

1-2006

Small Hop-diameter Sparse Spanners for Doubling Metrics

T.-H. Hubert Chan
Carnegie Mellon University

Anupam Gupta
Carnegie Mellon University

Follow this and additional works at: <http://repository.cmu.edu/compsci>

This Conference Proceeding is brought to you for free and open access by the School of Computer Science at Research Showcase @ CMU. It has been accepted for inclusion in Computer Science Department by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Small Hop-diameter Sparse Spanners for Doubling Metrics*

T-H. Hubert Chan

Anupam Gupta

Abstract

Given a metric $M = (V, d)$, a graph $G = (V, E)$ is a t -spanner for M if every pair of nodes in V has a “short” path (i.e., of length at most t times their actual distance) between them in the spanner. Furthermore, this spanner has a *hop diameter* bounded by D if every such short path also uses at most D edges. We consider the problem of constructing sparse $(1 + \varepsilon)$ -spanners with small hop diameter for metrics of low doubling dimension.

In this paper, we show that given any metric with constant doubling dimension k , and any $0 < \varepsilon < 1$, one can find a $(1 + \varepsilon)$ -spanner for the metric with nearly linear number of edges (i.e., only $O(n \log^* n + n\varepsilon^{-O(k)})$ edges) and a *constant* hop diameter, and also a $(1 + \varepsilon)$ -spanner with linear number of edges (i.e., only $n\varepsilon^{-O(k)}$ edges) which achieves a hop diameter that grows like the functional inverse of the Ackermann’s function. Moreover, we prove that such tradeoffs between the number of edges and the hop diameter are asymptotically optimal.

1 Introduction

The study of finite metrics and their properties has been a very fruitful area of research, with applications to many different problems: many commonly arising problems (e.g., clustering, near-neighbor finding, network routing, just to name a few) deal with sets of points on which a distance function has been defined, and one wants to store and process this metric in different ways.

However, metrics vary in their “complexity”: some metrics, like the well-understood Euclidean spaces, seem to be intrinsically simpler to manipulate than others (say ℓ_1 space), which in turn are simpler than arbitrary metrics. Hence, merely using the number of points n in a given metric space to quantify the performance of algorithms (e.g., running time, or quality of the output) seems too pessimistic. To this end, there has been much recent interest in defining a notion of “dimension”

for arbitrary metrics, and to devise algorithms whose behavior degrades gracefully as the dimension of the input metric increases.

One such notion of dimension has been that of the doubling dimension [5, 13, 11]. The *doubling dimension* of a metric $M = (V, d)$ is the minimum value k such that every ball B in the metric can be covered by 2^k balls of half the radius of B . This can be seen as a generalization of Euclidean dimension to arbitrary metric spaces; indeed, it is not difficult to see that \mathbb{R}^k equipped with any of the ℓ_p norms had doubling dimension $\Theta(k)$.

Apart from being a generalization of the ℓ_p notion of dimension, designing algorithms that only use the doubling properties (instead of the geometry of \mathbb{R}^k) has other advantages: the notion of doubling dimension is fairly resistant to small perturbations in the distances: for instance, if one takes a distance matrix of a set of points in ℓ_p^k and slightly changes some of the entries, then the doubling dimension does not change by much, but the metric may not remain isometrically embeddable in ℓ_p (into any number of dimensions). To this end, there has been much interest in understanding this notion of dimension, and in generalizing algorithms to adapt gracefully to the dimension of the input metric; see, e.g., [10, 14, 11, 22, 16, 15, 7, 12]

In this paper, we will focus on obtaining sparse representations of metrics: these are called *spanners*, and they have been studied extensively both for general and Euclidean metrics. Formally, a t -spanner for a metric $M = (V, d)$ is an undirected graph $G = (V, E)$ such that the distances according to d_G (the shortest-path metric of G) are close to the distances in d : i.e., $d(u, v) \leq d_G(u, v) \leq t d(u, v)$. Clearly, one can take a complete graph and obtain $t = 1$, and hence the quality of the spanner is typically measured by how few edges can G contain whilst maintaining a *stretch* of at most t . The notion of spanners has been widely studied for general metrics (see, e.g. [18, 3, 8]), and for geometric distances (see, e.g., [6, 21, 24, 4]).

Very recently, there have been good constructions of spanners for doubling metrics as well: given a metric with doubling dimension \dim , the results of Chan et al. [7], and independently, those of Har-Peled and Mendel [12] showed how to construct $(1 + \varepsilon)$ -spanners

*Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213. {hubert, anupamg} at cs.cmu.edu. This research was partly supported by the NSF grant CCR-0122581 (The ALADDIN project), the NSF CAREER award CCF-0448095, and by an Alfred P. Sloan Fellowship.

with $n(1 + 1/\varepsilon)^{O(\dim)}$ edges. (Here, as in the rest of the paper, $|V| = n$ is the number of points in the metric.)

Our Results. In this paper, we extend these results to find spanners that also have small *hop-diameter*. A t -spanner has hop-diameter D if every pair $u, v \in V$ are connected by some short path in G having length at most $t d(u, v)$, and there are at most D edges on this path. We prove essentially matching upper bounds as well as lower bounds in this paper.

THEOREM 1.1. (Upper Bound) *Given a metric $M = (V, d)$ with doubling dimension \dim and $n = |V|$, there exists a $(1 + \varepsilon)$ -spanner with $m + (2 + \frac{1}{\varepsilon})^{O(\dim)} n$ edges and hop diameter $O(\alpha(m, n))$, where α is the inverse of Ackermann’s function. Such a spanner can be constructed in $2^{O(\dim)} n \log n$ time.*

Note that the result above allows us to trade off the number of edges in the spanner with the hop-diameter: if we desire only a linear number of edges, then the hop-diameter goes as $\alpha(n)$, and as we increase the number of edges, the hop-diameter decreases. After proving this result (which turns out to be fairly straight-forward), we then turn to the lower bound and show that the trade-off in Theorem 1.1 is essentially tight.

