

Low Distortion Delaunay Embedding of Trees in Hyperbolic Plane

Rik Sarkar

Institut Für Informatik
Freie Universität Berlin, Germany.
sarkar@inf.fu-berlin.de

Abstract. This paper considers the problem of embedding trees into the hyperbolic plane. We show that any tree can be realized as the Delaunay graph of its embedded vertices. Particularly, a weighted tree can be embedded such that the weight on each edge is realized as the hyperbolic distance between its embedded vertices. Thus the embedding preserves the metric information of the tree along with its topology. Further, the distance distortion between non adjacent vertices can be made arbitrarily small – less than a $(1 + \varepsilon)$ factor for any given ε . Existing results on low distortion of embedding discrete metrics into trees carry over to hyperbolic metric through this result. The Delaunay character implies useful properties such as guaranteed greedy routing and realization as minimum spanning trees.

1 Introduction

Embedding given data into a standard space lets us use properties of the target space as additional structure in the original dataset, and brings to front information that is hard to detect in the raw input. If the target space is Euclidean, that allows us to visualize and treat the data geometrically. In general, if it is a metric space, the properties of the metric can aid in understanding the original data and answering queries. This approach has been found relevant to a variety of subjects such as data visualization, network analysis, routing, localization, machine learning, statistics, biology and many others.

Trees are an important class of data structures. They occur commonly in natural scenarios, therefore associating trees with geometric spaces can be of benefit in many domains. Realizing trees as Delaunay graphs lets us combine the structural properties of trees with those of Delaunay graphs as well as those of the ambient space.

In this paper, we show that this can always be achieved in the hyperbolic spaces. In particular, we tackle the question for weighted trees and show that any given edge weights can be realized exactly (upto a choice of unit) in the delaunay embedding, while keeping the overall distortion arbitrarily low.

1.1 Related Work

Trees in Euclidean Plane. Monma and Suri [15] address embedding minimum spanning trees in the Euclidean space. They analyze questions of perturbations to the embedding and their effect on the topology of the MST. Relevant to us, they consider the problem of which trees can be realized as minimum spanning tree in the Euclidean plane, and show that any tree with maximum vertex degree of 5 or less admits an embedding as a minimum spanning tree. The topic of distortion of a tree embedded in euclidean metric is analyzed in [13].

Low Distortion Metric Embeddings. This is an extensively studied subject. The goal is to embed one metric into a target metric space with minimal distortion. A small distortion of the original metric means much of the original information is preserved. Embedding into Euclidean and L_p spaces has been considered in a number of papers [1, 17, 3, ?] and is known to have many applications [2, 14].

Embedding metrics into trees will be relevant to our results in the following chapters. Fakcharoenphol et al. [9] showed that for every n -point metric, there is a distribution over embeddings in trees such that the expected distortion is bounded by $O(\log n)$. Dhamdhere et al. [6] show that if the original metric is a weighted graph, and we consider embedding into its spanning trees, there is a distribution over them with expected distortion of $O(\log^2 n)$. Elkin et al. [7] show that every weighted graph contains a spanning tree into which it can be embedded such that edges incur at most $O(\log^2 n \log \log n)$ distortion on average.

The hyperbolic metric behaves like a tree in many respects. One way to model this is the concept of δ -hyperbolic metric[11]. Chepoi and co-authors[4] show that a graph that has an n -node δ -hyperbolic graph admits an approximating tree of additive error $O(\delta \log n)$.

Hyperbolic Embeddings Embedding graphs into hyperbolic spaces can provide advantages that we do not get in Euclidean spaces or trees. Kleinberg [12] considers the problem of *greedy routing* in wireless networks. This is the method, where a node routes a message by forwarding it to the neighbor that is nearest to the destination. The idea in [12] is to embed a spanning tree of the network into the hyperbolic plane such that this routing always works successfully on the tree, and thereby on the original network. Eppstein and Goodrich [8] show a related method that uses small sized coordinates. Cvetkovski et al. [5] extend [12] to incorporate dynamic insertion of edges. Papadopoulos et al. [16] show that hyperbolic embedding can naturally give rise to scale free networks and as before, can support greedy routing. Zeng et al. [20] embed the universal covering space of a network into the hyperbolic plane as a Riemann surface and show that this can be used to easily find paths of different homotopy types in the network.

