

# On Triangulation of Simple Networks

[Extended Abstract]

Robert Krauthgamer  
IBM Almaden Research Center  
San Jose, CA 95120, USA  
robi@almaden.ibm.com

## ABSTRACT

Network triangulation is a method for estimating distances between nodes in the network, by letting every node measure its distance to a few beacon nodes, and deducing the distance between every two nodes  $x, y$  by using their measurements to their common beacons and applying the triangle inequality. Kleinberg, Slivkins and Wexler [FOCS 2004] initiated a theoretical study of triangulation in metric spaces, and Slivkins [PODC 2005] subsequently showed that metrics of bounded doubling dimension admit a triangulation that approximates arbitrarily well all pairwise distances using only  $O(\log n)$  beacons per point, where  $n$  is the number of points in the network. He then asked whether this term is necessary (for doubling metrics).

We provide the first lower bounds on the number of beacons required for a triangulation in some specific simple networks. In particular, these bounds (i) answer Slivkins' question positively, even for one-dimensional metrics, and (ii) prove that, up to constants, Slivkins' triangulation achieves an optimal number of beacons (as a function of the approximation guarantee and the doubling dimension).

## Categories and Subject Descriptors

C.2 [Computer-Communication Networks]: Network Architecture and Design; G.2 [Discrete Mathematics]: Applications

## General Terms

Algorithms, Theory

## Keywords

Network Triangulation, Distance Estimation, Doubling Metrics, Tree Metrics

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SPAA'07, June 9–11, 2007, San Diego, California, USA.  
Copyright 2007 ACM 978-1-59593-667-7/07/0006 ...\$5.00.

## 1. INTRODUCTION

The distance between nodes in a network such as the Internet, is of high importance in many networking scenarios and applications, raising the problem of estimating the distance between nodes using only limited resources, such as preprocessing and local storage.<sup>1</sup> Recent networking research indicates that severe violations to the triangle inequality are not very common, and hence the network can be modeled by a finite metric space  $(X, d)$  with  $n = |X|$  points.<sup>2</sup> Indeed, the following “triangulation” approach, suggested by Guyton and Schwartz [8], has proven to be successful empirically, see e.g. [3]: let every node measure (precompute) its distance to a few “beacon” nodes, and then the distance between  $x$  and  $y$  is estimated (from above) by the minimum of  $d(x, b) + d(b, y)$  over the beacons  $b$ . Alternatively,  $d(x, y)$  can be estimated from below by the maximum of  $|d(x, b) - d(y, b)|$  over the beacons  $b$ . Clearly, the efficiency of this distributed approach decreases as the number of beacons grows.

Kleinberg, Slivkins and Wexler [11] initiated a theoretical study of network triangulation, formally defined as follows. A *triangulation* of a finite metric  $(X, d)$  is a labeling of each point  $x \in X$  with a subset  $S_x \subseteq X$  (i.e., it is a mapping  $X \mapsto 2^X$ ). We say that the triangulation has *order*  $s$  if  $\max_{x \in X} |S_x| \leq s$ . For two points  $x, y \in X$ , the triangulation induces an upper bound and a lower bound on  $d(x, y)$ :

$$D^+(x, y) = \min_{b \in S_x \cap S_y} [d(x, b) + d(b, y)]$$

$$D^-(x, y) = \max_{b \in S_x \cap S_y} |d(x, b) - d(b, y)|$$

The triangulation is called an  $(\varepsilon, \delta)$ -*triangulation* if

$$D^+(x, y)/D^-(x, y) \leq 1 + \delta$$

holds for all but an  $\varepsilon$ -fraction of the pairs  $x, y \in X$ .<sup>3</sup> We shall say that  $\varepsilon$  is the *slack* of the triangulation, and  $\delta$  is the accuracy. Clearly, some metrics do not have a low-order triangulation; for instance, in an equilateral metric (i.e.,  $d(x, y) = 1$  for all  $x, y \in X$ ), every distance that is not directly measured will have a lower estimate of 0.

<sup>1</sup>Distances in the network often represent node to node latencies.

<sup>2</sup>Formally,  $X$  is a set of points, endowed with a positive, symmetric, distance function  $d(\cdot, \cdot)$  that satisfies the triangle inequality.

<sup>3</sup>If we let the label of  $x$  contain, for each beacon  $b \in S_x$ , a unique identifier of  $b$  together with  $d(x, b)$ , then the estimates  $D^-(x, y), D^+(x, y)$  could be computed from directly the labels of  $x$  and  $y$ .

We digress momentarily to recall the family of doubling metrics, which plays a crucial role in network triangulation. The *doubling dimension* of a metric space  $(X, d)$  is defined as the smallest  $k > 0$  such that every ball in the metric can be covered by  $2^k$  balls of half the radius [2, 7]. A metric satisfying this property will also be called  $2^k$ -doubling, or just doubling if  $k = O(1)$ . It is easy to see that subsets of  $\mathbb{R}^k$  have doubling dimension  $O(k)$  (under any  $L^p$  norm). Furthermore, it includes the family of metrics with bounded growth rate of balls (see [7]), which was suggested as a model for the Internet in several recent studies, such as [17, 15, 21]. This notion was suggested in [7] (based on a classical notion of [2], and inspired by the approach in [10]) as an intrinsic notion of metric dimension that affects algorithmic tractability. Indeed, it was subsequently studied in the context of metric embeddings, Nearest Neighbor Search, the Traveling Salesman Problem, and low-stretch routing, see e.g. [7, 12, 20, 14, 19, 1].

Kleinberg et al. [11] gave an analytic explanation for the empirical success of triangulation, by proving (among other results) that every  $2^k$ -doubling metric admits an  $(\varepsilon, \delta)$ -triangulation whose order is independent of  $n$ .<sup>4</sup> Slivkins [19] showed (among other results) that metrics of doubling dimension  $k$  admit a  $(0, \delta)$ -triangulation of order  $\delta^{-O(k)} \log^2 n$ . and in [18] he further improved it (among other results) to a  $(0, \delta)$ -triangulation of order  $\delta^{-O(k)} \log n$ , and asked [Section 6 of the full version] whether the  $O(\log n)$  term is necessary (for doubling metrics).

We provide the first lower bounds on the order required for a triangulation of some specific simple networks. In particular, these bounds (i) answer Slivkins' question positively, even for one-dimensional metrics, and (ii) prove that, up to constants, Slivkins' triangulation achieves an optimal number of beacons (as a function of the approximation guarantee and the doubling dimension). As we discuss in more detail in Section 1.3, the lower bounds known previously for  $(\varepsilon, \delta)$ -triangulation in doubling metrics were very limited, e.g., only for  $\varepsilon = 0$  and extremely large diameter.

## 1.1 Results

We first show (in Section 2) the following relatively simple lower bound (for every constant  $\delta < 1$ ):

- A  $(0, \delta)$ -triangulation of the  $n$ -cycle must have order at least  $\Omega(\log n)$ .

Clearly, the cycle (ring) metric is one-dimensional, and hence the lower bound resolves Slivkins' question in the positive, even in extremely simple metrics. Our main technical result generalizes the above lower bound as follows (see Section 3 for the precise requirements, e.g. that  $n$  is sufficiently large):

- An  $(\varepsilon, \delta)$ -triangulation of certain simple  $2^k$ -doubling  $n$ -point metrics must have order of the form  $\delta^{-\Omega(k)} \log \frac{1}{\varepsilon}$ .

