$ (or nearly $O(n)$) edges. Second, its *weight* $w(G) = \sum_{e \in E} w(e)$ should not be much greater than the weight $w(\text{MST}(S))$ of the minimum spanning tree $\text{MST}(S)$ of S . Third, its *diameter* $\Lambda = \Lambda(G)$ should be small, i.e., for every pair of points $x, y \in S$ there should exist a path P in G that contains at most Λ edges and has weight $w(P) = \sum_{e \in E(P)} w(e) \leq (1+\epsilon)\|x-y\|$. Fourth, its *maximum degree* (henceforth, *degree*) $\Delta(G)$ should be small.

In a seminal STOC'95 paper, Arya et al. [4] have devised a construction of $(1+\epsilon)$ -spanners with lightness² $O(\log^2 n)$, diameter $O(\log n)$ and constant degree. They have also devised a construction of $(1+\epsilon)$ -spanners with $O(n)$ (respectively, $O(n \log^* n)$) edges and diameter $O(\alpha(n))$ (resp., at most $O(1)$). However, in the latter construction the resulting spanners may have arbitrarily large (i.e., at least $\Omega(n)$) lightness and degree. There are also a few other known constructions of $(1+\epsilon)$ -spanners. Das and Narasimhan [18] devised a construction with constant degree and lightness, but the diameter may be arbitrarily large. There is also another construction by Arya et al. [4] that guarantees that both the diameter and the lightness are $O(\log n)$, but the degree may be arbitrarily large. While these constructions address some important practical scenarios, they certainly do not address all of them. In particular, they fail to address situations in which we are prepared to compromise on one of the parameters, but cannot afford this parameter to be arbitrarily large.

In this paper we devise a novel *unified* construction that trades between degree, diameter and weight gracefully. For a positive integer k , our construction provides a $(1+\epsilon)$ -spanner with degree $O(k)$, diameter $O(\log_k n + \alpha(k))$, lightness $O(k \log_k n \log n)$, and $O(n)$ edges. Also, we can improve the bound on the diameter from $O(\log_k n + \alpha(k))$ to $O(\log_k n)$, at the expense of increasing the number of edges from $O(n)$ to $O(n \log^* n)$. Note that for $k = O(1)$ our tradeoff gives rise to degree $O(1)$, diameter $O(\log n)$ and lightness $O(\log^2 n)$, which is one of the aforementioned results of [4]. Also, for $k = n^{1/\alpha(n)}$ it gives rise to a spanner with degree $O(n^{1/\alpha(n)})$, diameter $O(\alpha(n))$ and lightness $O(n^{1/\alpha(n)}(\log n)\alpha(n))$. In the corresponding result from [4] the spanner has the same number of edges and diameter, but its lightness and degree may be arbitrarily large.

In addition, we can achieve lightness $o(\log^2 n)$ at the expense of increasing the diameter. Specifically, for a parameter k the second variant of our construction provides a $(1+\epsilon)$ -spanner with degree $O(1)$, diameter $O(k \log_k n)$, and lightness $O(\log_k n \log n)$. For example, for $k = \log^\delta n$, for an arbitrarily small constant

¹ The notion “spanner” was coined by Peleg and Ullman [33], who have also introduced spanners for general graphs.

² For convenience, we will henceforth refer to the normalized notion of weight $\Psi(G) = \frac{w(G)}{w(\text{MST}(S))}$, which we call *lightness*.

$\delta > 0$, we get a $(1 + \epsilon)$ -spanner with degree $O(1)$, diameter $O(\log^{1+\delta} n)$ and lightness $O(\frac{\log^2 n}{\log \log n})$. (See Table 1 in Appendix B for a concise comparison of previous and new results.)

Our unified construction can be implemented in $O(n \log n)$ time. This matches the state-of-the-art running time of the aforementioned constructions [4]. Moreover, our construction can be implemented in linear time if the spread³ of the point set is bounded by a polynomial in n . The linear time implementation of our construction is based on a recent work of Chan [11] which shows that well-separated pair decompositions can be constructed in linear time for point sets with polynomially bounded spread. The result of [11] holds in the real-RAM model, assuming that the floor function is available and that the word size is at least logarithmic. The same assumptions are required for implementing our construction in linear time.

Note that in any construction of spanners with degree $O(k)$, the diameter is $\Omega(\log_k n)$. Also, Chan and Gupta [12] showed that any $(1 + \epsilon)$ -spanner with $O(n)$ edges must have diameter $\Omega(\alpha(n))$. Consequently, our upper bound of $O(\log_k n + \alpha(k))$ on the diameter is tight under the constraints that the degree is $O(k)$ and the number of edges is $O(n)$. If we allow $O(n \log^* n)$ edges in the spanner, then our bound on the diameter is reduced to $O(\log_k n)$, which is again tight under the constraint that the degree is $O(k)$.

In addition, Dinitz et al. [20] showed that for any construction of spanners, if the diameter is at most $O(\log_k n)$, then the lightness is at least $\Omega(k \log_k n)$ and vice versa, if the lightness is at most $O(\log_k n)$, the diameter is at least $\Omega(k \log_k n)$. This lower bound implies that the bound on lightness in both our tradeoffs cannot possibly be improved by more than a factor of $\log n$. The same slack of $\log n$ is present in the result of [4] that guarantees lightness $O(\log^2 n)$, diameter $O(\log n)$ and constant degree.

Euclidean Spanners for Random Point Sets. For random point sets in the d -dimensional unit cube (henceforth, unit cube), we “shave” a factor of $\log n$ from the lightness bound in both our tradeoffs, and show that the first (respectively, second) variant of our construction achieves maximum degree $O(k)$ (resp., $O(1)$), diameter $O(\log_k n + \alpha(k))$ (resp., $O(k \log_k n)$) and lightness that is with high probability (henceforth, w.h.p.) $O(k \log_k n)$ (resp., $O(\log_k n)$). Note that for $k = O(1)$ both these tradeoffs give rise to degree $O(1)$, diameter $O(\log n)$ and lightness (w.h.p.) $O(\log n)$. In addition to these tradeoffs, we can get a $(1 + \epsilon)$ -spanner with diameter $O(\log n)$ and lightness (w.h.p.) $O(1)$. Finally, under the aforementioned assumptions these constructions can be implemented in randomized linear time.

Spanners for Doubling Metrics. Spanners for doubling metrics⁴ have received much attention in recent years (see, e.g., [13, 12, 23]). In particular, Chan et al. [13] showed that for any doubling metric (X, δ) there exists a $(1 + \epsilon)$ -spanner with constant maximum degree. In addition, Chan and Gupta [12] devised a construction of $(1 + \epsilon)$ -spanners for doubling metrics that achieves the optimal tradeoff between the number of edges and the diameter. We present a single

³ The *Spread* of a point set is the ratio between the largest pairwise distance and the smallest pairwise distance.

⁴ The *doubling dimension* of a metric (X, δ) is the smallest value ζ such that every ball B in the metric can be covered by at most 2^ζ balls of half the radius of B . The metric (X, δ) is called *doubling* if its doubling dimension ζ is constant.

construction of $O(1)$ -spanners for doubling metrics that achieves the optimal tradeoff between the degree, the diameter and the number of edges in the entire range of parameters. Specifically, for a parameter k , our construction provides an $O(1)$ -spanner with maximum degree $O(k)$, diameter $O(\log_k n + \alpha(k))$ and $O(n)$ edges. Also, we can improve the bound on the diameter from $O(\log_k n + \alpha(k))$ to $O(\log_k n)$, at the expense of increasing the number of edges from $O(n)$ to $O(n \log^* n)$. More generally, we can achieve the same optimal tradeoff between the number of edges and the diameter as the spanners of [12] do, while also having the optimal maximum degree. The drawback is, however, that the stretch of our spanners is $O(1)$ rather than $1 + \epsilon$.

Spanners for Tree Metrics. We denote by ϑ_n the metric induced by n points v_1, v_2, \dots, v_n lying on the x -axis with coordinates $1, 2, \dots, n$, respectively. In a classical STOC'82 paper [38], Yao showed that there exists a 1-spanner⁵ $G = (V, E)$ for ϑ_n with $O(n)$ edges and diameter $O(\alpha(n))$, and that this is tight. Chazelle [15] extended the result of [38] to arbitrary tree metrics. Other proofs of Chazelle's result appeared in [2, 8, 37]. The problem is also closely related to the well-studied problem of computing partial-sums [36, 38, 16, 32].

In all constructions [38, 15, 2, 8, 37] of 1-spanners for tree metrics, the degree and lightness of the resulting spanner may be arbitrarily large. Moreover, the constraint that the diameter is $O(\alpha(n))$ implies that the degree must be $n^{\Omega(1/\alpha(n))}$. A similar lower bound on lightness follows from the result of [20].

En route to our tradeoffs for Euclidean spanners, we have extended the results of [38, 15, 2, 8, 37] and devised a construction that achieves the *optimal* (up to constant factors) *tradeoff between all involved parameters*. Specifically, consider an n -vertex tree T of degree $\Delta(T)$, and let k be a positive integer. Our construction provides a 1-spanner for the metric M_T induced by T with $O(n)$ edges, degree $O(k + \Delta(T))$, diameter $O(\log_k n + \alpha(k))$, and lightness $O(k \log_k n)$. We can also get a spanner with $O(n \log^* n)$ edges, diameter $O(\log_k n)$, and the same degree and lightness as above. For the complementary range of diameter, another variant of our construction provides a 1-spanner with $O(n)$ edges, degree $O(\Delta(T))$, diameter $O(k \log_k n)$ and lightness $O(\log_k n)$. As was mentioned above, both these tradeoffs are optimal up to constant factors. In addition, this construction of 1-spanners can be implemented in linear time.

We show that this general tradeoff between various parameters of 1-spanners for tree metrics is useful for deriving new results (and improving existing results) in the context of Euclidean spanners and spanners for doubling metrics. We anticipate that this tradeoff will be found useful in the context of partial sums problems, and for other applications.

Structure of the Paper. In Sect. 2 we describe our construction of 1-spanners for tree metrics. Therein we start (Sect. 2.1) with outlining our basic scheme. We proceed (Sect. 2.2) with describing our 1-dimensional construction. In Sect. 2.3 we extend this construction to general tree metrics. In Sect. 3 we derive our results for Euclidean spanners and spanners for doubling metrics. Due to space constraints, we leave all issues of running time out of this extended abstract, and relegate the figures and Table 1, as well as some proofs, to Appendices. Also, a detailed overview of our and previous techniques appears in Appendix A.

⁵ The graph G is said to be a *1-spanner* of ϑ_n if for every pair of distinct vertices $v_i, v_j \in V$, the distance between them in G is equal to their distance $\|i - j\|$ in ϑ_n . Yao stated this problem in terms of partial sums.

Preliminaries. Let G be a spanning subgraph of a metric space $M = (V, dist)$. The *stretch* between two vertices $u, v \in V$ is defined as $\frac{dist_G(u,v)}{dist(u,v)}$. We say that G is a t -*spanner* for M if the maximum stretch taken over all pairs of points in V is at most t . For any two vertices u, v in a tree T , their (weighted) distance in T is denoted by $dist_T(u, v)$. The tree metric M_T induced by a tree T is defined as $M_T = (V(T), dist_T)$. The *size* of a tree T , denoted $|T|$, is the number of vertices in T . For a positive integer n , we denote the set $\{1, 2, \dots, n\}$ by $[n]$.

2 1-Spanners for Tree Metrics

2.1 The Basic Scheme

Consider an arbitrary n -vertex (weighted) rooted tree (T, rt) , and let M_T be the tree metric induced by T . Clearly, T is both a 1-spanner and an MST of M_T , but its diameter may be huge. We would like to reduce the diameter of this 1-spanner by adding to it some edges. On the other hand, the number of edges of the resulting spanner should still be linear in n . Moreover, the lightness and the maximum degree of the resulting spanner should also be reasonably small.

Let H be a spanning subgraph of M_T . The *monotone distance* between any two points u and v in H is defined as the minimum number of hops in a 1-spanner path in H connecting them. Two points in M_T are called *comparable* if one is an ancestor of the other in the underlying tree T . The *monotone diameter* (respectively, *comparable monotone diameter*) of H , denoted $\Lambda(H)$ (resp., $\bar{\Lambda}(H)$), is defined as the maximum monotone distance in H between any two points (resp., any two comparable points) in M_T . Observe that if any two *comparable* points are connected via a 1-spanner path that consists of at most h hops, then any two *arbitrary* points are connected via a 1-spanner path that consists of at most $2h$ hops. Consequently, $\bar{\Lambda}(H) \leq \Lambda(H) \leq 2\bar{\Lambda}(H)$. We henceforth restrict our attention to comparable monotone diameter in the sequel.

Let k be a fixed parameter. The first ingredient of the algorithm is to select a set of $O(k)$ *cut-vertices* whose removal from T partitions it into a collection of subtrees of size $O(n/k)$ each. As will become clear in the sequel, we also require this set to satisfy several additional properties. Having selected the cut-vertices, the next step of the algorithm is to connect the cut-vertices via $O(k)$ edges, so that the monotone distance between any pair of comparable cut-vertices will be small. (This phase does not involve a recursive call of the algorithm.) Finally, the algorithm calls itself recursively for each of the subtrees.