THEOREM 1.2. (Lower Bound) *For any $\varepsilon > 0$, there are infinitely many integers n such that there exists a metric M induced by n points on the real line, for which any $(1 + \varepsilon)$ -spanner for M with at most m edges has hop diameter at least $\Omega(\alpha(m, n))$.*

Our Techniques and Related Work. The upper bound Theorem 1.1 generalizes a result of Arya et al. [4] for Euclidean spaces. Indeed, the proof of our result is not difficult given previously known techniques. The basic idea is to first construct a *net-tree* representing a sequence of nested nets of the metric space: this is fairly standard, and has been used earlier, e.g., in [22, 16, 7]. A nearly-linear-time construction of net-trees is given by Har-Peled and Mendel [12]. A second phase then adds some more edges in order to “short-cut” paths in this net tree which have too many hops. The techniques we use are based on those originally used by Yao [25] for range queries on the line, and on the extensions to trees due to Chazelle [9]. As pointed out by Arya et al. [4], a similar construction was given by Alon and Schieber [2].

To the best of our knowledge, there are no previously known lower bounds which show metrics with low doubling (or Euclidean) dimension that require many edges in order to get low hop-diameter $(1 + \varepsilon)$ -spanners. We first consider lower bounds for binary “hierarchically well-separated” trees (HSTs), where the length of

an edge from each node to its child node is much smaller than that to its parent node: this well-separation ensures that low-stretch paths must be “well-behaved”: i.e., the low-stretch path between vertices in any subtree cannot escape the subtree, thus allowing us to reason about them. Our lower bound result for line metrics then follows from the fact that binary HSTs with large separation embed into the real line with small distortion. We note that the lower bounds for the range-query problem given by Yao [25], and Alon and Scheiber [2], while inspiring our work, directly apply to our problem only for the case $\varepsilon = 0$; i.e., for the case where we are not allowed to introduce any further stretch in the second, “short-cutting” phase. Thus Theorem 1.2 can be seen as generalizing Yao’s lower bound proof to all $\varepsilon > 0$.

Other Related Work. Previously known algorithms to obtain low-stretch spanners for doubling metrics [7, 12] have a hop-diameter of $\Omega(\log \Delta)$, where Δ is the diameter of the metric; in fact, there are constructions of spanners in Chan et al. [7] with $(1 + 1/\varepsilon)^{\dim}$ degree (i.e., constant degree for doubling metrics), and it trivially follows that such spanners must have a hop diameter of $\Omega(\log \Delta)$. Abraham et al. [1] study compact routing on Euclidean metrics, and their construction also essentially gives a $(1 + \varepsilon)$ -spanner with $O_\varepsilon(n)$ edges that has hop diameter $O(\log \Delta)$ with high probability.

Low-stretch Spanners with small hop-diameter are potentially useful in network routing protocols. For example, many wireless ad-hoc networks find paths that minimize hop count [20, 17, 19]. Our results may be useful in such situations to build sparse networks admitting paths having few hops and low stretch simultaneously.

2 Preliminaries and Notation

We consider a finite metric $M = (V, d)$ where $|V| = n$. A metric has *doubling dimension* [11] at most k if for every $R > 0$, every ball of radius R can be covered by 2^k balls of radius $R/2$.

DEFINITION 2.1. ((1 + ε)-spanner) *Let (V, d) be a finite metric. Suppose $G = (V, E)$ is an undirected graph such that each edge $\{u, v\} \in E$ has weight $d(u, v)$, and $d_G(u, v)$ is the length of the shortest path between vertices u and v in G . The graph G , or equivalently, the set E of edges, is a $(1 + \varepsilon)$ -spanner for (V, d) if for all pairs u and v , $d_G(u, v)/d(u, v) \leq 1 + \varepsilon$.*

A $(1 + \varepsilon)$ -path in the metric $M = (V, d)$ between u and v is one with length at most $(1 + \varepsilon)d(u, v)$. Thus a $(1 + \varepsilon)$ -spanner is a subgraph $G = (V, E)$ that contains a $(1 + \varepsilon)$ path for each pair of nodes in V .

DEFINITION 2.2. (Hop Diameter) *A $(1 + \varepsilon)$ -spanner is said to have hop diameter at most D if for every*

pair of nodes, there exists a $(1 + \varepsilon)$ -path in the spanner between them having at most D edges or hops.

We would use hierarchical trees to analyze the properties of spanners.

DEFINITION 2.3. (Hierarchical Tree) A hierarchical tree for a set V is a pair (T, φ) , where T is a rooted tree, and φ is a labeling function $\varphi : T \rightarrow V$ that labels each node of T with an element in V , such that the following conditions hold.

1. Every leaf is at the same depth from the root.
2. The function φ restricted to the leaves of T is a bijection into V .
3. If u is an internal node of T , then there exists a child v of u such that $\varphi(v) = \varphi(u)$. This implies that the nodes mapped by φ to any $x \in V$ form a connected subtree of T .

We also use *net trees*; these are similar to those defined by Har-Peled and Mendel [12], who showed how to construct them in $2^{O(k)}n \log n$ time.

DEFINITION 2.4. (Net-Tree) A net tree for a metric (V, d) is a hierarchical tree (T, φ) for the set V such that the following conditions hold.

1. Let N_i be the set of nodes of T that have height i . (The leaves have height 0.) Suppose δ is the minimum pairwise distance in (V, d) . Let $0 < r_0 < \delta/2$, and $r_{i+1} = 2r_i$, for $i \geq 0$. Then, for $i \geq 0$, $\varphi(N_{i+1})$ is an r_{i+1} -net for $\varphi(N_i)$.
2. Let node $u \in N_i$, and its parent node be p_u . Then, $d(\varphi(u), \varphi(p_u)) \leq r_{i+1}$.

In [7], it was shown that any metric with bounded doubling dimension admits a sparse spanner. This result can be rephrased in terms of net trees in the following theorems.

THEOREM 2.1. ([7]) Given a finite metric $M = (V, d)$ with doubling dimension bounded by \dim . Let $\varepsilon > 0$ and (T, φ) be any net tree for M . For each $i \geq 0$, let

$$E_i := \{\{u, v\} \mid u, v \in \varphi(N_i), d(u, v) \leq (4 + \frac{32}{\varepsilon}) \cdot r_i\}.$$

(Here the parameters N_i, r_i are as in Definition 2.4.) Then $\widehat{E} := \cup_i E_i$ forms a $(1 + \varepsilon)$ -spanner for (V, d) , with the number of edges being $|\widehat{E}| \leq (2 + \frac{1}{\varepsilon})^{O(\dim)} |V|$.