In particular, for  $(0, \delta)$ -triangulation we obtain a lower bound of  $\delta^{-\Omega(k)} \log n$ , which matches the upper bound  $\delta^{-O(k)} \log n$  shown by Slivkins [18].

Finally, we present (in Section 4) a simple lower bound for trees and ultrametrics (for all  $\varepsilon \leq 1/9$  and  $\delta < 1$ ):

<sup>4</sup>In fact, they show that following very simple randomized construction succeeds with high probability: choose uniformly at random a subset  $S \subset X$  whose cardinality is some function of  $k, \varepsilon$ , and  $\delta$ , and let all nodes use this set  $S$  as their beacon set.

- An  $(\varepsilon, \delta)$ -triangulation of certain simple tree metrics and ultrametrics on  $n$  points must have order at least  $(\Omega(n))^{1/1+\delta}$ .
- An  $(\varepsilon, \delta)$ -triangulation of certain doubling ultrametrics on  $n$  points must have order at least  $\Omega(\log \frac{1}{\varepsilon})$ .

## 1.2 Techniques

The main difficulty in proving the lower bounds is to set up the accounting method. Our method is illustrated in great clarity in the first lower bound, for the cycle, because its proof is rather short (but perhaps tricky). The main idea goes roughly as follows: if for two points  $x, y$  a beacon  $b \in S_x \cap S_y$  provides a good upper estimate (namely,  $d(x, b) + d(b, y) \leq 2d(x, y)$ ) then the distance of  $b$  from at least one of the two points, say  $x$ , must be in the range  $[\frac{1}{2}d(x, y), \frac{3}{2}d(x, y)]$ , and we can “blame”  $x$  for this event. The core technical idea is then a subtle averaging argument that shows that some points  $x$  must be “blamed”  $\Omega(\log n)$  “distinct” times, namely, for many integers  $i \in \{1, 2, \dots, \log n\}$  the point  $x$  is “blamed” when  $d(x, y) = 2^i$ . Hence, this  $x$  has a beacon  $b$  with  $d(x, b) \in [\frac{1}{2}2^i, \frac{3}{2}2^i]$ , and we can thus conclude that  $|S_x| \geq \Omega(\log n)$ .

The proof for  $2^k$ -doubling metrics is a generalization of the our basic technique mentioned above, although the core idea is cluttered by the additional technical complications. A natural approach is to take a  $k$ -dimensional torus, but Slivkins et al. [11, Full version, Theorem 2.4] show that, for the grid in  $\mathbb{R}^k$  endowed with the  $\ell_1$  norm, there is a  $(\varepsilon, 0)$ -triangulation whose order depends only on  $\varepsilon$  and  $k$ . A similar problem arises with the  $\ell_\infty$  norm, namely, there are many “different” paths between any two points such as  $x = (1, \dots, 1)$  and  $y = (k, \dots, k)$ , and hence relatively few beacons suffice to obtain an accurate estimate  $D^+(x, y)$ . We thus use the  $\ell_2$ -norm, and we need to prove that a point that is close to the shortest path between  $x$  and  $y$ , is not also close to the shortest path between  $x$  and  $y' = (k, \dots, k, (1 - \delta)k)$ .

The results for trees and ultrametrics use once again the accounting technique established above for the  $n$ -cycle, except that they use the lower estimate  $D^-(x, y)$ .

## 1.3 Related work

In the *distance labeling problem*, the goal is to assign labels to the points of an input metric space  $(X, d)$ , so that the distance between every pair of points can be computed (exactly or approximately), only from their labels, and the goal is to have the maximum size of a label be as small as possible. We shall assume that the smallest distance in the metric is 1, and denote the largest pairwise distance by  $\Delta$ . Peleg [16] initiated the study of this problem. He showed that tree metrics have an exact distance labeling with  $O(\log n \cdot \log(n + \Delta))$ -bit labels, and that for every  $\kappa \geq 1$ , every graph has  $8\kappa$ -approximate distance labeling with  $O(\log n \cdot \log \Delta \cdot \kappa \cdot n^{1/\kappa})$ -bit labels. Additional results, e.g. for planar graphs and doubling/hyperbolic metrics, are given in [6, 4, 7, 20, 18, 5, 13].

A  $(0, \delta)$ -triangulation of order  $s$  immediately implies a  $(1 + \delta)$ -approximate distance labeling scheme with  $O(s \log(n + \Delta))$ -bit labels, as follows: The label of a point  $x$  records, for all points  $z \in S_x$ , a unique identifier of  $z$  and the distance  $d(x, z)$ .<sup>5</sup> The distance between two points  $x, y$  can

<sup>5</sup>In fact, Slivkins [18] points out that the distance can be

then be approximated by computing  $D^+(x, y)$  (or alternatively  $D^-(x, y)$ ). Notice that it suffices that the only upper estimate is always good, namely,  $D^+(x, y) \leq (1 + \delta)d(x, y)$  for all pairs  $x, y \in X$  (or alternatively, that only the lower estimate  $D^-(x, y)$  is always good). Indeed, this is essentially the approach employed for trees in [16], for planar graphs in [6], and for doubling metrics in [18]; in the first two cases, the triangulation is constructed by finding in the graph a vertex separator (of size 1 or  $O(\sqrt{n})$ , respectively) and proceeding recursively.

Consequently, distance labeling lower bounds imply triangulation lower bounds as follows: if  $(1 + \delta)$ -approximate distance labeling require at least  $B$  bits per label, then  $(0, \delta)$ -triangulation must have order  $\Omega(B/\log(n + \Delta))$ .<sup>6</sup> There are two known lower bounds for approximate distance labeling in metrics of doubling dimension  $k$ : Talwar [20] shows that the labels require  $2^{\Omega(k)}$  bits, and Har-Peled and Mendel [9] show that labels require  $\Omega(\log n)(\log \log \Delta - \log \log n)$  bits, even if  $k = 1$ . But neither of these implies a nontrivial lower bound on triangulation in metrics of doubling dimension  $O(1)$  and diameter  $n^{O(1)}$ . More generally, distance labeling requires only one of the upper and lower estimates to (always) have good accuracy, and hence distance labeling lower bounds may lead to poor triangulation lower bounds (compare e.g. our result for trees with the  $\Omega(\log^2 n)$  bits lower bound known for exact distance labeling [6]). Furthermore, distance labeling lower bounds all work (inherently) by counting arguments, and therefore do not apply to any explicit input metric (unlike our lower bounds, which apply e.g. to the  $n$ -cycle).

Finally, we point out that there are other “generic” methods to obtain a distance labeling scheme, such as to embed the input metric into  $\ell_\infty^m$ , or into an ultrametric (or use a collection of such embeddings); see e.g. [7, 20, 13]. Often, a basic building block in these embeddings is the distance to a subset of the points  $A \subseteq X$  (e.g. each coordinate in the embedding is a map of the form  $x \mapsto d(x, A)$ ), which is dramatically different from a triangulation when  $|A|$  is large.

## 1.4 Preliminaries

For an integer  $k$ , let  $[k] = \{0, 1, 2, \dots, k - 1\}$ . We make no attempt to optimize the constants.

## 2. ONE-DIMENSIONAL METRICS

The following theorem shows that even in simple networks, the order has to be  $\Omega(\log n)$ , and hence Slivkins’ upper bound is tight in the case where  $\delta$  and the doubling dimension  $k$  are constant. After proving the theorem, we discuss some interesting aspects of the proof and further extensions of it.