We insert all edges of the original tree T into our final spanner H . These edges connect between cut-vertices and subtrees in the spanner. We remark that the spanner contains no other edges that connect between cut-vertices and subtrees, or between different subtrees.

2.2 1-Dimensional Spaces

In this section we devise an optimal construction of 1-spanners for ϑ_n . (See Sect. 1 for its definition.) Our argument extends easily to any 1-dimensional space.

Denote by P_n the path $(v_1, v_2), (v_2, v_3), \dots, (v_{n-1}, v_n)$ that induces the metric ϑ_n . We remark that the edges of P_n belong to all spanners that we construct.

Selecting the Cut-Vertices. The task of selecting the cut-vertices in the 1-dimensional case is trivial. (We assume for simplicity that n is an integer power of k .) In addition to the two endpoints v_1 and v_n of the path, we select the $k-1$ points r_1, r_2, \dots, r_{k-1} to be cut-vertices, where for each $i \in [k-1]$, $r_i = v_{i(n/k)}$. Indeed, by removing the $k+1$ cut-vertices $r_0 = v_1, r_1, \dots, r_{k-1}, r_k = v_n$ from the path (along with their incident edges), we are left with k intervals I_1, I_2, \dots, I_k of length at most n/k each. The two endpoints v_1 and v_n of the path are called the *sentinels*, and they play a special role in the construction. (See Fig. 1 in Appendix B for an illustration for the case $k=2$.)

1-Spanners with Low Diameter. In this section we devise a construction $H_k(n)$ of 1-spanners for ϑ_n with comparable monotone diameter $\bar{\Lambda}(n) = \bar{\Lambda}(H_k(n))$ in the range $\Omega(\alpha(n)) = \bar{\Lambda}(n) = O(\log n)$.

First, the algorithm connects the $k+1$ cut-vertices $r_0 = v_1, r_1, \dots, r_{k-1}, r_k = v_n$ via one of the aforementioned constructions of 1-spanners from [38, 15, 2, 8, 37] (henceforth, *list-spanner*). In other words, $O(k)$ edges are added between cut-vertices to guarantee that the monotone distance between any two cut-vertices⁶ will be at most $O(\alpha(k))$. Then the algorithm connects each of the two sentinels to all other k cut-vertices. Finally, the algorithm calls itself recursively for each of the intervals I_1, I_2, \dots, I_k . At the bottom level of the recursion, i.e., when $n \leq k$, the algorithm uses the list-spanner to connect all points, and also connects both sentinels v_1 and v_n to all the other $n-2$ points. (See Fig. 2 in Appendix B for an illustration.)

Denote by $E(n)$ the number of edges in $H_k(n)$, excluding edges of P_n . Clearly, $E(n)$ satisfies the recurrence $E(n) \leq O(k) + kE(n/k)$, with the base condition $E(n) = O(n)$, for $n \leq k$, yielding $E(n) = O(n)$. Denote by $\Delta(n)$ the maximum degree of a vertex in $H_k(n)$, excluding edges of P_n . Clearly, $\Delta(n)$ satisfies the recurrence $\Delta(n) \leq \max\{k, \Delta(n/k)\}$, with the base condition $\Delta(n) \leq n-1$, for $n \leq k$, yielding $\Delta(n) \leq k$. Including edges of P_n , the number of edges increases by $n-1$ units, and the maximum degree increases by at most two units.

To bound the weight $w(n) = w(H_k(n))$ of $H_k(n)$, first note that at most $O(k)$ edges are added between cut-vertices. Each of these edges has weight at most $n-1$. The total weight of all edges within an interval I_i is at most $w(n/k)$. Observe also that $w(P_n) = n-1$. Hence $w(n)$ satisfies the recurrence $w(n) \leq O(nk) + kw(n/k)$, with the base condition $w(n) = O(n^2)$, for $n \leq k$. It follows that $w(n) = O(nk \log_k n) = O(k \log_k n)w(MST(\vartheta_n))$.

Next, we show that the comparable monotone diameter $\bar{\Lambda}(n)$ of $H_k(n)$ is at most $O(\log_k n + \alpha(k))$. The *monotone radius* $R(n)$ of $H_k(n)$ is defined as the maximum monotone distance in $H_k(n)$ between one of the sentinels (either v_1 or v_n) and some other point in ϑ_n . Let v_j be a point in ϑ_n , and let i be the index such that $i(n/k) \leq j < (i+1)(n/k)$. Then the 1-spanner path Π in $H_k(n)$ connecting the sentinel v_1 and the point v_j will start with the two edges $(v_1, v_{i(n/k)})$, $(v_{i(n/k)}, v_{i(n/k)+1})$. The point $v_{i(n/k)+1}$ is a sentinel of the i th interval I_i . Hence, the path Π will continue recursively, from $v_{i(n/k)+1}$ to v_j . It follows that the monotone radius $R(n)$ satisfies the recurrence $R(n) \leq 2 + R(n/k)$, with the base condition $R(n) = 1$, for $n \leq k$, yielding $R(n) = O(\log_k n)$. Finally, it is easy to verify that $\bar{\Lambda}(n)$ satisfies the recurrence $\bar{\Lambda}(n) \leq \max\{\bar{\Lambda}(n/k), O(\alpha(k)) + 2R(n/k)\}$, with the base condition $\bar{\Lambda}(n) = O(\alpha(n))$, for $n \leq k$. Hence $\bar{\Lambda}(n) = O(\log_k n + \alpha(k))$.

⁶ In the 1-dimensional case any two points are comparable.

Theorem 1. *For any n -point 1-dimensional space and a parameter k , there exists a 1-spanner with $O(n)$ edges, maximum degree at most $k + 2$, diameter $O(\log_k n + \alpha(k))$ and lightness $O(k \log_k n)$.*

1-Spanners with High Diameter. In this section we devise a construction $H'_k(n)$ of 1-spanners for ϑ_n with comparable monotone diameter $\bar{A}'(n) = \bar{A}(H'_k(n))$ in the range $\bar{A}'(n) = \Omega(\log n)$.

The algorithm connects the $k + 1$ cut-vertices $r_0 = v_1, r_1, \dots, r_{k-1}, r_k = v_n$ via a path of length k , i.e., it adds the edges $(r_0, r_1), (r_1, r_2), \dots, (r_{k-1}, r_k)$ into the spanner. In addition, it calls itself recursively for each of the intervals I_1, I_2, \dots, I_k . At the bottom level of the recursion, i.e., when $n \leq k$, the algorithm adds no additional edges to the spanner. (See Fig. 1 and Fig. 2 for an illustration.)

We denote by $\Delta'(n)$ the maximum degree of a vertex in $H'_k(n)$, excluding edges of P_n . Clearly, $\Delta'(n)$ satisfies the recurrence $\Delta'(n) \leq \max\{2, \Delta'(n/k)\}$, with the base condition $\Delta'(k) = 0$, for $n \leq k$, yielding $\Delta'(n) \leq 2$. Including edges of P_n , the maximum degree increases by at most two units, and so $\Delta(H'_k(n)) \leq 4$. Consequently, the number of edges in $H'_k(n)$ is no greater than $2n$.

To bound the weight $w'(n) = w(H'_k(n))$ of $H'_k(n)$, first note that the path connecting all $k + 1$ cut-vertices has weight $n - 1$. Observe also that $w(P_n) = n - 1$. Thus $w'(n)$ satisfies the recurrence $w'(n) \leq 2(n - 1) + kw'(n/k)$, with the base condition $w'(n) \leq n - 1$, for $n \leq k$, yielding $w'(n) = O(n \log_k n) = O(\log_k n)w(MST(\vartheta_n))$.

Note that the monotone radius $R'(n)$ of $H'_k(n)$ satisfies the recurrence $R'(n) \leq k + R'(n/k)$, with the base condition $R'(n) \leq n - 1$, for $n \leq k$. Hence, $R'(n) = O(k \log_k n)$. It is easy to verify that the comparable monotone diameter $\bar{A}'(n) = \bar{A}(H'_k(n))$ of $H'_k(n)$ satisfies the recurrence $\bar{A}'(n) \leq \max\{\bar{A}'(n/k), k + 2R'(n/k)\}$, with the base condition $\bar{A}'(n) \leq n - 1$, for $n \leq k$, and so $\bar{A}'(n) = O(k \log_k n)$.

Finally, we remark that the spanner $H'_k(n)$ is a planar graph.

Theorem 2. *For any n -point 1-dimensional space and a parameter k , there exists a 1-spanner with maximum degree 4, diameter $O(k \log_k n)$ and lightness $O(\log_k n)$. This 1-spanner is a planar graph.*

2.3 General Tree Metrics

In this section we extend the constructions of Sect. 2.2 to general tree metrics.

Selecting the Cut-Vertices. In this section we present a procedure for selecting, given a tree T , a subset of $O(k)$ vertices whose removal from the tree partitions it into subtrees of size $O(|T|/k)$ each.

Let (T, rt) be a rooted tree. For an inner vertex v in T with $ch(v)$ children, we denote its children by $c_1(v), c_2(v), \dots, c_{ch(v)}(v)$. Suppose without loss of generality that the size of the subtree $T_{c_1(v)}$ of v is no smaller than the size of any other subtree of v , i.e., $|T_{c_1(v)}| \geq |T_{c_2(v)}|, |T_{c_3(v)}|, \dots, |T_{c_{ch(v)}(v)}|$. We say that the vertex $c_1(v)$ (respectively, the subtree $T_{c_1(v)}$) is the *left-most* child (resp., subtree) of v . Also, an edge in T is called *left-most* if it connects a vertex v in T and its left-most child $c_1(v)$. We denote by $P(v) = (v, c_1(v), \dots, l(v))$ the path of left-most edges leading down from v to the left-most vertex $l(v)$ in the subtree T_v of T rooted at v . A vertex v in T is called *d-balanced*, for $d \geq 1$, or simply

balanced if d is clear from the context, if $|T_{c_1(v)}| \leq |T| - d$. The first balanced vertex along $P(v)$ is denoted by $B(v)$.

Next, we present the Procedure CV that accepts as input a rooted tree (T, rt) and a parameter d , and returns as output a subset of $V(T)$. If $|T| < 2d$, the procedure returns the empty set \emptyset . Otherwise, for each child $c_i(b)$ of the first balanced vertex $b = B(rt)$ along $P(rt)$, $i \in [ch(b)]$, the procedure recursively constructs the subset $C_i = CV((T_{c_i(b)}, c_i(b)), d)$, and then returns $\bigcup_{i=1}^{ch(b)} C_i \cup \{b\}$.

Let (T, rt) be an n -vertex rooted tree, and let d be a fixed parameter. Next, we analyze the properties of the set $C = CV((T, rt), d)$ of *cut-vertices*.

For a tree τ , the root $rt(\tau)$ of τ and its left-most vertex $l(\tau)$ are called the *sentinels* of τ . Similarly to the 1-dimensional case, we add the two sentinels $rt(T)$ and $l(T)$ of the original tree T to the set C of cut-vertices. From now on we refer to the appended set $\tilde{C} = C \cup \{rt(T), l(T)\}$ as the set of *cut-vertices*. Observe that the subset \tilde{C} induces a tree $\tilde{Q} = \mathcal{Q}(T, \tilde{C})$ over \tilde{C} in the natural way: a vertex $v \in \tilde{C}$ is defined to be a child of its closest ancestor in T that belongs to \tilde{C} . We denote by $T \setminus \tilde{C}$ the forest obtained from T by removing all vertices in \tilde{C} , along with the edges that are incident to them. The next proposition summarizes the properties of the set \tilde{C} of cut-vertices. Refer to Appendix C for its proof.

Proposition 1. 1) For $n \geq 2d$, $|\tilde{C}| \leq (n/d) + 1$. 2) The size of any tree in $T \setminus \tilde{C}$ is smaller than $2d$. 3) $\tilde{Q} = \mathcal{Q}(T, \tilde{C})$ is a spanning tree of \tilde{C} rooted at $rt(T)$, with $\Delta(\tilde{Q}) \leq \Delta(T)$. 4) For any tree T' in $T \setminus \tilde{C}$, only the two sentinels of T' are incident to a vertex in \tilde{C} . Also, $rt(T')$ is incident only to its parent in T and $l(T')$ is incident only to its left-most child in T , unless it is a leaf in T .

Intuitively, part (4) of this proposition shows that the Procedure CV “slices” the tree in a “path-like” fashion, i.e., in a way analogous to the partition of ϑ_n into intervals. (See Fig. 4 in Appendix B for an illustration.)

1-Spanners with Low Diameter. Consider an n -vertex (weighted) tree T , and let M_T be the tree metric induced by T . In this section we devise a construction $\mathcal{H}_k(n)$ of 1-spanners for M_T with comparable monotone diameter $\bar{\Delta}(n) = \bar{\Delta}(\mathcal{H}_k(n))$ in the range $\Omega(\alpha(n)) = \bar{\Delta}(n) = O(\log n)$. Both in this construction and in the one with high diameter presented in the sequel, all edges of the original tree T are added to the spanner.

