

Distributed Spanner Construction in Doubling Metric Spaces

Mirela Damian¹, Saurav Pandit², and Sriram Pemmaraju²

¹ Department of Computer Science, Villanova University, Villanova, PA 19085
mirela.damian@villanova.edu.

² Department of Computer Science, The University of Iowa, Iowa City, IA 52242-1419
{spandit, sriram}@cs.uiowa.edu.

Abstract. This paper presents a distributed algorithm that runs on an n -node unit ball graph (UBG) G residing in a metric space of constant doubling dimension, and constructs, for any $\varepsilon > 0$, a $(1 + \varepsilon)$ -spanner H of G with maximum degree bounded above by a constant. In addition, we show that H is “lightweight”, in the following sense. Let Δ denote the aspect ratio of G , that is, the ratio of the length of a longest edge in G to the length of a shortest edge in G . The total weight of H is bounded above by $O(\log \Delta) \cdot wt(MST)$, where MST denotes a minimum spanning tree of the metric space. Finally, we show that H satisfies the so called *leapfrog property*, an immediate implication being that, for the special case of Euclidean metric spaces with fixed dimension, the weight of H is bounded above by $O(wt(MST))$. Thus, the current result subsumes the results of the authors in PODC 2006 that apply to Euclidean metric spaces, and extends these results to metric spaces with constant doubling dimension.

1 Introduction

A *unit ball graph* (UBG) is a graph whose vertices reside in some metric space and whose edges connect pairs of vertices at distance at most 1. The *doubling dimension* of a metric space is the smallest ρ such that any ball in this metric space can be covered by 2^ρ balls of half the radius. It is easy to verify that the d -dimensional Euclidean space, equipped with any of the L_p norms, has doubling dimension $\Theta(d)$. If ρ is a fixed constant (independent of the size of the UBG), then we call the UBG a *doubling UBG*. A t -spanner of a graph G is a spanning subgraph H of G such that, for all pairs of vertices $u, v \in V$, the length of a shortest uv -path in H is at most t times the length of a shortest uv -path in G . In this paper we present a distributed algorithm for constructing a low-weight, $(1 + \varepsilon)$ -spanner of bounded degree for doubling UBGs.

Precisely stated, our result is this: for any fixed $\varepsilon > 0$, our algorithm runs in $O(\log^* n)$ communication rounds on an n -node doubling UBG G , to construct a $(1 + \varepsilon)$ -spanner H of G with maximum degree bounded above by a constant. This constant depends on ε and ρ , the doubling dimension of the metric space in which G resides. Recall that $\log^* n = \min\{t \mid \log^{(t)} n \leq 2\}$, where $\log^{(0)} n = n$ and $\log^{(i)} n = \log(\log^{(i-1)} n)$ for any positive integer i . In addition, we show

that H is “lightweight,” in the following sense. Let Δ denote the aspect ratio of G , that is, the ratio of the length of a longest edge in G to the length of a shortest edge in G . We show that the total weight of H is bounded above by $O(\log \Delta) \cdot wt(MST)$, where MST denotes a minimum spanning tree of G . Thus we obtain a spanner that provides an $O(\log \Delta)$ -approximation to a spanner of G of minimum weight. Finally, we also show that H satisfies the so called *leapfrog property* [8], which informally says that any uv -path in H that does not include $\{u, v\}$ must have length greater than $\{u, v\}$ by a constant factor. An immediate implication of this property is that, for the special case of Euclidean metric spaces with fixed dimension, the weight of H is bounded above by $O(wt(MST))$ [7]. Thus, our current result subsumes the results in [6] that apply to Euclidean metric spaces, and extends these results to metric spaces with constant doubling dimension.

1.1 Topology Control

Our result is motivated by the *topology control* problem in wireless ad-hoc networks. For an overview of topology control, see the survey by Rajaraman [16]. Since an ad-hoc network does not come with fixed infrastructure, there is no topology to start with and informally speaking, the topology control problem is one of selecting neighbors for each node so that the resulting topology has a number of useful properties such as sparseness, small weight, or maximum vertex degree bounded above by a constant. Most topology control protocols that provide worst case guarantees on the quality of the topology assume that the network is modeled by a unit disk graph (UDG) (see [14] for a recent example). The results in this paper apply to the more general model of doubling unit ball graphs (UBG). Doubling metric spaces have received a great deal of attention recently [4,11,12,13,17], partly because they are thought to capture real-world phenomena such as latencies in peer-to-peer networks and in the Internet. Also, doubling metrics are robust in the sense that the doubling dimension is roughly preserved under distortion (see Proposition 3 in [17]). Thus distorted versions of low dimensional Euclidean space also have small doubling dimension. Consequently, doubling UBGs can model wireless networks in which nodes have non-uniform transmission ranges or have erroneous perception of distances to other nodes. Finally, doubling metrics imply the following “bounded growth” phenomenon that seems to be characteristic of large scale wireless ad-hoc and sensor networks: the number of nodes that are far away from each other and yet are all in the vicinity of a particular node, is small. In other words, no node can have an arbitrarily large independent set in its neighborhood.

1.2 Net Trees

Let (V, d) be a metric space with $|V| = n$ and doubling dimension ρ . In a recent paper, Chan, Gupta, Maggs, and Zhou [2] show how to construct, via a sequential, polynomial-time algorithm, a $(1 + \varepsilon)$ -spanner of (V, d) with maximum degree bounded above by $(\frac{1}{\varepsilon})^{O(\rho)}$. We will refer to this algorithm as the CGMZ

algorithm. The problem of constructing a spanner for a metric space can be thought of as a special case of our problem, in which the given UBG is a complete graph. Underlying the result in [2] is the notion of *net trees*, independently proposed by Har-Peled and Mendel [10]. Let $B(u, r)$ denote the ball of radius r centered at point u . A subset $U \subseteq V$ is an r -net of V if it satisfies two properties:

- r*-packing: For every u and v in U , $d(u, v) > r$.
- r*-covering: The union $\cup_{u \in U} B(u, r)$ covers V .