1.2 Our Contributions

The choice of the target space determines the properties we can expect to obtain from the embedding. In this paper, we use the hyperbolic plane as the standard target space and show that this works remarkably well for all kinds of trees –

weighted as well as unweighted, and induces many desirable features. We present the important ideas and theorems in the main body of the paper. Details of proofs can be found in the full version of the paper online [18].

Delaunay embedding. Delaunay graphs are known to have many useful properties. If a graph can be treated as the Delaunay graph in some space, we can expect to leverage some of these properties for our purposes. With this motivation, we take up the question of realizing trees as Delaunay graphs, and show that the vertices of any tree can be embedded in the hyperbolic plane such that their Delaunay graph is the original tree. This leads to two immediate consequences: Realization of any tree as the minimum spanning tree of its vertices and guaranteed delivery by greedy routing. The first overcomes the constant degree bound of Euclidean case, the second implies the result of [12]. The embedding in [12] is a specific Delaunay embedding based on tiling of \mathbb{H} by ideal regular d -gons, where d is the max degree. We generalize this to not depend on a tiling, which gives greater flexibility needed to create the low distortion embeddings in sections 4 and 5.

Weighted Trees. Suppose the given tree has weights or lengths defined for its edges. This induces a metric and we refer to the input as a *metric tree*. We show that the vertices of any such metric tree has a Delaunay embedding in the hyperbolic plane, such that each Delaunay edge has a length that is the multiple of the edge’s prescribed length by a constant. As before, the properties implied by Delaunay embedding (MST and greedy routing) are retained.

Thus, the embedding not only preserves the topology of the tree, it also preserves the geometry – the metric information from the tree. The effect of the constant is negligible: since it is same globally, once it has been computed we can scale our unit of length to eliminate the factor. Our definition of distortion is oblivious to such global scaling.

Distortion of Metric Trees. While we preserve edge lengths precisely, the overall metric may not be preserved. The hyperbolic distance between v_i and v_j will be less than or equal to the path length between the two vertices in the tree. We show that this distortion can be kept arbitrarily low. Given any ε , we show it is possible to do the Delaunay embedding such that the distortion is less than $1 + \varepsilon$. Thus the tree can be seen as a hyperbolic spanner of the embedded vertices.

This implies distortion bounds for embedding of metrics in general. The results in [9, 6, 7] show low distortion bounds (probabilistic and average) for embedding arbitrary n point metrics into trees. It is therefore possible to embed arbitrary metrics into hyperbolic metric with the same bounds, using the tree embedding as an intermediate step.

2 Basics and notations

This section introduces some basic facts and notations about the geometry of hyperbolic plane that will be used later. We denote the hyperbolic plane by \mathbb{H} . For a more detailed exposition of these ideas, see [10]. We present the ideas in

terms of the plane \mathbb{H} , but they generalize directly to higher dimensions. We use $|\cdot|_{\mathbb{H}}$ to denote distances in hyperbolic metric. Correspondingly, we use $|\cdot|_T$ to denote distances in the metric of the input tree.

In hyperbolic geometry, the axioms of euclidean geometry are all true, except the parallel axiom, which is replaced by the *Hyperbolic Axiom*: *There is a line l and point P not on l , such that there are at least 2 different lines through P parallel to l .*

This leads to the general property that given a line, there are an infinite number of different lines parallel to it through an outside point. Parallel lines can be of different types. A pair of parallel lines are said to be *limiting parallel* if they approach each-other asymptotically without intersecting. Such a pair does not admit a common line perpendicular to both, and the distance between the two does not have a minimum. A pair of parallel lines are called *divergent parallel* if they move away from each-other in both directions. They have a common segment perpendicular to both. This segment achieves the minimum distance between the two lines.




Fig. 1. Line $m \perp l$ and rays $m_{-\alpha}$ and $m_{+\alpha}$ are limiting parallel to l . The angle α depends on the length of the segment PQ . Hyperbolic straight lines $m_{-\alpha}$ and $m_{+\alpha}$ look curved because our figure is in euclidean space.