Let k be a fixed parameter such that $4 \leq k \leq n/2 - 1$, and set $d = n/k$. (Observe that $d > 2$.) To select the set \tilde{C} of cut-vertices, we invoke the procedure CV on the input (T, rt) and d . Set $C = CV((T, rt), d)$ and $\tilde{C} = C \cup \{rt(T), l(T)\}$. Since $k \geq 4$, it holds that $2d = 2n/k < n$. Denote the subtrees in the forest $T \setminus \tilde{C}$ by T_1, T_2, \dots, T_p . By Proposition 1, $|\tilde{C}| \leq k + 1$, and each tree T_i in $T \setminus \tilde{C}$ has size less than $2n/k$. Observe that $\sum_{i=1}^p |T_i| = n - |\tilde{C}| \geq n - k - 1$, implying that $p \geq (n - k - 1)/(2n/k) \geq k/4$.

To connect the set \tilde{C} of cut-vertices, the algorithm first constructs the tree $\tilde{Q} = \mathcal{Q}(T, \tilde{C})$. Note that \tilde{Q} inherits the tree structure of T , i.e., for any two points u and v in \tilde{C} , u is an ancestor of v in \tilde{Q} iff it is its ancestor in T . Consequently, any 1-spanner path in \tilde{Q} between two arbitrary comparable⁷ points is also a 1-spanner path for them in the original tree T . The algorithm proceeds by building

⁷ This may not hold true for two points that are not comparable, as their least common ancestor may not belong to \tilde{Q} .

a 1-spanner for \tilde{Q} via one of the aforementioned generalized constructions from [15, 2, 8, 37] (henceforth, *tree-spanner*). In other words, $O(k)$ edges between cut-vertices are added to the spanner $\mathcal{H}_k(n)$ to guarantee that the monotone distance in the spanner between any two comparable cut-vertices is $O(\alpha(k))$. Then the algorithm adds to the spanner $\mathcal{H}_k(n)$ edges that connect each of the two sentinels to all other cut-vertices. (In fact, the leaf $l(T)$ needs not be connected to all cut-vertices, but rather only to those which are its ancestors in T .) Finally, the algorithm calls itself recursively for each of the subtrees T_1, T_2, \dots, T_p . At the bottom level of the recursion, i.e., when $n < 2k + 2$, the algorithm uses the tree-spanner to connect all points, and, in addition, it adds to the spanner edges that connect both sentinels $rt(T)$ and $l(T)$ to all the other points.

The properties of the spanner $\mathcal{H}_k(n)$ are summarized in the next theorem.

Theorem 3 (See proof in Appendix D). *For any tree metric M_T and a parameter k , there exists a 1-spanner $\mathcal{H}_k(n)$ with $O(n)$ edges, maximum degree at most $\Delta(T) + 2k$, diameter $O(\log_k n + \alpha(k))$ and lightness $O(k \log_k n)$.*

We remark that the maximum degree $\Delta(\mathcal{H})$ of the spanner $\mathcal{H} = \mathcal{H}_k(n)$ cannot be in general smaller than the maximum degree $\Delta(T)$ of the original tree. Indeed, consider a unit weight star T with edge set $\{(rt, v_1), (rt, v_2), \dots, (rt, v_{n-1})\}$. Obviously, any spanner \mathcal{H} for M_T with $\Delta(\mathcal{H}) < n - 1$ distorts the distance between the root rt and some other vertex.

1-Spanners with High Diameter. The next theorem gives our construction of 1-spanners for M_T with comparable monotone diameter in the range $\Omega(\log n)$.

Theorem 4 (See proof in Appendix E). *For any tree metric M_T and a parameter k , there exists a 1-spanner with $O(n)$ edges, degree at most $2\Delta(T)$, diameter $O(k \log_k n)$ and lightness $O(\log_k n)$. This 1-spanner is a planar graph.*

3 Euclidean Spanners

In this section we demonstrate that our 1-spanners for tree metrics can be used for constructing Euclidean spanners and spanners for doubling metrics.

We start with employing the Dumbbell Theorem of [4] in conjunction with our 1-spanners for tree metrics to construct our Euclidean spanners.

Theorem 5. (“Dumbbell Theorem”, Theorem 2 in [4]) *Given a set S of n points in \mathbb{R}^d and a parameter $\epsilon > 0$, there exists a forest \mathcal{D} consisting of $O(1)$ rooted binary trees of size $O(n)$ each, having the following properties: 1) For each tree in \mathcal{D} , there is a 1-1 correspondence between the leaves of this tree and the points of S . 2) Each internal vertex in the tree has a unique representative point, which can be selected arbitrarily from the points in any of its descendant leaves. 3) Given any two points $u, v \in S$, there is a tree in \mathcal{D} , so that the path formed by walking from representative to representative along the unique path in that tree between these vertices, is a $(1 + \epsilon)$ -spanner path for u and v .*

For each dumbbell tree in \mathcal{D} , we use the following representative assignment from [4]. Leaf labels are propagated up the tree. An internal vertex chooses to itself one of the propagated labels and propagates the other one up the tree. Each label is used at most twice, once at a leaf and once at an internal vertex. Any

label assignment induces a weight function over the edges of the dumbbell tree in the obvious way. (The weight of an edge is set to be the Euclidean distance between the representatives corresponding to the two endpoints of that edge.) Arya et al. [4] proved that the lightness of dumbbell trees is always $O(\log n)$, regardless of which representative assignment is chosen for the internal vertices.

Next, we describe our construction of Euclidean spanners with diameter in the range $\Omega(\alpha(n)) = \Lambda = O(\log n)$. For each dumbbell tree $DT_i \in \mathcal{D}$, denote by M_i the $O(n)$ -point tree metric induced by DT_i . To obtain our construction of $(1 + \epsilon)$ -spanners with low diameter, we set $k = n^{1/\Lambda}$, and build the 1-spanner construction $\mathcal{H}^i = \mathcal{H}_k^i(O(n))$ of Theorem 3 for each of the tree metrics M_i . Then we translate each \mathcal{H}^i to be a spanning subgraph $\hat{\mathcal{H}}^i$ of S in the obvious way. Let $\mathcal{E}_k(n)$ be the spanner obtained from the union of all the graphs $\hat{\mathcal{H}}^i$.

It is easy to see that the number of edges in $\mathcal{E}_k(n)$ is $O(n)$.

Next, we show that $\Delta(\mathcal{E}_k(n)) = O(\log_k n + \alpha(k))$. By the Dumbbell Theorem, for any pair of points $u, v \in S$, there exists a dumbbell tree DT_i , so that the unique path $P_{u,v}$ between u and v in DT_i is a $(1 + \epsilon)$ -spanner path. Theorem 3 implies that there is a 1-spanner path P in \mathcal{H}^i between u and v that consists of $O(\log_k n + \alpha(k))$ hops. By the triangle inequality, the weight of the corresponding translated path \hat{P} in $\hat{\mathcal{H}}^i$ is no greater than the weight of $P_{u,v}$. Hence, \hat{P} is a $(1 + \epsilon)$ -spanner path for u and v that consists of $O(\log_k n + \alpha(k))$ hops.

We proceed by showing that $\Delta(\mathcal{E}_k(n)) = O(k)$. Since each dumbbell tree DT_i is binary, theorem 3 implies that $\Delta(\mathcal{H}^i) = O(k)$. Recall that each label is used at most twice in DT_i , and so $\Delta(\hat{\mathcal{H}}^i) \leq 2\Delta(\mathcal{H}^i) = O(k)$. The union of $O(1)$ such graphs will also have maximum degree $O(k)$.

Finally, we argue that the lightness $\Psi(\mathcal{E}_k(n))$ of $\mathcal{E}_k(n)$ is $O(k \log_k n \log n)$. Consider a dumbbell tree DT_i . Recall that the lightness of all dumbbell trees is $O(\log n)$, and so $w(DT_i) = O(\log n)w(MST(S))$. By Theorem 3, the weight $w(\mathcal{H}^i)$ of \mathcal{H}^i is at most $O(k \log_k n)w(DT_i)$. By the triangle inequality, the weight of each edge in $\hat{\mathcal{H}}^i$ is no greater than the corresponding weight in \mathcal{H}^i , implying that $w(\hat{\mathcal{H}}^i) \leq w(\mathcal{H}^i) \leq O(k \log_k n \log n)w(MST(S))$. The union of $O(1)$ such graphs will also have weight $O(k \log_k n \log n)w(MST(S))$.

To obtain our construction of Euclidean spanners in the range $\Lambda = \Omega(\log n)$, we use our 1-spanners for tree metrics from Theorem 4 instead of Theorem 3.

Corollary 1. *For any set S of n points in \mathbb{R}^d , any $\epsilon > 0$ and a parameter k , there exists a $(1 + \epsilon)$ -spanner with $O(n)$ edges, maximum degree $O(k)$, diameter $O(\log_k n + \alpha(k))$ and lightness $O(k \log_k n \log n)$. There also exists a $(1 + \epsilon)$ -spanner with degree $O(1)$, diameter $O(k \log_k n)$ and lightness $O(\log_k n \log n)$.*

In Appendix G we show that the lightness of well-separated pair constructions for random point sets in the unit cube is (w.h.p.) $O(1)$. Also, the lightness of well-separated pair constructions provides an asymptotic upper bound on the lightness of dumbbell trees. We derive the following result as a corollary.

Corollary 2. *For any set S of n points that are chosen independently and uniformly at random from the unit cube, any $\epsilon > 0$ and a parameter k , there exists a $(1 + \epsilon)$ -spanner with $O(n)$ edges, maximum degree $O(k)$, diameter $O(\log_k n + \alpha(k))$ and lightness (w.h.p.) $O(k \log_k n)$. There also exists a $(1 + \epsilon)$ -spanner with maximum degree $O(1)$, diameter $O(k \log_k n)$ and lightness (w.h.p.) $O(\log_k n)$.*

Chan et al. [13] showed that for any doubling metric (X, δ) there exists a $(1 + \epsilon)$ -spanner with constant maximum degree. On the way to this result they

proved the following lemma, which we employ in conjunction with our 1-spanners for tree metrics to construct our spanners for doubling metrics.

Lemma 1 (Lemma 3.1 in [13]). *For any doubling metric (X, δ) , there exists a collection \mathcal{T} of $m = O(1)$ spanning trees for (X, δ) , $\mathcal{T} = \{\tau_1, \tau_2, \dots, \tau_m\}$, that satisfies the following two properties: i) For each index $i \in [m]$, the maximum degree $\Delta(\tau_i)$ of the tree τ_i is constant, i.e., $\Delta(\tau_i) = O(1)$. ii) For each pair of points $x, y \in X$ there exists an index $i \in [m]$ such that $\text{dist}_{\tau_i}(x, y) = O(1)\delta(x, y)$.*

To obtain our spanners for doubling metrics we start with constructing the collection $\mathcal{T} = \{\tau_1, \tau_2, \dots, \tau_m\}$ of spanning trees with properties listed in Lemma 1. Next, we apply Theorem 3 with some parameter k to construct a 1-spanner $\mathcal{Z}^i = \mathcal{Z}_k^i(n)$ for the tree metric induced by the i th tree τ_i in \mathcal{T} , for each $i \in [m]$. Our spanner \mathcal{Z} is set to be the union of all the 1-spanners \mathcal{Z}_i , i.e., $\mathcal{Z} = \bigcup_{i=1}^m \mathcal{Z}_i$. We summarize the properties of the resulting spanner \mathcal{Z} in the next statement.

Corollary 3 (See proof in Appendix F). *For any n -point doubling metric (X, δ) and a parameter k , there exists an $O(1)$ -spanner \mathcal{Z} with $O(n)$ edges, maximum degree $O(k)$ and diameter $O(\log_k n + \alpha(k))$.*

Acknowledgments: We are grateful to Sunil Arya, David Mount and Michiel Smid for helpful discussions.