Such nets always exist for any $r > 0$, and can be easily computed using a greedy algorithm. Assume without loss of generality that the largest pairwise distance in V is exactly 1 (this can be achieved by appropriate scaling). Let α , with $\sqrt{1 + \varepsilon} \leq \alpha$, and $\gamma = \frac{2\alpha}{\alpha - 1} (1 + \frac{4\alpha}{\varepsilon})$ be constants (we use the fact that $\sqrt{1 + \varepsilon} \leq \alpha$ in the proof of Lemma 7). Let h be the smallest positive integer such every pairwise distance is greater than $\frac{1}{\alpha^h}$. Let $r_0 = \frac{1}{\alpha^h}$ and let $r_i = \alpha \cdot r_{i-1}$, for $i > 0$. A *net tree* is a sequence of subsets $\langle V_0, V_1, V_2, \dots, V_h \rangle$, such that $V_0 = V$ and V_i is an r_i -net of V_{i-1} , for $i > 0$. Note that every V_i , including V_0 , is an r_i -packing. Also note that V_h , which is a 1-net of V_{h-1} , is a singleton, since the maximum separation between any pair of points is 1. To view the sequence $\langle V_0, V_1, V_2, \dots, V_h \rangle$ as a tree, let $i(v) = \max\{i \mid v \in V_i\}$ for each $v \in V$. Then, for each $v \in V$, $i(v) + 1$ copies of v appear as nodes in the tree. These are denoted $(0, v), (1, v), \dots, (i(v), v)$, where (i, v) represents the occurrence of v in V_i . For each $0 \leq i < i(v)$, the parent of node (i, v) is $(i+1, v)$. Node $(i(v), v)$ has no parent and is the root of the net tree, if $i(v) = h$; otherwise, vertex $v \notin V_{i(v)+1}$ and there is some vertex $u \in V_{i(v)+1}$ such that $B(u, r_{i(v)+1})$ contains v . Arbitrarily pick one such u and let $(i(v) + 1, u)$ be the parent of $(i(v), v)$. Informally speaking,




Fig. 1. A net tree with six levels

higher levels in the net tree (leaves are at level 0) represent the structure of V at lower resolution. Figure 1 shows an example of a net tree with 6 levels. Below we present the CGMZ algorithm [2]. For any two points $u, v \in V$, we use $d(u, v)$ to denote the distance between u and v in the underlying metric space.

The CGMZ Algorithm

1. Build a net tree $\langle V_0, V_1, \dots, V_h \rangle$ of V .
2. Let $\gamma = \frac{2\alpha}{\alpha-1} \left(1 + \frac{4\alpha}{\varepsilon}\right)$. Construct the edge sets

$$E_0 = \{\{u, v\} \in V_0 \times V_0 \mid d(u, v) \leq \gamma \cdot r_0\},$$

and

$$E_i = \{\{u, v\} \in V_i \times V_i \mid \gamma \cdot r_{i-1} < d(u, v) \leq \gamma \cdot r_i\},$$

for each $i = 1, \dots, h$ and let $\widehat{E} = \cup_i E_i$.

3. Replace some edges in \widehat{E} by other edges to obtain a new edge set \widetilde{E} .

Chan and coauthors [2] work with the version of the algorithm for $\alpha = 2$. They show that the graph $H = (V, \widehat{E})$ obtained after Step (2) is a $(1 + \varepsilon)$ -spanner of the metric space and has linear number of edges, but may not satisfy the bounded degree requirement. Short paths in H can be obtained from the net tree in a natural manner. A uv -path in H whose length is at most $(1 + \varepsilon) \cdot d(u, v)$ can be obtained by traveling up the net tree from the leaf u and from the leaf v until some level i is reached, such that the ancestors of u and v at level i are connected by an edge in H . In Step (3), a subset of the edges in \widehat{E} is considered and each edge in this subset is replaced by at most one new edge. This step, which will be described in detail in Section 2.2, redistributes the edges so that all vertex-degrees are bounded above by a constant. The techniques used by Chan and coauthors for bounding vertex degrees play a critical role in this paper as well. In [6] we also describe an algorithm for constructing a bounded-degree $(1 + \varepsilon)$ -spanner for Euclidean UBGs, but our results rely on purely geometric arguments to bound the vertex degree of the constructed spanner. Chan and coauthors [2] obtain the following theorem.

Theorem 1. [Chan, Gupta, Maggs, Zhou] *Let (V, d) be a finite metric with doubling dimension bounded by ρ . For any $\varepsilon > 0$, there is a $(1 + \varepsilon)$ -spanner for (V, d) , with maximum degree bounded above by $\left(\frac{1}{\varepsilon}\right)^{O(\rho)}$.*

Our algorithm is a modification of the CGMZ algorithm [2] that takes into account the fact that pairs of points separated by a distance greater than 1 are not connected by an edge and therefore such edges cannot be used in the spanner. A high level view of our algorithm is that it uses a slightly modified version of the CGMZ algorithm and constructs a graph H that may contain some *virtual edges*, that is, edges of length more than 1. H has all the desired properties with respect to the input UBG G . Subsequently, we show how to replace each virtual edge in H by at most one *real edge*, that is, an edge of length at most 1. The resulting graph is a $(1 + \varepsilon)$ -spanner of G with degree bounded above by a constant.

To obtain a distributed implementation of the above idea in $O(\log^* n)$ rounds, we use an algorithm due to Kuhn, Moscibroda, and Wattenhofer [13]. For a given doubling n -node UBG G , the algorithm in [13] deterministically computes a $(1, O(1))$ -network decomposition, that is, a partition of G into clusters such

that each cluster has diameter 1 and the resulting cluster graph has chromatic number $O(1)$. We use the same algorithm to compute a net tree. After computing the net tree, we require a constant number of additional rounds to construct the spanner.

2 Spanners for Doubling UBGs

Let (V, d) be a metric space with doubling dimension ρ . Let $G = (V, E)$ be the UBG induced by this metric space. Thus, for all $u, v \in V, u \neq v, \{u, v\} \in E$ if and only if $d(u, v) \leq 1$. For a fixed $\varepsilon > 0$, let the quantities h, r_i, α and γ be defined as in Section 1.2. Run Steps (1) and (2) of the CGMZ Algorithm to construct a set of edges \widehat{E} . Let $H = (V, \widehat{E})$. Note that V_h may not be a singleton since V may contain points whose pairwise distance is more than 1. So the sequence $\langle V_0, V_1, \dots, V_h \rangle$ should be viewed as a forest of net trees, rooted at points in V_h . Recall that $\widehat{E} = \cup_{i=0}^h E_i$ and further recall that for $i > 0, E_i$ consists of edges connecting all pairs of points $u, v \in V$ such that $d(u, v) \in (\gamma \cdot r_{i-1}, \gamma \cdot r_i]$. Note that there are values of i for which the right endpoint of the interval $(\gamma \cdot r_{i-1}, \gamma \cdot r_i]$ may be greater than 1 and for such values of i, E_i may contain edges that are not in E . Thus H is not necessarily a subgraph of G . Let $\delta = \lceil \log_\alpha \gamma \rceil$. It is easy to verify that for $0 \leq i \leq h - \delta, E_i \subseteq E$; for $i = h - \delta + 1$, the edge-set E_i may contain some edges in E and some edges not in E ; and for $i > h - \delta + 1$, all edges in E_i are outside E . We call edges in H that also belong to E , *real edges*. Any edge in H that is not real is a *virtual edge*. Clearly, a spanner for G may not contain virtual edges, however virtual edges in H do carry important proximity information that will provide clues on how to replace them with real edges.