Given a line l and a point P outside, there is always a point Q on l such that $PQ \perp l$. Through P , there are always rays $m_{+\alpha}$ and $m_{-\alpha}$ that are limiting parallel to l in the two directions. The angle between PQ and $m_{+\alpha}$ (or symmetrically, the angle between PQ and $m_{-\alpha}$) is called the angle of parallelism and represented by $\Pi(PQ)$. The Bolyai and Lobachevsky formula gives its value in terms of the length $|PQ|_{\mathbb{H}}$:

$$\tan \frac{\Pi(PQ)}{2} = e^{-|PQ|_{\mathbb{H}}/k}, \quad (1)$$

where k is a constant for the hyperbolic plane in consideration. Note that $\Pi(PQ)$ is always less than $\pi/2$ radians, since PQ cannot be perpendicular to a ray that is limiting parallel to l . Given l and P , the limiting parallel rays and the angle $\Pi(PQ)$ are unique. A ray that creates a larger angle with PQ will be divergent parallel, while a ray with smaller angle will intersect l .

The region bounded by rays $m_{-\alpha}$ and $m_{+\alpha}$ and containing the ray $m = \overrightarrow{PQ}$ will be important to our discussion. Let us refer to it as a closed cone $\bar{C}(\overrightarrow{PQ}, \alpha)$. We are particularly interested in the set $\bar{C}(\overrightarrow{PQ}, \alpha) \setminus \{P\}$, we call it a cone of angle α at P (or rooted at P) and denote it by $C(\overrightarrow{PQ}, \alpha)$.

The usual Euclidean axioms of betweenness, incidence and angles hold in the hyperbolic case. Therefore two cones $C(m, \alpha)$ and $C(n, \beta)$ at P do not intersect if and only if the angle between m and n is greater than $\alpha + \beta$. We say such pairs of cones are *disjoint*.

Observation 1 *Given any finite integer d we can always construct d mutually disjoint cones at P by taking d different rays and cones of suitably small angles around them.*

Our goal is to compute an embedding function $\Phi : V \rightarrow \mathbb{H}$, where $V = v_0, v_1, v_2, \dots$ is the set of vertices of the tree. To abbreviate notations, we write $\varphi_i = \Phi(v_i)$. We sometimes abbreviate our notations for cones as $C_{ij}^\alpha = C(\overrightarrow{\varphi_i \varphi_j}, \alpha)$.

We consider distortion over some set W of pairs of distinct vertices in question. For example, W can be the set of edges in the tree, or it can be the set of all pairs of vertices.

Distortion. We define the *contraction factor* over W as $\delta_c = \max_{(i,j) \in W} \frac{|v_i v_j|_{\mathcal{T}}}{|\varphi_i \varphi_j|_{\mathbb{H}}}$, and similarly the *expansion factor* as $\delta_e = \max_{(i,j) \in W} \frac{|\varphi_i \varphi_j|_{\mathbb{H}}}{|v_i v_j|_{\mathcal{T}}}$. The *distortion* is defined as $\delta = \delta_c \cdot \delta_e$.

Observe that if the embedding globally scales all distances for pairs in W by the same factor, then $\delta = 1$.

We consider Voronoi diagrams in \mathbb{H} . Given a finite set of vertices $v_0, v_1, \dots \subset \mathbb{H}$, the *Voronoi cell* of v_i , denoted $\mathcal{V}(v_i)$ is the set of points whose distance to v_i is not larger than the distance to v_j for any $j \neq i$.

Definition 1. Delaunay Graph. *Given a set of vertices in \mathbb{H} their Delaunay graph is one where a pair of vertices are neighbors if their Voronoi cells intersect.*

As with the euclidean case, this delaunay graph contains the MST:

Theorem 2. *For vertices in \mathbb{H} :*

1. *Any minimum spanning tree is contained in the Delaunay graph.*
2. *If the Delaunay graph of a set of vertices is a tree, then it is the minimum spanning tree.*

Delaunay Embedding of Graphs: Given a graph G , its Delaunay embedding in \mathbb{H} is an embedding of the vertices such that their Delaunay graph is G .

