Circular Partitions with Applications to Visualization and Embeddings

Krzysztof Onak[*]
Massachusetts Institute of Technology
konak@mit.edu

Anastasios Sidiropoulos
Massachusetts Institute of Technology
tasos@mit.edu

ABSTRACT

We introduce a hierarchical partitioning scheme of the Euclidean plane, called *circular partitions*. Such a partition consists of a hierarchy of convex polygons, each having small aspect ratio, and satisfying specified volume constraints. We apply these partitions to obtain a natural extension of the popular Treemap visualization method. Our proposed algorithm is not constrained in using only rectangles, and can achieve provably better guarantees on the aspect ratio of the constructed polygons.

Under relaxed conditions, we can also construct circular partitions in higher-dimensional spaces. We use these relaxed partitions to obtain improved approximation algorithms for embedding ultrametrics into d-dimensional Euclidean space. In particular, we give a polylog(Δ)-approximation algorithm for embedding n-point ultrametrics into $\mathbb{R}^d$ with minimum distortion (Δ denotes the spread of the metric). The previously best-known approximation ratio for this problem was polynomial in n [2]. This is the first algorithm for embedding a non-trivial family of weighted graph metrics into a space of constant dimension that achieves polylogarithmic approximation ratio.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems; I.3.m [**Computer Graphics**]: Miscellaneous

General Terms

Algorithms, Theory

Keywords

embeddings, approximation algorithms, ultrametrics, visualization, TreeMap

[*]Supported in part by NSF grant 0514771.

1. INTRODUCTION

The visualization of hierarchical structures is a fundamental problem in graph drawing, and computer graphics in general. One of the most successful practical algorithms for this problem, that has attracted a lot of attention over the past years, is Treemap [27]. More precisely, one is given a hierarchy of elements represented as a rooted tree with positive weights on its leaves. The weight of each internal vertex is the sum of the weights of the leaves in its subtree. Treemap assigns a rectangle to each vertex such that:

- the area of the rectangle is equal to the weight of the vertex;
- the rectangles of the children of each internal vertex v are disjoint, and are contained inside the rectangle of v.

We propose a natural extension of Treemap, based on a novel hierarchical partitioning scheme of the plane called *circular partitions*, and we give provable guarantees for its performance.

Another geometric approach for visualizing hierarchies can be obtained by low-distortion metric embeddings into low-dimensional spaces. In this visualization scenario, the natural metric analog of a hierarchy is an *ultrametric* (see later in this section for definitions). The goal is to compute an embedding of a given ultrametric into $\mathbb{R}^d$, for some fixed $d \geq 2$, minimizing the distortion (see [2]). Using a simple relaxation of circular partitions, we derive an interesting connection between this problem and Treemap. Our results improve the best-known approximation guarantee for embedding ultrametrics into $\mathbb{R}^d$.

1.1 Our Results

An Extension of Treemap.

The most important goal of the plane partition computed by Treemap is the minimization of the aspect ratio of each rectangle. However, it is easy to construct instances where the aspect ratio of any such rectangular assignment is unbounded. For example, consider a tree with a root and two leaves, where the first leaf has weight 1, and the second has weight L. The optimal aspect ratio of Treemap in this case is unbounded as $L \to \infty$. This simple observation leads to the following natural question:

> Is there a hierarchical partitioning of the plane into convex polygons that achieves aspect ratio independent of the weights?

We answer this question in the affirmative. More precisely, we present an algorithm that given an n-vertex tree of depth d, outputs a partitioning into convex polygons, each having aspect ratio $O(\mathrm{poly}(d, \log n))$.

We remark that the problem of modifying Treemap so that it uses only sets of small aspect ratio has been considered in [10, 4, 3, 30]. However, our work provides the first provable guarantees on the aspect ratio.

Figure 1 depicts partitions computed by our algorithm on synthetic hierarchical data. It would be interesting to compare our algorithm with existing implementations of Treemap, on real data.

Furthermore, if it is required that all polygons assigned to vertices of the tree be rectangles, we show that it is possible to construct a relaxed partition with small aspect ratio, that we call a *rectangular partition with slack*. The difference from the standard partition is that the area of the rectangle assigned to an internal vertex can exceed the sum of the areas of the rectangles assigned to its children by a factor of at most $1 + \epsilon$.

Embedding Ultrametrics into $\mathbb{R}^d$.

It turns out that the notion of rectangular partitions easily generalizes to any dimension. We call the generalized partitions *hyperrectangular partitions with slack*.

Surprisingly, both hyperrectangular partitions with slack and circular partitions can be used to achieve embeddings of ultrametrics into Euclidean spaces. In particular, using the hyperrectangular partitioning scheme, we obtain a polynomial-time polylog(Δ)-approximation algorithm for the problem of embedding ultrametrics into ℓ_2^d with minimum distortion. This is an exponential improvement over the previously best known algorithm for this problem, which has approximation ratio $n^{O(1)}$ [2]. In fact, all previously known approximation algorithms for related problems have polynomial approximation ratios (see Related Work for a detailed discussion).

1.2 Related Work

Treemap.

The *Treemap* algorithm was proposed by Shneiderman [27], and its first efficient implementation was given by Johnson and Shneiderman [17]. There have been several improvements of the original algorithm. Bruls et al. [10] proposed a variant of Treemap that heuristically tries to minimize the aspect ratio of the resulting rectangles. Shneiderman and Wattenberg [28] have proposed a modified algorithm that minimizes the aspect ratio while preserving certain ordering constraints of the rectangles of the children of each vertex. The quality of the representation of a partition has been further improved by van Wijk and van de Wetering [31], who developed a method for displaying the rectangles using more intuitive shading.

Voronoi treemaps [4, 3] are probably the most closely related to ours. The algorithm is not limited to output a partitioning of the plane into rectangles, but is allowed to output arbitrary, even nonconvex objects. Partitioning of an area is done as follows. First a set S of points that correspond to subtrees is placed within the area. Then, each point of the area is assigned to the closest point in S, where the distance function is modified for each point p in S according to the weight of the subtree corresponding to p. An iterative process is used to optimize the placement of points, and the size of an area assigned to a point may slightly differ from the expected. A version of Voronoi treemaps provides a partitioning into polygons. As opposed to the partitioning scheme discussed in this paper, Voronoi treemaps are not known to give any theoretical guarantees on aspect ratios of computed areas.

Another proposed extension of Treemap to non-rectangular objects are *circular treemaps* [35], which use circles instead of rectangles. Circular treemaps are visually appealing, and nicely display nesting, but a lot of space may be wasted in the process of partitioning a circle into smaller circles.

Extensions of Treemap for visualization in 3-dimensional space have been considered by Rekimoto and Green [25], Bladh et al. [9], and Bladh et al. [8]. A variant of Treemap that constructs radial partitions was proposed by Stasko et al. [29].

The Treemap algorithm has been used to visualize a wide range of hierarchical data, including stock portfolios [18], news items [34], blogs [33], business data [32], tennis matches [16], photo collections [7], and file-system usage [27, 35].

Shneiderman maintains a webpage [26] that describes the history of his invention. It gives an overview of applications and proposed extensions to his original idea.

Approximation Algorithms for Metric Embeddings.