2.1 Properties of H

We will now prove some important properties of H . Let d_H be the distance metric induced by shortest paths in H . Specifically, we will show that H satisfies the following three properties:

- (1) For every $\{u, v\} \in E, d_H(u, v) \leq (1 + \varepsilon) \cdot d(u, v)$ (Lemma 4).
- (2) Edges of H can be oriented in such a way that the out-degree of H is bounded by $(\frac{1}{\varepsilon})^{O(\rho)}$ (Lemma 5).
- (3) The weight of H is $wt(H) = O(\log \Delta) \cdot (\frac{1}{\varepsilon})^{O(\rho)} \cdot wt(MST)$ (Lemma 6).

Property (1) implies that H is connected, since G is assumed to be connected. Property (2) implies that H has a linear number of edges, though it does not imply that H has bounded maximum degree. In Section 2.2 we describe a method to alter H so as to bound the in-degree of H as well, while maintaining all the properties listed above. The proofs of these properties are based on some intermediate results, that we now establish. Proofs of Lemma 4 and Lemma 5 are similar to those in [3]. The next observation follows immediately from the definition of the doubling dimension of a metric space.

Proposition 1. *If (X, d) is a metric with doubling dimension ρ and $Y \subseteq X$ is a subset of points with aspect ratio Δ , then $|Y| \leq 2^{\rho \lceil \log_2 \Delta \rceil}$.*

For any point $u \in V_i$, let $N_i(u) = \{v \in V_i \mid \{u, v\} \in E_i\}$ denote the set of points connected to u by edges in E_i . We now show an upper bound on the size of $N_i(u)$.

Lemma 1. *For each $u \in V_i$, $|N_i(u)| \leq (\frac{1}{\varepsilon})^{O(\rho)}$.*

Proof. That the aspect ratio of $N_i(u)$ is bounded by 2γ follows from two observations: (1) any two points in $N_i(u)$ are more than distance r_i apart, and (2) any point in $N_i(u)$ is at distance at most $\gamma \cdot r_i$ from u and therefore, by using the triangle inequality, any two points in $N_i(u)$ are at most $2\gamma \cdot r_i$ apart. Then Proposition 1 implies the lemma.

Lemma 2. *Suppose $u, v \in V_i$ and $d(u, v) \leq \gamma \cdot r_i$. Then $\{u, v\} \in \widehat{E}$.*

Proof. If $\gamma \cdot r_{i-1} < d(u, v) \leq \gamma \cdot r_i$, then by definition of E_i , $\{u, v\} \in E_i$. Otherwise, (a) $d(u, v) \leq \gamma \cdot r_0$ or (b) for some $j < i$, $\gamma \cdot r_{j-1} < d(u, v) \leq \gamma \cdot r_j$. Since $V_i \subseteq V_j$ for all $0 \leq j \leq i$, in case (a), $\{u, v\} \in E_0$ and in case (b), $\{u, v\} \in E_j$.

Lemma 3. *For each $u \in V$ and for each i , there exists $v \in V_i$ such that $d_H(u, v) \leq \frac{\alpha}{\alpha-1} \cdot r_i$.*

Proof. The proof is by induction on i . For $i = 0$, $u \in V_0 = V$ and $d_H(u, u) = 0 < \frac{\alpha}{\alpha-1} \cdot r_0$, proving this case true. For $i > 0$, apply the inductive hypothesis to infer that there exists $w \in V_{i-1}$ such that $d_H(u, w) \leq \frac{\alpha}{\alpha-1} \cdot r_{i-1}$. Furthermore, since V_i is an r_i -net of V_{i-1} , there exists $v \in V_i \subseteq V_{i-1}$ such that $d(w, v) \leq r_i \leq \gamma \cdot r_{i-1}$. This along with Lemma 2 shows that $\{w, v\} \in \widehat{E}$ and therefore $d_H(w, v) = d(w, v) \leq r_i$. By the triangle inequality we have that $d_H(u, v) \leq d_H(u, w) + d_H(w, v) \leq \frac{\alpha}{\alpha-1} \cdot r_{i-1} + r_i = \frac{\alpha}{\alpha-1} \cdot r_i$.

In addition to proving the existence of a vertex v at each level i , Lemma 3 implies a certain path from vertex u to $v \in V_i$. Start from node $(0, u)$ in the tree (that is, the copy of u corresponding to a leaf) and follow the path through a sequence of parents, until a level- i node (i, v) is reached. Lemma 3 shows that the distance in H along this path is at most $\frac{\alpha}{\alpha-1} \cdot r_i$.

Lemma 4. [Property 1] *For any edge $\{u, v\} \in E$, $d_H(u, v) \leq (1 + \varepsilon) \cdot d(u, v)$.*

Proof. For ease of presentation, let $\lambda = \frac{\alpha}{\alpha-1}$. Let q be the smallest integer such that $\frac{4\lambda}{\alpha^q} \leq \varepsilon < \frac{8\lambda}{\alpha^q}$. Thus $q = \lceil \log_\alpha \frac{4\lambda}{\varepsilon} \rceil$. Let i be such that $r_i \leq d(u, v) < r_{i+1}$, and assume first that $i \leq q - 1$. Then $d(u, v) < \alpha^q \cdot r_0 \leq \frac{8\lambda}{\varepsilon} \cdot r_0 \leq \gamma r_0$, since $\gamma = 2\lambda (1 + \frac{4\alpha}{\varepsilon}) > \frac{8\lambda}{\varepsilon}$. Also since both u and v belong to V_0 , by Lemma 2, we have that $\{u, v\} \in \widehat{E}$. This implies that $d_H(u, v) = d(u, v)$, proving the lemma true for this case. Assume now that $i \geq q$ and let $s = i - q \geq 0$. Note that $r_i = \alpha^q \cdot r_s$. By Lemma 3, there exist $x, y \in V_s$ such that $d_H(u, x) \leq \lambda \cdot r_s$ and $d_H(y, v) \leq \lambda \cdot r_s$. By the triangle inequality,

$$\begin{aligned}
 d(x, y) &\leq d(x, u) + d(u, v) + d(v, y) \\
 &\leq \lambda \cdot r_s + d(u, v) + \lambda \cdot r_s && (d(x, u) \leq d_H(x, u), d(v, y) \leq d_H(v, y)) \\
 &< \lambda \cdot r_s + \alpha \cdot r_i + \lambda \cdot r_s && (\text{since } d(u, v) < r_{i+1}) \\
 &= r_s(2\lambda + \alpha \cdot \alpha^q) && (\text{since } r_i = \alpha^q \cdot r_s) \\
 &\leq r_s(2\lambda + \alpha \frac{8\lambda}{\varepsilon}) \\
 &= \gamma \cdot r_s
 \end{aligned}$$