References

1. Agarwal, P. K., Wang, Y., Yin, P.: Lower bound for sparse Euclidean spanners. In: Proc. of 16th SODA, pp. 670–671 (2005).
2. Alon, N., Schieber, B.: Optimal preprocessing for answering on-line product queries. Manuscript (1987).
3. Althöfer, I., Das, G., Dobkin, D. P., Joseph, D., Soares, J.: On sparse spanners of weighted graphs. *Discrete & Computational Geometry* 9, 81–100 (1993).
4. Arya, S., Das, G., Mount, D. M., Salowe, J. S., Smid, M. H. M.: Euclidean spanners: short, thin, and lanky. In: Proc. of 27th STOC, pp. 489–498 (1995).
5. Arya, S., Mount, D. M., Netanyahu, N. S., Silverman, R., Wu, A. Y.: An optimal algorithm for approximate nearest neighbor searching in fixed dimensions. *J. ACM* 45(6), 819–923 (1998).
6. Arya, S., Smid, M. H. M.: Efficient construction of a bounded degree spanner with low weight. *Algorithmica* 17(1), 33–54 (1997).
7. de Berg, M., Cheong, O., van Kreveld, M., Overmars, M.: *Computational Geometry: Algorithms and Applications*, third edition. Springer-Verlag, Heidelberg (2008).
8. Bodlaender, H. L., Tel, G., Santoro, N.: Trade-offs in non-reversing diameter. *Nord. J. Comput.* 1(1), 111–134 (1994).
9. Callahan, P. B., Kosaraju, S. R.: A decomposition of multi-dimensional point-sets with applications to k -nearest-neighbors and n -body potential fields. In: Proc. of 24th STOC, pp. 546–556 (1992).
10. Callahan, P. B., Kosaraju, S. R.: Faster algorithms for some geometric graph problems in higher dimensions. In: Proc. of 4th SODA, pp. 291–300 (1993).
11. Chan, T. M.: Well-separated pair decomposition in linear time?. *Inf. Process. Lett.* 107(5), 138–141 (2008).
12. Chan, H. T.-H., Gupta, A.: Small hop-diameter sparse spanners for doubling metrics. In: Proc. of 17th SODA, pp. 70–78 (2006).
13. Chan, H. T.-H., Gupta, A., Maggs, B. M., Zhou, S.: On hierarchical routing in doubling metrics. In: Proc. of 16th SODA, pp. 762–771 (2005).

14. Chandra, B.: Constructing sparse spanners for most graphs in higher dimensions. *Inf. Process. Lett.* 51(6), 289–294 (1994).
15. Chazelle, B.: Computing on a free tree via complexity-preserving mappings. *Algorithmica* 2, 337–361 (1987).
16. Chazelle, B., Rosenberg, B.: The complexity of computing partial sums off-line. *Int. J. Comput. Geom. Appl.* 1, 33–45 (1991).
17. Chew, L. P.: There is a planar graph almost as good as the complete graph. In: *Proc. of 2nd SOCG*, pp. 169–177 (1986).
18. Das, G., Narasimhan, G.: A fast algorithm for constructing sparse Euclidean spanners. In: *Proc. of 10th SOCG*, pp. 132–139 (1994).
19. Das, G., Narasimhan, G., Salowe, J. S.: A new way to weigh malnourished Euclidean graphs. In: *Proc. of 6th SODA*, pp. 215–222 (1995).
20. Dinitz, Y., Elkin, M., Solomon, S.: Shallow-low-light trees, and tight lower bounds for Euclidean spanners. In: *Proc. of 49th FOCS*, pp. 519–528 (2008).
21. Fischer, J., Har-Peled, S.: Dynamic well-separated pair decomposition made easy. In: *Proc. of 17th CCCG*, pp. 235–238 (2005).
22. Frederickson, G. N.: A data structure for dynamically maintaining rooted trees. In: *Proc. of 4th SODA*, pp. 175–184 (1993).
23. Gottlieb, L., Roditty, L.: An optimal dynamic spanner for doubling metric spaces. In: *Proc. of 16th ESA*, pp. 478–489 (2008).
24. Gudmundsson, J., Levkopoulos, C., Narasimhan, G., Smid, M. H. M.: Approximate distance oracles for geometric graphs. In: *Proc. of 13th SODA*, pp. 828–837 (2002).
25. Gudmundsson, J., Levkopoulos, C., Narasimhan, G., Smid, M. H. M.: Approximate distance oracles for geometric spanners. *ACM Transactions on Algorithms* 4(1), (2008).
26. Gudmundsson, J., Narasimhan, G., Smid, M. H. M.: Fast pruning of geometric spanners. In: *Proc. of 22nd STACS*, pp. 508–520 (2005).
27. Hassin, Y., Peleg, D.: Sparse communication networks and efficient routing in the plane. In: *Proc. of 19th PODC*, pp. 41–50 (2000).
28. Keil, J. M., Gutwin, C. A.: Classes of graphs which approximate the complete Euclidean graph. *Discrete & Computational Geometry* 7, 13–28 (1992).
29. Lenhof, H. P., Salowe, J. S., Wrege, D. E.: New methods to mix shortest-path and minimum spanning trees. *Manuscript* (1994).
30. Mansour, Y., Peleg, D.: An approximation algorithm for min-cost network design. *DIMACS Series in Discr. Math and TCS* 53, 97–106 (2000).
31. Narasimhan, G., Smid, M.: *Geometric Spanner Networks*. Cambridge University Press (2007).
32. Pătrașcu, M., Demaine, E. D.: Tight bounds for the partial-sums problem. In: *Proc. of 15th SODA*, pp. 20–29 (2004).
33. Peleg, D., Ullman, J. D.: An optimal synchronizer for the hypercube. *SIAM J. Comput.* 18(4), 740–747 (1989).
34. Rao, S., Smith, W. D.: Approximating geometrical graphs via “spanners” and “banyans”. In: *Proc. of 30th STOC*, pp. 540–550 (1998).
35. Smid, M. H. M. Private communication.
36. Tarjan, R. E.: Applications of path compression on balanced trees. *J. ACM* 26(4), 690–715 (1979).
37. Thorup, M.: Parallel shortcutting of rooted trees. *J. Algorithms* 23(1), 139–159 (1997).
38. Yao, A. C.: Space-time tradeoff for answering range queries. In: *Proc. of 14th STOC*, pp. 128–136 (1982).

Appendix

A An Overview of Our and Previous Techniques

The starting point for our construction is the construction of Arya et al. [4] that achieves diameter $O(\log n)$, lightness $O(\log^2 n)$ and constant degree. The construction of [4] is built in two stages. First, a construction for the 1-dimensional case is devised. Then the 1-dimensional construction is extended to arbitrary constant dimension. For 1-dimensional spaces Arya et al. [4] start with devising a construction of 1-spanners with diameter, lightness and degree all bounded by $O(\log n)$. This construction is quite simple; it is essentially a flattened version of a deterministic skip-list. Next, by a more involved argument they show that the degree can be reduced to $O(1)$, at the expense of increasing the stretch parameter from 1 to $1 + \epsilon$. Finally, the generalization of their construction to point sets in the plane (or, more generally, to \mathbb{R}^d) is far more involved. Specifically, to this end Arya et al. [4] employed two main tools. The first one is the *dumbbell trees*, the theory of which was developed by Arya et al. in the same paper [4]. (See also Chapter 11 of [31].) The second one is the bottom-up clustering technique that was developed by Frederickson [22] for topology trees. Roughly speaking, the *Dumbbell Theorem* of [4] states that for every point set S , one can construct a forest \mathcal{D} of $O(1)$ dumbbell trees in which there exists a tree $T \in \mathcal{D}$ for every pair x, y of points from S , such that the distance $dist_T(x, y)$ between x and y in T at most $(1 + \epsilon)$ times their Euclidean distance $\|x - y\|$. Arya et al. employ Frederickson's clustering technique on each of these $O(1)$ trees to obtain their ultimate spanner.

Similarly to [4], we start with devising a construction of 1-spanners for the 1-dimensional case. However, our construction achieves both diameter and lightness at most $O(\log n)$, in conjunction with degree at most 4. (Note that [4] paid for decreasing the degree from $O(\log n)$ to $O(1)$ by increasing the stretch of the spanner from 1 to $1 + \epsilon$. Our construction achieves stretch 1 in conjunction with logarithmic diameter and lightness, and with constant degree.) Moreover, our construction is far more general, as it provides the *entire suite* of all possible values of diameter, lightness and degree, and it is *optimal up to constant factors* in the entire range of parameters. We then proceed to extending it to arbitrary tree metrics. Finally, we employ the dumbbell trees of Arya et al. [4]. Specifically, we construct our 1-spanners for the metrics induced by each of these dumbbell trees, and return their union as our ultimate spanner. As a result we obtain a *unified* construction of Euclidean spanners that achieves near-optimal tradeoffs in the entire range of parameters. We remark that it is unclear whether the construction of Arya et al. [4] can be extended to provide additional combinations between diameter and lightness other than $O(\log n)$ and $O(\log^2 n)$, respectively; roughly speaking, the logarithms there come from the number of levels in Frederickson's topology trees. In particular, the construction of Arya et al. [4] that achieves diameter $O(\alpha(n))$ and arbitrarily large lightness and degree is based on completely different ideas. On the other hand, our construction yields a stronger result (diameter $O(\alpha(n))$, lightness and degree $n^{O(1/\alpha(n))}$), and this result is obtained by substituting a different parameter into one of our tradeoffs. Moreover, our construction is much simpler and much more modular than that of [4]. In

particular, it does not employ Frederickson’s bottom-up clustering technique, but rather constructs 1-spanners for dumbbell trees directly.

Also, our construction of 1-spanners for tree metrics (that we use for dumbbell trees) is fundamentally different from the previous constructions due to [38, 15, 2, 8, 37]. In particular, the techniques of [15, 2, 8, 37] for generalizing constructions of 1-spanners from 1-dimensional metrics to general tree metrics ensure that the diameter of the resulting spanners is not (much) greater than the diameter in the 1-dimensional case. However, the degree and/or lightness of spanners for tree metrics that are obtained by these techniques may be arbitrarily large. To overcome this obstacle we adapt the techniques of [15, 2, 8] to our purposes. Next, we overview this adaptation. A central ingredient in the generalization techniques of [15, 2, 8] is a tree decomposition procedure. Given an n -vertex rooted tree (T, rt) and a parameter k , this procedure computes a set C of $O(k)$ cut-vertices. This set has the property that removing all vertices of C from the tree T decomposes T into a collection \mathcal{F} of trees, so that each tree $\tau \in \mathcal{F}$ contains $O(n/k)$ vertices. This decomposition induces a tree $\mathcal{Q} = \mathcal{Q}(\tau, C)$ over the vertex set $C \cup \{rt\}$ in a natural way: a cut-vertex $w \in C$ is defined to be a child of its closest ancestor in T which is a cut-vertex. For our purposes, it is crucial that the degree of the tree \mathcal{Q} will not be (much) greater than the degree of T . In addition, it is essential that each tree $\tau \in \mathcal{F}$ will be incident to at most $O(1)$ cut-vertices. We devise a novel decomposition procedure that guarantees these two basic properties. Intuitively, our decomposition procedure “slices” the tree in a “path-like” fashion. This path-like nature of our decomposition enables us to keep the degree and lightness of our construction for tree metrics (essentially) as small as in the 1-dimensional case.

B Table and Figures

	[4]	[4]	New	New	New	New	New	New
	I,II		I	I	I	II	II	II
k	1		$\log^\delta n$	$2^{\sqrt{\log n}}$	$n^{1/\alpha(n)}$	$\log^\delta n$	$2^{\sqrt{\log n}}$	n^ζ
Δ	1	n	$\log^\delta n$	$2^{\sqrt{\log n}}$	$n^{1/\alpha(n)}$	1	1	1
A	$\log n$	$\alpha(n)$	$\frac{\log n}{\log \log n}$	$\sqrt{\log n}$	$\alpha(n)$	$\log^{1+\delta} n$	$2^{O(\sqrt{\log n})}$	n^ζ
Ψ	$\log^2 n$	n	$\log^{2+\delta} n$	$2^{O(\sqrt{\log n})}$	$n^{O(1/\alpha(n))}$	$\frac{\log^2 n}{\log \log n}$	$\log^{3/2} n$	$\log n$

Table 1. A concise comparison of previous and new results. Each column corresponds to a set of parameters that can be achieved simultaneously. For each column the first row indicates whether the result is new or due to [4]. (The first column is due to [4], but can also be achieved from both our tradeoffs.) For new results, the second row indicates whether it is obtained by the first (I) or the second (II) tradeoff. (The first tradeoff is degree $O(k)$, diameter $O(\log_k n + \alpha(k))$, and lightness $O(k \log_k n \log n)$. The second tradeoff is degree $O(1)$, diameter $O(k \log_k n)$ and lightness $O(\log_k n \log n)$.) The third row indicates the value of k that is substituted in the corresponding tradeoff. The next three rows indicate the resulting degree (Δ), diameter (A) and lightness (Ψ). The number of edges used in all constructions is $O(n)$. To save space, the O notation is omitted everywhere except for the exponents. The letters δ and ζ stand for arbitrarily small positive constants.




Fig. 1. The construction for $k = 2$. Only the first level of the recursion is illustrated. (Path-edges are not depicted in the figure.) The cut-vertex $r_1 = v_{n/2}$ is connected via edges to the two sentinels v_1 and v_n . The construction proceeds recursively for each of the two intervals I_1 and I_2 .




Fig. 2. The constructions $H_k(n)$ and $H'_k(n)$ for a general k . Only the first level of the recursion is illustrated. (Path-edges are not depicted in the figure.) For $H_k(n)$, all the cut-vertices are connected via the list-spanner, and, in addition, each of the two sentinels is connected to all other k cut-vertices. For $H'_k(n)$, each cut-vertex r_{i-1} is connected to the next cut-vertex r_i in line, $i \in [k]$.

C Proof of Proposition 1

In this section we prove Proposition 1 from Sect. 2.3.

For convenience, we define $n_i = |T_{c_i(b)}|$, for each $i \in [ch(b)]$.

The first assertion of Proposition 1 follows immediately from the next lemma.