The problem of embedding ultrametrics into $\mathbb{R}^2$ has been shown to be NP-complete in [2]. The same paper gives an $O(n^{1/3})$-approximation algorithm for this problem. They also extended the algorithm for embedding ultrametrics into $\mathbb{R}^d$, obtaining a $(n^{\frac{1}{d}-\Theta(\frac{1}{d^2})})$-approximation. Badoiu *et al.* [12] gave an $O(1)$-approximation algorithm for embedding subsets of the 3-dimensional sphere into $\mathbb{R}^2$.

Recently, it has been shown by Matoušek and Sidiropoulos [23] that for any $d \geq 2$, minimum distortion embedding of general metrics into $\mathbb{R}^d$ is NP-hard to approximate within $\Omega(n^{1/(17d)})$. This result implies that restricting our attention to special classes of metrics, such as ultrametrics, is in general necessary in order to obtain a poly-logarithmic approximation ratio.

For the case of embedding into the line, Badoiu *et al.* [12] have given a $O(\sqrt{n})$-approximation for embedding unweighted graphs, and a $O(n^{1/3})$-approximation for embedding unweighted trees. For weighted metrics, it has been shown by Badoiu *et al.* [11] that there exist $0 < \alpha < \beta < 1$ such that embedding general trees into the line is $\Omega(n^\alpha)$-hard to approximate, while there exists a $O(n^\beta)$-approximation for the same problem. For embedding general metrics into the line, there exists an $O(\Delta^{3/4} n^{1-\epsilon})$-approximation algorithm, for some $\epsilon > 0$ [11].

Although the guarantees on the distortion on the above positive results are better when the optimal distortion is small, in terms of approximation ratio our algorithm for embedding ultrametrics into $\mathbb{R}^d$ is the first one that achieves sub-polynomial approximation guarantee for embedding a non-trivial family of graphs metrics into a space of constant dimension.

Approximation algorithms for embeddings into high-dimensional spaces have also been considered. In particular, it has been shown by Linial *et al.* [21] that there exists a polynomial-time algorithm for computing an optimal embedding of a metric space into ℓ_2. Lee *et al.* [20] gave a

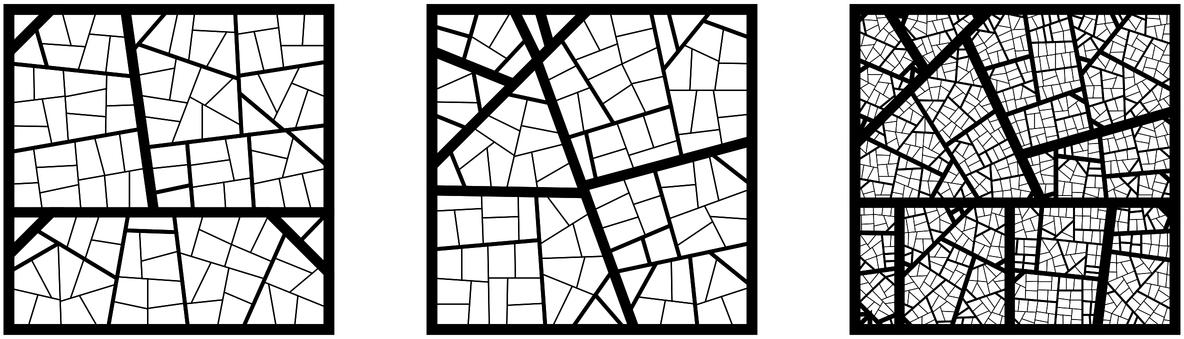

Figure 1: Hierarchical partitions computed by the modified Treemap algorithm on synthetic data. Thicker boundaries correspond to higher levels of the partition.

$O(1)$-approximation algorithm for embedding trees into ℓ_p.

For the case of embedding into ℓ_1, Avis and Deza [1] have shown that it is NP-hard to decide whether a given metric space embeds isometrically (i.e. with distortion 1). Interestingly, it has been shown by Malitz and Malitz [22] (see also Edmonds [14]) that deciding isometric embedding into 2-dimensional ℓ_1 can be done in polynomial time, while Edmonds [14] has shown that it is NP-hard for 3-dimensional ℓ_1.

The question of approximating the minimum distortion has also been investigated under the requirement that the embedding is a bijection (cf. [19, 24, 15, 13]).

1.3 Definitions and Notation

Metrics and Embeddings.

The spread of a metric space is the ratio of the diameter to its minimum distance.

An embedding of metric space $M = (X, D)$ into a metric space $M' = (X', D')$ is a mapping $f : X \to X'$. The distortion of such an embedding is defined as $\max_{x,y \in X} \frac{D'(f(x),f(y))}{D(x,y)} \cdot \max_{z,w \in X} \frac{D(z,w)}{D'(f(z),f(w))}$.

An ultrametric $M = (X, D)$ is a metric space that can be realized as the shortest-path metric over the leaves of a rooted weighted tree T, such that the distance between the root and any leaf is the same. Equivalently, M is an ultrametric iff for any $x, y, z \in X$, $D(x, z) \le \max\{D(x, y), D(y, z)\}$.

Geometry and Aspect Ratios.

For a set $A \subset \mathbb{R}^d$, let $\mathrm{Vol}(A)$, and $\mathrm{diam}(A)$ denote the d-dimensional volume, and the diameter of A, respectively. Let also $\mathrm{int}(A)$, $\mathrm{cl}(A)$, and ∂A, be the interior, the closure, and the boundary of A respectively. We define the *aspect ratio* of a polygon A to be $\lambda(A) = \frac{\mathrm{diam}(A)^2}{\mathrm{Vol}(A)}$.

For a d-dimensional hyperrectangle R of sides $s_1, s_2, \ldots, s_d \in \mathbb{R}_+$, the *rectangular aspect ratio* $\lambda_{\mathrm{rect}}(R)$ of R equals $\frac{\max_i s_i}{\min_i s_i}$. It can easily be shown that for 2-dimensional rectangles, the aspect ratio and the rectangular aspect ratio are within a constant factor.

2. HIERARCHICAL CIRCULAR PARTITIONS OF THE PLANE

We show an algorithm that constructs a partition of the plane that reflects properties of a tree with weights $w(\cdot)$ assigned to its vertices. There is a 1-to-1 correspondence between the polygons in the partition and the vertices of the tree, and each polygon has volume equal to the weight of the corresponding vertex.

Throughout the paper, we will refer to this partition as *hierarchical circular partition*. We call it "hierarchical" because if a vertex v is a descendant of another vertex u, then the polygon corresponding to v is contained inside the polygon corresponding to u. Furthermore, if two vertices are not in the ancestor-descendant relation in the tree, the interiors of the polygons corresponding to these two vertices are disjoint. The term "circular" is used because we require all the polygons to have small aspect ratio. Intuitively, if a polygon has small aspect ratio, it is close to a circle. The main technical difficulty that we face is showing that the aspect ratios of all polygons in our partition are small.

A formal specification of all the desired properties of such a partition follows. We write $\mathcal{P}(S)$ to denote the power set of S, i.e., the set of all subsets of S.

Definition 1. (γ-Hierarchical Circular Partition) Let $T = (V, E)$ be a rooted tree with n leaves, and depth d. Let $w : V \to \mathbb{R}_{\ge 0}$ be a function such that for any internal vertex $v \in V(T)$, with children $u_1, \ldots, u_k$, $w(v) \ge \sum_{i=1}^{k} w(u_i)$. Then, for some $\gamma > 0$, a *γ-hierarchical circular partition* for (T, w) is a mapping $f : V(T) \to \mathcal{P}(\mathbb{R}^2)$, such that:

- For each $v \in V(T)$, $f(v)$ is a convex polygon in $\mathbb{R}^2$ with $\lambda(f(v)) \le \gamma$.
- For each $v \in V(T)$, $\mathrm{Vol}(f(v)) = w(v)$.
- For each $u, v \in V(T)$, such that u is the parent of v in T, $f(v) \subseteq f(u)$.
- For each $u, v \in V(T)$, such that u is not an ancestor of v, and v is not an ancestor of u, $\mathrm{int}(f(u)) \cap \mathrm{int}(f(v)) = \emptyset$.