Hence, by Lemma 2, $\{x, y\} \in \widehat{E}$ and therefore $d_H(x, y) = d(x, y)$. Using the triangle inequality again, we get

$$\begin{aligned}
 d_H(u, v) &\leq d_H(u, x) + d_H(x, y) + d_H(y, v) \\
 &\leq 2\lambda \cdot r_s + d(x, y) \\
 &\leq 4\lambda \cdot r_s + d(u, v) && (\text{from the upper bound derivation of } d(x, y)) \\
 &\leq (1 + \frac{4\lambda}{\alpha^q}) \cdot d(u, v) && (\text{since } r_i = \alpha^q \cdot r_s \leq d(u, v)) \\
 &\leq (1 + \varepsilon) \cdot d(u, v)
 \end{aligned}$$

This completes the proof.

Lemma 4 also identifies a uv -path in H of length at most $(1 + \varepsilon) \cdot d(u, v)$. Simply follow the sequence of parents, starting at the node $(0, u)$ in the tree and similarly, starting at the node $(0, v)$. At a certain level (denoted s in the proof), the ancestor of u and the ancestor of v at that level are connected by an edge in H .

We now prove Property (2) of H . Recall the notation: for each point u , $i(v) = \max\{i \mid v \in V_i\}$. For each edge $\{u, v\} \in \widehat{E}$, direct $\{u, v\}$ from u to v , if $i(u) < i(v)$. If $i(u) = i(v)$, pick an arbitrary orientation. This edge orientation is identical to the one used in [2]. Call the resulting digraph \vec{H} .

Lemma 5. [Property 2] *The out-degree of \vec{H} is bounded above by $(\frac{1}{\varepsilon})^{O(\rho)}$.*

Proof. Let $\{u, v\} \in \widehat{E}$ be an arbitrary edge directed from u to v , and let i be such that $\{u, v\} \in E_i$. Then $d(u, v) \leq \gamma \cdot r_i$. Now note that $r_{i+\delta} = \alpha^\delta \cdot r_i \geq \gamma \cdot r_i$ (recall that $\delta = \lceil \log_\alpha \gamma \rceil$). This, along with the fact that $V_{i+\delta}$ is an $r_{i+\delta}$ -net, implies that it is not possible for both u and v to exist in $V_{i+\delta}$. Since $i(u) \leq i(v)$ (by our assumption), it follows that $i(u) \leq i + \delta$. On the other hand, $u \in V_i$ and so $i(u) \geq i$.

Summarizing, we have that $i(u) - \delta \leq i \leq i(u)$. This tells us that there are at most $\delta + 1 = O(\log_\alpha \gamma)$ values of i for which E_i may contain an edge outgoing from u . For each such i , by Lemma 1 there are at most $|N_i(u)| \leq (\frac{1}{\varepsilon})^{O(\rho)}$ edges in E_i outgoing from u . It follows that the total number of edges in \widehat{E} outgoing from u is $(\frac{1}{\varepsilon})^{O(\rho)} \cdot O(\log_\alpha \gamma) = (\frac{1}{\varepsilon})^{O(\rho)}$.

We now prove Property (3) of H , showing that H has bounded weight.

Lemma 6. [Property 3] *The total weight of H is $wt(H) = O(\log \Delta) \cdot (\frac{1}{\varepsilon})^{O(\rho)}$. $wt(MST)$, where MST is a minimum spanning tree of V , and Δ is the aspect ratio of G .*

Proof. We show that, for each i , $wt(E_i) = \left(\frac{1}{\varepsilon}\right)^{O(\rho)} \cdot wt(MST)$. This along with the fact that there are $h + 1 = \log_{\alpha} \frac{1}{r_0} + 1 = O(\log_{\alpha} \Delta)$ levels i , proves the claim of the lemma.

Let $U_i \subseteq V_i$ be the points in V_i incident to edges in E_i , and let $t = |U_i|$. Recall that any edge $\{u, v\} \in E_i$ satisfies $r_i < d(u, v) \leq \gamma \cdot r_i$. Thus, any spanning tree of a set of points containing U_i has weight at least $(t - 1) \cdot r_i$, implying that $wt(MST) \geq (t - 1) \cdot r_i$. Also note that the weight of E_i is bounded by $\sum_{u \in U_i} |N_i(u)| \cdot \gamma \cdot r_i \leq \left(\frac{1}{\varepsilon}\right)^{O(\rho)} \cdot t \cdot \gamma \cdot r_i$, using the upper bound on $|N_i(u)|$ given by Lemma 1. Using the lower bound on $wt(MST)$, we see that the weight of E_i is bounded above by $\left(\frac{1}{\varepsilon}\right)^{O(\rho)} \cdot \gamma \cdot (wt(MST) + r_i)$. Summing this expression over all E_i , yields the upper bound claimed in the lemma.

2.2 Altering H for Bounded Degree

In this section we show how to modify H so as to bound the degree of each vertex by a constant. Lemma 5 shows that an oriented version of H , namely \vec{H} , has bounded out-degree. Next we describe a method that carefully replaces some directed edges in \vec{H} by others so as to guarantee constant bound on the in-degree as well, without increasing the out-degree. The replacement procedure is similar to the one used in [2], slightly adjusted to work with UBGs. Assume without loss of generality that $\varepsilon \leq \frac{1}{2}$; otherwise, if $\varepsilon > \frac{1}{2}$, we proceed with $\varepsilon = \frac{1}{2}$. We use the fact that $\varepsilon \leq \frac{1}{2}$ in the proof of Lemma 9. Let ℓ be the smallest positive integer such that $\frac{1}{\alpha^{\ell-1}} \leq \varepsilon$. Thus $\ell = O(\log_{\alpha} \frac{1}{\varepsilon})$.

Edge Replacement Procedure. Let u be an arbitrary point in V and let $M(u, i)$ be the set of all vertices $v \in V_i$ such that $\{v, u\}$ is an edge in E_i directed from v to u in \vec{H} . Let $I(u) = \langle i_1, i_2, \dots \rangle$ be the increasing sequence of all indices i_k for which $M(u, i_k)$ is nonempty. For $1 \leq k \leq \ell$, we do not disturb any of the edges from points in $M(u, i_k)$ to u . For each $k > \ell$ such that $i_k \leq h - \delta - 2$, edges $\{v, u\}$ connecting $v \in M(u, i_k)$ to u are replaced by other edges. Specifically, an edge $\{v, u\}$, with $v \in M(u, i_k)$, is replaced by an edge $\{v, w\}$, where w is an arbitrary vertex in $M(u, i_{k-\ell})$. The replacement can be equivalently viewed as happening in either H or its oriented version \vec{H} . In \vec{H} , we replace the directed edge (v, u) by the directed edge (v, w) . In the next two lemmas, our arguments will use \vec{H} or H , as convenient.