Lemma 2. For $n \geq 2d$, $|C| \leq (n/d) - 1$.

Proof. The proof is by induction on $n = |T|$. Throughout the proof we use the observation that for $n < 2d$, it holds that $C = \emptyset$.

Basis: $2d \leq n < 3d$. Fix an index $i \in [ch(b)]$. Since b is balanced, we have

$$n_i \leq n_1 \leq n - d < 2d,$$




Fig. 3. A rooted tree (T, rt) with $n = 18$ vertices $v_1 = rt, v_2, \dots, v_{18}$. The first 6-balanced vertex along $P(rt)$ is v_2 . The procedure CV on input (T, rt) and $d = 6$ returns the subset $\{v_2, v_8\}$.




Fig. 4. A “path-like” decomposition of the tree T into subtrees $T^{(1)}, T^{(2)}, \dots, T^{(6)}$. The two cut-vertices in the figure are depicted by white dots, whereas the twelve sentinels of the subtrees $T^{(1)}, T^{(2)}, \dots, T^{(6)}$ are depicted by black dots. Similarly to the 1-dimensional case, each subtree $T^{(i)}$ is incident to at most two cut-vertices. Edges in T that connect sentinels of subtrees with cut-vertices are depicted by dashed lines.

implying that $C_i = \emptyset$. It follows that $C = \bigcup_{i=1}^{ch(b)} C_i \cup \{b\} = \{b\}$, and so $|C| = 1 \leq (n/d) - 1$.

Induction Step: We assume the correctness of the statement for all smaller values of n and prove it for n . Let I be the set of all indices i in $[ch(b)]$ for which $n_i \geq 2d$. Observe that for each $i \in [ch(b)] \setminus I$, $C_i = \emptyset$, and by induction hypothesis, for each $i \in I$, $|C_i| \leq (n_i/d) - 1$. By construction, the vertex sets $C_1, C_2, \dots, C_{ch(b)}$ and $\{b\}$ are pairwise disjoint, and $C = \bigcup_{i=1}^{ch(b)} C_i \cup \{b\}$. Hence

$$|C| = \sum_{i=1}^{ch(b)} |C_i| + |\{b\}| = \sum_{i \in I} |C_i| + 1 \leq \sum_{i \in I} ((n_i/d) - 1) + 1. \quad (1)$$

The analysis splits into three cases depending on the size of $|I|$.

Case 1: $|I| = 0$. Equation (1) yields $|C| \leq 1 \leq (n/d) - 1$.

Case 2: $|I| = 1$. By construction, $n_1 \geq n_i$, for each $i \in [ch(b)]$, implying that $I = \{1\}$. Since b is balanced, $n_1 \leq n - d$, and so (1) yields $|C| \leq (n_1/d) - 1 + 1 \leq (n/d) - 1$.

Case 3: $|I| \geq 2$. Clearly, $\sum_{i \in I} n_i \leq n - 1$, and so (1) yields

$$|C| \leq \sum_{i \in I} ((n_i/d) - 1) + 1 = \sum_{i \in I} (n_i/d) - |I| + 1 \leq (n-1)/d - 2 + 1 \leq (n/d) - 1.$$

□

We denote by $T \setminus C$ the forest obtained from T by removing all vertices in C along with the edges that are incident to them. The second assertion of Proposition 1 follows immediately from the next lemma.

Lemma 3. *The size of any tree in the forest $T \setminus C$ is smaller than $2d$.*

Proof. We start by proving the following claim.

Claim. Let \overline{T}_b be the tree obtained from T by removing the subtree T_b . Then $|\overline{T}_b| < d$.

Proof. If $b = rt$, then \overline{T}_b is empty and the result is immediate. Otherwise, consider the parent $\pi(b)$ of b in T . Since b is the first balanced vertex along $P(rt)$, $\pi(b)$ is non-balanced, and so $|T_b| = |T_{c_1(\pi(b))}| > n - d$. Hence $|\overline{T}_b| = n - |T_b| < d$, which completes the proof of Claim C. □

The proof of the lemma proceeds by induction on $n = |T|$. The basis $n < 2d$ is trivial.

Induction Step: We assume the correctness of the statement for all smaller values of n and prove it for n . In this case, C is non-empty, and so the size of any tree in the forest $T \setminus C$ is strictly smaller than n . Consider a tree T' in the forest $T \setminus C$. Observe that for $n \geq 2d$, we have $T \setminus C = \bigcup_{i=1}^{ch(b)} (T_{c_i(b)} \setminus C_i) \cup \{\overline{T}_b\}$. Consequently, either $T' = \overline{T}_b$, or it belongs to the forest $T_{c_i(b)} \setminus C_i$, for some index $i \in [ch(b)]$. In the former case, the size bound follows from Claim C, whereas in the latter case it follows from the induction hypothesis. Lemma 3 follows. □

Observe that the subset C induces a forest $Q = \mathcal{Q}(T, C)$ over C in the natural way: a vertex $v \in C$ is defined to be a child of its closest ancestor in T that belongs to C . Observe that for $n < 2d$, $C = \emptyset$, and so $Q = \emptyset$. The third assertion of Proposition 1 follows from the following lemma.

Lemma 4. For $n \geq 2d$, \mathcal{Q} is a spanning tree of C rooted at $b = B(rt)$, with $\Delta(\mathcal{Q}) \leq \Delta(T)$.

Proof. We prove by induction on $n = |T|$, $n \geq 2d$, the next stronger statement: \mathcal{Q} is a spanning tree of C rooted at $b = B(rt)$, such that for each vertex v in C , the number of children $ch_{\mathcal{Q}}(v)$ of v in \mathcal{Q} is no greater than the corresponding number $ch(v)$ in T .

Basis: $2d \leq n < 3d$. In this case $C = \{b\}$, and so \mathcal{Q} consists of a single root vertex b .

Induction Step: We assume the correctness of the statement for all smaller values of n and prove it for n . Let I be the set of all indices i in $[ch(b)]$ for which $n_i \geq 2d$, and write $I = \{i_1, i_2, \dots, i_I\}$. Note that for each $i \in [ch(b)] \setminus I$, $C_i = \emptyset$, and so $\mathcal{Q}(T_{c_i(b)}, C_i)$ is an empty tree. By the induction hypothesis, for each $i \in I$, $\mathcal{Q}_i = \mathcal{Q}(T_{c_i(b)}, C_i)$ is a spanning tree of C_i rooted at $b_i = B(c_i(b))$ in which the number of children of each vertex is no greater than the corresponding number in $T_{c_i(b)}$. By construction, the only children of b in \mathcal{Q} are the roots $b_{i_1}, b_{i_2}, \dots, b_{i_I}$ of the non-empty trees $\mathcal{Q}_{i_1}, \mathcal{Q}_{i_2}, \dots, \mathcal{Q}_{i_I}$, respectively, and so $ch_{\mathcal{Q}}(b) = |I| \leq ch(b)$. Also, b has no parent in \mathcal{Q} , and so it is the root of \mathcal{Q} . \square

The next lemma implies the fourth assertion of Proposition 1. The proof of this lemma follows similar lines as those in the proof of Lemma 3, and is omitted.

Lemma 5. For any tree T' in $T \setminus C$, no other vertex in T' other than its two sentinels is incident to a vertex in C . Moreover, $rt(T')$ is incident only to its parent in T , unless it is the root of T , and $l(T')$ is incident only to its left-most child in T , unless it is a leaf in T .

Lemmas 2, 3, 4 and 5 imply Proposition 1.

D Proof of Theorem 3

In this section we prove Theorem 3 from Sect. 2.3.

We denote by $E(n)$ the number of edges in $\mathcal{H}_k(n)$, excluding edges of T . Clearly, $E(n)$ satisfies the recurrence $E(n) \leq O(k) + \sum_{i=1}^p E(|T_i|)$, with the base condition $E(n) = O(n)$, for $n < 2k + 2$. Recall that for each $i \in [p]$, $|T_i| \leq 2n/k$, and $p \geq k/4$. Also, $\sum_{i=1}^p |T_i| = n - |\tilde{C}| \leq n - 2$. It is easy to verify that $E(n) = O(n)$. Denote by $\Delta(n)$ the maximum degree of a vertex in $\mathcal{H}_k(n)$, excluding edges of T . Since $|\tilde{C}| \leq k + 1$, $\Delta(n)$ satisfies the recurrence $\Delta(n) \leq \max\{k, \Delta(2n/k)\}$, with the base condition $\Delta(n) \leq 2k$, for $n < 2k + 2$, yielding $\Delta(n) \leq 2k$. Including edges of the tree T , the number of edges increases by $n - 1$ units and the maximum degree increases by at most $\Delta(T)$ units.

Next, we show that the weight $w(\mathcal{H}_k(n))$ of the spanner $\mathcal{H}_k(n)$ satisfies $w(\mathcal{H}_k(n)) = O(k \log_k n)w(T)$. To this end, we extend the notion of *load* defined in [1]⁸ for 1-dimensional spaces to general tree metrics. Consider an edge $e' = (v, w)$ connecting two arbitrary points in M_T , and a tree-edge $e \in E(T)$. The edge e' is said to *load* e if the unique path in T between the endpoints v and w of e' traverses e . For a spanning subgraph H of M_T , the number of edges $e' \in E(H)$ that load a tree-edge e is called the load of e by H and it is denoted

⁸ Agarwal et al. [1] used a slightly different notion called *covering*. The notion of load as defined above was introduced in [20], but the two notions are very close.

$\chi(e) = \chi_H(e)$. The load of H , $\chi(H)$, is the maximum load of a tree-edge by H . By double-counting,

$$\begin{aligned} w(H) &= \sum_{e' \in E(H)} w(e') = \sum_{e' \in E(H)} \sum_{\{e \in E(T) : e \text{ loaded by } e'\}} w(e) \\ &= \sum_{e \in E(T)} \chi_H(e) w(e) \leq \chi(H) \sum_{e \in E(T)} w(e) = \chi(H) w(T), \end{aligned} \quad (2)$$

implying that $w(H)/w(T) \leq \chi(H)$. Thus it suffices to provide an upper bound of $O(k \log_k n)$ on the load $\chi(n) = \chi(\mathcal{H}_k(n))$ of $\mathcal{H}_k(n)$. Since $\mathcal{H}_k(n)$ contains only $O(k)$ edges that connect cut-vertices, each subtree in the forest $T \setminus \tilde{C}$ is loaded by at most $O(k)$ such edges. In addition, $\mathcal{H}_k(n)$ contains all edges of the original tree T . These edges contribute an additional unit of load to each subtree in the forest $T \setminus \tilde{C}$. Hence $\chi(n)$ satisfies the recurrence $\chi(n) \leq O(k) + \chi(2n/k)$, with the base condition $\chi(n) = O(n)$, for $n < 2k + 2$, yielding $\chi(n) = O(k \log_k n)$.

Finally, we show that $\bar{A}(n) = \bar{A}(\mathcal{H}_k(n)) = O(\log_k n + \alpha(k))$. The *leaf radius* $\hat{R}(n)$ of $\mathcal{H}_k(n)$ is defined as the maximum monotone distance between the leftmost vertex $l(T)$ in T and one of its ancestors in T . By Proposition 1, similarly to the 1-dimensional case, $\hat{R}(n)$ satisfies the recurrence $\hat{R}(n) \leq 2 + \hat{R}(2n/k)$, with the base condition $\hat{R}(n) = 1$, for $n < 2k + 2$, yielding $\hat{R}(n) = O(\log_k n)$. Similarly, we define the *root radius* $\check{R}(n)$ as the maximum monotone distance between the root $rt(T)$ of T and some other point in T . By the same argument we get $\check{R}(n) = O(\log_k n)$. Applying again Proposition 1 and reasoning similar to the 1-dimensional case, we get that $\bar{A}(n) \leq \max\{\bar{A}(2n/k), O(\alpha(k)) + \hat{R}(2n/k) + \check{R}(2n/k)\}$, with the base condition $\bar{A}(n) = O(\alpha(n))$, for $n < 2k + 2$. Hence, $\bar{A}(n) = O(\log_k n + \alpha(k))$, and we are done.

E Proof of Theorem 4

In this section we devise a construction $\mathcal{H}'_k(n)$ of 1-spanners for M_T with comparable monotone diameter $\bar{A}'(n) = \bar{A}(\mathcal{H}'_k(n))$ in the range $\Omega(\log n)$. For a parameter k , the spanner $\mathcal{H}'_k(n)$ has $O(n)$ edges, maximum degree at most $2\Delta(T)$, diameter $O(k \log_k n)$ and lightness $O(\log_k n)$, proving Theorem 4 of Sect. 2.3.

The algorithm starts with constructing the tree $\tilde{\mathcal{Q}} = \mathcal{Q}(T, \tilde{C})$ that spans the set \tilde{C} of cut-vertices. All edges of $\tilde{\mathcal{Q}}$ are inserted into $\mathcal{H}'_k(n)$. Observe that the depth of $\tilde{\mathcal{Q}}$ is at most k , implying that any two comparable cut-vertices are connected via a 1-spanner path in $\tilde{\mathcal{Q}}$ that consists of at most k hops. Since $\tilde{\mathcal{Q}}$ inherits the tree structure of T , this path is also a 1-spanner path in T . Then the algorithm calls itself recursively for each of the subtrees T_1, T_2, \dots, T_p . At the bottom level of the recursion, i.e., when $n < 2k + 2$, the algorithm adds no additional edges to the spanner.