2.1 Existence of a good cut

The main component of a proof that hierarchical circular partitions with good properties exist will be the following lemma. It shows that there is always a way to cut a polygon into two smaller polygons of required volumes so that the aspect ratios of the new polygons are bounded. The proof of the lemma is long and consists of a case analysis.

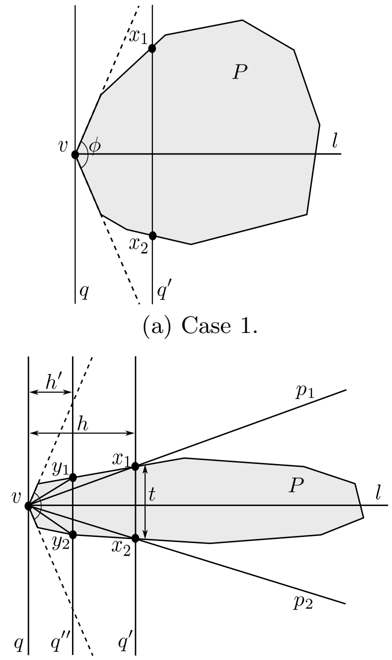

(a) Case 1.

(b) Case 1.2.

Figure 2: Partitioning P into P_1, and P_2, when $\alpha \leq 1/k^2$.

Lemma 1 (Circular Cut). *Let $P \subset \mathbb{R}^2$ be a convex polygon with k vertices, and aspect ratio $\lambda(P)$, and let $a \in (0, 1/2]$. Then, P can be partitioned into two convex polygons P_1, and P_2, such that*

- *Each of the P_1, and P_2 has at most $k + 1$ vertices.*
- $\mathrm{Vol}(P_1) = a \cdot \mathrm{Vol}(P)$, *and* $\mathrm{Vol}(P_2) = (1 - a) \cdot \mathrm{Vol}(P)$.
- *The aspect ratio of each of the P_1, P_2 is at most* $\max\{\lambda(P_1), \lambda(P_2)\} \leq \max\left\{\lambda(P)\left(1 + \frac{6}{k}\right), k^8\right\}$.

Proof. We distinguish between the following two cases.

Case 1: $a \leq 1/k^2$. Let ϕ be the smallest angle of P, and let v be a vertex of P, incident to an angle ϕ. Since P has k vertices, we have

$$\phi \leq \pi\left(1 - \frac{2}{k}\right)$$

Let l be the bisector of ϕ, and let q be the line normal to l. Let S be the halfplane with boundary q, such that $S \cap P = v$. Consider the translation S' of S, such that

$$\mathrm{Vol}(S' \cap P) = a \cdot \mathrm{Vol}(P)$$

Let also q' be the boundary of S'. We define $P_1 = S' \cap P$, and $P_2 = \mathrm{cl}(P \setminus S')$. Clearly, P_1, and P_2 are convex polygons with at most $k+1$ vertices each, such that $\mathrm{Vol}(P_1) = a \cdot \mathrm{Vol}(P)$, and $\mathrm{Vol}(P_2) = (1 - a) \cdot \mathrm{Vol}(P)$. Therefore, it remains to bound the aspect ratios of P_1, and P_2.

Since $P_2 \subset P$, we have

$$\begin{aligned}
\lambda(P_2) &= \frac{\mathrm{diam}(P_2)^2}{\mathrm{Vol}(P_2)} \leq \frac{\mathrm{diam}(P)^2}{(1-a)\cdot \mathrm{Vol}(P)} = \frac{\lambda(P)}{1-a} \\
&< \lambda(P)\,(1+2a) < \lambda(P)\left(1 + \frac{2}{k^2}\right) \\
&< \lambda(P)\left(1 + \frac{1}{k}\right).
\end{aligned}$$

We next bound $\lambda(P_1)$. Let x_1, x_2 be the two points where q' intersects ∂P, and let t be the distance between x_1, and x_2. Let h be the distance between the lines q and q'. Figure 2(a) depicts the arrangement. We distinguish between the following cases.

Case 1.1: $t \geq h/k^2$. Since P is convex, the triangle vx_1x_2 is contained in P_1. Therefore, $\mathrm{Vol}(P_1) \geq h \cdot t/2 \geq h^2/(2k^2)$. On the other hand, since S' is normal to the bisector of the angle of v, it follows that P_1 is contained inside a rectangle of width h, and height H, with

$$\begin{aligned}
H &\leq 2 \cdot h \cdot \tan(\phi/2) \leq 2 \cdot h \cdot \tan\left(\frac{\pi(1 - 2/k)}{2}\right) \\
&\leq 2 \cdot h/\tan(\pi/k) \leq 2 \cdot h \cdot k/\pi
\end{aligned}$$

Thus, $\mathrm{diam}(P_1) < h(1 + 2 \cdot k/\pi)$. It follows that

$$\lambda(P_1) = \frac{\mathrm{diam}(P_1)^2}{\mathrm{Vol}(P_1)} < \frac{(h + 2 \cdot h \cdot k/\pi)^2}{h^2/(2k^2)} < k^5$$

Case 1.2: $t < h/k^2$. Let p_1 be the line passing through v, and x_1, and let p_2 be the line passing through v, and x_2. Let γ be the angle between p_1, and p_2. Observe that P_2 is contained between p_1 and p_2. Therefore, there exist a point $u \in P_2$, such that

$$\frac{\gamma}{2\pi}\pi\|u - v\|_2^2 \geq \mathrm{Vol}(P_2)$$

It follows that $\mathrm{diam}(P)^2 \geq \|u - v\|_2^2 \geq \frac{2}{\gamma}(1 - a)\,\mathrm{Vol}(P)$. Therefore,

$$\lambda(P) = \frac{\mathrm{diam}(P)^2}{\mathrm{Vol}(P)} \geq \frac{2}{\gamma}\,(1-a) \geq \frac{2}{\gamma}\left(1 - \frac{1}{k^2}\right)$$

We now give an upper bound on the diameter of P_1. Assume w.l.o.g. that $\|v - x_2\|_2 \geq \|v - x_1\|_2$, and let $R = \|v - x_2\|_2$. Consider a line q'', parallel to q, that lies between q and q'. Let h' be the distance between q and q''. The line q'' intersects ∂P_1 on two points y_1, y_2 (see Figure 2(b)). We will show that $\|y_1 - y_2\|_2 \leq 2t$. Assume for the sake of contradiction, that $\|y_1 - y_2\|_2 > 2t$. Let g_1 be the line passing through y_1, and x_1, and let g_2 be the line passing through y_2, and x_2. Observe that since $\|y_1 - y_2\|_2 > \|x_1 - x_2\|_2$, it follows that g_1, and g_2 intersect at a point w, such that P_2 is contained in the triangle x_1x_2w. Observe that the polygon $vy_1x_1x_2y_2$ is contained in P_1. If $h' \geq h/2$, then the volume of the triangle vy_1y_2 is greater or equal to the volume of the triangle x_1x_2w. Therefore, $\mathrm{Vol}(P_1) \geq \mathrm{Vol}(P_2)$, contradicting the fact that $a \leq 1/k^2$. If on the other hand $h' < h/2$, then the volume of the quadrilateral $y_1x_1x_2y_2$, is greater than the volume of the triangle x_1x_2w, implying that $\mathrm{Vol}(P_1) \geq$

$\mathrm{Vol}(P_2)$, a contradiction. Therefore, we obtain that $\|y_1 - y_2\|_2 \le 2t$.