Let \tilde{E} be the resulting set of edges. By our construction, $|\tilde{E}| \leq |\hat{E}|$. An important observation here is that the replacement procedure above is carried out only for edges in E_i , with $i \leq h - \delta - 2$ (that is, only edges of length no greater than $1/\alpha^2$). This is to ensure that only real edges get replaced and no virtual edges get added, a guarantee that is shown in the following lemma.

Lemma 7. $\tilde{E} \setminus \hat{E}$ contains no virtual edges.

Proof. Let $\{v, u\}$ be an edge that gets replaced by $\{v, w\}$, with $v \in M(u, i_k)$ and $w \in M(u, i_{k-\ell})$. Recall that $k > \ell$ and $i_k \leq h - \delta - 2$. Using the definitions of

E_{i_k} and $E_{i_{k-\ell}}$ and the fact that $\frac{1}{\alpha^{\ell-1}} \leq \varepsilon$, it follows that $d(w, u) \leq \varepsilon \cdot d(v, u)$. By the triangle inequality, $d(v, w) \leq d(v, u) + d(w, u) \leq (1 + \varepsilon)d(v, u)$. Now note that $d(v, u) \leq 1/\alpha^2$. This is because edges in E_{i_k} have length no greater than $\gamma \cdot r_{i_k} \leq 1/\alpha^2$, for any $i_k \leq h - \delta - 2$. Therefore $d(v, w) \leq (1 + \varepsilon)/\alpha^2 \leq 1$, for any $\alpha^2 \geq (1 + \varepsilon)$.

Let $J = (V, \tilde{E})$. First we show that J indeed has bounded degree (Lemma 8). Second we show that the metric distance d_J induced by shortest paths in J is a good approximation of d_H (Lemma 9). A consequence of this is that J remains connected, and maintains spanner paths between endpoints of real edges.

Lemma 8. *Every vertex in $J = (V, \tilde{E})$ has degree bounded by $(\frac{1}{\varepsilon})^{O(\rho)}$.*

Proof. Let A be the maximum out-degree of a vertex of \vec{H} . By Lemma 5, $A \leq (\frac{1}{\varepsilon})^{O(\rho)}$. Let B be the largest of $|N_i(u)|$, for all i and all u . By Lemma 1, $B \leq (\frac{1}{\varepsilon})^{O(\rho)}$. The edge-replacement procedure replaces a directed edge (v, u) by a directed edge (v, w) . So the out-degrees of vertices remain unchanged by the edge-replacement procedure, and continue to be bounded above by $(\frac{1}{\varepsilon})^{O(\rho)}$. Thus, we can simply focus on the in-degrees of vertices. We bound these by accounting for the in-degree of an arbitrary vertex x with respect to old edges (in $\tilde{E} \cap \hat{E}$) and with respect to new edges (in $\tilde{E} \setminus \hat{E}$); we show that both in-degrees are bounded above by $(\frac{1}{\varepsilon})^{O(\rho)}$.

In-degree of x with respect to $\tilde{E} \cap \hat{E}$. Out of the edges in \vec{H} that come into x , at most $B(\ell + \delta + 2)$ remain in \tilde{E} . More specifically, at most B edges at each of the first ℓ levels i_1, i_2, \dots, i_ℓ in $I(x)$, plus at most B edges in each of E_i , $i = h - \delta - 1, h - \delta, \dots, h$, remain in \tilde{E} . Any other edge directed into x gets replaced by an edge not incident to x . We end this case by noting that $B(\ell + \delta + 2) = (\frac{1}{\varepsilon})^{O(\rho)}$.

In-degree of x with respect to $\tilde{E} \setminus \hat{E}$. Vertex x has a new in-coming edge whenever it plays the role of w in the edge-replacement procedure. Recall that in the edge-replacement procedure, w and v are both in-neighbors of u . For each edge (w, u) , there are at most B edges (v, u) directed into u that may get replaced by (v, w) . Furthermore, there are A edges (w, u) outgoing from w . This gives an upper bound of $AB = (\frac{1}{\varepsilon})^{O(\rho)}$ on the in-degree of x .

It remains to show that d_J is a good approximation of d_H . Intuition for this is provided by the proof of Lemma 7. In that proof, it is shown that when $\{v, w\}$ replaces $\{v, u\}$, $d(w, u) \leq \varepsilon \cdot d(v, u)$ and $d(v, w) \leq (1 + \varepsilon) \cdot d(v, u)$. Thus, if the path $\langle v, w, u \rangle$ existed in \tilde{E} , this path would have length at most $(1 + 2\varepsilon) \cdot d(v, u)$. However, edge $\{w, u\}$ may not exist in \tilde{E} , since it may itself have been replaced. Thus a shortest path from w to u in \tilde{E} may be longer than $d(w, u)$. However, since $d(w, u) \leq \varepsilon \cdot d(v, u)$, the extra cost of replacing $\{w, u\}$ is marginal and the eventual sum of all of these lengths is still bounded above by $(1 + 2\varepsilon) \cdot d(v, u)$. Thus we have the following lemma, whose proof we skip (due to space restrictions).

Lemma 9. $d_J \leq (1 + 2\varepsilon)d_H$.

2.3 Eliminating Virtual Edges

The only impediment in having $J = (V, \tilde{E})$ serve as a spanner for the input UBG G is the presence of virtual edges in J . Recall that these are edges of length greater than 1 and clearly do not exist in G . In this section we show that there exist real edges that can take over the role of virtual edges in J , without violating the properties J is expected to have.