THEOREM 2.2. Consider the construction in Theorem 2.1. For any x, y in V , the spanner \widehat{E} contains a $(1 + \varepsilon)$ -path of the following form. If x_0 and y_0 are the leaf nodes in T with $\varphi(x_0) = x$ and $\varphi(y_0) = y$, and x_i and y_i are the ancestors of x_0 and y_0 at height $i \geq 1$, then there exists i^* such that the path

$$x = \varphi(x_0), \varphi(x_1), \dots, \varphi(x_{i^*}), \\ \varphi(y_{i^*}), \dots, \varphi(y_1), \varphi(y_0) = y$$

is a $(1 + \varepsilon)$ -path (after removing repeated vertices).

3 Adding few extra edges to ensure small hop diameter

Since our spanner in Theorem 2.1 has $O_{k, \varepsilon}(n)$ edges, it is optimal (with respect to n) in terms of the sparsity achieved while preserving shortest path distance. It is easy to check that the number of hops in a $(1 + \varepsilon)$ -path obtained in Theorem 2.2 is $\Theta(\log \Delta)$, where Δ is the *aspect ratio* of the metric (V, d) (i.e., the ratio of the maximum to the minimum pairwise distances). Indeed, the net tree (T, φ) has a height of $\Theta(\log \Delta)$, and in general, a $(1 + \varepsilon)$ -path can have $\Omega(\log \Delta)$ hops.

Before we begin in earnest to investigate how many extra edges are required in order to achieve small hop diameter, let us make a simple observation. For each node u in the tree T , let L_u be the set of leaves under u . For each node u , suppose we add an edge between $\varphi(u)$ and every point in $\varphi(L_u)$. Since the tree has $O(\log \Delta)$ levels, the number of extra edges added is $O(n \log \Delta)$, while the hop diameter of the augmented spanner is at most 3. In the next section, we will build on this idea to show how one can reduce the number of additional edges to $O(n \log n)$ (independent of the aspect ratio Δ) and achieve the same hop-diameter.

3.1 A Warm-up: Obtaining $O(\log n)$ hop-diameter

Notice that Theorem 2.1 holds for any net tree (T, φ) . Hence, by choosing a net tree more carefully, we could possibly improve the trade-off between the hop diameter of the spanner and its size. Indeed, we show in the next theorem that we can improve the parameter $\log \Delta$ to $\log n$ in both cases. (Note that since the metric is doubling, $\log \Delta = \Omega(\log n)$.)

THEOREM 3.1. Suppose (V, d) is a finite metric, where $|V| = n$. Then, there exists a net tree (T, φ) from which the spanner \widehat{E} constructed in the manner described in Theorem 2.1 has the following properties.

1. The hop diameter of the spanner \widehat{E} is $O(\log n)$.
2. It is possible to add $n(\lfloor \log_2 n \rfloor - 1)$ extra edges such that for all leaves $u \in N_0$ in T and any ancestor v of u , there is an edge between $\varphi(u)$ and $\varphi(v)$. (Hence, the hop diameter of the spanner can be reduced to 3.)

Proof. We describe a way to construct a net tree (T, φ) . Let N_0 be the set of leaves for which there is a one-one correspondence φ onto V .

Suppose we have obtained the set N_i of nodes of height i . We would be done if $|N_i| = 1$. Otherwise,

we would obtain an r_{i+1} -net for $\varphi(N_i)$ in the following way. We show a way to greedily construct a net for a set. Start with a list L initially containing all the nodes in N_i , ordered such that a node containing more leaves in its subtree would appear earlier.

As long as the list L is not empty, we repeat the following process. Remove the first node u in the remaining list, form a new node $v \in N_{i+1}$ such that $\varphi(v) := \varphi(u)$ and set the parent of u to be v . For each node w in the remaining list L such that $d(\varphi(w), \varphi(v)) \leq r_{i+1}$, remove w from list L and set the parent of w to be v .

CLAIM 3.1. *For each $x \in N_0$, let A_x be the set of its ancestors in T . Then, $|\varphi(A_x)| \leq \lfloor \log_2 n \rfloor + 1$. In particular, $|\varphi(A_x) \setminus \{\varphi(x)\}| \leq \lfloor \log_2 n \rfloor$.*

Proof. Let a_i be the ancestor of z in N_i . Suppose there exists i such that $\varphi(a_i) \neq \varphi(a_{i+1})$. It follows that the node a_i must have a sibling c , for which $\varphi(c) = \varphi(a_{i+1})$, whose subtree contains at least as many leaves as the subtree at a_i does. Hence, the subtree at a_{i+1} contains at least twice as many leaves as a_i does. Thus there can be at most $\lfloor \log_2 n \rfloor$ values of i for which $\varphi(a_i) \neq \varphi(a_{i+1})$. ■

For the first part of the theorem, it follows that the $(1+\varepsilon)$ -path guaranteed in Theorem 2.1 has the number hops at most $2\lfloor \log_2 n \rfloor + 1$.

For the second part of the theorem, for every $z \in N_0$, we add an edge between $\varphi(z)$ and every point in $\varphi(A_z) \setminus \{\varphi(z)\}$. Note that $|\varphi(A_z) \setminus \{\varphi(z)\}| \leq \lfloor \log_2 n \rfloor$. Suppose y is the lowest ancestor of z such that $\varphi(z) \neq \varphi(y)$, and suppose x is the ancestor of z that is also the child of y . Then, observe that the spanner \hat{E} already includes the edge between $\varphi(y)$ and $\varphi(x) = \varphi(z)$. Hence, for each vertex z , we actually only need to add at most $\lfloor \log_2 n \rfloor - 1$ extra edges. The $(1+\varepsilon)$ -path in Theorem 2.1 can be reduced to $x = \varphi(x_0), \varphi(x_{i^*}), \varphi(y_{i^*}), \varphi(y_0) = y$, which has 3 hops. ■

In the following section, we will investigate the tradeoff between the hop-diameter of a spanner and the number of edges, this time using any given net tree instead.

3.2 The General Upper Bound In this section, we assume that the given metric (V, d) has doubling dimension bounded by k . Given a net tree (T, φ) for the metric, suppose E_T is the spanner obtained in Theorem 2.1. Note that E_T is dependent on the stretch parameter ε . However, for ease of notation, we would leave out the dependency on ε throughout this session.