It now follows that any point $u \in P_1$ is at distance at most $2t$ from the line segment vx_2. Thus,

$$\begin{aligned}\mathrm{diam}(P_1) &= \max_{u,u' \in P_1} \|u - u'\|_2 \\ &\le \max_{u,u' \in P_1} \{2t + \|v - x_2\|_2 + 2t\} \\ &\le R + 4t \le R\left(1 + \frac{4}{k^2}\right).\end{aligned}$$

Let x^* be the point on the line segment x_1x_2, that is closest to v. Since $R \ge h$, we have

$$\begin{aligned}\mathrm{Vol}(P_1) &\ge \frac{\gamma}{2\pi}\pi\|v - x^*\|_2^2 \ge \frac{\gamma}{2}(R-t)^2 \\ &\ge \frac{\gamma}{2}R^2\left(1 - \frac{1}{k^2}\right).\end{aligned}$$

Therefore,

$$\begin{aligned}\lambda(P_1) &= \frac{\mathrm{diam}(P_1)^2}{\mathrm{Vol}(P_1)} \le \frac{2}{\gamma} \cdot \frac{(1+4/k^2)^2}{1-1/k^2} \\ &\le \lambda(P)\frac{(1+4/k^2)^2}{(1-1/k^2)^2} \le \lambda(P) \cdot (1+6/k^2)^2 \\ &\le \lambda(P) \cdot (1+2/k)^2 \le \lambda(P) \cdot (1+6/k)\end{aligned}$$

Case 2: $a > 1/k^2$.

Case 2.1: $\lambda(P) \le k^6$. We pick an arbitrary half-plane H, such that $\mathrm{Vol}(P \cap H) = a \cdot \mathrm{Vol}(P)$. We set $P_1 = P \cap H$, and $P_2 = \mathrm{cl}(P \setminus H)$. Clearly, we have

$$\lambda(P_1) = \frac{\mathrm{diam}(P_1)^2}{\mathrm{Vol}(P_1)} \le \frac{\mathrm{diam}(P)^2}{a \cdot \mathrm{Vol}(P)} \le k^2 \cdot \lambda(P) \le k^8$$

and

$$\begin{aligned}\lambda(P_2) &= \frac{\mathrm{diam}(P_2)^2}{\mathrm{Vol}(P_2)} \le \frac{\mathrm{diam}(P)^2}{(1-a) \cdot \mathrm{Vol}(P)} \\ &\le 2 \cdot \lambda(P) \le 2 \cdot k^6 < k^7\end{aligned}$$

Case 2.2: $\lambda(P) > k^6$. Pick points $v_1, v_2 \in P$, such that $\|v_1 - v_2\|_2 = \mathrm{diam}(P)$. Let ρ be the line passing through v_1, and v_2. Let also ν_1, and ν_2, be the lines normal to ρ, passing through v_1, and v_2 respectively. Note that P is contained between ν_1, and ν_2.

For each $z \in [0, \mathrm{diam}(P)]$, let $\nu(z)$ be a line normal to ρ that is at distance z from ν_1, and at distance $\mathrm{diam}(P) - z$ from ν_2. Define $f(z)$ to be the length of the intersection of P with $\nu(z)$. Observe that

$$\mathrm{Vol}(P) = \int_{z=0}^{\mathrm{diam}(P)} f(z)dz$$

Pick $s_1, s_2 \in [0, \mathrm{diam}(P)]$, so that

$$a \cdot \mathrm{Vol}(P) = \int_{z=0}^{s_1} f(z)dz = \int_{z=\mathrm{diam}(P)-s_2}^{\mathrm{diam}(P)} f(z)dz$$

Let Q_1 be the part of P that is contained between ν_1, and $\nu(s_1)$. Similarly, let Q_2 be the part of P that is contained between $\nu(\mathrm{diam}(P) - s_2)$, and ν_2. Clearly, both Q_1, and Q_2 are convex polygons with at most $k+1$ vertices.

First, we will show that

$$\min\left\{\frac{\mathrm{Vol}(Q_1)}{s_1}, \frac{\mathrm{Vol}(Q_2)}{s_2}\right\} \le \frac{\mathrm{Vol}(P)}{\mathrm{diam}(P)}$$

Assume for the sake of contradiction that $\frac{\mathrm{Vol}(Q_1)}{s_1} > \frac{\mathrm{Vol}(P)}{\mathrm{diam}(P)}$, and $\frac{\mathrm{Vol}(Q_2)}{s_2} > \frac{\mathrm{Vol}(P)}{\mathrm{diam}(P)}$. It follows that there exist $z_1 \in [0, s_1]$, and $z_2 \in [\mathrm{diam}(P) - s_2]$, such that $f(z_1) > \frac{\mathrm{Vol}(P)}{\mathrm{diam}(P)}$, and $f(z_2) > \frac{\mathrm{Vol}(P)}{\mathrm{diam}(P)}$. Since P is convex, f is a bitonic function. Therefore, for each $z \in [z_1, z_2]$, $f(z) > \frac{\mathrm{Vol}(P)}{\mathrm{diam}(P)}$. It follows that

$$\begin{aligned}\mathrm{Vol}(P) &= \mathrm{Vol}(Q_1) + \mathrm{Vol}(Q_2) \\ &\quad + \mathrm{Vol}(P \setminus (Q_1 \cup Q_2)) \\ &> \frac{\mathrm{Vol}(P)}{\mathrm{diam}(P)} \cdot \mathrm{diam}(P),\end{aligned}$$

a contradiction.

We can therefore assume w.l.o.g. that

$$\frac{\mathrm{Vol}(Q_1)}{s_1} \le \frac{\mathrm{Vol}(P)}{\mathrm{diam}(P)}$$

Note that this implies

$$s_1 \ge a \cdot \mathrm{diam}(P)$$

We set $P_1 = Q_1$, and $P_2 = P \setminus Q_1$. It remains to bound $\lambda(P_1)$, and $\lambda(P_2)$.