3 Delaunay Embedding of Trees

In this section, we describe the basic construct of embedding a tree as a Delaunay graph. The Delaunay graph automatically has the minimum spanning tree embedding and greedy embedding property. The basic idea is not new, and has been used in [12, 5, 19]. But we wish to make it more general and write in terms of cones rooted at the embedded vertices. This makes it easier to handle weighted trees and low distortion embedding in following sections.

Reorganizing equation 1, we have that given rays $m_{+\alpha}$ and $m_{-\alpha}$ through P , the distance $|PQ|_{\mathbb{H}}$ to a point Q on m so that the line $l \perp PQ$ at Q is limiting parallel to $m_{-\alpha}$ and $m_{+\alpha}$ is given by:

$$|PQ|_{\mathbb{H}} = -k \ln \left(\tan \frac{\alpha}{2} \right). \quad (2)$$

Since α is always less than $\pi/2$ radians, we have $\tan \frac{\alpha}{2} < 1$. Therefore $|PQ|_{\mathbb{H}}$ is positive and monotone decreasing in α . This means in particular, that if $|PQ|_{\mathbb{H}}$ is larger, then l is limiting parallel to the bounding lines of a smaller cone, and therefore divergent parallel to $m_{-\alpha}$ and $m_{+\alpha}$, and fully contained in $C(m, \alpha)$.

We construct a function Φ that embeds the vertices of a tree T into \mathbb{H} . Function Φ is designed such that T is the Delaunay graph of the embedded vertices.

Lemma 1. *Given two cones $C(\overrightarrow{PR}, \alpha)$ and $C(\overrightarrow{RP}, \beta)$, and $\gamma = \min(\alpha, \beta)$, if $|PR|_{\mathbb{H}} \geq -2k \ln \left(\tan \frac{\gamma}{2} \right)$ the perpendicular bisector of the segment PR lies in the intersection $C(\overrightarrow{PR}, \gamma) \cap C(\overrightarrow{RP}, \gamma)$ which is contained in the intersection $C(\overrightarrow{PR}, \alpha) \cap C(\overrightarrow{RP}, \beta)$.*

Proof. Suppose without loss of generality that $\gamma = \alpha$. Let us say that the perpendicular bisector l intersects PR at its midpoint Q .

Then $|PQ|_{\mathbb{H}} \geq -k \ln \left(\tan \frac{\alpha}{2} \right)$ and therefore $l \subset C(\overrightarrow{PR}, \alpha)$. Symmetrically, $|RQ|_{\mathbb{H}} \geq -k \ln \left(\tan \frac{\alpha}{2} \right) \implies l \subset C(\overrightarrow{RP}, \alpha)$. Since $\alpha \leq \beta$, we have $C(\overrightarrow{RP}, \alpha) \subseteq C(\overrightarrow{RP}, \beta)$. Therefore $l \subset C(\overrightarrow{PR}, \alpha) \cap C(\overrightarrow{RP}, \alpha)$ and $l \subset C(\overrightarrow{PR}, \alpha) \cap C(\overrightarrow{RP}, \beta)$. \square

This means in general we can consider the smaller of the two angles and consider the two cones to be of this same angle. See Figure 2. We say a segment PR is Delaunay for angle γ if it satisfies the conditions of the lemma for $\gamma = \alpha = \beta$.

Now we describe embedding an edge into a given cone. The goal is to embed in a way such that the Voronoi cell of one embedded vertex is completely contained inside a cone at the other vertex, and the edge is realized as the Delaunay edge between the embedded vertices, and the two cones can have the same angle. Again, see Figure 2.

Edge embedding function Φ . Given an edge $v_i v_j$ and an angle α , we select a cone $C(m, \alpha)$ at point P , and embed as follows. Vertex v_i is embedded at $\varphi_i = P$. We select a point R on the ray m such that $|PR|_{\mathbb{H}} \geq -2k \ln \left(\tan \frac{\alpha}{2} \right)$, and embed v_j as $\varphi_j = R$. We call this a Delaunay embedding of the edge for angle α .

This construction satisfies the conditions of lemma 1 with $\gamma = \alpha = \beta$, and therefore cones of angle α at φ_i and φ_j contain the perpendicular bisector of $\varphi_i \varphi_j$.