**THEOREM 2.1.** *For every  $n \geq 2$  there is an  $n$ -point one-dimensional metric, in which every  $(0, \delta)$ -triangulation with  $\delta < 1$  has order  $\Omega(\log n)$ .*

**PROOF.** We assume, without loss of generality, that  $n$  is a power of 2. Let  $(X, d)$  be the  $n$ -point cycle metric, namely, encoded (approximately) using  $O(\log \frac{1}{\delta})$  bits for the mantissa and  $O(\log \log \Delta)$  bits for the exponent, yielding a total of  $O(s \log(n + \frac{1}{\delta} + \log \Delta))$  bits per label.

<sup>6</sup>Using the approximate encoding of Slivkins mentioned above, it also follows that  $(0, \delta/2)$ -triangulation must have order  $\Omega(B/\log(n + \frac{1}{\delta} + \log \Delta))$ .

$X = [n]$  and  $d(x, y) = \min\{|x - y|, n - |x - y|\}$ . It is easy to verify that this metric space has doubling dimensions  $\alpha = 1$ .

Set  $M = (\log n) - 3$  and assume towards contradiction that  $(X, d)$  has a  $(0, \delta)$ -triangulation  $x \mapsto S_x$  of order  $k \leq M/8$ . For  $x \in X$  and  $j \in [M]$ , let  $A(x, j)$  be the event that  $S_x$  contains a beacon  $b$  for which  $2^{j-1} \leq d(x, b) < 2^j$ . Fix  $x \in X$ , and observe that a single  $b \in S_x$  can contribute to at most two of the events  $A(x, 0), \dots, A(x, M - 1)$ , and hence

$$\sum_{j \in [M]} \mathbf{1}_{A(x, j)} \leq 2|S_x| \leq M/4.$$

It follows that for  $j' \in [M]$  chosen uniformly at random,

$$\Pr_{j'}[A(x, j')] \leq 1/4. \quad (1)$$

The last inequality holds for every point  $x \in X$ , and therefore also if  $x \in X$  is chosen at random (as long as it is independent of  $j'$ ).

Now choose uniformly at random  $j' \in [M]$  and  $x' \in X$ , and define  $y' = (x' + 2^{j'}) \bmod n$ . By (1), we have that  $\Pr_{j', x'}[A(x', j')] \leq 1/4$ . Observe that  $y' = (x' + 2^{j'}) \bmod n$  is distributed uniformly over  $X$ , and furthermore that  $y'$  and  $j'$  are independent. We can thus apply (1) again to get that  $\Pr_{j', x'}[A(y', j')] \leq 1/4$ . Using a union bound, we then get

$$\Pr_{j', x'}[A(x', j') \vee A(y', j')] \leq 1/2. \quad (2)$$

It follows by the probabilistic method that there exist  $j \in [M]$  and  $x \in X$  such that for  $y = (x + 2^j) \bmod n$ , neither  $A(x, j)$  nor  $A(y, j)$  occurs. We claim that  $D^+(x, y) \geq 2d(x, y)$ ; this will complete the proof, since clearly  $D^-(x, y) \leq d(x, y)$  and therefore this cannot be a  $(0, \delta)$ -triangulation for  $\delta < 1$ .

It remains to prove the claim. First notice that  $d(x, y) = 2^j$  (because  $j \leq M \leq (\log n) - 1$ ). Now consider  $b \in S_x \cap S_y$ . At least one of  $d(x, b)$  and  $d(y, b)$  must be greater or equal to  $\frac{1}{2}d(x, y) = 2^{j-1}$ . In the case  $d(x, b) \geq 2^{j-1}$ , the fact that  $A(x, j)$  does not occur implies that  $d(x, b) \geq 2^{j+1}$ . In the case  $d(y, b) \geq 2^{j-1}$ , the fact that  $A(y, j)$  does not occur implies that  $d(y, b) \geq 2^{j+1}$ . In either case, we get that  $d(x, b) + d(y, b) \geq 2^{j+1} = 2d(x, y)$ . Since the last inequality holds for all  $b \in S_x \cap S_y$ , we conclude that  $D^+(x, y) \geq 2d(x, y)$ , as claimed.  $\square$

### Extensions..

Theorem 2.1 has several immediate extensions as follows.

1. The lower bound shown by the proof is actually  $\log n - O(1)$ .
2. The proof only uses the fact that  $D^+(x, y) \leq (1 + \delta)d(x, y)$  for all  $x, y \in X$ .
3. The same  $\Omega(\log n)$  lower bound applies also for the  $n$ -path. Indeed, if the  $n$ -path has a triangulation of order  $s$  satisfying  $D^+(x, y) \leq (1 + \delta)d(x, y)$  for all  $x, y$ , then the  $2n$ -cycle has a triangulation of order  $2s + 4$  and having similar accuracy guarantees, constructed as follows. For each  $z \in \{0, \frac{n}{2}, n, \frac{3n}{2}\}$ , take the order  $s$  triangulation we assumed exists for the  $n$ -path  $\{z, z + 1, \dots, z + n - 1\}$  (this handles all  $d(x, y) < n/2$ ), and let every point in the cycle also have the four additional beacons  $0, \frac{n}{2}, n, \frac{3n}{2}$  (this handles all  $n/2 \leq d(x, y) \leq n$ ).

4. The proof does not seem to extend to a lower bound of  $\Omega(\delta^{-1} \log n)$ . In fact, for the  $n$ -cycle there is a triangulation of order  $O(\log n)$  in which  $D^+(x, y) = d(x, y)$  for all  $x, y \in [n]$  (i.e., the upper estimate is exact with zero slack): Every  $x \in [n]$  maintains as beacons the “nearby” multiples of every power of 2, namely

$$S_x = \bigcup_{j \in [\log n]} \{b \in [n] : b \in 2^j \mathbb{Z} \text{ and } d(x, b) < 2^{j+1}\}.$$

It is easy to verify that  $|S_x| \leq 4 \log n$ ; in addition, for every two points  $x, y \in [n]$ , letting  $j \in [\log n]$  be such that  $2^j \leq d(x, y) < 2^{j+1}$ , the shortest path between  $x$  and  $y$  contains a point  $b \in [n]$  which is a multiple of  $2^j$  and thus  $b \in S_x \cap S_y$ , implying that  $D^+(x, y) = d(x, y)$ .<sup>7</sup>

The proof of Theorem 2.1 extends also to triangulation with  $\varepsilon$ -slack, as follows.