By the convexity of P, $\mathrm{Vol}(P) \ge \max_{z \in [0, \mathrm{diam}(P)]} f(z) \cdot \mathrm{diam}(P)/2$. Since $\lambda(P) > k^6$, it follows that

$$\max_{z \in [0, \mathrm{diam}(P)]} f(z) < \frac{2}{k^6} \cdot \mathrm{diam}(P).$$

This implies that P is contained inside a rectangle with one edge of length $\mathrm{diam}(P)$ parallel to ρ, and one edge of length $\frac{4}{k^6} \cdot \mathrm{diam}(P)$ normal to ρ. Thus,

$$\mathrm{diam}(P_1) \le s_1 + \frac{4}{k^6} \cdot \mathrm{diam}(P).$$

Let σ_1, σ_2 be the two points where $\nu(s_1)$ intersects ∂P. Let ζ_1, ζ_2, be the lines passing through v_1, and σ_1, σ_2 respectively. Let also σ_1', and σ_2', be the points where ζ_1, and ζ_2 respectively intersect ν_2 (see Figure 3). By the convexity of P and P_1, we have

$$\begin{aligned}\mathrm{Vol}(P_1) &\ge \mathrm{Vol}(v_1\sigma_1\sigma_2) \\ &= \left(\frac{s_1}{\mathrm{diam}(P)}\right)^2 \cdot \mathrm{Vol}(v_1\sigma_1'\sigma_2') \\ &\ge \left(\frac{s_1}{\mathrm{diam}(P)}\right)^2 \cdot \mathrm{Vol}(P).\end{aligned}$$

Since $\mathrm{Vol}(P_1) = \alpha \cdot \mathrm{Vol}(P)$, it follows that

$$s_1 \le \sqrt{\alpha} \cdot \mathrm{diam}(P).$$

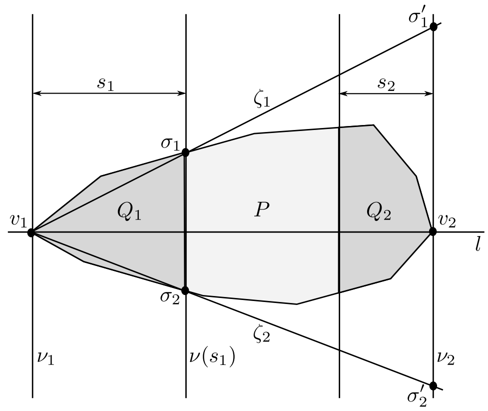

Figure 3: Partitioning P into P_1, and P_2, when $\alpha > 1/k^2$: Case 2.2.

Therefore,

$$\begin{aligned}\lambda(P_1) &= \frac{\mathrm{diam}(P_1)^2}{\mathrm{Vol}(P_1)}\\ &\leq \frac{(s_1 + 4\cdot \mathrm{diam}(P)/k^6)^2}{\mathrm{Vol}(P_1)}\\ &\leq \frac{(\sqrt{\alpha}\cdot \mathrm{diam}(P) + 4\cdot \mathrm{diam}(P)/k^6)^2}{\alpha\cdot \mathrm{Vol}(P)}\\ &< \frac{\mathrm{diam}(P)}{\sqrt{\mathrm{Vol}(P)}}\cdot (1+4/k^4)^2\\ &\leq \lambda(P)\cdot(1+8/k^4+16/k^{16})\\ &\leq \lambda(P)\cdot(1+1/k)\end{aligned}$$

Since f is bitonic, it follows that

$$\min_{z\in[s_1,\mathrm{diam}(P)-s_2]} f(z) \geq$$

$$\min\{\max_{z\in[0,s_1]} f(z), \max_{z\in[\mathrm{diam}(P)-s_2,\mathrm{diam}(P)]} f(z)\}$$

Therefore,

$$\frac{\mathrm{Vol}(P_2)}{\mathrm{diam}(P)-s_1} \geq \frac{\mathrm{Vol}(P_1)}{s_1}$$

We have

$$\mathrm{diam}(P_2) \leq \mathrm{diam}(P) - s_1 + \frac{4}{k^6}\cdot \mathrm{diam}(P)$$

Thus,

$$\begin{aligned}\lambda(P_2) &= \frac{\mathrm{diam}(P_2)^2}{\mathrm{Vol}(P_2)}\\ &\leq \frac{(\mathrm{diam}(P)(1+4/k^6)-s_1)^2}{(1-a)\cdot \mathrm{Vol}(P)}\\ &\leq \lambda(P)\cdot\left(\frac{1+4/k^6-a}{\sqrt{1-a}}\right)^2\\ &\leq \lambda(P)\cdot(1+4\cdot\sqrt{2}/k^6)^2\\ &\leq \lambda(P)\cdot(1+1/k^2)^2\\ &\leq \lambda(P)\cdot(1+3/k^2)\\ &\leq \lambda(P)\cdot(1+1/k)\end{aligned}$$

This concludes the proof. □

2.2 Circular partitions

Now we have all the necessary tools to prove that for any tree T, there exists a γ-hierarchical circular partition with γ polynomial in the depth of T and the logarithm of the number of leaves in T. Initially, we transform T into an equivalent balanced binary tree. For a binary tree, at each internal vertex we can split the polygon corresponding to it into two polygons corresponding to its children with a single cut. To determine the cut, we use Lemma 1, which yields that the aspect ratios of all the polygons will be bounded.

Lemma 2. (Existence of Hierarchical Circular Partitions) *Let $T = (V, E)$ be a rooted tree with n leaves, and depth d. Let $w : V \to \mathbb{R}_{\geq 0}$ be a function such that for any interval vertex $v \in V(T)$, with children $u_1, \ldots, u_k$, $w(v) \geq \sum_{i=1}^{k} w(u_i)$. Then, there exists an $O\left((d\cdot \lg n)^{17}\right)$-hierarchical circular partition for (T, w).*

Proof. Let r be the root of T. We first construct a *binary* tree $T' = (V, E)$, such that $V(T) \subseteq V(T')$, and for each $u, v \in V(T)$, if u is an ancestor of v in T, then u is also an ancestor of v in T'. Clearly, this can be done as follows: For each non-leaf vertex $v \in V(T)$, we replace the set of edges connecting u with its children by a balanced binary tree of depth at most $\lceil \lg n\rceil$. The resulting tree has depth $d' \leq d\cdot\lceil \lg n\rceil$. We define weights w' of nodes in T' as follows. For each node $v \in V(T)$, we set $w'(v) = w(v)$. For each other node $v \in V(T')\setminus V(T)$, that was added to T' as a result of replacing the edges adjacent to a vertex u by a balanced binary tree, we set the value $w'(v)$ to be the sum of the weights of the children of u that are below v in T'. Note that for any node $v \in V(T')$, the sum of the weights of its children in T' is at most $w'(v)$.

We will define inductively a hierarchical circular partition f, starting from r. We set $f(r)$ to be a square in $\mathbb{R}^2$ of volume $w(r)$. Consider now a non-leaf vertex $v \in V(T')$ such that $f(v)$ has already been defined. The volume of the polygon $f(v)$ is $w'(v)$. Let t be the sum of the weights of the children of v in T'. Let P be the polygon obtained by uniform shrinking of $f(v)$ by a factor of $\sqrt{t/w'(v)}$ with any point inside $f(v)$ being a fixed point of the transformation. The volume of P equals t. If v has exactly one child u in T', then we simply set $f(u) = P$. Otherwise, let u_1, u_2 be the children of v in T'. Let $a = \frac{w'(u_1)}{w'(u_1)+w'(u_2)}$. Applying Lemma 1, we partition $f(v)$ into two convex polygons P_1, and P_2, such that $\mathrm{Vol}(P_1) = a\cdot \mathrm{Vol}(f(v)) = w'(u_1)$, and $\mathrm{Vol}(P_2) = (1-a)\cdot \mathrm{Vol}(f(v)) = w'(u_2)$. Moreover, we have $\max\{\lambda(P_1), \lambda(P_2)\} \leq \max\left\{\lambda(f(v))\left(1+\frac{6}{k}\right), k^8\right\}$. We set $f(u_1) = P_1$, and $f(u_2) = P_2$.