The approach we use is similar to that used by Arya et al. [4] for Euclidean metrics, which is a subclass of doubling metrics. Instead of using net trees, they worked with “dumbbell trees”, which have similar properties. Applying a construction from [9, 2] to “shortcut” edges in the net-tree, we can show that one can add few extra edges to E_T in order to achieve small hop diameter. Moreover, as shown in [2], this can be done in $O(n \log n)$ time.

We first consider how to add extra edges to a tree such that every pair of nodes has a path with a small number of hops between them.

DEFINITION 3.1. *Define $g(m, n)$ to be the minimum i such that for any tree metric with vertex set V , where $|V| = n$, there exists a spanner P with m edges that preserves all pairwise distances exactly, and for any pair of points, there is a shortest path in P with i hops.*

LEMMA 3.1. *Suppose a metric (V, d) with n points has a net tree (T, φ) , and suppose E_T is the $(1+\varepsilon)$ -spanner obtained in Theorem 2.1. Then, it is possible to add m extra edges to E_T such that the hop diameter of the new spanner is at most $2g(m, n) + 1$.*

Proof. Suppose u is an internal node of T that has a child v such that $\varphi(u) = \varphi(v)$. We contract the edge $\{u, v\}$ by merging the two nodes u and v , and renaming the new node v' such that $\varphi(v') = \varphi(v)$. We repeat the process to obtain the resulting tree (T', φ) . Note that (T', φ) is a tree with V as its vertex set, and is no longer a net tree or a hierarchical tree. However, observe that if u is an ancestor of v in T , then $\varphi(u)$ is an ancestor of $\varphi(v)$ in T' .

Consider the tree T' with unit weights on its edges. By the definition of g , there is a spanner F on T' that preserves all pairwise distances such that for every pair of nodes, there is a shortest path with at most $g(m, n)$ hops. We add the following set of edges to the spanner E_T .
$$E_F := \{\{\varphi(a), \varphi(b)\} : \{a, b\} \in F\}.$$

Suppose x and y points in V , x_0 and y_0 are the leaf nodes in T such that $\varphi(x_0) = x$ and $\varphi(y_0) = y$, and x_i and y_i are the ancestors in T at height i for x_0 and y_0 respectively. By Theorem 2.2, there exists i^* such that the following points form a $(1+\varepsilon)$ -path P_0 , after removing repeated points.

$$\begin{aligned} x &= \varphi(x_0), \varphi(x_1), \dots, \varphi(x_{i^*}), \\ \varphi(y_{i^*}), \dots, \varphi(y_1), \varphi(y_0) &= y \end{aligned}$$

Suppose x_{i^*} and y_{i^*} are contracted to \hat{x} and \hat{y} respectively in T' . By the choice of F , there exist at most $g(m, n) - 1$ intermediate vertices $\{v_i\}_{i=1}^k$ on the

path from x_0 to \hat{x} in T' such that $\{x_0, v_1\}$, $\{v_i, v_{i+1}\}$ ($1 \leq i < k$) and $\{v_k, \hat{x}\}$ are in F . Hence, we have a path with at most $g(m, n)$ hops from x to $\varphi(\hat{x})$: $x = \varphi(x_0), \varphi(v_1), \varphi(v_2), \dots, \varphi(v_k), \varphi(\hat{x})$. Since this sequence of points is a subsequence of $\varphi(x_0), \varphi(x_1), \dots, \varphi(x_{i^*})$, it follows this length of this path is at most that of the sub-path from $\varphi(x_0)$ to $\varphi(x_{i^*})$ in P_0 .

Similarly, there is a path with at most $g(m, n)$ hops from $\varphi(\hat{y})$ to y whose length is at most that of the corresponding sub-path in P_0 . Hence, there is a $(1 + \varepsilon)$ -path with at most $2g(m, n) + 1$ hops from x to y in the spanner $E_T \cup E_F$. ■

THEOREM 3.2. (Chazelle [9]) For $m \geq 2n$, $g(m, n) = O(\alpha(m, n))$, where α is the functional inverse of Ackermann's function.

DEFINITION 3.2. (Ackermann's function [23]) Let $A(i, j)$ be a function defined for integers $i, j \geq 0$ as the following.

$$\begin{aligned} A(0, j) &= 2j && \text{for } j \geq 0 \\ A(i, 0) &= 0, A(i, 1) = 2 && \text{for } i \geq 1 \\ A(i, j) &= A(i-1, A(i, j-1)) && \text{for } i \geq 1, j \geq 2 \end{aligned}$$

Define the function α as $\alpha(m, n) = \min\{i \mid i \geq 1, A(i, 4\lceil m/n \rceil) > \log_2 n\}$.

From Lemma 3.1 and Theorem 3.2, we obtain the following theorem.

THEOREM 3.3. Suppose a metric (V, d) with n points has a net tree (T, φ) , and suppose E_T is the $(1 + \varepsilon)$ -spanner obtained in Theorem 2.1. Then, it is possible to add m extra edges to E_T such that the hop diameter of the new spanner is at most $O(\alpha(m, n))$.

Observing that $A(2, 4 \log^* n) > \log_2 n$, we have the following corollary.

COROLLARY 3.1. Suppose a metric (V, d) with n points has a net tree (T, φ) , and suppose E_T is the $(1 + \varepsilon)$ -spanner obtained in Theorem 2.1. Then, it is possible to add $n \log^* n$ extra edges to E_T such that the hop diameter of the new spanner is $O(1)$.

4 Lower Bound

We now show that the trade-off between the size of the spanner and its hop diameter obtained in Theorem 1.1 is essentially optimal.

THEOREM 4.1. For any $\varepsilon > 0$, for infinitely many integers n , there exists a metric M induced by n points on the real line such that any $(1 + \varepsilon)$ -spanner with m edges on the metric M has hop diameter $\Omega(\alpha(m, n))$.

Our general approach first consider a family of metrics, each of which induced by some binary ‘‘hierarchically well-separated tree’’ (HST). We define a function $G(i, j)$ that is a variant of the Ackermann's function such that if a metric from the family contains $n \geq G(i, j)$ points, then any spanner on the metric with hop diameter bounded by $i + 1$ must have more than $\Omega(jn)$ edges. The relationship between $G(i, j)$ and the Ackermann's function is used to obtain the lower bound for HSTs. The proof technique we used is an extension of that used in Yao's paper [25]. Our lower bound result for line metrics then follows from the fact that binary HSTs with large separation embed into the real line with small distortion.