**THEOREM 2.2.** *For every  $n \geq 2$  there is an  $n$ -point one-dimensional metric, in which every  $(\varepsilon, \delta)$ -triangulation with  $\delta < 1$  has order  $\Omega(\log \frac{1}{\varepsilon})$ .*

**PROOF.** The proof is quite similar to that of Theorem 2.1, and we only point out the differences. First, we define  $A(x, j)$  to be the event that  $S_x$  contains a beacon  $b$  for which  $2^{j-2} \leq d(x, b) < 2^{j+1}$ . Second, we let  $j'$  be chosen uniformly at random from  $\{\log(\varepsilon n) + 3, \dots, M\}$ ; we consider now a triangulation of order at most  $\frac{1}{12}(\log \frac{1}{\varepsilon} - 6)$ , and thus we can still derive (1). Third, we define  $y' = (x' + r) \bmod n$ , where  $r$  is chosen uniformly from  $\{2^{j'-1}, \dots, 2^{j'}\}$ ; since  $y'$  is still uniformly distributed and independent of  $j'$ , and we can still derive (2). By averaging, it follows that there exists  $j \in \{\log(\varepsilon n) + 3, \dots, M\}$ , for which

$$\Pr_{x', r} [A(x', j) \vee A(y', j)] \leq 1/2. \quad (3)$$

It follows that there are at least  $\frac{1}{2} \cdot n \cdot 2^{j-1} \geq 2\varepsilon n^2$  pairs of points  $(x, y)$  for which neither  $A(x, j)$  nor  $A(y, j)$  occurs; the latter still implies that  $D^+(x, y) \geq 2^{j+1} \geq 2d(x, y)$ , implying that this is not an  $(\varepsilon, \delta)$ -triangulation.  $\square$

### 3. DOUBLING METRICS

In this section we give lower bounds on the order required for a triangulation of certain simple metrics that have low doubling dimension. We start with some preparatory lemmas about the Euclidean space  $\ell_2^k$  (Section 3.1). We then use these lemmas to prove lower bounds for  $(0, \delta)$ -triangulation (Section 3.2) and for  $(\varepsilon, \delta)$ -triangulation (Section 3.3).

#### 3.1 Grid in Euclidean spaces

We will need the following three technical lemmas regarding  $\ell_2^k$ , the  $k$ -dimensional Euclidean metric. For the sake of exposition, the proofs of these lemmas are relegated to Appendix A. We use  $\|\cdot\|$  to denote the  $\ell_2$ -norm. As usual, the distance between a point  $x \in X$  and set  $A \subseteq X$  is defined to be  $d(x, A) = \inf_{p \in A} \|x - p\|$ .

<sup>7</sup>On the other hand, it can be shown that  $(0, \delta)$ -triangulation of the  $n$ -cycle requires order  $\Omega(1/\delta)$ , by using arguments similar to Theorem 2.1, but this time applied to  $D^-$ . This is in contrast to the  $n$ -path, in which two beacons suffice to obtain an exact lower estimate  $D^-$  for all vertex pairs.

The first two lemmas limit the geometric configurations of two points  $x, y$  and a beacon  $b$  such that  $b$  provides a good estimate on the distance between  $x$  and  $y$ . The third lemma provides an estimate on the number of integer-grid points in a ball, compared to a ball with the same center and a smaller radius. Let us briefly motivate their intended use: Suppose a beacon  $b$  gives a good upper estimate for  $\|x - y\|$ , and assume that  $b$  is closer to  $y$  than to  $x$ . Then by the Euclidean geometry, it is clear that  $b$  must reside close to the line segment connecting  $x$  with  $y$ . The first lemma puts this differently — the point  $y$  must reside close to the line connecting  $x$  and  $b$ , and furthermore, to its segment that goes between  $b$  and  $x' = b - (x - b)$  (i.e.  $x'$  is the mirror of  $x$  with respect to  $b$ ). The second lemma then says that all these potential positions for  $y$  can be covered by a few balls of a relatively small radius. In our application,  $y$  must be in an integer-grid point, and hence the third lemma would come in handy to upper bound the total number of such points  $y$  (for a given  $x$  and  $b$ ).

**LEMMA 3.1.** *Fix two points  $x, b \in \mathbb{R}_2^k$  and  $0 < \delta < 1/2$ . Then there exists a line segment  $L$  of length  $\|x - b\|$ , such that every point  $y \in \mathbb{R}^k$  satisfying*

$$\|x - b\| + \|b - y\| \leq (1 + \delta)\|x - y\| \text{ and } \|x - b\| \geq \|y - b\|, \quad (4)$$

*is within distance  $5\sqrt{\delta}\|x - b\|$  from  $L$ , i.e.  $d(y, L) \leq 5\sqrt{\delta}\|x - b\|$ .*

**LEMMA 3.2.** *Fix two points  $x, b \in \mathbb{R}_2^k$  and  $0 < \delta < 1/2$ , and let  $Y \subset \mathbb{R}^k$  be finite and non-empty. Then there exist  $1/\delta$  points  $y_1, \dots, y_{1/\delta} \in Y$  such that every  $y \in Y$  satisfying (4) is contained in  $\bigcup_{t=1}^{1/\delta} B(y_t, 12\sqrt{\delta})$ .*

**LEMMA 3.3.** *Fix  $k \geq 2$ . Let  $B(x, r) = \{z \in \mathbb{R}^k : \|z - x\| \leq r\}$  be the closed ball of radius  $r > 0$  around  $x \in \mathbb{R}^k$ , and let  $\hat{B}(x, r) = \mathbb{Z}^k \cap B(x, r)$ . Then for every  $x \in \mathbb{R}^k$ ,  $\alpha \geq 2$ , and  $r \geq 6\sqrt{k}$ ,*

$$(\alpha - 1/4)^k \leq \frac{|\hat{B}(x, \alpha r)|}{|\hat{B}(x, r)|} \leq (\alpha + 1/2)^k.$$

#### 3.2 Zero-slack triangulation

**THEOREM 3.4.** *For every  $k \geq 2$ ,  $\delta < 1/2$ , and  $n \geq (k/\delta)^k$ , there is a  $2^{O(k)}$ -doubling metric on  $n$  points in which every  $(0, \delta)$ -triangulation has order at least  $(c\delta)^{-k/2+1} \cdot \log(n^{1/k})$ , for some universal constant  $c > 0$ .*

**Remark.** It may seem that this bound has an extra factor of  $1/k$  compared to the one stated in Section 1.1, but modifying the constant  $c$  can easily absorb that factor.

**PROOF.** Let  $(X, d)$  be an  $\ell_2$ -like metric on the  $k$ -dimensional discrete torus (or an  $\ell_2$ -product of  $k$  cycles), formally defined to be as follows. Let  $m = n^{1/k}$  and assume, without loss of generality, it is integral and a power of 2. Let  $X = [m]^k$  and define  $d_i(x, y) = \min\{|x_i - y_i|, m - |x_i - y_i|\}$  and  $d(x, y) = [\sum_i (d_i(x, y))^2]^{1/2}$ . It is easy to verify that this metric space has doubling dimension  $O(k)$  (similar to  $\mathbb{R}^k$ ). As usual, let  $B(x, r) = \{y \in X : d(y, x) \leq r\}$  denote the closed ball (in  $X$ ) of radius  $r$  centered at  $x$ .

Set  $M = \log m$ . Assume towards contradiction that  $(X, d)$  has a  $(0, \delta)$ -triangulation  $x \mapsto S_x$  of order  $\frac{\delta}{64k} \cdot (24\sqrt{\delta})^{-k} \log n$ .

For  $x, y \in X$ , define  $A(x, y)$  to be the event that  $S_x$  contains a beacon  $b$  for which  $d(x, b) + d(b, y) \leq (1 + \delta)d(x, y)$  and  $d(b, x) \geq d(b, y)$ . We immediately get the following claim.