We would like to bound $\lambda(f(v))$, for each $v \in V(T)$. Since $f(r)$ is a square, we have that $\lambda(f(r)) = 2$. Consider now $v \in V(T')$. Let t be the distance between r and v in T'. Let p be the path from r to v in T', with $p = v_0, v_2, \ldots, v_t$, where $v_0 = r$, and $v_t = v$. Observe that for each $i \in \{0, \ldots, t\}$, $f(v_i)$ is a convex polygon with at most $i+4$ vertices. It follows by Lemma 1, that for each $i \in \{1, \ldots, t\}$,

$$\lambda(f(v_i)) \leq \max\left\{(i+3)^8, \lambda(f(v_{i-1}))\cdot\left(1+\frac{6}{i+3}\right)\right\}.$$

Hence, we have

$$\begin{aligned}\lambda(f(v_i)) &\leq (t+3)^8 \cdot \prod_{j=3}^{t+3}\left(1+\frac{6}{j}\right)\\ &= (t+3)^8 \cdot \frac{\prod_{j=3}^{t+3}(j+6)}{\prod_{j=3}^{t+3} j}\\ &\leq (t+3)^8 \cdot (t+9)^6 \leq (t+9)^{14}.\end{aligned}$$

□

2.3 Implementation Remark

The proof of Lemma 1 is constructive and shows how to efficiently compute a good cut. Nevertheless, from the practical perspective, a natural heuristic to consider is to always compute the best cut. This is how the circular partitions in Picture 1 were computed.

3. PARTITIONS WITH SLACK

In this section, we show that if we allow small distortion of the volumes at each level of the tree, then there exists a partition of a hypercube into hyperrectangles (d-dimensional rectangles) of small aspect ratio. For each internal node, the hyperrectangles assigned to its children, may have volumes shrunken by a factor in the range $[1-\epsilon, 1]$ with respect to the volume assigned to their parent.

In the algorithm, we always use cuts perpendicular to the longest side of a hyperrectangle. We try to balance the weights of the children assigned to each resulting hyperrectangle. If this is possible, the two resulting hyperrectangles also have small aspect ratios. Otherwise, one child must have large weight. Therefore, we can maintain small aspect ratios by slightly shrinking the volume of its hyperrectangle, and using the resulting empty space to improve the aspect ratio of the other, small hyperrectangle.

Definition 2. (Hierarchical Hyperrectangular Partition with Slack) Let $T=(V,E)$ be a rooted tree with n leaves, and depth d. Let $w : V \to \mathbb{R}_{\geq 0}$ be a function such that for any internal vertex $v \in V(T)$, with children $u_1, \ldots, u_k$, $w(v) \geq \sum_{i=1}^{k} w(u_i)$. Then a *$\gamma$-hierarchical hyperrectangular partition with ϵ-slack* for (T, w) is a mapping $f : V(T) \to \mathcal{P}(\mathbb{R}^d)$, for some $d \geq 2$, such that:

- For each $v \in V(T)$, $f(v)$ is a d-dimensional hyperrectangle with $\lambda_{\text{rect}}(f(v)) \leq \gamma$.
- For the root r of T, $\text{Vol}(f(r)) = w(r)$.
- For each $u, v \in V(T)$, such that u is the parent of v in T, $f(v) \subseteq f(u)$, and
$$(1-\epsilon)\frac{\text{Vol}(f(u))}{w(u)} \leq \frac{\text{Vol}(f(v))}{w(v)} \leq \frac{\text{Vol}(f(u))}{w(u)}.$$
- For each $u, v \in V(T)$, such that u is not an ancestor of v, and v is not an ancestor of u, $\text{int}(f(u)) \cap \text{int}(f(v)) = \emptyset$.

Lemma 3. *Let $\epsilon \in (0, 1/3)$, and let $d \geq 2$. Let $T=(V,E)$ be a rooted tree of depth t. Let $w : V \to \mathbb{R}_{\geq 0}$ be a function such that for any interval vertex $v \in V(T)$, with children $u_1, \ldots, u_k$, $w(v) \geq \sum_{i=1}^{k} w(u_i)$. Then, there exists a $1/\epsilon$-hierarchical hyperrectangular partition $f : V \to \mathcal{P}(\mathbb{R}^d)$ for (T, w) with ϵ-slack.*

Proof. We create a mapping f such that for each $u \in V$, $f(u)$ is a hyperrectangle. We start from a hypercube of volume $w(r)$, where r is the root of the the tree. We fix $f(x)$ to be this hypercube. Its rectangular aspect ratio is 1.

We show by induction how to construct f and prove that the rectangular aspect ratio of each $f(u)$ is at most $1/\epsilon$. This implies that the (standard) aspect ratio of each $f(u)$ is at most $\sqrt{d}/\epsilon$.

For each $f(u)$, we define $w'_v = \frac{\text{Vol}(f(u))}{w(u)} \cdot w(v)$ for each child v of u in T. Then we shrink $f(u)$ so that the volume of the shrunken hyperrectangle R is exactly equal to the sum of w'_v over the children v of u.

Whenever we want to subdivide a hyperrectangle R of rectangular aspect ratio at most $1/\epsilon$ among a subset S of at least two children of u, we do what follows. We split S with a cut which is perpendicular to the longest side of R. Let $s \in S$ be the child in S of the largest w'_s. There are two cases.

- If $w'_s / \sum_{v \in S} w'_v \leq 1-\epsilon$, then we can split S into two sets S_1 and S_2 each of weight which is at most an $1-\epsilon$ fraction of the total weight of S. Then we split R with a cut which is perpendicular to the longest cut, so that we create two hyperrectangles R_1 and R_2 of volume proportional to the total weight of S_1 and S_2, respectively. All sides but the longest are preserved in the new hyperrectangles, and the length of the initially longest side becomes an at least ϵ fraction of the original value. This implies that if the rectangular aspect ratio of R_1 or R_2 increases with respect to the ratio of R, then it cannot be greater than $1/\epsilon$.
- The second case is when $w'_s / \sum_{v \in S} w'_v > 1-\epsilon$, i.e., there is a very heavy element in S. In this case, we must be more careful to avoid assigning a bad hyperrectangle. We first split R into two hyperrectangles R_1 and R_2 with a cut perpendicular to the longest side, so that $\text{Vol}(R_1) = (1-\epsilon)\text{Vol}(R)$ and $\text{Vol}(R_2) = \epsilon \text{Vol}(R)$. The rectangular aspect ratio of both R_1 and R_2 is at most $1/\epsilon$. We set $f(s)$ to be R_1. This means that we assign to s a hyperrectangle of volume smaller by a factor of at most $1-\epsilon$ than what is implied by the weight of s. To the other elements we assign R_2 uniformly shrunken so that its volume equals $\sum_{x \in S \setminus \{s\}} w'_x$. The shrunken R_2 is a subset of the initial R_2. We proceed with it recursively, until S has only one element.

□

4. EMBEDDING ULTRAMETRICS INTO $\mathbb{R}^d$

In this section, we give an approximation algorithm for embedding ultrametrics into $\mathbb{R}^d$. Before we describe the algorithm, we define *α-hierarchical well separated trees* (HST), introduced by Bartal [5]. For some $\alpha > 1$, an α-HST is a rooted tree T, with all the leaves on the same level. For each vertex v there is an associated label $l(v) > 0$, such that for each child u of v in T, $l(v) = \alpha \cdot l(u)$. The metric space that corresponds to the HST T is defined on the leaves of T, and the distance between leaves x, y is equal to the label of the nearest common ancestor of x, and y in T.