Let $\{u, v\} \in E$ be an arbitrary edge and let i be such that $r_i \leq d(u, v) < r_{i+1}$. Let q be as in the proof of Lemma 4: the smallest integer such that $\frac{\alpha}{\alpha-1} \cdot \frac{4}{\alpha^q} \leq \varepsilon < \frac{\alpha}{\alpha-1} \cdot \frac{8}{\alpha^q}$. As mentioned before, the proof of Lemma 4 implies a certain uv -path of length at most $(1 + \varepsilon) \cdot d(u, v)$ in $H = (V, \hat{E})$. If $i \leq q - 1$, this path is just the edge $\{u, v\}$, because $\{u, v\}$ is guaranteed to exist in \hat{E} . The Edge Replacement Procedure (Section 2.2) ensures that only real edges are replaced, and each real edge is replaced by a path consisting only of real edges. This along with Lemma 9 ensures that even in \tilde{E} there is a uv -path of length at most $(1 + 2\varepsilon) \cdot d(u, v)$, consisting of real edges only. If $i \geq q$, the uv -path in H implied by Lemma 4 may have more than one edge. Let $s = i - q$ and (s, u^*) (respectively, (s, v^*)) be the level- s ancestor of the leaf $(0, u)$ (respectively, the leaf $(0, v)$) in the net tree $\langle V_0, V_1, \dots, V_h \rangle$. Then the edge $\{u^*, v^*\}$ is guaranteed to be present in \hat{E} and the uv -path implied by Lemma 4 starts at $(0, u)$, goes up to the net tree via parents to (s, u^*) , then to (s, v^*) , and then follows the unique path down the tree from (s, v^*) to $(0, v)$. It is easy to check that of all the edges in this path, only $\{u^*, v^*\}$ may be virtual. Specifically, when the edge $\{u, v\}$ is long enough to guarantee that $i \geq h - \delta + 1 + q$, then $s = i - q \geq h - \delta + 1$ and the edge $\{u^*, v^*\}$ may belong to E_s . Recall that for $j \geq h - \delta + 1$, edges in E_j may not be real and in particular $\{u^*, v^*\}$ may be a virtual edge. Since the uv -path implied by Lemma 4 passes through edge $\{u^*, v^*\}$, one has to be careful in replacing $\{u^*, v^*\}$ by a real edge. Our virtual edge replacement procedure is given below.

For any node (i, v) in the net tree, let $T(i, v)$ denote the set of all vertices $u \in V$, such that the subtree of the net tree rooted at (i, v) contains a copy of u . In other words, $T(i, v) = \{u \in V \mid (i, v) \text{ is an ancestor of } (j, u) \text{ for some } j \leq i\}$.

Virtual Edge Replacement Procedure. For a virtual edge $\{u, v\} \in E_i$, if there is a real edge $\{x, y\}$ already in the spanner H , with $x \in T(i, u)$ and $y \in T(i, v)$, then simply delete $\{u, v\}$. Similarly, if there is no such real edge $\{x, y\}$ in the input graph G with $x \in T(i, u)$ and $y \in T(i, v)$ then simply delete $\{u, v\}$. Otherwise, find a real edge $\{x, y\} \in E$, $x \in T(i, u)$ and $y \in T(i, v)$, and replace $\{u, v\}$ by $\{x, y\}$.

The reason why this replacement procedure works can be intuitively explained as follows. A virtual edge $\{u, v\} \in E_i$ is important for pairs of vertices $\{a, b\}$, with $a \in T(i, u)$ and $b \in T(i, v)$, for which all ab -paths of length at most $(1 + \varepsilon) \cdot d(a, b)$ pass through $\{u, v\}$. Replacing $\{u, v\}$ by $\{x, y\}$ provides the following alternate ab -path that is short enough: starting at the leaf a , go up the tree rooted at (i, u) via parents until an ancestor common to a and x is reached, then come down to x , take edge $\{x, y\}$, go up the tree rooted at (i, v) until an ancestor common to




Fig. 2. A short ab -path passes through virtual edge $\{u, v\}$. After replacing virtual edge $\{u, v\}$ by real edge $\{x, y\}$, there is a short ab -path through $\{x, y\}$.

b and y is reached, and finally go down to b . Figure 2 illustrates this alternate path. Note that this entire path consists only of real edges.

We finally state our main result. Let G' be the graph obtained from J by replacing virtual edges using the Virtual Edge Replacement Procedure.

Theorem 2. $G' = (V, E')$ is a $(1 + \varepsilon)$ -spanner of G with degree bounded above by $(\frac{1}{\varepsilon})^{O(\rho)}$ and weight bounded above by $O(\log \Delta) \cdot (\frac{1}{\varepsilon})^{O(\rho)} \cdot wt(MST)$.

A proof similar to that of Lemma 4 can be used to show the spanner property of G' . The fact that G' is lightweight simply follows from the fact that a virtual edge of length greater than 1 in J , either gets eliminated, or gets replaced by at most one real edge of length at most 1 in G' . The constant degree bound follows from the observation that, for a vertex x to acquire a new incident edge, there is an ancestor of x in the net tree at level $h - \delta + 1$ or higher, that loses an incident edge at that level. There are a constant number of such ancestors and from Lemma 1, we know that any vertex has a constant number of incident edges at any particular level.

We conclude this section with a summary of our algorithm.

Algorithm SPANNER($(V, d), \varepsilon$)

Let $\sqrt{1 + \varepsilon} < \alpha$ be a constant, $\gamma = \frac{2\alpha}{\alpha - 1} (1 + \frac{4\alpha}{\varepsilon})$, and $\delta = \lceil \log_\alpha \gamma \rceil$.

Let h be the smallest such that $\frac{1}{\alpha^h}$ is smaller than the minimum inter-point distance.

Let $r_0 = \frac{1}{\alpha^h}$ and let $r_i = \alpha \cdot r_{i-1}$, for all $i > 0$.

Constructing a linear size $(1 + \varepsilon)$ -spanner $H = (V, \hat{E})$.

1. Construct the net tree $\langle V_0, V_1, \dots, V_h \rangle$.

[Let $i(u) = \max\{i \mid u \in V_i\}$.]

2. Construct the sets

$$E_0 = \{\{u, v\} \in V_0 \times V_0 \mid d(u, v) \leq \gamma \cdot r_0\},$$

$$E_i = \{\{u, v\} \in V_i \times V_i \mid \gamma \cdot r_{i-1} < d(u, v) \leq \gamma \cdot r_i\}, \text{ for } 1 \leq i \leq h.$$

[Let $\hat{E} = \cup_i E_i$ and $H = (V, \hat{E})$.]

Replacing edges to obtain a constant degree bound.

3. Orient each edge $\{u, v\} \in \hat{E}$ from u to v if $i(u) \leq i(v)$, breaking ties arbitrarily.

[Let $M(u, i)$ denote the set of vertices $v \in V_i$, with $\{v, u\} \in \hat{E}$.]

4. For each $u \in V$, construct the increasing sequence $I(u) = \langle i_1, i_2, \dots \rangle$ of all i_k with $M(u, i_k) \neq \emptyset$. [Let ℓ be the smallest integer with $\frac{1}{\alpha^{\ell-1}} \leq \varepsilon$.]
 5. For each $u \in V$ and each $i_k \in I(u)$, with $k > \ell$ and $i_k \leq h - \delta - 2$, do
 6. Replace directed edge (v, u) , $v \in M(u, i_k)$ by edge (v, w) , for arbitrary $w \in M(u, i_{k-\ell})$.
- [Let $J = (V, \tilde{E})$ be the resulting graph, with distance metric d_J .]