Similarly to the proof of Theorem 3 we get that the number of edges in $\mathcal{H}'_k(n)$ is $O(n)$. Denote by $\Delta'(n)$ the maximum degree of a vertex in $\mathcal{H}'_k(n)$, excluding edges of T . By the third assertion of Proposition 1, $\Delta(\tilde{\mathcal{Q}}) \leq \Delta(T)$, and so $\Delta'(n)$ satisfies the recurrence $\Delta'(n) \leq \max\{\Delta(T), \Delta'(2n/k)\}$, with the base condition $\Delta'(n) = 0$, for $n < 2k + 2$, yielding $\Delta'(n) \leq \Delta(T)$. It follows that the maximum degree $\Delta(\mathcal{H}'_k(n))$ of $\mathcal{H}'_k(n)$ is at most $2\Delta(T)$.

Next, we show that the load $\chi'(n) = \chi(\mathcal{H}'_k(n))$ of $\mathcal{H}'_k(n)$ is $O(\log_k n)$, which, by (2), implies that $\Psi(\mathcal{H}'_k(n)) = O(\log_k n)$. By the fourth assertion of Proposition

1, each subtree T_i in the forest $T \setminus \tilde{C}$ is loaded by at most one edge in \tilde{Q} , namely, the edge connecting the parent of $rt(T_i)$ in T and the left-most child of $l(T_i)$ in T , if exists. In addition, $\mathcal{H}'_k(n)$ contains all edges of the original tree T . These edges contribute an additional unit of load to each subtree in $T \setminus \tilde{C}$. Hence $\chi'(n)$ satisfies the recurrence $\chi'(n) \leq 2 + \chi'(2n/k)$, with the base condition $\chi'(n) \leq 1$, for $n < 2k + 2$, yielding $\chi'(n) = O(\log_k n)$.

By Proposition 1, similarly to the 1-dimensional case, the leaf radius $\hat{R}'(n)$ of $\mathcal{H}'_k(n)$ satisfies the recurrence $\hat{R}'(n) \leq k + \hat{R}'(2n/k)$, with the base condition $\hat{R}'(n) \leq n - 1$, for $n < 2k + 2$, yielding $\hat{R}'(n) = O(k \log_k n)$. Similarly, we get that $\tilde{R}'(n) = O(k \log_k n)$. Applying Proposition 1 and reasoning similar to the 1-dimensional case, we get that the comparable monotone diameter $\bar{A}'(n) = \bar{A}(\mathcal{H}'_k(n))$ of $\mathcal{H}'_k(n)$ satisfies the following recurrence $\bar{A}'(n) \leq \max\{\bar{A}'(2n/k), k + \tilde{R}'(2n/k) + \hat{R}'(2n/k)\}$, with the base condition $\bar{A}'(n) \leq n - 1$, for $n < 2k + 2$, yielding $\bar{A}'(n) = O(k \log_k n)$.

Finally, we remark that $\mathcal{H}'_k(n)$ is a planar graph, which completes the proof.

F Proof of Corollary 3

In this section we prove Corollary 3 from Sect. 3.

First, it is easy to see that the number of edges in \mathcal{Z} is $O(n)$.

To argue that \mathcal{Z} is an $O(1)$ -spanner consider a pair of points $x, y \in X$. By Lemma 1, there exists an index $i \in [m]$, such that $dist_{\tau_i}(x, y) = O(1)\delta(x, y)$. Also, since \mathcal{Z}^i is a 1-spanner for the metric induced by τ_i , it follows that $dist_{\mathcal{Z}^i}(x, y) = dist_{\tau_i}(x, y)$. Finally, since $\mathcal{Z}^i \subseteq \mathcal{Z}$, we conclude that $dist_{\mathcal{Z}}(x, y) \leq dist_{\mathcal{Z}^i}(x, y) \leq dist_{\tau_i}(x, y) = O(1)\delta(x, y)$. Observe also that $\Lambda(\mathcal{Z}^i) = O(\log_k n + \alpha(k))$, and so there is a path between x and y in \mathcal{Z}^i that consists of at most $\Lambda(\mathcal{Z}^i)$ edges and has length at most $dist_{\mathcal{Z}^i}(x, y)$. Hence, $\Lambda(\mathcal{Z}) = O(\log_k n + \alpha(k))$.

By Theorem 3, the maximum degree of each 1-spanner \mathcal{Z}^i satisfies $\Delta(\mathcal{Z}^i) \leq \Delta(\tau_i) + 2k$. By Lemma 1, for all $i \in [m]$, $\Delta(\tau_i) = O(1)$. Hence $\Delta(\mathcal{Z}^i) = O(k)$. Since $m = O(1)$, it follows that $\Delta(\mathcal{Z}) \leq \sum_{i=1}^m \Delta(\mathcal{Z}^i) \leq mO(k) = O(k)$.

G Well-Separated Pair Constructions for Random Points

In this section we show that for any set \mathcal{S} of points that are chosen independently and uniformly at random from the unit square, the lightness of well-separated pair constructions is (w.h.p.) $O(1)$. Our argument extends to any constant dimension.

The next lemma from [31] provides a lower bound on the weight of $MST(\mathcal{S})$.

Lemma 6 (Lemma 15.1.6 in [31]). *For a set \mathcal{S} of n points that are chosen independently and uniformly at random from the unit square, there are constants $c > 0$ and $0 < \rho < 1$, such that $Pr(w(MST(\mathcal{S})) < c\sqrt{n}) \leq \rho^n$.*

The following statement shows that the lightness of well-separated pair constructions for \mathcal{S} is (w.h.p.) $O(1)$.

Proposition 2. *For any set \mathcal{S} of n points in the unit square, the weight of well-separated pair constructions is $O(\sqrt{n})$.*

Our proof of Proposition 2 employs a simple, self-contained, combinatorial argument for bounding the weight of the WSPD-spanner due to Callahan and Kosaraju [9, 10] *directly*. (See Sect. G.1 for the definition of the WSPD-spanner.) In particular, our argument is independent of the theory of dumbbell trees due to Arya et al. [4]. Before we prove Proposition 2, we provide (Sect. G.1) the relevant background and introduce some notation. The proof of Proposition 2 appears in Sect. G.2.

Arya et al. [4] devised a well-separated pair construction of $(1 + \epsilon)$ -spanners with both diameter and lightness at most $O(\log n)$. In addition, Lenhof et al. [29] showed that there exist point sets for which any well-separated pair construction must admit lightness at least $\Omega(\log n)$. While this existential bound holds true in the worst-case scenario, our probabilistic upper bound of $O(1)$ on the lightness of well separated pair constructions for random point sets implies that on average one can do much better.

Corollary 4. *For any set S of n points that are chosen independently and uniformly at random from the unit cube, there is a $(1 + \epsilon)$ -spanner with $O(n)$ edges, diameter $O(\log n)$ and lightness (w.h.p.) $O(1)$.*

Remark: After communicating this result to Michiel Smid, he [35] pointed out the following alternative argument for obtaining this bound on lightness. First, Chandra [14] showed that for random point sets in the unit cube, any edge set that satisfies the gap property has lightness (w.h.p.) $O(1)$. Second, consider the edge set E of the dumbbell trees of [4]. As shown in [4] this set can be partitioned into $E = E' \cup E''$, such that E' satisfies the *gap property* and $w(E'') = O(w(E'))$. Finally, use the observation that the lightness of well-separated pair constructions is asymptotically equal to that of dumbbell trees. On the other hand, our proof employs a simple, self-contained, combinatorial argument for analyzing the lightness of well-separated pair constructions *directly*. Hence we believe that our approach is advantageous, since, in particular, it does not take a detour through the heavy dumbbell trees machinery of [4].

G.1 Background and Notation

In what follows, let $s > 0$ be a real fixed number.

We say that two point sets in the plane A and B are *well-separated* with respect to s if A and B can be enclosed in two circles of radius r , such that the distance between the two circles is at least sr . The number s is called the *separation ratio* of A and B . A *well-separated pair decomposition* (WSPD) for a point set P in the plane with respect to s is a set $\{\{A_1, B_1\}, \{A_2, B_2\}, \dots, \{A_m, B_m\}\}$ of pairs of nonempty subsets of P , for some integer m , such that: 1) For each $i \in [m]$, A_i and B_i are well-separated with respect to s , 2) For any two distinct points p and q of P , there is exactly one index i in $[m]$, such that either $p \in A_i$ and $q \in B_i$, or $p \in B_i$ and $q \in A_i$.

Next, we describe a well-known algorithm due to Callahan and Kosaraju [9] for computing a WSPD for P with respect to s . The algorithm consists of two phases. In the first phase, we construct a *split tree*, that is, a tree that corresponds to a hierarchical decomposition of P into rectangles of bounded aspect ratio, where rectangles serve as vertices of the tree, each being split into smaller rectangles as long as it contains more than one point of P . Observe that

the split tree does not depend on s . In the second phase, we employ the split tree to construct the WSPD itself.

There are many variants of a split tree, and we outline below the fair split tree due to Callahan and Kosaraju [9]. Place a smallest-possible rectangle $R(P)$ about the point set P . The root of the fair split tree is $R(P)$. Choose the longer side of $R(P)$ and divide it into two equal parts, thus splitting $R(P)$ into two smaller rectangles of equal size, R_l and R_r . The left and right subtrees of the root $R(P)$ are the fair split trees that are constructed recursively for the point sets $R_l \cap P$ and $R_r \cap P$, respectively. This recursive process is repeated until a single point remains, in which case the split tree consists of just a single vertex that stores this point. Following Arya et al. [4], we consider a fair split tree in an ideal form, henceforth the *idealized box split tree*. In this tree rectangles are squares, each split recursively into four identical squares of half the side length. In other words, the idealized box split tree is a *quadtrees*. (Refer to Chapter 14 of [7] for the definition of quadtree.) While actual constructions will be performed using the fair split tree or other closely related variants (see, e.g., the *compressed quadtrees* of [21] and [11], and the *balanced box-decomposition tree* of [5]), the idealized box split tree provides a clean and elegant way of conceptualizing the fair split tree in all its variants for purposes of analysis.




Fig. 5. An illustration of a typical internal vertex v in \mathcal{T} . The vertex v has four children v_1, v_2, v_3 and v_4 , each being a square of half the side length $side(v)/2$. Each child v_i of v , $i \in [4]$, has four children of its own (unless it is a leaf), of side length $side(v)/2^2$ each, and so on. In the illustration only the four children of v_2 are depicted.

Consider the idealized box split tree $\mathcal{T} = \mathcal{T}(P)$ that is constructed for P . We identify each vertex v in the tree \mathcal{T} with the square in the plane corresponding to it. For example, the root $rt = rt(\mathcal{T})$ of \mathcal{T} is identified with the smallest-possible square $R(P)$ about the point set P . Thus referring to, e.g., the side length of a vertex v in \mathcal{T} , is well-defined. Suppose without loss of generality that the sides of the square rt are parallel to the x and y axes. Consequently, each vertex v of \mathcal{T} is a square whose sides are parallel to the x and y axes. Denote the four sides of v by $North(v)$, $South(v)$, $East(v)$ and $West(v)$, with $North(v)$ and $South(v)$ (respectively, $East(v)$ and $West(v)$) being parallel to the x -axis (resp., y -axis). Denote the four children of an internal vertex v in \mathcal{T} by v_1, v_2, v_3 and v_4 , each being a square of half the side length of v , where v_1, v_2, v_3 and v_4 are the North-Eastern, South-Eastern, South-Western and North-Western parts of v , respectively. The side length of v is denoted by $side(v)$. Notice that each child of rt has side length $\frac{side(rt)}{2}$, each grandchild of rt has side

length $\frac{side(rt)}{2^2}$, etc. More generally, a vertex v in \mathcal{T} of level⁹ $L(v)$, $0 \leq L(v) \leq depth(\mathcal{T})$, has side length $side(v) = \frac{side(rt)}{2^{L(v)}}$. (See Fig. 5 for an illustration.) Define $P(v) = v \cap P$. The vertex v is called *empty* if $P(v) = \emptyset$. Otherwise, it is *non-empty*. The depth of a vertex v in \mathcal{T} is defined as the depth of the subtree \mathcal{T}_v of \mathcal{T} rooted at v . For any two vertices u and v in \mathcal{T} , we denote by $dist(u, v)$ the distance of closest approach between u and v , i.e., the minimum distance between a point lying on the boundary of u and a point lying on the boundary of v . Also, we denote by $distMax(P(u), P(v))$ the maximum distance between a point in $P(u)$ and a point in $P(v)$. Clearly, $distMax(P(u), P(v))$ is no smaller than $dist(u, v)$. On the other hand, it is bounded from above by the distance of furthest approach between u and v , i.e., the maximum distance between a point lying on the boundary of u and a point lying on the boundary of v , which is, in turn, bounded from above by $dist(u, v) + 2\sqrt{2} \max\{side(u), side(v)\}$. Thus, $dist(u, v) \leq distMax(P(u), P(v)) \leq dist(u, v) + 2\sqrt{2} \max\{side(u), side(v)\}$.