Now we extend the embedding to the entire tree. The idea is to embed such that each edge is Delaunay for a suitable angle, and the corresponding cones rooted at a vertex are disjoint.

Definition 2. *Tree Embedding Function Φ .* The vertices of a tree are embedded in a way that allows an assignment of an angle $\theta(e)$ to each edge e such that:

1. Cones determined by θ over all edges incident on any vertex are disjoint.




Fig. 2. Embedding an edge v_0v_1 . (a) Shows the symmetry in the embedding – it is the same from the point of view of both end points. The Voronoi edge is the dashed line. (b) shows a different view, which is more akin to the view from point P . This view gives the intuition that we can embed any number of cones of small enough angles at P , and by symmetry, same is true at point R .

2. For any edge v_iv_j , its embedding $\varphi_i\varphi_j$ is Delaunay for angle $\theta(v_iv_j)$.

The following algorithm describes the construction of such a Φ . Without loss of generality, we treat the input as a rooted tree.

Algorithm: Construction of Φ for T . Embed the root v_0 to an arbitrary point φ_0 . If v_0 has d children, they are Delaunay embedded individually into d disjoint cones at φ_0 . We embed all other vertices inductively as follows. Suppose v_j is a child of v_i , and has been embedded in a cone of angle α at φ_i . The children of v_j are Delaunay embedded in mutually disjoint cones that are also disjoint from the cone $C(\overrightarrow{\varphi_j\varphi_i}, \alpha)$.

By construction, this algorithm produces an embedding that satisfies definition 2. The construction works for infinite trees as well.

Lemma 2. If v_i is the parent of v_{i+1} , and Φ embeds the edge v_iv_{i+1} in the cone $C(\overrightarrow{\varphi_i\varphi_{i+1}}, \alpha)$ then

1. The Voronoi cells of all nodes in the subtree rooted at φ_{i+1} are contained in $C(\overrightarrow{\varphi_i\varphi_{i+1}}, \alpha)$.
2. The Voronoi cell of any node not in the subtree rooted at φ_{i+1} are contained in the cone $C(\overrightarrow{\varphi_{i+1}\varphi_i}, \alpha)$.

The hyperbolic plane contains many mutually parallel lines. Voronoi edges can be aligned to such lines, therefore, the Voronoi diagram consists of disjoint lines, carving out many disjoint half planes. The consequence is that Voronoi cells of nodes from different subtrees do not intersect. The only pairs of Voronoi cells that can intersect correspond to pairs of nodes that form edges of the tree. Therefore:

Theorem 3. The function Φ embeds T as a Delaunay graph in \mathbb{H} .

This directly implies that T is embedded as the minimum spanning tree of its vertices, by Theorem 2.

Greedy Embedding. Greedy routing is a well studied problem in wireless networks. Delaunay graphs are significant to this problem as well:




Fig. 3. Embedding a tree, by using cones of the type shown in Fig. 2 for each edge. 1, 2, 3... are the vertices of the tree. The rays form the cones. The thick edges are the edges of the tree.

Theorem 4. *A Delaunay embedding guarantees delivery by greedy routing.*

This implies the result of [12]. In fact the embedding of [12] is a specific instance of Delaunay embedding.

Edge Insertion. It is easy to adapt the algorithm to allow dynamic insertion into trees. At every vertex, we select the cones such that there is always enough angle left over to create more cones. This can always be done, since we can make the new cone small enough to not cover the entire available space. A new edge can be inserted into a new cone created in this space. This implies the edge insertion result of [5].

4 Delaunay Embedding of Metric Trees

Now we show a stronger embedding. Suppose $T = (V, E, w)$ is a weighted tree, where $w : E \rightarrow \mathbb{R}$ is the weight or length function on the edges. The goal is to realize the weight $w(v_i v_j)$ on each edge $v_i v_j$ as the length $|\varphi_i \varphi_j|_{\mathbb{H}}$ of the edge in the Delaunay embedding of the tree. For now, we are interested only in the distortion of individual edges of the tree, and show that in that sense there is an embedding with no distortion. The general distortion case is handled in the next section.