REMARK 4.1. For technical reasons, we assume that a spanner contains a self-loop for every point. Since any spanner must contain a linear number of edges, this assumption does not affect the asymptotic lower bound.

Construction of the family of HST metrics. For $k \geq 0$, let M_k be the metric induced by the 2^k leaves of the weighted complete binary tree T_k defined as follows. Let $\beta > 0$ be the separation parameter for the HST. The tree T_k is a binary tree containing 2^k leaves such that for each internal node u at height $h \geq 1$, the distance from u to any of the leaves in the subtree rooted at u is β^{h-1} .

The following proposition follows from the construction of the metrics M_k .

PROPOSITION 4.1. Let the HST metric M_k be defined as above.

- (a) Suppose M_k is constructed with separation $\beta \geq 100(1 + \varepsilon)$. Let U be the subset of points corresponding to the leaves of T_k which are the descendants of some internal node. Then, any $(1 + \varepsilon)$ -path between points in U cannot contain any point outside U .
- (b) Consider T_k and suppose $h \leq k$. Suppose T' is the tree obtained from T_k by replacing each subtree rooted at an internal node of height h by a leaf whose distance from the root is the same as before, i.e., β^{h-1} . Then, T' is isomorphic to T_{k-h} .
- (c) For every $k \geq 0$, the metric M_k with expansion $\beta \geq 4$ has doubling dimension at most 2.

We will use Proposition 4.1(a) crucially in our analysis. Unless otherwise stated, we assume the HST metric M_k is always constructed with separation β large enough such that the statement holds.

We prove the following theorem that states the lower bound result for the HST metrics.

THEOREM 4.2. For each integer $k \geq 1$ and any $\varepsilon > 0$, there exists an HST metric M_k with large enough

separation β such that any $(1 + \varepsilon)$ -spanner on M_k with at most m edges has hop diameter at least $\Omega(\alpha(m, n))$.

We observe that HST metrics with large separation embed into the real line with small distortion in the following claim.

CLAIM 4.1. *For each integer $k \geq 1$ and any $\rho > 0$, for sufficiently large $\beta > 0$, the HST metric M_k with separation β embeds into the real line with distortion at most $1 + \rho$.*

Proof. We embed the leaves associated with M_k into the real line in their natural ordering, i.e. leaves in the subtree rooted at some internal node are clustered together in the line. The distance between embedded points is the same as that between them in the tree. Such an embedding does not contract distances.

Consider the expansion of the distance between a pair of leaves whose lowest common ancestor is at height r . Hence, their distance in the tree is $2\beta^r$. Observe that their embedded distance is at most $2 \cdot \{2^r + 2^{r-1}\beta + \dots + 2\beta^{r-1} + \beta^r\}$. Hence, the distortion is at most

$$\begin{aligned} \frac{2^r + 2^{r-1}\beta + \dots + 2\beta^{r-1} + \beta^r}{\beta^r} &= \frac{2^r}{\beta^r} \cdot \frac{(\beta/2)^r - 1}{\beta/2 - 1} + 1 \\ &\leq \frac{1}{\beta/2 - 1} + 1, \end{aligned}$$

which is at most $1 + \rho$ for $\beta \geq 2(1 + \frac{1}{\rho})$. \blacksquare

Hence, the main theorem of this section follows from Theorem 4.2 and Claim 4.1.

Proof of Theorem 4.1: Suppose $n = 2^k$ is a power of two. We construct a line metric M with n points. Let $\varepsilon' = 2\varepsilon$ and $\rho > 0$ be small enough such that $(1 + \varepsilon)(1 + \rho) \leq 1 + \varepsilon'$. Suppose the HST metric M_k has large enough separation β such that by Theorem 4.2, any $(1 + \varepsilon')$ -spanner for M_k with m edges has hop diameter $\Omega(\alpha(m, n))$, and by Claim 4.1, M_k embeds into some line metric M with distortion at most $1 + \rho$.

Suppose P is a $(1 + \varepsilon)$ -spanner for metric M with m edges and hop diameter at most D . Since $(1 + \varepsilon)(1 + \rho) \leq 1 + \varepsilon'$, it follows spanner P corresponds to a $(1 + \varepsilon')$ -spanner in M_k with m edges and hop diameter at most D . Therefore, $D = \Omega(\alpha(m, n))$. \blacksquare

The rest of the section obtains the lower bound result for the HST metrics. We define a variant of the Ackermann's function.

DEFINITION 4.1. *Define the function $G(i, j)$, for $i \geq 0, j \geq 0$ to be:*

$$\begin{aligned} G(0, 0) &= 0, G(0, j) = 2^{\lceil \log_2 j \rceil}; & j \geq 1 \\ G(i, 0) &= 0, G(i, 1) = 1; & i \geq 1 \\ G(i, j) &= G(i, j-1)G(i-1, 4G(i, j-1)); & i \geq 1, j \geq 2 \end{aligned}$$

PROPOSITION 4.2. *Suppose $G(i, j)$ is the function defined as above.*

- (a) *For all $i \geq 0, j \geq 1$, $G(i, j)$ is a power of two.*
- (b) *For $j \geq 1, j \leq G(0, j) \leq 2j$.*

We next prove the main technical lemma for the lower bound for the HST metrics.

LEMMA 4.1. *Suppose $2^k \geq G(i, j)$, where $i \geq 0$ and $j \geq 1$; suppose $\varepsilon > 0$ and the HST metric M_k has large enough separation β . Suppose X is a subset of M_k such that $|X| = n \geq 1$. Let $\rho = n/2^k$. Then, any $(1 + \varepsilon)$ -spanner for X with hop diameter at most $i + 1$ must have more than $\frac{1}{4}\rho j n$ edges.*

Proof. We prove the result by induction on the lexicographical order of (i, j) .

Base cases. For $i = 0, j \geq 1$, any spanner with hop diameter 1 on n points must have exactly $\frac{1}{2}n(n-1) + n$ edges, recalling that we require that a spanner must contain a self-loop for each point. Hence, observing that $j \leq G(0, j) \leq 2^k$ from Proposition 4.2, we conclude that such a spanner cannot have the number of edges less than $\frac{1}{4}\rho j n \leq \frac{1}{4}n^2 < \frac{1}{2}n(n-1) + n$.