**CLAIM 3.5.** *Let  $x, y \in X$ . If neither  $A(x, y)$  nor  $A(y, x)$  occurs, then  $D^+(x, y) > (1 + \delta)d(x, y)$ .*

**PROOF OF CLAIM 3.5.** Fix  $x, y \in X$  and assume to the contrary that  $D^+(x, y) \leq (1 + \delta)d(x, y)$ . Let  $b \in S_x \cap S_y$  such that  $D^+(x, y) = d(x, b) + d(b, y)$ . Now either  $d(b, x) \geq d(b, y)$  or  $d(b, x) \leq d(b, y)$ , and hence at least one of  $A(x, y)$  and  $A(y, x)$  must occur.  $\square$

Let  $\mu$  be the following distribution: Choose uniformly at random  $x \in X$  and  $j \in \{M/2, \dots, M-1\}$ , and then choose uniformly at random  $y \in B(x, 2^j) \setminus B(x, 2^{j-1})$ . Observe that  $y$  is uniformly distributed over  $X$  (because balls of same radius  $2^j$  have same cardinality), and furthermore  $y$  and  $j$  are independent (because that last observation is true even when conditioned on  $j$ ). The following claim is now proved using the technical machinery developed in Section 3.1. As explained earlier, the main idea is that for every given  $x$  and  $b$ , there are only a few points  $y$  such that  $b$  would be a “useful” beacon for the pair  $x, y$ .

**CLAIM 3.6.**  $\Pr_\mu[A(x, y)] \leq 1/4$ .

**PROOF OF CLAIM 3.6.** We are actually going to prove the following stronger claim: For all  $x' \in X$ ,

$$\Pr_\mu[A(x, y)|x = x'] \leq 1/4.$$

For  $x, y \in X$  and  $b \in S_x$ , let  $A(x, y, b)$  be the event that both  $\{d(x, b) + d(b, y) \leq (1 + \delta)d(x, y)$  and  $d(b, x) \geq d(b, y)\}$ . By definition  $A(x, y) = \bigcup_{b \in S_x} A(x, y, b)$  and hence  $\Pr_\mu[A(x, y)|x = x'] \leq \sum_{b \in S_{x'}} \Pr[A(x, y, b)|x = x']$ .

Now fix  $x' \in X$  and a beacon  $b' \in S_{x'}$ , and let us bound the probability (under the conditional distribution  $\mu|x = x'$ ) of the event  $A(x', y, b')$ . Observe that when this event occurs, we have  $\frac{1}{1+\delta}d(b', x') \leq d(x, y) \leq 2d(b', x')$ , and also  $2^{j-1} < d(x, y) \leq 2^j$ , which together restrict  $j$  to the interval  $[\log d(b', x') - 1, \log d(b', x') - 2]$ . Since  $j$  is an integer chosen independently of  $x = x'$ , the probability of choosing  $j$  in that interval is at most  $8/M$ .

Suppose now that  $j$  is chosen as one of these four values, and consider the randomness in choosing  $y \in B(x', 2^j) \setminus B(x', 2^{j-1})$ . Observe that the ball  $B(x', 2^j)$  in  $X$  looks like (i.e. is isometric to) a finite subset of  $\ell_2^k$ , and thus by Lemma 3.2, there is a collection of  $1/\delta$  points  $y_1, \dots, y_{1/\delta} \in X$  such that every  $y \in X$  satisfying  $A(x', y, b')$  is contained in  $\bigcup_{t=1}^{1/\delta} B(y_t, 6\sqrt{\delta}d(x', b'))$ . Thus, the probability that choosing  $y \in B(x', 2^j) \setminus B(x', 2^{j-1})$  would satisfy  $A(x', y, b')$  is at most

$$\begin{aligned} \frac{\sum_{t=1}^{1/\delta} |B(y_t, 6\sqrt{\delta}d(x', b'))|}{|B(x', 2^j) \setminus B(x', 2^{j-1})|} &\leq \frac{1}{\delta} \cdot \frac{|B(y_1, 6\sqrt{\delta} \cdot 2^{j+1})|}{|B(x', 2^j)|/2} \\ &\leq \frac{2}{\delta} \cdot (24\sqrt{\delta})^k. \end{aligned}$$

where the last inequality follows from Lemma 3.3 (which is applicable because  $6\sqrt{\delta} \cdot 2^{j+1} \geq 6\sqrt{\delta}m \geq 6\sqrt{k}$ ).

Altogether, we conclude that

$$\Pr_\mu[A(x, y)|x = x'] \leq |S_{x'}| \cdot \frac{8}{M} \cdot \frac{2}{\delta} (24\sqrt{\delta})^k \leq 1/4,$$

which proves the claim.  $\square$

We can now complete the proof of Theorem 3.4. Since  $x$  and  $y$  has symmetric roles in the distribution  $\mu$ , we can use Claim 3.6 twice and apply a union bound, to get  $\Pr_\mu[A(x, y) \vee A(y, x)] \leq 1/2$ . It follows by the probabilistic method that there exist  $x, y \in X$  such that neither  $A(x, j)$  nor  $A(y, j)$  occurs. It then follows from claim 3.5 that for these two points,  $D^+(x, y) > (1 + \delta)d(x, y)$ . This completes the proof of Theorem 3.4, since clearly  $D^-(x, y) \leq d(x, y)$  and therefore this is not a  $(0, \delta)$ -triangulation.

### 3.3 Triangulation with slack

**THEOREM 3.7.** *For every  $k \geq 2$ ,  $\delta < 1/2$ ,  $\varepsilon < 1/(3k)^k$ , and  $n \geq (k/\delta)^k$ , there is a  $2^{O(k)}$ -doubling metric on  $n$  points in which every  $(\varepsilon, \delta)$ -triangulation has order at least  $(c\delta)^{-k/2+1} \cdot \log(\varepsilon^{-1/2k})$ , for some universal constant  $c > 0$ .*

**PROOF.** We modify the proof of Theorem 3.4 (in an analogous fashion to the proof of Theorem 2.2). First, we change the distribution  $\mu$  so that  $j$  is chosen at random from  $\{M - \log_{\varepsilon_1} \frac{1}{\varepsilon}, \dots, M-1\}$ , where  $\varepsilon_1 = \varepsilon^{1/2k}$ . Changing the proof of Claim 3.6 accordingly, we get that

$$\Pr_\mu[A(x, y)|x = x'] \leq |S_{x'}| \cdot \frac{4}{\log(1/\varepsilon_1)} \cdot \frac{2}{\delta} (24\sqrt{\delta})^k \leq 1/4,$$

where the last inequality follows from assuming that the triangulation has a low order.

Using this claim, we have  $\Pr_\mu[A(x, y) \vee A(y, x)] \leq 1/2$ , and by averaging we get that there exists  $j'$  such that

$$\Pr_\mu[A(x, y) \vee A(y, x)|j = j'] \leq 1/2.$$

Observe that choosing  $(x, y)$  at random from the distribution on the LHS, namely, from  $(\mu|j = j')$ , is actually a uniform choice over the set  $P = \{(x, y) : x \in X, y \in B(x, 2^{j'}) \setminus B(x, 2^{j'-1})\}$ . Thus, for at least half the pairs  $(x, y)$  in  $P$ , neither  $A(x, y)$  nor  $A(y, x)$  occurs, implying by Claim 3.5 that  $D^+(x, y) > (1 + \delta)d(x, y)$ . We conclude that the slack of this triangulation is at least  $\frac{1}{2}|P|/n^2$ .