Let $M = (X, D)$ be the given ultrametric. After scaling M, we can assume that the minimum distance is 1, and the diameter is Δ. It is known, and easy to see that for any

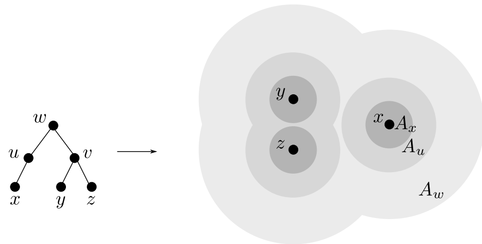

Figure 4: The sets A_v, for a non-contracting embedding of an HST.

$\alpha > 1$, M can be embedded into an α-HST, with distortion α (cf. [6]). Given M, we initially compute an embedding of M into a 2-HST T, with distortion 2. Let $M' = (X, D')$ be the metric space corresponding to T. Any embedding of M' into $\mathbb{R}^d$ with distortion c', is clearly also an embedding of M into $\mathbb{R}^d$ with distortion at most $c = O(c')$. It therefore suffices to embed of M' into $\mathbb{R}^d$.

4.1 A Lower Bound

We will now briefly describe a lower bound given in [2] on the optimal distortion. Consider a non-contracting embedding ϕ of M' into $\mathbb{R}^d$. For each $v \in V(T)$ we define a set $A_v \subset \mathbb{R}^d$ as follows. For a leaf v of T, let A_v be a ball of radius $1/2$ around $\phi(v)$ in $\mathbb{R}^d$. For a non-leaf vertex v, with children $u_1, \ldots, u_k$, let A_v be the Minkowski sum of $\bigcup_{i=1}^k A_{u_i}$ with a d-dimensional ball of radius $l(v)$. By the non-contraction of ϕ it follows that for each pair of vertices x, y that are on the same level of T, $\mathrm{int}(A_x) \cap \mathrm{int}(A_y) = \emptyset$. Therefore, by using the Brunn-Minkowski inequality we can derive a lower bound on $\mathrm{Vol}(A_v)$, for each $v \in V(T)$ (see Figure 4 for an example).

Following [2], we define a function $C : V(T) \to \mathbb{R}$, which (up-to scaling factors) corresponds to the volume of A_v. For any $r > 0$ let $V_d(r)$ be the volume of a d-dimensional ball or radius r, $V_d(r) = \frac{\pi^{d/2} r^d}{\Gamma(1+d/2)}$. Formally, if v is a leaf of T, we set $C(v) = V_d(1/2)$. Otherwise, for an internal vertex v, with children $u_1, \ldots, u_k$, we set

$$C(v) = \sum_{i=1}^{k} \left((C(u_i))^{1/d} + (V_d(l(v)/4))^{1/d} \right)^d$$

Intuitively, a large value of $C(v)$ implies that the volume of A_v should be large for some vertex v. This in turn can be translated via an isoperimetric argument to a lower bound on the distance between the images of two points with ancestor v. The above intuition is formalized in the following Lemma, that has been shown in [2]. For any $V > 0$ let $\rho_d(V)$ be the radius of a d-dimensional ball that has volume V, i.e. $\rho_d(V) = \left(\frac{V \cdot \Gamma(1+d/2)}{\pi^{d/2}} \right)^{1/d}$.

Lemma 4 ([2], Corollary 1). *Let v be some non-leaf vertex of T, and let ϕ be a non-contracting embedding of M into $\mathbb{R}^d$, under the ℓ_2 norm, with distortion c'. Then, $c' \geq \frac{\rho_d(C(v))}{l(v)} - 1$.*

4.2 The Algorithm

We are now ready to describe the embedding f of M' into $\mathbb{R}^d$. The intuition behind our algorithm is as follows. The lower bound given by Lemma 4 implies that an embedding is nearly-optimal if it results in sets A_v with small aspect ratio. Our approach, however, is essentially reversed. We first compute a hierarchical partition of $\mathbb{R}^d$ into sets with small aspect ratio. The sets in the lower level of the partition would roughly correspond to balls around the images of the points in our embedding. Therefore, given the hierarchical partition we will be able to easily obtain the embedding.

More precisely, the algorithm works as follows. Initially, we compute the values $C(v)$, for each vertex v of the HST T. Then, using Lemma 3, we compute a $(\log \Delta)$-hierarchical hyperrectangular partition g for (T, C) (i.e. with weight assignment $w(v) = C(v)$). We further define a mapping $g' : V(T) \to \mathcal{P}(\mathbb{R}^d)$ by slightly modifying g as follows. Starting from the root of T, we traverse all the vertices of T. When we visit a vertex u, and we shrink uniformly all the hyperrectangles of the vertices in the subtree rooted at u, by a factor of $1 - 1/\log \Delta$, with the center of the hyperrectangle of u being the fixed point in the transformation. Let $g' : V(T) \to \mathcal{P}(\mathbb{R}^d)$ be the resulting mapping. Observe that for each $v \in V(T)$, $\mathrm{Vol}(g'(v)) \geq (1 - 1/\log \Delta)^{\log \Delta} \mathrm{Vol}(g(v)) = \Omega(\mathrm{Vol}(g(v))$, and that $\lambda_{\mathrm{rect}}(g(v')) = \lambda_{\mathrm{rect}}(g(v))$. For each point $x \in X$, let v_x be the leaf of T corresponding to x. Having computed g', we simply set $f(x)$ to be the center of the hyperrectangle $g'(v_x)$. It remains to bound the distortion of f.

Lemma 5. *The expansion of f is $O(\log \Delta \cdot c')$.*

Proof. Consider points $x, y \in X'$, and let v_x, v_y, be the leafs of T that correspond to x, and y respectively. Let v be the nearest common ancestor of v_x, and v_y, in T. We have $D'(x, y) = l(v)$. By Lemma 3, it follows that in the partition g' computed by the algorithm, v is mapped to a hyperrectangle $g(v') \subset \mathbb{R}^d$, with $\lambda_{\mathrm{rect}}(g'(v)) \leq \log \Delta$. Note that $f(x) \in g'(v_x)$, $f(y) \in g'(v_y)$, and also $g'(v_x) \subseteq g'(v)$, $g'(v_y) \subseteq g'(v)$. Since $\mathrm{Vol}(g'(v)) \leq \mathrm{Vol}(g(v)) \leq C(v)$, we have $\|f(x) - f(y)\|_2 \leq \mathrm{diam}(g'(v)) \leq \mathrm{diam}(g(v)) \leq d \cdot \log \Delta \cdot (C(v))^{1/d}$. Therefore, by Lemma 4, we obtain that $\|f(x) - f(y)\|_2 = O(c' \cdot l(v) \cdot \log \Delta) = O(\log \Delta \cdot c' \cdot D'(x, y))$. □

Lemma 6. *The contraction of f is $O(\log^{O(1)} \Delta)$.*

Proof. Since the depth of T is $\log \Delta$, it follows that for each vertex $u \in V(T)$, $\mathrm{Vol}(g'(u)) = \Omega(\mathrm{Vol}(g(u))) = \Omega((1 - 1/\log \Delta)^{\log \Delta} C(u)) = \Omega(C(u))$. Consider points $x, y \in X'$, and let $v_x, v_y \in V(T)$ be the leafs of T corresponding to x, y respectively. Let v be the nearest common ancestor of v_x, and v_y in T. We will consider the following two cases for v:

Case 1: v is the parent of v_x, and v_y in T. Since the minimum distance in M' is 1, it follows that $D'(x, y) = 1$. By the construction, $f(x)$ is the center of $g'(v_x)$. Let t be the distance between $f(x)$, and $\partial g'(v_x)$. Since $\lambda_{\mathrm{rect}}(g'(v_x)) \leq \log \Delta$, we have

$$t \geq \frac{(\mathrm{Vol}(g'(v_x)))^{1/d}}{\log \Delta} = \frac{\Omega((C(v_x)^{1/d})}{\log \Delta} = \Omega(1/\log \Delta).$$

Thus, $\|f(x) - f(y)\|_2 \geq t = \Omega(D(x, y)/\log \Delta)$.