For $i \geq 1, j = 1$, we observe that any spanner on n points must have at least n edges. Hence, the number of edges in a spanner cannot be less than $\frac{1}{4}\rho n \leq \frac{1}{4}n < n$.

Inductive Step. Suppose X is a subset of M_k such that $2^k \geq G(i, j)$ for some $i \geq 1$ and $j \geq 2$, where $|X| = n$ and $\rho = n/2^k$. For contradiction's sake, assume there is a $(1 + \varepsilon)$ -spanner E with hop diameter $i + 1$ for X such that $|E| \leq \frac{1}{4}\rho j n$.

Let I be the indexing set for the subtrees of T_k , each rooted at some internal node and containing exactly $G(i, j-1)$ leaves. Observing that $G(i, j-1)$ is a power of 2 from Proposition 4.2, it follows that

$$\begin{aligned} |I| &= 2^k / G(i, j-1) \geq G(i, j) / G(i, j-1) \\ &= G(i-1, 4G(i, j-1)). \end{aligned}$$

For each $s \in I$, let V_s be the set of leaves contained in the corresponding sub-tree. Let us also define:

- $E_s^1 := \{\{u, v\} \in E : u, v \in V_s\}$, for each $s \in I$, and $E^1 := \cup_{s \in I} E_s^1$.
- $E^2 := \{\{u, v\} \in E : u \in V_s, v \in V_t, s \neq t\}$.

We describe the high level idea to obtain a contradiction. Suppose for each $s \in I$, we replace the subtree containing V_s by a leaf in the same manner as Proposition 4.1(b), then we would obtain a tree T' which is isomorphic to $T_{\hat{k}}$, where $2^{\hat{k}} = |I| \geq G(i-1, 4G(i, j-1))$.

Let $X_s := X \cap V_s$ and $J := \{s \in I : |X_s| \geq 1\}$. Identifying each X_s 's with the corresponding leaf in the modified tree T' , consider the submetric of M_k induced by the non-empty X_s 's, whose point set we write as $X' := \{X_s : s \in J\}$. Hence, X_s is a subset of metric M_k , as well as a point in metric X' .

Define $E' := \{\{X_s, X_t\} : \{u, v\} \in E^2, u \in X_s, v \in X_t\}$. Observe that E' is a $(1 + \varepsilon)$ -spanner for X' with hop diameter at most $i + 1$. Since we wish to apply the induction hypothesis, we need to show that the size of E' is small. Moreover, since $|I| \geq G(i - 1, 4G(i, j - 1))$, the induction hypothesis can only say about spanners of hop diameter at most i . To resolve this issue, we would remove some points in X' and modify the spanner appropriately such that its hop diameter is at most i . First observing that $|E'| \leq |E^2|$, it suffices to show that $|E^2|$ is small.

CLAIM 4.2. $|E^2| < \frac{1}{4}\rho n$.

Proof. Let $|X_s| = n_s$ and $\rho_s = n_s/G(i, j - 1)$. Observe from Proposition 4.1(a) that for each $s \in I$, any $(1 + \varepsilon)$ -path between vertices inside X_s cannot go outside X_s . Hence, for $n_s \geq 1$, it follows E_s^1 is a spanner for X_s having hop diameter at most $i + 1$. Applying the induction hypothesis for $(i, j - 1)$, we have for each s , $|E_s^1| > \frac{1}{4}\rho_s(j - 1)n_s$. Summing over $s \in I$, we have

$$|E^1| > \sum_{s \in I} \frac{1}{4}\rho_s(j - 1)n_s \geq \frac{1}{4} \cdot \frac{j - 1}{G(i, j - 1)} \sum_{s \in I} n_s^2.$$

Observing that $\sum_{s \in I} n_s = n$ and the fact that $x \mapsto x^2$ is a convex function, the last term is minimized when all n_s 's are equal. Hence,

$$|E^1| > \frac{j - 1}{4G(i, j - 1)} \cdot |I| \cdot \left(\frac{n}{|I|}\right)^2 = \frac{1}{4}(j - 1)\rho n.$$

Since there are at most $\frac{1}{4}\rho j n$ edges in total, it follows that $|E^2| < \frac{1}{4}\rho n$. ■

Next, we describe a procedure that removes some points from X' and modify E' to obtain a spanner with hop diameter at most i . Note that the points from X' are indexed by J . The procedure labels the removed points *bad*.

1. Place the index set J in a list L in an arbitrary order.
2. Consider each element s in list L according to the ordering:
 - (a) If there exists an element t appearing after s in the list L such that any $(1 + \varepsilon)$ -path in E' between X_s and X_t takes at least $i + 1$ hops,
 - (i) Label s *bad* and remove it from list L .
 - (ii) Modify E' so that if X_p is a point in list L closest to X_s , every edge incident on X_s will now be incident on X_p , i.e., X_s and X_p are merged.
 - (b) Move on to the next element in list L .

Any two remaining points certainly have a $(1 + \varepsilon)$ -path with at most i hops; otherwise, the one appearing earlier in the list would have been removed. Moreover, observe in step (ii) of the procedure that X_s and X_p are equidistant from any other X_q 's in the list. Hence, the length of any $(1 + \varepsilon)$ -path for two points still in the list does not increase. Moreover, since we have merged X_s with X_p , the number of hops for any $(1 + \varepsilon)$ -path cannot increase.

Let B be the set of $s \in J$ that are labelled bad. Let $R := J - B$ be the set of remaining indices. Let \widehat{E} be the modified edge set. It follows that \widehat{E} is a spanner with hop diameter at most i for $\widehat{X} := \{X_s : s \in R\}$. However, we need to show that not too many bad points are removed.

CLAIM 4.3. $\sum_{s \in R} |X_s| \geq \frac{1}{2}n$.