It remains to estimate  $|P|$ . Using Lemma 3.3 twice, we get that

$$|B(x, 2^{j'}) \setminus B(x, 2^{j'-1})| \geq \frac{|B(x, 2^{j'})|}{2} \geq \frac{|B(x, m\sqrt{k})|}{2(2m\sqrt{k}/2^{j'})^k}.$$

Since  $B(x, m\sqrt{k}) = X$  and  $j' \geq \log(\varepsilon_1 m)$ , we have that  $|P| \geq \frac{1}{2}(\varepsilon_1/2\sqrt{k})^k n^2 > 2\varepsilon n^2$ , and hence the slack of the triangulation is bigger than  $\varepsilon$ .  $\square$

We remark that a requirement of the form  $\varepsilon < 1/(3k)^k$  seems necessary in the current proof. Without such an assumption, a ball containing about  $\varepsilon$ -fraction of the points in the metric has radius bigger than  $m$ , and is thus not isomorphic to a similar ball in  $\mathbb{Z}_2^k$ . It is plausible that this requirement can be eliminated by extending our proof to a suitable net on the unit sphere in  $\ell_2^k$  (essentially by modifying the lemmas in Section 3.1 accordingly).

## 4. TREE METRICS AND ULTRAMETRICS

This section first shows a strong lower bound for triangulation of tree metrics and ultrametrics. We then use ideas from this proof to show an  $\Omega(\log \frac{1}{\varepsilon})$  lower bound for  $(\varepsilon, \delta)$ -triangulation of one-dimensional ultrametrics. In particular, the latter gives another proof to Theorem 2.2.

**THEOREM 4.1.** *There is a family of tree metrics (and similarly ultrametrics) on  $n$  points, in which every  $(1/9, \delta)$ -triangulation has order  $\Omega(n)^{1/(1+\delta)}$ .*

**PROOF.** Let  $T$  be a complete binary tree of height  $k \geq 1$ . Fix  $\varepsilon \leq 1/9$  and assume that the shortest-path metric  $d(\cdot, \cdot)$  of  $T$  has an  $(\varepsilon, \delta)$ -triangulation of order  $m$ . We let  $n = 2^{k+1} - 1$  denote the number of vertices in  $T$ , and  $L$  denote the set of leaves in  $T$ . Notice that  $|L| = 2^k > n/2$ . Let the vertices of  $T$  at depth  $q = \lceil k/(1+\delta) \rceil$  be denoted by  $z_1, \dots, z_{2^q}$ . For  $j = 1, \dots, 2^q$ , we denote that subtree of  $T$  rooted under  $z_j$  by  $T_j$ . For a leaf  $v \in L$ , define  $j(v)$  as the value  $j \in \{1, \dots, 2^q\}$  such that  $v \in T_j$  (i.e.,  $v$  is a descendant of  $z_j$ ). Roughly speaking, we will show that for most leaves  $v \in L$  the beacon set has a nonempty intersection with many (a constant fraction) of the  $2^q$  trees  $T_j$ .

Let  $L = L_1 \cup L_2$  be a partition of the leaves according to the child of the root which they are a descendant of, and note that  $|L_1| = |L_2| = |L|/2$ . Let  $L^* = (L_1 \times L_2) \cup (L_2 \times L_1)$  and consider a random leaf pair  $(x, y) \in L^*$ . First, observe that the  $(\varepsilon, \delta)$ -triangulation satisfies

$$\Pr_{(x,y) \in L^*} [D^-(x, y) < d(x, y)/(1+\delta)] \leq \varepsilon n^2 / |L^*| < 8\varepsilon. \quad (5)$$

Second, fixing (conditioning on)  $x$ , we know that  $S_x$  contains vertices from at most  $|S_x| \leq m$  subtrees among  $T_1, \dots, T_{2^q}$ , and that the point  $y$  is distributed uniformly over  $|L|/2$  leaves, and furthermore  $T_{j(y)}$  is distributed uniformly over  $2^{q-1}$  of subtrees among  $T_1, \dots, T_{2^q}$ , we get that

$$\Pr_{(x,y) \in L^*} [S_x \cap T_{j(y)} \neq \emptyset] \leq m/2^{q-1}. \quad (6)$$

Applying a similar argument to  $S_y$ , we also get

$$\Pr_{(x,y) \in L^*} [S_y \cap T_{j(x)} \neq \emptyset] \leq m/2^{q-1}. \quad (7)$$

We now claim that  $S_x \cap T_{j(y)} = S_y \cap T_{j(x)} = \emptyset$  imply that the lower estimate is “wrong”, i.e.,  $D^-(x, y) < d(x, y)/(1+\delta)$ . This will complete the proof of the theorem, since the claim together with (5), (6) and (7) imply that

$$\begin{aligned} 1 - 2m/2^{q-1} &\leq \Pr_{(x,y) \in L^*} [S_x \cap T_{j(y)} = S_y \cap T_{j(x)} = \emptyset] \\ &\leq \Pr_{(x,y) \in L^*} [D^-(x, y) < d(x, y)/(1+\delta)] \\ &\leq 8\varepsilon, \end{aligned}$$

and by rearranging we get  $m \geq \Omega(2^q) > \Omega((n/2)^{1/(1+\delta)})$ , which indeed proves the theorem.

It remains to prove the above claim. Fix  $(x, y) \in L^*$  with  $S_x \cap T_{j(y)} = S_y \cap T_{j(x)} = \emptyset$ . Notice that  $d(x, y) = 2k$  by definition of  $L^*$ , and let us show that for every beacon  $b \in S_x \cap S_y$  we have  $|d(x, b) - d(y, b)| < 2q$ . (The formal argument below can be replaced by a case-analysis, but note that  $b$  need not be a leaf.) Indeed, by our assumption  $b \notin T_{j(x)} \cup T_{j(y)}$  and  $j(x) \neq j(y)$ . Let  $w$  be the median vertex (in  $T$ ) for the triple  $\{x, y, b\}$ ;<sup>8</sup> we then get

$$\begin{aligned} d(x, b) - d(y, b) &= d(x, w) - d(y, w) \\ &= d(x, y) - 2d(y, w) \\ &= 2k - 2d(y, w). \end{aligned} \quad (8)$$

<sup>8</sup>In a tree, the median of three vertices is the (unique) vertex that lies on the shortest path between every two of the three given vertices.

Since  $w$  lies on the shortest path between  $b, x \notin T_{j(y)}$ , it must be the case that also  $w \notin T_{j(y)}$ . Therefore,  $d(y, w) > k - q$ , and using (8),  $d(x, b) - d(y, b) < 2q$ . A similar argument shows that  $d(y, b) - d(x, b) < 2q$ , and we obtain, as desired, that  $|d(x, b) - d(y, b)| < 2q$ . We conclude from the above that  $D^-(x, y) < 2q = d(x, y)/(1+\delta)$ , which proves the claim, and thus the theorem for trees.

The above proof immediately extends to ultrametrics, since it only uses the leaves of the tree, and they all reside at the same depth from the root.  $\square$

We remark that the above proof only uses the fact that  $D^-(x, y) \geq d(x, y)/(1+\delta)$  for  $\varepsilon$ -fraction of the pairs  $(x, y)$ .

**THEOREM 4.2.** *There is a family of  $n$ -point ultrametrics of doubling dimension 1 in which every  $(\varepsilon, \frac{3}{4})$ -triangulation has order  $\Omega(\log \frac{1}{\varepsilon})$ .*

**PROOF.** Let  $T$  be a complete binary tree of height  $k \geq 1$ , where the edges at depth  $k$  (incident at the leaves) have length 1 and edges at depth  $1 \leq j < k$  have length  $2^{k-j-1}$ . Let  $L$  be the set of leaves in  $T$ , and let  $n = |L| = 2^k$ . The shortest-path metric  $d(\cdot, \cdot)$  in the tree clearly induces an ultrametric on the leaves  $L$ . Furthermore, this ultrametric  $(L, d)$  has doubling dimension 1, because a ball in the metric corresponds to (the leaves in) a subtree in  $T$ . Observe that if the least common ancestor of  $x, y \in L$  is at depth  $j \in \{0, 1, \dots, k\}$  from the root then  $d(x, y) = 2(1 + 2^0 + 2^1 + \dots + 2^{k-j-2}) = 2^{k-j}$ .