In the previous section we saw that each edge has to be embedded to a minimum length depending on the cone in which it is embedded. Based on this idea, we proceed as follows.

For each edge e , we compute a minimum length $L(e)$ needed to embed it as a Delaunay edge along with its neighboring edges. This gives a minimum scaling factor $\eta(e) = L(e)/w(e)$ for each edge. Next we compute $\eta_{\max} = \max_{e \in E} \eta(e)$ as the worst scaling needed at any edge. Then the tree is embedded such that each edge has length $\eta_{\max} \cdot w(e)$, that is, each edge is scaled by the same amount. Therefore, for unit length η_{\max} the Delaunay graph realizes the metric tree.

Algorithm: Computation of η_{\max} . Let us write $d(v)$ for the degree of a vertex v . The algorithm executes 5 passes over the tree:

1. Select for each vertex a maximum cone angle $\mu(v_i) < 2\pi/d(v_i)$.
2. Select for each edge $v_i v_j$ the maximum cone angle as $\alpha_{ij} = \min(\mu(v_i), \mu(v_j))$.
3. Compute for each edge the minimum required length $L(v_i v_j) = -2k \ln \left(\tan \frac{\alpha_{ij}}{2} \right)$.
4. Compute for each edge the minimum scaling factor $\eta(v_i v_j) = L(v_i v_j)/w(v_i v_j)$.

5. Compute the max value of η over all edges : $\eta_{\max} = \max_{v_i v_j \in E} \eta(v_i v_j)$.

Algorithm: Embedding function Φ_w of metric tree T . The function is a special case of Φ , and proceeds the same way. The cone construction and embedding steps are made more specific as follows:

1. At vertex φ_i with parent φ_h , we create $d(v_i) - 1$ disjoint cones of angle $\mu(v_i)$, disjoint from cone $C(\overrightarrow{\varphi_i \varphi_h}, \mu(v_i))$.
2. A child v_j of v_i is embedded in the cone $C(\overrightarrow{\varphi_i \varphi_j}, \mu(v_i))$ such that $|\varphi_i \varphi_j|_{\mathbb{H}} = \eta_{\max} \cdot w(v_i v_j)$.

For any edge $v_i v_j$ in E , the function Φ_w embeds it in cones $C(\overrightarrow{\varphi_i \varphi_j}, \mu(v_i))$ and $C(\overrightarrow{\varphi_j \varphi_i}, \mu(v_j))$ respectively.

The need to compute η_{\max} implies that algorithm works only on finite trees. But if constant bounds are known for the maximum vertex degree and the minimum edge length, then η_{\max} can be computed beforehand and the algorithm applied to infinite trees.

Lemma 3. $|\varphi_i \varphi_j|_{\mathbb{H}} \geq -2k \ln \left(\tan \frac{\alpha_{ij}}{2} \right)$.

Proof. By the definition of Φ_w , we have $|\varphi_i \varphi_j|_{\mathbb{H}} = \eta_{\max} \cdot w(v_i v_j) \geq \eta(v_i v_j) \cdot w(v_i v_j) \geq L(v_i v_j) \geq -2k \ln \left(\tan \frac{\alpha_{ij}}{2} \right)$ \square

Therefore, the embedding satisfies definition 2 and:

Theorem 5. *The embedding Φ_w is a Delaunay embedding of a metric tree with distortion 1 over the set of edges of the tree.*

Proof. That it is a Delaunay embedding follows from definition 2 and Theorem 3. Since all edges are scaled by a factor η_{\max} , we have $\delta_c = 1/\eta_{\max}$ and $\delta_e = \eta_{\max}$, and therefore distortion $\delta = \delta_c \cdot \delta_e = 1$ over the set of edges. \square

5 Delaunay Embedding with $(1 + \varepsilon)$ Distortion: Hyperbolic Spanner

In this section, we address the question of reducing the overall distortion to arbitrarily close to 1. Thus, in such an embedding, the tree acts as hyperbolic spanner of the embedded points. This embedding also implies that known results that show existence of *(tree, embedding)* pairs translate to existence of embeddings in the hyperbolic plane.