Case 2: v is not the parent of v_x, and v_y in T. Let u_x be the child of v, that lies on the path from v to v_x, in T. Let γ be the distance between x, and $\partial g'(u_x)$. By the construction of g' we have $\|f(x) - f(y)\|_2 \geq \gamma = \Omega((C(u_x))^{1/d}/\log^{O(1)} \Delta) = \Omega(l(u_x)/\log^{O(1)} \Delta) = \Omega(D(x, y)/\log^{O(1)} \Delta)$. □

Combining lemmas 6, and 5, we obtain the main result of the section.

THEOREM 1. *For any fixed $d \geq 2$, there exists a polynomial-time,* $\mathrm{polylog}(\Delta)$*-approximation algorithm, for the problem of embedding ultrametrics into $\mathbb{R}^d$ with minimum distortion.*

Acknowledgments

The authors wish to thank Roberto Tamassia for providing pointers to the applicability of circular partitions in visualization.

5. REFERENCES

[1] D. Avis and M. Deza. The cut cone, l_1 embedability, complexity and multicommodity flows. *Networks*, 21:595–617, 1991.

[2] M. Badoiu, J. Chuzhoy, P. Indyk, and A. Sidiropoulos. Embedding ultrametrics into low-dimensional spaces. In *Proceedings of the 22nd ACM Symposium on Computational Geometry*, 2006.

[3] M. Balzer and O. Deussen. Voronoi treemaps. In *INFOVIS*, page 7, 2005.

[4] M. Balzer, O. Deussen, and C. Lewerentz. Voronoi treemaps for the visualization of software metrics. In *SOFTVIS*, pages 165–172, 2005.

[5] Y. Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. *Annual Symposium on Foundations of Computer Science*, 1996.

[6] Y. Bartal and M. Mendel. Dimension reduction for ultrametrics. In *SODA '04: Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 664–665, Philadelphia, PA, USA, 2004. Society for Industrial and Applied Mathematics.

[7] B. B. Bederson, B. Shneiderman, and M. Wattenberg. Ordered and quantum treemaps: Making effective use of 2d space to display hierarchies. *ACM Trans. Graph.*, 21(4):833–854, 2002.

[8] T. Bladh, D. A. Carr, and M. Kljun. The effect of animated transitions on user navigation in 3d tree-maps. In *IV*, pages 297–305, 2005.

[9] T. Bladh, D. A. Carr, and J. Scholl. Extending tree-maps to three dimensions: A comparative study. In *APCHI*, pages 50–59, 2004.

[10] D. M. Bruls, C. Huizing, and J. J. van Wijk. Squarified treemaps. In *W. de Leeuw, R. van Liere (eds.), Data Visualization 2000, Proceedings of the joint Eurographics and IEEE TCVG Symposium on Visualization*, pages 33–42. Springer, 2000.

[11] M. Bădoiu, J. Chuzhoy, P. Indyk, and A. Sidiropoulos. Low-distortion embeddings of general metrics into the line. *Annual ACM Symposium on Theory of Computing*, 2005.

[12] M. Bădoiu, K. Dhamdhere, A. Gupta, Y. Rabinovich, H. Raecke, R. Ravi, and A. Sidiropoulos. Approximation algorithms for low-distortion embeddings into low-dimensional spaces. *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, 2005.

[13] M. Cary, A. Rudra, and A. Sabharwal. On the hardness of embeddings between two finite metrics. In *ICALP*, pages 1412–1423, 2005.

[14] J. Edmonds. Embedding into ℓ_∞^2 is easy embedding into ℓ_∞^3 is np-complete. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms*, 2007.

[15] A. Hall and C. H. Papadimitriou. Approximating the distortion. In *APPROX-RANDOM*, pages 111–122, 2005.

[16] L. Jin and D. C. Banks. Tennisviewer: A browser for competition trees. *IEEE Computer Graphics and Applications*, 17(4):63–65, 1997.

[17] B. Johnson and B. Shneiderman. Tree-maps: a space-filling approach to the visualization ofhierarchical information structures. In *Proc., IEEE Conference on Visualization*, pages 284–291, 1991.

[18] W.-A. Jungmeister and D. Turo. Adapting treemaps to stock portfolio visualization. Technical Report UMCP-CSD CS-TR-2996, College Park, Maryland 20742, U.S.A., 1992.

[19] C. Kenyon, Y. Rabani, and A. Sinclair. Low distortion maps between point sets. *Annual ACM Symposium on Theory of Computing*, 2004.

[20] J. R. Lee, A. Naor, and Y. Peres. Trees and markov convexity. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, 2006.

[21] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Proceedings of 35th Annual IEEE Symposium on Foundations of Computer Science*, pages 577–591, 1994.

[22] S. Malitz and J. Malitz. A bounded compactness theorem for l_1-embeddings of metric spaces in the plane. *In Discrete Comput. Geom.*, 8:373–385, 1992.

[23] J. Matoušek and A. Sidiropoulos. On the computational near-optimality of random projection. *Manuscript*, 2008.

[24] C. Papadimitriou and S. Safra. The complexity of low-distortion embeddings between point sets. *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 112–118, 2005.

[25] J. Rekimoto and M. Green. The information cube: Using transparency in 3d information visualization. In *Proceedings of the Third Annual Workshop on Information Technologies and Systems (WITS'93)*, pages 125–132, 1993.

[26] B. Shneiderman. Treemaps for space-constrained visualization of hierarchies. Available at `http://www.cs.umd.edu/hcil/treemap-history/index.shtml`.

[27] B. Shneiderman. Tree visualization with tree-maps: 2-d space-filling approach. *ACM Trans. Graph.*, 11(1):92–99, 1992.

[28] B. Shneiderman and M. Wattenberg. Ordered treemap layouts. In *IEEE Symposium on Information Visualization, INFOVIS '01*, pages 73–78, 2001.

[29] J. Stasko, R. Catrambone, M. Guzdial, and K. McDonald. An evaluation of space-filling information visualizations for depicting hierarchical structures. *International Journal of Human-Computer Studies*, 53(5):663–694, 2000.

[30] D. Turo and B. Johnson. Improving the visualization of hierarchies with treemaps: Design issues and experimentation. *IEEE Visualization*, pages 124–131, 1992.

[31] J. J. van Wijk and H. van de Wetering. Cushion

treemaps: Visualization of hierarchical information. In *INFOVIS*, pages 73–78, 1999.

[32] R. Vliegen, J. J. van Wijk, and E.-J. van der Linden. Visualizing business data with generalized treemaps. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):789–796, 2006.

[33] S. Wan. Blog treemap visualizer (`http://www.samuelwan.com/information/archives/000159.html`).

[34] M. Weskamp. Newsmap. http://marumushi.com/apps/newsmap/newsmap.cfm.

[35] K. Wetzel. Using Circular Treemaps to visualize disk usage. Available at `http://lip.sourceforge.net/ctreemap.html`.