Set  $m = \log \frac{1}{\varepsilon} - 3$ , and assume that the ultrametric  $(L, d)$  has an  $(\varepsilon, \frac{3}{4})$ -triangulation of order  $m/4$ . For  $x \in L$  and  $j \in [m]$ , let  $A(x, j)$  be the event that  $S_x$  contains a beacon  $b \in L$  whose distance from  $x$  is exactly  $2^{k-j}$ . Clearly, for every  $x \in L$ , every beacon in  $S_x$  contributes to an event  $A(x, j)$  for at most one  $j$ , and thus if we choose  $j \in [m]$  uniformly at random then

$$\Pr_j [A(x, j)] \leq \frac{|S_x|}{m} \leq 1/4. \quad (9)$$

Now consider a leaf pair  $(x, y) \in L \times L$  chosen at random as follows: (i) choose  $x$  uniformly at random from  $L$ ; (ii) choose uniformly at random  $j \in [m]$ ; and (iii) choose  $y$  uniformly at random from all points whose distance from  $x$  is exactly  $2^{k-j}$  (i.e., the least common ancestor of  $x$  and  $y$  is at depth  $j - 1$ ).

Since  $x$  and  $j$  are chosen independently, we have by (9) that  $\Pr_{x,j,y} [A(x, j)] \leq 1/4$ . Observe that  $j$  is independent of  $y$  (because  $y$  is uniformly distributed over  $L$  even when  $j$  is given), and hence we can apply (9) again to get  $\Pr_{x,j,y} [A(y, j)] \leq 1/4$ . Altogether, we have

$$\Pr_{x,j,y} [A(x, j) \vee A(y, j)] \leq 1/2. \quad (10)$$

We now claim that if neither  $A(x, j)$  nor  $A(y, j)$  occurs then  $D^+(x, y) \geq 4d(x, y)$  and  $D^-(x, y) = 0$ . Indeed, consider  $b \in S_x \cap S_y$ . Let  $z$  denote the least common ancestor of  $x$  and  $y$ , and notice that if  $b$  is a descendant of  $z$ , then exactly one of  $d(x, b)$  or  $d(y, b)$  must be equal to  $d(x, y) = 2^{k-j}$ , contradicting our assumption that neither  $A(x, j)$  nor  $A(y, j)$  occurs. Thus,  $b$  is not a descendant of  $z$ , and thus  $d(x, b) = d(y, b) \geq 2^{k-j+1}$ . It follows that  $D^-(x, y) = 0$  and  $D^+(x, y) \geq 4d(x, y)$ .

We are now in a position to complete the proof for ultrametrics. By averaging, it follows from (10) that there exists

$j = j_0$  for which  $\Pr_{x,j,y}[A(x,j) \vee A(y,j) \mid j = j_0] \leq 1/2$ . In other words, there are at least  $\frac{1}{2}n \cdot n/2^{j_0+1} \geq 2\epsilon n^2$  distinct leaf pairs  $(x,y) \in L \times L$  for which neither  $A(x,j)$  nor  $A(y,j)$  occurs, and by the claim above,  $D^+(x,y) \geq 4d(x,y)$  and  $D^-(x,y) = 0$ , implying that this is not a  $(\epsilon, \frac{3}{4})$ -triangulation.  $\square$

We remark that the above proof can use either the fact that  $D^-(x,y) \geq d(x,y)/(1+\delta)$  for  $\epsilon$ -fraction of the pairs  $(x,y)$  or that  $D^+(x,y) \leq (1+\delta)d(x,y)$  for  $\epsilon$ -fraction of the pairs  $(x,y)$ .

## 5. CONCLUDING REMARKS

Our results indicate that known constructions may potentially be improved (such as refining the triangulation of [11] and/or its analysis, or extending the  $(0,\delta)$ -triangulation of [18] to  $\epsilon$ -slack). They also suggest that certain quantitative tradeoffs between accuracy and slack are unavoidable. In addition, they expose inherent limitations of standard triangulation, where one wishes to obtain an upper estimate  $D^+(x,y)$  and a lower estimates  $D^-(x,y)$  that are both accurate within  $1+\delta$  (in which case, one obtains a ‘‘certificate’’ for the estimates’ accuracy). A possible workaround is to relax the methodology and obtain only an upper estimate or only a lower estimate; for particular network topologies (e.g., tree metrics or doubling metrics) such an estimate may provably give a good approximation to the true distance.

## Acknowledgments

I thank Alex Andoni and Alex Slivkins for useful discussions of certain technical specifics at different stages of the work.

## 6. REFERENCES

- [1] I. Abraham, C. Gavoille, A. V. Goldberg, and D. Malkhi. Routing in networks with low doubling dimension. In *26th IEEE International Conference on Distributed Computing Systems (ICDCS 2006)*, page 75. IEEE Computer Society, 2006.
- [2] P. Assouad. Plongements lipschitziens dans  $\mathbf{R}^n$ . *Bull. Soc. Math. France*, 111(4):429–448, 1983.
- [3] P. Francis, S. Jamin, C. Jin, Y. Jin, D. Raz, Y. Shavitt, and L. Zhang. IDMaps: A global internet host distance estimation service. *IEEE/ACM Trans. Netw.*, 9(5):525–540, 2001.
- [4] C. Gavoille, M. Katz, N. A. Katz, C. Paul, and D. Peleg. Approximate distance labeling schemes. In *Proceedings of the 9th Annual European Symposium on Algorithms*, pages 476–487. Springer-Verlag, 2001.
- [5] C. Gavoille and O. Ly. Distance labeling in hyperbolic graphs. In *16th International Symposium on Algorithms and Computation (ISAAC)*, volume 3827 of *Lecture Notes in Computer Science*, pages 1071–1079. Springer, 2005.
- [6] C. Gavoille, D. Peleg, S. Pérennes, and R. Raz. Distance labeling in graphs. *Journal of Algorithms*, 53(1):85–112, 2004.
- [7] A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *44th Annual IEEE Symposium on Foundations of Computer Science*, pages 534–543, Oct. 2003.
- [8] J. D. Guyton and M. F. Schwartz. Locating nearby copies of replicated internet servers. In *Proceedings of SIGCOMM ’95*, pages 288–298, New York, NY, USA, 1995. ACM Press.
- [9] S. Har-Peled and M. Mendel. Fast construction of nets in low dimensional metrics, and their applications. *SIAM J. Comput.*, 35(5):1148–1184, 2006.
- [10] D. Karger and M. Ruhl. Finding nearest neighbors in growth-restricted metrics. In *34th Annual ACM Symposium on the Theory of Computing*, pages 63–66, 2002.
- [11] J. M. Kleinberg, A. Slivkins, and T. Wexler. Triangulation and embedding using small sets of beacons. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 444–453, 2004.
- [12] R. Krauthgamer and J. R. Lee. Navigating nets: Simple algorithms for proximity search. In *15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 791–801, Jan. 2004.
- [13] R. Krauthgamer and J. R. Lee. Algorithms on negatively curved spaces. In *Proceedings of 47th FOCS*, 2006. To appear.
- [14] R. Krauthgamer, J. R. Lee, M. Mendel, and A. Naor. Measured descent: A new embedding method for finite metrics. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 434–443. IEEE, Oct. 2004.
- [15] T. S. E. Ng and H. Zhang. Predicting internet network distance with coordinates-based approaches. In *INFOCOM*, volume 1, pages 170–179, 2002.
- [16] D. Peleg. Proximity-preserving labeling schemes. *Journal of Graph Theory*, 33(3):167–176, 2000.
- [17] C. G. Plaxton, R. Rajaraman, and A. W. Richa. Accessing nearby copies of replicated objects in a distributed environment. *Theory Comput. Syst.*, 32(3):241–280, 1999.
- [18] A. Slivkins. Distance estimation and object location via rings of neighbors. In *Proceedings of the 24th annual ACM symposium on Principles of distributed computing*, pages 41–50. ACM Press, 2005.
- [19] A. Slivkins. Distributed approaches to triangulation and embedding. In *16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 640–649, 2005.
- [20] K. 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.
- [21] B. Y. Zhao, L. Huang, S. C. Rhea, J. Stribling, A. D. Joseph, and J. Kubiatowicz. Tapestry: A global-scale overlay for rapid service deployment. *IEEE Journal on Selected Areas in Communications*, 22(1):41–53, 2004.