To compute the WSPD of P , we use a simple recursive algorithm which consists of the two procedures below. (This algorithm is essentially taken from Callahan and Kosaraju [9].) We initially invoke Procedure 1 below by making the call $WSPD(rt(\mathcal{T}))$, where $rt(\mathcal{T}) = R(P)$. The output returned by this call is the WSPD for P . We omit the proof of correctness, which resembles that of [9]. Notice that for any pair of vertices u and v in \mathcal{T} , both calls $WSPD(u, v)$ and $WSPD(u)$ return sets of well-separated pairs of P . In what follows we write $WSPD(u, v)$ and $WSPD(u)$ to refer to the sets that are returned by these calls (rather than to the calls themselves).

Procedure 1 $WSPD(u)$:

- 1: **if** $|P(u)| \leq 1$ **then**
 - 2: return \emptyset
 - 3: **end if**
 - 4: return $\bigcup_{1 \leq i \leq 4} WSPD(u_i) \cup \bigcup_{1 \leq i < j \leq 4} WSPD(u_i, u_j)$
-

Procedure 2 $WSPD(u, v)$:

- 1: **if** $P(u) = \emptyset$ or $P(v) = \emptyset$ **then**
 - 2: return \emptyset
 - 3: **end if**
 - 4: **if** $P(u)$ and $P(v)$ are well-separated **then**
 - 5: return $\{\{P(u), P(v)\}\}$
 - 6: **end if**
 - 7: **if** $side(u) \geq side(v)$ **then**
 - 8: return $\bigcup_{1 \leq i \leq 4} WSPD(u_i, v)$
 - 9: **else**
 - 10: return $\bigcup_{1 \leq i \leq 4} WSPD(u, v_i)$
 - 11: **end if**
-

A *representative assignment* for the split tree $\mathcal{T} = \mathcal{T}(P)$ is a mapping φ between vertices of \mathcal{T} and points of P , sending each vertex v in \mathcal{T} to a point $\varphi(v)$ in $P(v)$. The point $\varphi(v)$ is called the *representative* of v under the mapping φ . We say that a pair (A, B) of nonempty sets of P *belongs* to \mathcal{T} , if there are two

⁹ The *level* of a vertex in a rooted tree is its unweighted distance from the root.

vertices u and v in \mathcal{T} , such that $A = P(u)$ and $B = P(v)$. Given a representative assignment φ , there is a natural correspondence between a well-separated pair $\{P(u), P(v)\}$ that belongs to \mathcal{T} and the edge $(\varphi(u), \varphi(v))$ connecting the representatives of u and v under φ . In the same way, there is a natural correspondence between a set S of well-separated pairs of P that belong to \mathcal{T} and the edge set $E_\varphi(S)$, where $E_\varphi(S) = \{(\varphi(u), \varphi(v)) \mid \{P(u), P(v)\} \in S\}$. The weight $w(H)$ of an edge set H is defined as the sum $\sum_{e=(u,v) \in H} w(u, v)$ of all edge weights in it, where $w(u, v) = \|u - v\|$. Callahan and Kosaraju [10] showed that for any representative assignment φ , the edge set $E^* = E_\varphi(WSPD(P))$ that corresponds to $WSPD(P) = WSPD(rt(\mathcal{T}))$ constitutes a $(1 + \epsilon)$ -spanner (with $O(n)$ edges), henceforth the *WSPD-spanner* of P , where ϵ is an arbitrarily small constant depending on s . (It can be easily shown that $\epsilon \leq \frac{8}{s-4}$.)

G.2 Proof of Proposition 2

In this section we prove Proposition 2.

Let P be an arbitrary set of n points in the plane, and let $\mathcal{T} = \mathcal{T}(P)$ and $WSPD(P) = WSPD(rt(\mathcal{T}))$ be the idealized box split tree and the WSPD that are constructed for it, respectively. Also, fix an arbitrary representative assignment φ for \mathcal{T} . Next, we show that the weight $w(E^*)$ of the WSPD-spanner $E^* = E_\varphi(WSPD(P))$ is at most $c^* \text{side}(rt(\mathcal{T}))\sqrt{n}$, where c^* is a sufficiently large constant that depends only on s . (We do not try to optimize the constant c^* .) In particular, for a point set P in the unit square we have $\text{side}(rt(\mathcal{T})) = 1$, thus proving Proposition 2.

Observe that for any two vertices u and v in \mathcal{T} , both $WSPD(u, v)$ and $WSPD(u)$ are sets of well-separated pairs of P that belong to \mathcal{T} . Henceforth, we write $W(u, v)$ (respectively, $W(u)$) as a shortcut for $w(E_\varphi(WSPD(u, v)))$ (resp., $w(E_\varphi(WSPD(u)))$).

Lemma 7. *Let c be some constant, such that $1/2 \leq c \leq \sqrt{2}$. For any pair u, v of vertices in \mathcal{T} for which $\text{dist}(u, v) = c \max\{\text{side}(u), \text{side}(v)\}$, $W(u, v) \leq \alpha \max\{\text{side}(u), \text{side}(v)\}$, where $\alpha = \alpha_s$ is a sufficiently large constant that depends only on s .*

Proof. From standard packing arguments, it follows that $|WSPD(u, v)| \leq \tilde{\alpha}$, where $\tilde{\alpha} = \tilde{\alpha}_s$ is a sufficiently large constant that depends only on s . For each pair $\{P(x), P(y)\}$ in $WSPD(u, v)$, the weight $\|\varphi(x) - \varphi(y)\|$ of the corresponding edge $(\varphi(x), \varphi(y))$ is at most $\text{dist}(u, v) + 2\sqrt{2} \max\{\text{side}(u), \text{side}(v)\} = (c + 2\sqrt{2}) \max\{\text{side}(u), \text{side}(v)\}$. Define $\alpha = \tilde{\alpha}(c + 2\sqrt{2})$. It follows that $W(u, v) \leq \tilde{\alpha}(c + 2\sqrt{2}) \max\{\text{side}(u), \text{side}(v)\} = \alpha \max\{\text{side}(u), \text{side}(v)\}$. \square



Fig. 6. a) Two diagonal vertices w and z . b) Two adjacent vertices x and y .

We say that two vertices u and v in \mathcal{T} of the same level are *diagonal* if their boundaries intersect at a single point. (See Fig. 6.a for an illustration.) For example, for any vertex v in \mathcal{T} , its two children v_1 and v_3 are diagonal. Consider two diagonal vertices u and v in \mathcal{T} . Since by definition they are at the same level in \mathcal{T} , it holds that $side(u) = side(v)$. Also, notice that $dist(u, v) = 0$ and $\|\varphi(u) - \varphi(v)\| \leq distMax(P(u), P(v)) \leq 2\sqrt{2}side(u)$.

Lemma 8. *For any two diagonal vertices u and v in \mathcal{T} , $W(u, v) \leq \beta side(u)$, where $\beta = \beta_s$ is a sufficiently large constant that depends only on s .*

Proof. The proof is by induction on the sum $h = depth(u) + depth(v)$ of depths of u and v .

Basis: $h = 0$. In this case both u and v are leaves, and so each one of them contains at most one point. If either u or v is empty, then $WSPD(u, v) = \emptyset$, and so $W(u, v) = 0 < \beta side(u)$. Otherwise, $WSPD(u, v) = \{\{P(u), P(v)\}\}$, and so $W(u, v) = \|\varphi(u) - \varphi(v)\| \leq 2\sqrt{2}side(u) < \beta side(u)$.

Induction Step: We assume the correctness of the statement for all smaller values of h , and prove it for h . If either u or v is empty, then $WSPD(u, v) = \emptyset$, and so $W(u, v) = 0 < \beta side(u)$. Otherwise, if $P(u)$ and $P(v)$ are well-separated then $WSPD(u, v) = \{\{P(u), P(v)\}\}$, and so $W(u, v) = \|\varphi(u) - \varphi(v)\| \leq 2\sqrt{2}side(u) < \beta side(u)$. We henceforth assume that $P(u)$ and $P(v)$ are not well-separated. In this case $WSPD(u, v) = \bigcup_{1 \leq i \leq 4} WSPD(u_i, v)$, and so $W(u, v) = \sum_{1 \leq i \leq 4} W(u_i, v)$. Since u and v are diagonal, the intersection of u and exactly one child of u consists of a single point, whereas all the other children of u are disjoint from v . Suppose without loss of generality that the child of u that intersects v is u_1 . (See Fig. 7.a for an illustration.) Observe that $dist(u_2, v) = dist(u_4, v) = \frac{1}{2}side(v)$ and $dist(u_3, v) = \frac{1}{\sqrt{2}}side(v)$. Hence, by Lemma 7, for each $2 \leq i \leq 4$, $W(u_i, v) \leq \alpha side(v)$. Next, we bound $W(u_1, v)$. If u_1 is empty, then $WSPD(u_1, v) = \emptyset$, and so $W(u_1, v) = 0$. Also, if $P(u_1)$ and $P(v)$ are well-separated, then $WSPD(u_1, v) = \{\{P(u_1), P(v)\}\}$, and so $W(u_1, v) = \|\varphi(u_1) - \varphi(v)\| \leq 2\sqrt{2}side(v)$. Otherwise, $WSPD(u_1, v) = \bigcup_{1 \leq i \leq 4} WSPD(u_1, v_i)$, and so $W(u_1, v) = \sum_{1 \leq i \leq 4} W(u_1, v_i)$. Observe that $dist(u_1, v_2) = dist(u_1, v_4) = side(u_1)$ and $dist(u_1, v_1) = \sqrt{2}side(u_1)$. Hence, by Lemma 7, for each $i \neq 3$, $W(u_1, v_i) \leq \alpha side(u_1) = \frac{\alpha}{2}side(u)$. Notice that u_1 and v_3 are diagonal, and so by the induction hypothesis, $W(u_1, v_3) \leq \beta side(u_1) = \frac{\beta}{2}side(u)$. Set $\beta = 9\alpha$. Altogether,

$$\begin{aligned} W(u, v) &\leq 3\alpha side(v) + \sum_{1 \leq i \leq 4, i \neq 3} W(u_1, v_i) + W(u_1, v_3) \\ &\leq 3\alpha side(v) + 3\frac{\alpha}{2}side(u) + \frac{\beta}{2}side(u) = \beta side(u). \end{aligned}$$

□

We say that two vertices u and v in \mathcal{T} of the same level are *adjacent* if their boundaries intersect at a single side. (See Fig. 6.b for an illustration.) For example, for any vertex v in \mathcal{T} , its two children v_1 and v_2 are adjacent. Consider two adjacent vertices u and v in \mathcal{T} . Since they are at the same level in \mathcal{T} , we have $side(u) = side(v)$. Also, notice that $dist(u, v) = 0$ and $\|\varphi(u) - \varphi(v)\| \leq distMax(P(u), P(v)) \leq \sqrt{5}side(u)$. For a vertex v in \mathcal{T} , define $N(v) = |P(v)|$.

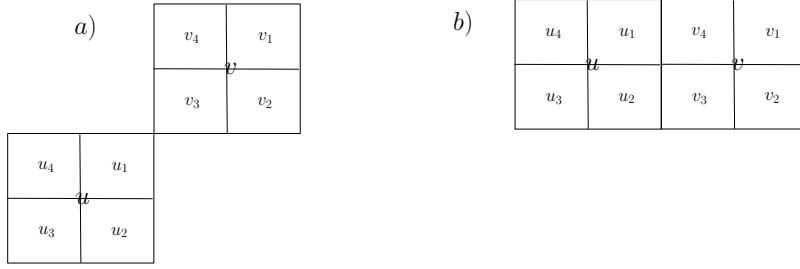


Fig. 7. a) An illustration of two diagonal vertices u and v . The vertex u_1 is diagonal to both u_3 and v_3 . The vertices u_1 and v_3 intersect at the same point as their respective parents u and v do. b) An illustration of two adjacent vertices u and v . The vertex u_1 is adjacent to u_2 , u_4 and v_4 . The vertices u_1 and v_4 intersect at a single side, which is the upper half of the side at which their respective parents u and v intersect.

Lemma 9. *For any two adjacent vertices u and v in \mathcal{T} such that $N(u) + N(v) \geq 1$, $W(u, v) \leq \gamma \text{side}(u) \log(N(u) + N(v))$, where $\gamma = \gamma_s$ is a sufficiently large constant that depends only on s .*

Proof. The proof is by induction on the sum $h = \text{depth}(u) + \text{depth}(v)$ of depths of u and v .

Basis: $h = 0$. In this case both u and v are leaves, and so each one of them contains at most one point. If either u or v is empty, then $WSPD(u, v) = \emptyset$, and so $W(u, v) = 0 = \gamma \text{side}(u) \log 1$. Otherwise, $WSPD(u, v) = \{\{P(u), P(v)\}\}$, and so $W(u, v) = \|\varphi(u) - \varphi(v)\| \leq \sqrt{5} \text{side}(u) < \gamma \text{side}(u) \log 2$.