Proof. For each $s \in B$, there exists $t \in J$ such that any $(1 + \varepsilon)$ -path between X_s and X_t in E' has at least $i + 1$ hops. Fix $b \in X_t$ and consider any $a \in X_s$, observe that there is a $(1 + \varepsilon)$ -path $P: a = v_0, v_1, \dots, v_l = b$ in E such that $l \leq i + 1$. For each v , let $\varphi(v)$ be the unique X_q that contains it. Then, it follows there is a $(1 + \varepsilon)$ -path $P': X_s = \varphi(v_0), \varphi(v_1), \dots, \varphi(v_l) = X_t$, after removing redundant X_q 's. Hence, $l = i + 1$ and there are no redundant X_q 's, otherwise there would be a $(1 + \varepsilon)$ -path from X_s to X_t with less than $i + 1$ hops. We associate $a \in X_s$ with the edge $\{a, v_1\} \in E^2$.

It follows for each $s \in B$ and each $a \in X_s$, there exists some edge $\{a, v\} \in E^2$. Each edge can be associated with at most two points in the bad X_s 's. Hence, we obtain the following.

$$\sum_{s \in B} |X_s| \leq 2|E^2| < \frac{1}{2}\rho n \leq \frac{1}{2}n,$$

where the middle inequality follows from Claim 4.2. Hence, it follows that $\sum_{s \in G} |X_s| \geq \frac{1}{2}n$. ■

We can now obtain a contradiction to the induction hypothesis for $(i-1, 4G(i, j-1))$, which states that if \hat{X} is a sub-metric of $T_{\hat{k}}$ such that $2^k \geq G(i-1, 4G(i, j-1))$ and $\hat{\rho} = |\hat{X}|/2^k$, then any $(1 + \varepsilon)$ -spanner for \hat{X} with hop diameter at most i must have more than $\frac{1}{4}\hat{\rho}(4G(i, j-1))|\hat{X}|$ edges.

Now, since for each $s \in R$, $|X_s| \leq G(i, j-1)$, it follows from Claim 4.3 that $|\hat{X}| = |R| \geq \frac{1}{2}n/G(i, j-1)$. Hence, $\hat{\rho} := |R|/|\hat{X}| \geq \frac{1}{2}\rho$. Moreover, $n = \rho G(i, j-1)|I| \leq 2|\hat{X}|G(i, j-1)$.

In conclusion, we have a subset \hat{X} in the metric $T_{\hat{k}}$ such that $2^k = |I| \geq G(i-1, 4G(i, j-1))$ and $\hat{\rho} = |\hat{X}|/|I| \geq \rho/2$. Moreover, \hat{E} is a $(1 + \varepsilon)$ -spanner for \hat{X} with hop diameter at most i and has the number of edges less than:

$$\frac{1}{4}\rho n \leq \frac{1}{4} \cdot (2\hat{\rho}) \cdot 2|\hat{X}|G(i, j-1) = \frac{1}{4}\hat{\rho}(4G(i, j-1))|\hat{X}|,$$

obtaining the desired contradiction against the induction hypothesis for $(i-1, 4G(i, j-1))$. This completes the inductive step of the proof. ■

If we substitute $\rho = 1$ in Lemma 4.1, we obtain the following corollary.

COROLLARY 4.1. *Suppose $n = 2^k \geq G(i, j)$, $j \geq 1$. Let $\varepsilon > 0$ and the HST metric M_k have large enough separation β . Then, any $(1 + \varepsilon)$ -spanner for M_k with hop diameter at most $i + 1$ must have more than $\frac{1}{4}jn$ edges.*

In order to get the desired lower bound on the hop diameter, we have to relate the function $G(i, j)$ to the Ackermann function $A(i, j)$.

DEFINITION 4.2. *Define the function $H(i, j)$, for $i \geq 0, j \geq 0$ to be:*

$$\begin{aligned} H(0, j) &= 8j^3 && \text{for } j \geq 0 \\ H(i, 0) &= 0, H(i, 1) = 8 && \text{for } i \geq 1 \\ H(i, j) &= H(i-1, H(i, j-1)) && \text{for } i \geq 1, j \geq 2 \end{aligned}$$

CLAIM 4.4. *Let $H(i, j)$ be as defined above.*

- (a) *For $i \geq 0, j \geq 0$, $H(i, j) \leq A(i+4, j+4) - 4$. In particular, $H(i, j) \leq A(i+4, j+4)$.*
- (b) *For $i \geq 0, j \geq 0$, $H(i, j) \geq 4j^2G(i, j)$. In particular, $H(i, j) \geq G(i, j)$.*

Proof. We prove both results by induction on the lexicographic order of (i, j) . Let us prove the claim of part (a) first.

Base cases. For $j \geq 0$, $H(0, j) = 8j^3 \leq A(4, j+4) - 4$. For $i \geq 1$, $H(i, 0) = 0 \leq A(i+4, 4) - 4$ and $H(i, 1) = 8 \leq A(i+4, 5) - 4$.

Inductive step. Suppose $i \geq 1, j \geq 2$. Then, using the induction hypothesis, we have

$$\begin{aligned} H(i, j) &= H(i-1, H(i, j-1)) \\ &\leq A(i+3, H(i, j-1) + 4) - 4 \\ &\leq A(i+3, A(i+4, j+3)) - 4 \\ &= A(i+4, j+4) - 4, \end{aligned}$$

which completes the inductive step of the first result.

We next prove the claim of part (b).

Base cases. For $j \geq 0$, $H(0, j) = 8j^3 \geq 4j^2G(0, j)$, by Proposition 4.2(b). For $i \geq 1$, $H(i, 0) \geq 8 \cdot 0^2G(i, 0)$, as both sides are zero; $H(i, 1) = 8 \geq 4 = 4G(i, 1)$.

Inductive step. Suppose $i \geq 1, j \geq 2$. Then, using the induction hypothesis, we have

$$\begin{aligned} H(i, j) &= H(i-2, H(i, j-1)) \\ &\geq 4H(i, j-1)^2G(i-1, H(i, j-1)) \\ &\geq 4H(i, j-1)^2G(i-1, 4(j-1)^2G(i, j-1)) \end{aligned}$$

Observe that since $i \geq 1$ and $j \geq 2$, $H(i, j-1) \geq 2^{j-1} \geq j$. Hence, $H(i, j) \geq 4j^2G(i-1, 4G(i, j-1)) = 4j^2G(i, j)$, completing the induction step of the second result. ■

The following claim describes some properties of the Ackermann's function and its functional inverse.