## APPENDIX

### A. SOME PROOFS

PROOF OF LEMMA 3.1. Let  $L$  be the line segment between  $b$  and  $z = 2b - x$ , and thus its length is  $\|z - b\| = \|b - x\|$ . (Note that the midpoint of the line segment between  $x$  and  $z$  is  $b$ .)

Now consider a point  $y$  as in the proposition’s premise. The case  $\|x - b\| \geq \|x - y\|$  is easy: it follows that  $\|b - y\| \leq \delta\|x - y\|$ , and we’re done by the fact that  $b \in L$ .

We thus assume henceforth that  $\|x - b\| < \|x - y\|$ . Let  $P$  be the orthogonal projection operator from the line going through  $x$  and  $b$  onto the line going through  $x$  and  $y$ . Define  $b' = P(b)$ , and since  $b$  is closer to  $x$  than to  $y$ ,

$$\|b' - x\| \geq \frac{1}{2}\|x - y\|. \quad (11)$$

Define  $\hat{y} = P^{-1}(y)$ , and note that  $\hat{y}$  is unique (as otherwise  $b - x$  is orthogonal to  $y - x$  and thus  $\|x - b\| < \|y - b\|$ ). By similarity of triangles and (11),

$$\|y - x\| \leq \|b - x\| \cdot \frac{\|y - x\|}{\|b' - x\|} \leq 2\|b - x\|,$$

and we get that  $\hat{y} \in L$ . Thus, to finish the proof, it remains to upper bound  $\|y - \hat{y}\|$ .

First, using again similarity of triangles and (11),

$$\|y - \hat{y}\| = \|b - b'\| \cdot \frac{\|y - x\|}{\|b' - x\|} \leq 2\|b - b'\|. \quad (12)$$

Second, we can show that

$$\|b - b'\| \leq \sqrt{5\delta} \cdot \|x - b'\|. \quad (13)$$

Indeed, if we assume to the contrary, then  $\|x - b\|^2 = \|x - b'\|^2 + \|b - b'\|^2 \geq (1 + 5\delta)\|x - b'\|^2$ , and we get that  $\|x - b\| + \|b - y\| > (1 + 2\delta)\|x - b'\| + \|y - b'\|$ . where the last inequality follows by This last inequality together with the triangle inequality  $\|x - b'\| + \|y - b'\| \geq \|x - y\|$  and (11), implies that  $\|x - b\| + \|b - y\| > (1 + \delta)\|x - y\|$ , which contradicts the assumption of the proposition.

We conclude from (12) and (13) that  $\|y - y'\| \leq \sqrt{20\delta}\|x - b'\| \leq \sqrt{20\delta}\|x - b\|$ , implying a similar upper bound on  $d(y, L)$ .  $\square$

**PROOF OF LEMMA 3.2.** Given  $x, b$ , and  $\delta$ , apply Lemma 3.1 and obtain the corresponding line segment  $L$  of length  $\|x - b\|$ . Let  $y'_0, y'_1, \dots, y'_{1/\delta}$  be equal-spaced points on  $L$  (including the endpoints of  $L$ ). Thus, every point  $y \in \mathbb{R}^k$  satisfying (4) is within distance  $6\sqrt{\delta}\|x - b\|$  from at least one point among  $y'_1, \dots, y'_{1/\delta}$ , and thus

$$y \in \bigcup_{t=1}^{1/\delta} B(y'_t, 6\sqrt{\delta}\|x - b\|). \quad (14)$$

For every  $t = 1, \dots, 1/\delta$ , let  $y_t$  be an arbitrary point in  $Y \cap B(y'_t, 6\sqrt{\delta}\|x - b\|)$  (if the latter set is empty, let  $y_t$  be an arbitrary point in  $Y$ ), and by the triangle inequality,

$$Y \cap B(y'_t, 6\sqrt{\delta}\|x - b\|) \subseteq B(y_t, 12\sqrt{\delta}\|x - b\|). \quad (15)$$

The lemma now follows from (14) and (15).  $\square$

**PROOF OF LEMMA 3.3.** For a set  $S \subset \mathbb{R}^k$ , let  $\text{vol}(S)$  denote the  $k$ -dimensional Euclidean volume (Lebesgue measure) of  $S$  (whenever it is defined). It is well-known that for every  $x \in \mathbb{Z}^k$ ,  $r > 0$  and  $\gamma > 1$ ,

$$\text{vol}(B(x, \gamma r)) = \gamma^k \cdot \text{vol}(B(x, r)).$$

Partition  $\mathbb{R}^k$  into unit cubes using the following equivalence relation:  $x \equiv y$  if for all  $i \in \{1, \dots, k\}$  we have  $\lfloor x_i \rfloor = \lfloor y_i \rfloor$ . Observe that each cube has the following 3 properties: (a) it contains exactly one point of  $\mathbb{Z}^k$ ; (b) it has diameter  $\sqrt{k}$ ; and (c) it has volume 1.

Now fix  $x \in \mathbb{Z}^k$ ,  $r \geq k$  and  $\alpha \geq 2$ . Using properties (a)-(c), every point in  $\hat{B}(x, r)$  gives rise to a distinct cube that has non-empty intersection with the ball  $B(x, r)$ , and each such cube is contained in  $B(x, r + \sqrt{k})$ . Thus,

$$|\hat{B}(x, r)| \leq \#\{\text{such cubes}\} \leq \text{vol}(B(x, r + \sqrt{k})).$$

Similarly, each cube that has a non-empty intersection with  $B(x, \alpha r - \sqrt{k})$  contains a distinct point of  $\hat{B}(x, \alpha r)$ , and thus

$$|\hat{B}(x, \alpha r)| \geq \#\{\text{such cubes}\} \geq \text{vol}(B(x, \alpha r - \sqrt{k})).$$

We conclude that

$$\frac{|\hat{B}(x, \alpha r)|}{|\hat{B}(x, r)|} \geq \frac{\text{vol}(B(x, \alpha r - \sqrt{k}))}{\text{vol}(B(x, r + \sqrt{k}))} \geq (\alpha - 1/4)^k.$$

A similar argument bounds  $\frac{|\hat{B}(x, \alpha r)|}{|\hat{B}(x, r)|}$  from above.  $\square$