Induction Step: We assume the correctness of the statement for all smaller values of h , and prove it for h . If either u or v is empty, then $WSPD(u, v) = \emptyset$, and so $W(u, v) = 0 \leq \gamma \text{side}(u) \log(N(u) + N(v))$. Otherwise, if $P(u)$ and $P(v)$ are well-separated then $WSPD(u, v) = \{\{P(u), P(v)\}\}$, and so $W(u, v) = \|\varphi(u) - \varphi(v)\| \leq \sqrt{5} \text{side}(u) < \gamma \text{side}(u) \log(N(u) + N(v))$. We henceforth assume that $P(u)$ and $P(v)$ are not well-separated. In this case $WSPD(u, v) = \bigcup_{1 \leq i \leq 4} WSPD(u_i, v)$, and so $W(u, v) = \sum_{1 \leq i \leq 4} W(u_i, v)$. Since u and v are adjacent, two adjacent children u_i and u_{i+1} of u intersect v , $i \in [3]$, each at a single side. Suppose without loss of generality that these children of u are u_1 and u_2 . (See Fig. 7.b for an illustration.) Note that $\text{dist}(u_3, v) = \text{dist}(u_4, v) = \frac{1}{2} \text{side}(v)$. By Lemma 7, $W(u_3, v), W(u_4, v) \leq \alpha \text{side}(v)$. Next, we bound $W(u_1, v)$. If u_1 is empty, then $WSPD(u_1, v) = \emptyset$, and so $W(u_1, v) = 0$. If $P(u_1)$ and $P(v)$ are well-separated, then $WSPD(u_1, v) = \{\{P(u_1), P(v)\}\}$, and so $W(u_1, v) = \|\varphi(u_1) - \varphi(v)\| \leq \sqrt{5} \text{side}(v)$. Otherwise, $WSPD(u_1, v) = \bigcup_{1 \leq i \leq 4} WSPD(u_1, v_i)$, and so $W(u_1, v) = \sum_{1 \leq i \leq 4} W(u_1, v_i)$. Note that $\text{dist}(u_1, v_1) = \text{dist}(u_1, v_2) = \text{side}(u_1)$. By Lemma 7, $W(u_1, v_1), W(u_1, v_2) \leq \alpha \text{side}(u_1) = \frac{\alpha}{2} \text{side}(u)$. Notice that u_1 and v_3 are diagonal, whereas u_1 and v_4 are adjacent. Hence, by Lemma 8, $W(u_1, v_3) \leq \beta \text{side}(u_1) = \frac{\beta}{2} \text{side}(u)$. Recall that u_1 is non-empty, and so $N(u_1) + N(v_4) \geq 1$. By the induction hypothesis, $W(u_1, v_4) \leq \gamma \text{side}(u_1) \log(N(u_1) + N(v_4)) = \frac{\gamma}{2} \text{side}(u) \log(N(u_1) + N(v_4))$. We get that

$$\begin{aligned} W(u_1, v) &\leq \alpha \text{side}(u) + \frac{\beta}{2} \text{side}(u) + \frac{\gamma}{2} \text{side}(u) \log(N(u_1) + N(v_4)) \\ &= \text{side}(u) \left(\alpha + \frac{\beta}{2} + \frac{\gamma}{2} \log(N(u_1) + N(v_4)) \right). \end{aligned}$$

Symmetrically, we get $W(u_2, v) \leq \text{side}(u) \left(\alpha + \frac{\beta}{2} + \frac{\gamma}{2} \log(N(u_2) + N(v_3)) \right)$. Note that $N(u_1) + N(u_2) \leq N(u)$ and $N(v_3) + N(v_4) \leq N(v)$, and so $2(N(u_1) + N(v_4))(N(u_2) + N(v_3)) \leq (N(u) + N(v))^2$. Set $\gamma = 2(4\alpha + \beta)$. Altogether,

$$\begin{aligned} W(u, v) &= \sum_{1 \leq i \leq 4} W(u_i, v) = [W(u_1, v) + W(u_2, v)] + [W(u_3, v) + W(u_4, v)] \\ &= \gamma \text{side}(u) \left(\frac{4\alpha + \beta}{\gamma} + \frac{1}{2} (\log(N(u_1) + N(v_4)) + \log(N(u_2) + N(v_3))) \right) \\ &= \gamma \text{side}(u) \left(\frac{1 + \log(N(u_1) + N(v_4)) + \log(N(u_2) + N(v_3))}{2} \right) \\ &\leq \gamma \text{side}(u) \log(N(u) + N(v)). \end{aligned}$$

□

We use the following claim to prove Lemma 10.

Claim. For any positive integers n_1, n_2, \dots, n_k, k and n , such that $\sum_{i=1}^k n_i = n$,

$$\sum_{i=1}^k \left(\sqrt{n_i} - \frac{\ln n_i}{8} \right) \leq f(n, k) = k \left(\sqrt{n/k} - \frac{\ln(n/k)}{8} \right).$$

Proof. The proof is by induction on k , for $k \in [n]$. The basis $k = 1$ is trivial.

Induction Step: We assume the correctness of the statement for all smaller values of k and prove it for k . By the induction hypothesis,

$$\sum_{i=1}^{k-1} \left(\sqrt{n_i} - \frac{\ln n_i}{8} \right) \leq f(n - n_k, k - 1) = (k - 1) \left(\sqrt{\frac{n - n_k}{k - 1}} - \frac{\ln \left(\frac{n - n_k}{k - 1} \right)}{8} \right).$$

It follows that

$$\sum_{i=1}^k \left(\sqrt{n_i} - \frac{\ln n_i}{8} \right) \leq (k - 1) \left(\sqrt{\frac{n - n_k}{k - 1}} - \frac{\ln \left(\frac{n - n_k}{k - 1} \right)}{8} \right) + \sqrt{n_k} - \frac{\ln n_k}{8}. \quad (3)$$

Define $g_{n,k}(x) = (k - 1) \left(\sqrt{\frac{n-x}{k-1}} - \frac{\ln \left(\frac{n-x}{k-1} \right)}{8} \right) + \sqrt{x} - \frac{\ln x}{8}$. Since $n_1, n_2, \dots, n_k \geq 1$ are positive integers and $\sum_{i=1}^k n_i = n$, we have that $1 \leq n_k \leq n - k + 1$. Hence, the maximum value of the function $g_{n,k}(x)$ in the range $1 \leq x \leq n - k + 1$ provides an upper bound on the right-hand side of (3). It is easy to verify that the function $g_{n,k}(x)$ in the range $1 \leq x \leq n - k + 1$ is maximized at $x = n/k$. Hence, in the range $1 \leq x \leq n - k + 1$, $g_{n,k}(x) \leq g_{n,k}(n/k) = k \left(\sqrt{n/k} - \frac{\ln(n/k)}{8} \right)$, and we are done. □

The next lemma implies that $w(E^*) = W(rt) \leq c^* \text{side}(rt) \left(\sqrt{n} - \frac{\ln(n)}{8} \right) \leq c^* \text{side}(rt) \sqrt{n}$. Hence, for any set of n points in the unit square, the weight of the WSPD-spanner is $O(\sqrt{n})$, thus proving Proposition 2.

Lemma 10. $W(u) \leq c^* \text{side}(u) \left(\sqrt{N(u)} - \frac{\ln(N(u))}{8} \right)$, for any non-empty vertex u in \mathcal{T} .

Proof. The proof is by induction on the depth $h = \text{depth}(u)$ of u . The basis $h = 0$ is trivial.

Induction Step: We assume the correctness of the statement for all smaller values of h , and prove it for h . First, suppose that $1 \leq N(u) < 20$. In this case, $|WSPD(u)| \leq c$, for a sufficiently large constant c . The weight of every edge in the edge set that corresponds to $WSPD(u)$ is at most $\sqrt{2}\text{side}(u)$, and so

$$W(u) \leq c\sqrt{2}\text{side}(u) < c^*\text{side}(u) \left(\sqrt{N(u)} - \frac{\ln(N(u))}{8} \right).$$

We henceforth assume that $N(u) \geq 20$. Hence,

$$WSPD(u) = \bigcup_{1 \leq i \leq 4} WSPD(u_i) \cup \bigcup_{1 \leq i < j \leq 4} WSPD(u_i, u_j),$$

and so

$$W(u) = \sum_{1 \leq i \leq 4} W(u_i) + \sum_{1 \leq i < j \leq 4} W(u_i, u_j). \quad (4)$$

To bound $W(u)$, we start with bounding the left sum $\sum_{1 \leq i \leq 4} W(u_i)$ in the right-hand side of (4). Denote by I the set of indices in $[4]$ for which $N(u_i) \geq 1$. By the induction hypothesis, for each $i \in I$, $W(u_i) \leq c^*\text{side}(u_i) \left(\sqrt{N(u_i)} - \frac{\ln(N(u_i))}{8} \right)$. Also, for each index $i \in [4] \setminus I$, we have $W(u_i) = 0$. It follows that

$$\begin{aligned} \sum_{1 \leq i \leq 4} W(u_i) &= \sum_{i \in I} W(u_i) \leq \sum_{i \in I} c^*\text{side}(u_i) \left(\sqrt{N(u_i)} - \frac{\ln(N(u_i))}{8} \right) \\ &= \frac{c^*}{2}\text{side}(u) \sum_{i \in I} \left(\sqrt{N(u_i)} - \frac{\ln(N(u_i))}{8} \right). \end{aligned} \quad (5)$$

Observe that $\sum_{i \in I} N(u_i) = N(u)$ and $1 \leq |I| \leq 4$. By Claim G.2,

$$\sum_{i \in I} \left(\sqrt{N(u_i)} - \frac{\ln(N(u_i))}{8} \right) \leq |I| \left(\sqrt{N(u)/|I|} - \frac{\ln(N(u)/|I|)}{8} \right).$$

Observe that the function $f_{N(u)}(x) = f(N(u), x) = x \left(\sqrt{N(u)/x} - \frac{\ln(N(u)/x)}{8} \right)$ is monotone increasing with x in the range $x > 0$. (The derivative $f'_{N(u)}(x)$ is strictly positive for all $x > 0$.) Since $|I| \leq 4$, we thus have

$$\begin{aligned} \sum_{i \in I} \left(\sqrt{N(u_i)} - \frac{\ln(N(u_i))}{8} \right) &\leq f(N(u), |I|) \leq 4 \left(\sqrt{N(u)/4} - \frac{\ln(N(u)/4)}{8} \right) \\ &= 2\sqrt{N(u)} - \frac{\ln(N(u)/4)}{2}. \end{aligned} \quad (6)$$

Plugging (6) into (5) yields

$$\sum_{1 \leq i \leq 4} W(u_i) \leq c^*\text{side}(u) \left(\sqrt{N(u)} - \frac{\ln(N(u)/4)}{4} \right). \quad (7)$$

We proceed with bounding the right sum $\sum_{1 \leq i < j \leq 4} W(u_i, u_j)$ in the right-hand side of (4). Observe that the two pairs (u_1, u_3) and (u_2, u_4) of children of u are

diagonal, whereas the four other pairs (u_1, u_2) , (u_1, u_4) , (u_2, u_3) and (u_3, u_4) are adjacent. By Lemma 8, $W(u_1, u_3), W(u_2, u_4) \leq \beta \text{side}(u_1) = \frac{\beta}{2} \text{side}(u)$. Consider a pair (u_i, u_j) among the four pairs of adjacent children of u . If both u_i and u_j are empty, then $W(u_i, u_j) = 0$. Otherwise, we have $N(u_i) + N(u_j) \geq 1$. Hence, by Lemma 9,

$$W(u_i, u_j) \leq \gamma \text{side}(u_i) \log(N(u_i) + N(u_j)) \leq \frac{\gamma}{2} \text{side}(u) \log(N(u)).$$

Recall that $\gamma = 2(4\alpha + \beta)$, and so $\beta \leq \frac{\gamma}{2}$. Altogether

$$\begin{aligned} \sum_{1 \leq i < j \leq 4} W(u_i, u_j) &\leq \beta \text{side}(u) + 2\gamma \text{side}(u) \log(N(u)) \\ &\leq \gamma \text{side}(u) \left(\frac{1}{2} + 2 \log(N(u)) \right) \leq 4\gamma \text{side}(u) \ln(N(u)). \end{aligned} \quad (8)$$

Plugging (7) and (8) into (4) yields

$$\begin{aligned} W(u) &= \sum_{1 \leq i \leq 4} W(u_i) + \sum_{1 \leq i < j \leq 4} W(u_i, u_j) \\ &\leq c^* \text{side}(u) \left(\sqrt{N(u)} - \frac{\ln(N(u)/4)}{4} \right) + 4\gamma \text{side}(u) \ln(N(u)). \end{aligned} \quad (9)$$

Observe that for a sufficiently large constant c^* and all $n \geq 20$, the right-hand side of (9) is at most $c^* \text{side}(u) \left(\sqrt{N(u)} - \frac{\ln(N(u))}{8} \right)$, and we are done. \square