Replacing virtual edges by real ones.

[Let $T(i, v) = \{x \in V \mid (i, v) \text{ is an ancestor of } (j, x) \text{ for some } j \leq i\}$.]

7. For each $i \geq h - \delta + 1$ and each virtual edge $\{u, v\} \in E_i$ do
 8. If there is a real edge $\{x, y\} \in \tilde{E}$, $x \in T(i, u)$ and $y \in T(i, v)$, then do nothing.
 9. Otherwise, if there is a real edge $\{x, y\} \in E$, with $x \in T(i, u)$ and $y \in T(i, v)$, replace $\{u, v\}$ by $\{x, y\}$.
- [Let E' be the set of resulting edges. Output is $G' = (V, E')$.]

3 Leapfrog Property

In Lemma 6, we showed that $H = (V, \hat{E})$ has total weight bounded above by $O(\log \Delta) \cdot (\frac{1}{\varepsilon})^{O(\rho)} \cdot wt(MST)$, where Δ is the aspect ratio of G . Thus, for fixed ε and constant doubling dimension ρ , the upper bound is within $O(\log \Delta)$ times the optimal value. In an attempt to show a bound that is within $O(1)$ times the optimal value, we use a tool that is widely used in the computational geometry literature [8,5,9]. In the context of building lightweight $(1 + \varepsilon)$ -spanners for Euclidean spaces, Das and Narasimhan [8] have shown that, if the set of edges in the spanner satisfy a property known as the *leapfrog property*, then the total weight of the spanner is bounded above by $O(wt(MST))$. Below we state the leapfrog property precisely.

Leapfrog Property. For any $t \geq t' > 1$, a set F of edges has the (t', t) -leapfrog property if, for every subset $S = \{\{u_1, v_1\}, \{u_2, v_2\}, \dots, \{u_s, v_s\}\}$ of F ,

$$t' \cdot d(u_1, v_1) < \sum_{i=2}^s d(u_i, v_i) + t \cdot \left(\sum_{i=1}^{s-1} d(v_i, u_{i+1}) + d(v_s, u_1) \right). \tag{1}$$

Informally, this definition says that, if there exists an edge between u_1 and v_1 , then any $u_1 v_1$ -path not including $\{u_1, v_1\}$ must have length greater than $t' \cdot d(u_1, v_1)$. Das and Narasimhan [8] show the following connection between the leapfrog property and the weight of the spanner.

Lemma 10. *Let $t \geq t' > 1$. If the line segments F in d -dimensional space satisfy the (t', t) -leapfrog property, then $wt(F) = O(wt(MST))$, where MST is a minimum spanning tree connecting the endpoints of line segments in F . The constant in the asymptotic notation depends on t, t' and d .*

It is well known that, if a spanner is built “greedily”, then the set of edges in the spanner satisfies the leapfrog property [8,5,9]. In [6] we showed that even a “relaxed” version of the greedy algorithm would ensure that the spanner edges

have the leapfrog property. This was critical to showing that the spanner constructed in a distributed manner for UBGs in Euclidean spaces [6] had total weight bounded above by $O(wt(MST))$. Here we ask if it is possible to do the same for UBGs in metric spaces with constant doubling dimension. In an attempt to answer this question we show that, using a variant of the SPANNER algorithm (outlined at the end of Section 2), we can build, for a given doubling UBG G , a $(1 + \varepsilon)$ -spanner with degree bounded above by a constant and with the (t, t') -leapfrog property, for some constants $t \geq t' > 1$. Note that this does not give us the desired $O(wt(MST))$ bound on the weight of the constructed spanner because we do not know if the equivalent of Lemma 10 holds for non-Euclidean metric spaces. The proof of Lemma 10 in [8] is quite geometric and does not suggest an approach to its generalization to metric spaces of constant doubling dimension.

To guarantee that the output spanner satisfies the (t', t) -leapfrog property, we need to make two modifications to the SPANNER algorithm. Let H_i denote the spanning subgraph of G induced by $E_0 \cup E_1 \cup \dots \cup E_i$.

1. We modify Step (2) of the algorithm and place an edge $\{u, v\}$ into E_i only if $\{u, v\} \in V_i \times V_i$, $\gamma \cdot r_{i-1} < d(u, v) \leq \gamma \cdot r_i$, and there is not already a uv -path of length at most $(1 + \varepsilon) \cdot d(u, v)$ in H_{i-1} .
2. Two edges $\{u, v\}$ and $\{u', v'\}$ in E_i are said to be *mutually redundant* if both of the following conditions hold:
 - (a) $d_{H_{i-1}}(v, u') + d(u', v') + d_{H_{i-1}}(v', u) \leq (1 + \varepsilon) \cdot d(u, v)$
 - (b) $d_{H_{i-1}}(v', u) + d(u, v) + d_{H_{i-1}}(v, u') \leq (1 + \varepsilon) \cdot d(u', v')$

Note that these conditions imply that $H_i \setminus \{u, v\}$ contains a uv -path of length at most $(1 + \varepsilon) \cdot d(u, v)$ and $H_i \setminus \{u', v'\}$ contains a $u'v'$ -path of length at most $(1 + \varepsilon) \cdot d(u', v')$. Thus, one of these can potentially be eliminated from H_i , without compromising the $(1 + \varepsilon)$ -spanner property of H_i . In fact, such mutually redundant pairs of edges need to be eliminated from H_i in order to show that H satisfies the leapfrog property.

These ideas and how they lead to a spanner that has the leapfrog property are discussed in detail in [6].

4 Distributed Implementation

In this section, we show that the SPANNER algorithm (Section 2) and its variant that ensures the leapfrog property (Section 3), both have distributed implementations that run in $O(\log^* n)$ rounds of communication. Here we focus on the SPANNER algorithm. It turns out that Step (1) of this algorithm takes $O(\log^* n)$ rounds, whereas the remaining steps take $O(1)$ additional rounds. We first examine Steps (2)-(9) of the algorithm.

It is easy to verify that in Steps (2)-(9), a node u needs to communicate only with other nodes that are either neighbors of u in G , or to which u is connected by a virtual edge. The main difficulty here is that the endpoints of a virtual edge $\{u, v\}$ may not be neighbors in the network. Consider a virtual edge $\{u, v\} \in E_i$.