CLAIM 4.5. *Suppose $a(x, j) := \min\{i \mid i \geq 1, A(i, j) > x\}$, where $x \geq 0, j \geq 0$.*

- (a) *For all $j \geq 0$, if $x \geq y \geq 0$, then $a(x, j) \geq a(y, j)$. In particular, $a(x, j) \geq a(\log_2 x, j)$.*
- (b) *For $k \geq 1$ and $x \geq 0$, $a(x, 4k+4) + 1 \geq a(x, 4k)$.*

Proof. The first statement follows trivially from the fact that the Ackermann's function $A(i, j)$ is monotone. For the proof of the second statement, suppose $i = a(x, 4k+4)$. Hence, $i \geq 1$ and $A(i, 4k+4) > x$. Observe that $A(i+1, 4k) = A(i, A(i+1, 4k-1))$ and $A(i+1, 4k-1) \geq 2^{4k-1} \geq 4k+4$, since $k \geq 1$ and $i \geq 1$. Hence, it follows that $A(i+1, 4k) \geq A(i, 4k+4) > x$ and thus $a(x, 4k) \leq a(x, 4k+4) + 1$, as required. ■

Hence, we can obtain the lower bound result for the HST metrics.

Proof of Theorem 4.2: Suppose E is a $(1 + \varepsilon)$ -spanner E for M_k . Let $j = \lceil \frac{4m}{n} \rceil$. Then, by Corollary 4.1, since $m \leq \frac{1}{4}jn$, if $G(i, j) \leq n$, the hop diameter of E is larger than $i + 1$. Hence, the hop diameter of E is at least

$$\begin{aligned}
& \min\{i + 1 \mid G(i, \lceil \frac{4m}{v} \rceil) > n\} \\
\geq & \min\{i + 1 \mid H(i, 4 \lceil \frac{m}{n} \rceil) > n\} && \text{(Claim 4.4(b))} \\
\geq & \min\{i + 1 \mid A(i + 4, 4 \lceil \frac{m}{n} \rceil + 4) > n\} && \text{(Claim 4.4(a))} \\
= & \min\{i \mid A(i, 4 \lceil \frac{m}{n} \rceil + 4) > n\} - 3 \\
= & a(n, 4 \lceil \frac{m}{n} \rceil + 4) - 3 \\
\geq & a(n, 4 \lceil \frac{m}{n} \rceil) - 4 && \text{(Claim 4.5(b))} \\
\geq & a(\log_2 n, 4 \lceil \frac{m}{n} \rceil) - 4 && \text{(Claim 4.5(a))} \\
= & \alpha(m, n) - 4,
\end{aligned}$$

which completes the proof. \blacksquare

References

- [1] Ittai Abraham and Dahlia Malkhi. Compact routing on euclidian metrics. In *PODC*, pages 141–149, 2004.
- [2] N. Alon and B. Schieber. Optimal preprocessing for answering on-line product queries. *Tech. report 71/87*, 1987.
- [3] Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete Comput. Geom.*, 9(1):81–100, 1993.
- [4] Sunil Arya, Gautam Das, David M. Mount, Jeffrey S. Salowe, and Michiel H. M. Smid. Euclidean spanners: short, thin, and lanky. *STOC*, pages 489–498, 1995.
- [5] Patrice Assouad. Plongements lipschitziens dans \mathbf{R}^n . *Bull. Soc. Math. France*, 111(4):429–448, 1983.
- [6] Paul B. Callahan and S. Rao Kosaraju. A decomposition of multidimensional point sets with applications to k -nearest-neighbors and n -body potential fields. *J. Assoc. Comput. Mach.*, 42(1):67–90, 1995.
- [7] Hubert T-H. Chan, Anupam Gupta, Bruce M. Maggs, and Shuheng Zhou. On hierarchical routing in doubling metrics. *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 762–771, 2005.
- [8] Barun Chandra, Gautam Das, Giri Narasimhan, and José Soares. New sparseness results on graph spanners. *Internat. J. Comput. Geom. Appl.*, 5(1-2):125–144, 1995. Eighth Annual ACM Symposium on Computational Geometry (Berlin, 1992).
- [9] Bernard Chazelle. Computing on a free tree via complexity-preserving mappings. In *Algorithmica 2*, pages 337–361, 1987.
- [10] K. L. Clarkson. Nearest neighbor queries in metric spaces. *Discrete Comput. Geom.*, 22(1):63–93, 1999.
- [11] Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 534–543, 2003.
- [12] Sariel Har-Peled and Manor Mendel. Fast construction of nets in low dimensional metrics, and their applications. *Symposium on Computational Geometry*, pages 150–158, 2005.
- [13] J. Heinonen. *Lectures on analysis on metric spaces*. Universitext. Springer-Verlag, New York, 2001.
- [14] David R. Karger and Matthias Ruhl. Finding nearest neighbors in growth-restricted metrics. In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computing*, pages 63–66, 2002.
- [15] Robert Krauthgamer and James R. Lee. The black-box complexity of nearest neighbor search. In *ICALP*, pages 858–869, 2004.
- [16] Robert Krauthgamer and James R. Lee. Navigating nets: Simple algorithms for proximity search. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 791–801, 2004.
- [17] Vincent D. Park and M. Scott Corson. A highly adaptive distributed routing algorithm for mobile wireless networks. In *INFOCOM*, pages 1405–1413, 1997.
- [18] D. Peleg and A. Schäffer. Graph spanners. In *J. Graph Theory 13*, pages 99–116, 1989.
- [19] Charles E. Perkins and Elizabeth M. Belding-Royer. Ad-hoc on-demand distance vector routing. In *WM-CSA*, pages 90–100, 1999.
- [20] Charles E. Perkins and Pravin Bhagwat. Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. In *SIGCOMM*, pages 234–244, 1994.
- [21] J.S. Salowe. Constructing multidimensional spanner graphs. In *Internat. J. Comput. Geom. Appl. 1*, pages 99–107, 1991.
- [22] Kunal Talwar. Bypassing the embedding: Algorithms for low-dimensional metrics. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 281–290, 2004.
- [23] R. E. Tarjan. Efficiency of a good but not linear set union algorithm. *Journal of ACM 22*, pages 215–225, 1975.
- [24] P.M. Vaidya. A sparse graph almost as good as the complete graph on points in k dimensions. In *Discrete Comput. Geom. 6*, pages 369–381, 1991.
- [25] Andrew C. Yao. Space-time tradeoff for answering range queries. In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing*, pages 128–136, 1982.