By definition of E_i , $d(u, v) \leq \gamma \cdot r_i \leq \gamma$. Even though the distance between u and v in the underlying metric space is bounded above by a constant, it is not necessary that the hop distance between u and v in G be similarly bounded above. Let us call a virtual edge $\{u, v\} \in E_i$, *useful*, if there exist $x \in T(i, u)$ and $y \in T(i, v)$ such that $\{x, y\} \in E$. Notice that only useful virtual edges need to be considered by our algorithm. If a virtual edge $\{u, v\}$ is not useful, then even though it is added in Step (2), it is eliminated in Steps (7)-(9). In the following lemma we show that the hop distance between endpoints of useful virtual edges is small.

Lemma 11. *The hop distance in G between the endpoints of any useful virtual edge $\{u, v\} \in E_i$ is at most $2(\frac{2\alpha}{\alpha-1} + 1)$.*

Proof. By definition of a useful virtual edge, there are points $x \in T(i, u)$ and $y \in T(i, v)$ such that $\{x, y\}$ is an edge in G . Thus a path in G between u and v is the following: start at node (i, u) in the net tree and travel down to a copy of x , follow the edge $\{x, y\}$, and then travel up to node (i, v) . Note that the edge $\{x, y\} \in E$, but it may not belong to \tilde{E} . The length of this path is at most $2(1 + \frac{1}{\alpha} + \frac{1}{\alpha^2} + \dots) + 1$, implying that $d_G(u, v) \leq \frac{2\alpha}{\alpha-1} + 1$. Now consider a shortest uv -path in G , say $\langle w_0 = u, w_1, \dots, w_{k+1} = v \rangle$. Because G is a UBG and due to the triangle inequality, $d(w_i, w_{i+2}) > 1$ for all $0 \leq i \leq k-1$ (otherwise w_{i+1} could be eliminated from the uv -path to obtain an even shorter uv -path in G). This yields a lower bound of $k/2$ on $d(u, v)$, and since $d_G(u, v) \geq d(u, v)$, we have that $d_G(u, v) \geq k/2$. Combining this with the upper bound of $\frac{2\alpha}{\alpha-1} + 1$, we obtain that $k \leq 2(\frac{2\alpha}{\alpha-1} + 1)$.

Thus, in Steps (2)-(9) of the SPANNER algorithm, a node only needs to communicate with nodes that are at most $O(\frac{1}{\alpha-1})$ hops away. This suggests a simple way of implementing Steps (2)-(9): after Step (1) is completed, each node u gathers neighborhood information and the values of $i(v)$ from all nodes v that are $O(\frac{1}{\alpha-1})$ hops away. After this, node u can do all of its computation with no further communication.

The fact that Step (1) can be implemented in $O(\log^* n)$ rounds of communication follows from a clever argument in [13]. Suppose that we have computed the set V_{i-1} . The computation of the set V_i , which is an r_i -net of V_{i-1} , reduces to a maximal independent set (MIS) computation on a degree-bounded graph. To see this, create a graph, say G_i , whose vertex set is V_{i-1} and whose edges connect any pair of vertices $u, v \in V_{i-1}$, if $d(u, v) \leq r_i$. Then it is easy to see that an MIS in G_i is an r_i -net of V_{i-1} . Furthermore, the fact that G_i has bounded degree follows from the fact that the underlying metric space has bounded doubling dimension. There is a well-known deterministic algorithm due to Linial [15] for computing an MIS, that runs in $O(\log^* n)$ communication rounds on graphs with bounded degree. Using this algorithm, one can compute the r_i -net V_i of V_{i-1} in $O(\log^* n)$ rounds. Since there are $h+1 = O(\log \Delta)$ such sets to compute, it seems like this approach will take $O(\log \Delta \cdot \log^* n)$ rounds. However, in [13] it is shown that in this algorithm, each node uses information only from nodes that are at most $O(\log^* n)$ hops away in G . Therefore, this algorithm has a $O(\log^* n)$ -round implementation in which each node u first gathers information

from nodes that are at most $O(\log^* n)$ hops away and then performs all steps of the SPANNER algorithm locally, using the collected information.

References

1. B. Awerbuch, A. Goldberg, M. Luby, and S. Plotkin. Network decomposition and locality in distributed computation. In *IEEE Symposium on Foundations of Computer Science*, pages 364–369, 1989.
2. Hubert T-H. Chan, Anupam Gupta, Bruce M. Maggs, and Shuheng Zhou. On hierarchical routing in doubling metrics. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 762–771, 2005.
3. T-H. Hubert Chan. Personal Communication, 2006.
4. T-H. Hubert Chan and Anupam Gupta. Small hop-diameter sparse spanners for doubling metrics. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 70–78, 2006.
5. A. Czumaj and H. Zhao. Fault-tolerant geometric spanners. *Discrete & Computational Geometry*, 32(2):207–230, 2004.
6. Mirela Damian, Saurav Pandit, and Sriram Pemmaraju. Local approximation schemes for topology control. In *PODC '06: Proceedings of the twenty-fifth annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, 2006.
7. G. Das, P. Heffernan, and G. Narasimhan. Optimally sparse spanners in 3-dimensional euclidean space. In *ACM Symposium on Computational Geometry*, pages 53–62, 1993.
8. G. Das and G. Narasimhan. A fast algorithm for constructing sparse euclidean spanners. *Int. J. Comput. Geometry Appl.*, 7(4):297–315, 1997.
9. J. Gudmundsson, C. Levkopoulos, and G. Narasimhan. Fast greedy algorithms for constructing sparse geometric spanners. *SIAM J. Comput.*, 31(5):1479–1500, 2002.
10. S. Har-Peled and M. Mendel. Fast construction of nets in low dimensional metrics, and their applications. In *SCG'05: Proceedings of the 21st annual symposium on Computational geometry*, pages 150–158, 2005.
11. R. Krauthgamer, A. Gupta, and J.R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *FOCS '03: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 534–543, 2003.
12. R. Krauthgamer and J.R. Lee. Navigating nets: simple algorithms for proximity search. In *SODA '04: Proceedings of the 15th annual ACM-SIAM symposium on Discrete algorithms*, pages 798–807, 2004.
13. Fabian Kuhn, Thomas Moscibroda, and Roger Wattenhofer. On the locality of bounded growth. In *PODC '05: Proceedings of the twenty-fourth annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 60–68, 2005.
14. Xiang-Yang Li and Yu Wang. Efficient construction of low weighted bounded degree planar spanner. *International Journal of Computational Geometry and Applications*, 14(1-2):69–84, 2004.
15. Nathan Linial. Locality in distributed graph algorithms. *SIAM J. Comput.*, 21(1):193–201, 1992.
16. R. Rajaraman. Topology control and routing in ad hoc networks: A survey. *SIGACT News*, 33:60–73, 2002.
17. K. Talwar. Bypassing the embedding: algorithms for low dimensional metrics. In *STOC'04: Proceedings of the 36th annual ACM symposium on Theory of computing*, pages 281–290, 2004.