Definition 3. β separated cones. *Suppose cones $C(\overrightarrow{\varphi_j \varphi_i}, \alpha)$ and $C(\overrightarrow{\varphi_j \varphi_x}, \gamma)$ adjacent with the same root φ_j . Then the cones are β separated if the two cones are an angle 2β apart. That is, for arbitrary points $p \in C(\overrightarrow{\varphi_j \varphi_i}, \alpha)$ and $q \in C(\overrightarrow{\varphi_j \varphi_x}, \gamma)$, the angle $\angle p \varphi_j q > 2\beta$.*

An embedding is globally β separated if every pair of adjacent cones used for embedding are β separated.

Lemma 4. *If cones $C(\overrightarrow{\varphi_j \varphi_i}, \alpha)$ and $C(\overrightarrow{\varphi_j \varphi_x}, \gamma)$ are β separated and $\varphi_r \in C(\overrightarrow{\varphi_j \varphi_i}, \alpha)$ and $\varphi_s \in C(\overrightarrow{\varphi_j \varphi_x}, \gamma)$ then there is a constant ν depending only on β such that $|\varphi_r \varphi_j|_{\mathbb{H}} + |\varphi_s \varphi_j|_{\mathbb{H}} > |\varphi_r \varphi_s|_{\mathbb{H}} > |\varphi_r \varphi_j|_{\mathbb{H}} + |\varphi_s \varphi_j|_{\mathbb{H}} - \nu$.*

Proof. Suppose l is the line limiting parallel to both the rays $\overrightarrow{\varphi_j\varphi_r}$ and $\overrightarrow{\varphi_j\varphi_s}$. Let us refer to the point φ_j as P , and the perpendicular on l as PQ . Then we know that $|PQ|_{\mathbb{H}} \leq -k \ln\left(\tan\frac{\beta}{2}\right)$.

Now, $\varphi_j, \varphi_r, \varphi_s$ are on the same side of l . Therefore, $\varphi_r\varphi_s$ intersects PQ , say at W . This implies that $PW < PQ$.

By triangle inequalities: $|\varphi_r\varphi_j|_{\mathbb{H}} + |\varphi_s\varphi_j|_{\mathbb{H}} > |\varphi_r\varphi_s|_{\mathbb{H}} > |\varphi_r\varphi_j|_{\mathbb{H}} + |\varphi_s\varphi_j|_{\mathbb{H}} - 2|PW|_{\mathbb{H}} > |\varphi_r\varphi_j|_{\mathbb{H}} + |\varphi_s\varphi_j|_{\mathbb{H}} - 2|PQ|_{\mathbb{H}}$. Substituting $\nu = 2|PQ|_{\mathbb{H}} \leq -2k \ln\left(\tan\frac{\beta}{2}\right)$, we get the result. \square

If we can embed a tree such that neighboring edges, that is, edges incident on the same vertex, are always Delaunay in β separated cones for some suitably small beta, then for such embeddings, we have the following theorem:

Theorem 6. *If all edges of T are scaled by a constant factor $\tau \geq \eta_{\max}$ such that each edge is longer than $\nu \frac{(1+\varepsilon)}{\varepsilon}$ and the Delaunay embedding of T is β separated, then the distortion over all vertex pairs is bounded by $1 + \varepsilon$.*

Proof. Let $v_i, v_{i+1}, v_{i+2}, \dots, v_{i+p}$ be the path in the tree between the two end points. For any vertex v_{i+j} on the path, observe that φ_i is contained in a cone $C(\overrightarrow{\varphi_{i+j}\varphi_{i+j-1}}, \alpha)$ – a consequence of lemma 2. This cone, by construction, is β separated from cone $C(\overrightarrow{\varphi_{i+j}\varphi_{i+j+1}}, \gamma)$ containing φ_{i+j+1} .

Thus we have: $|\varphi_i\varphi_{i+2}|_{\mathbb{H}} \geq |\varphi_i\varphi_{i+1}|_{\mathbb{H}} + |\varphi_{i+1}\varphi_{i+2}|_{\mathbb{H}} - \nu \geq |\varphi_i\varphi_{i+1}|_{\mathbb{H}} + |\varphi_{i+1}\varphi_{i+2}|_{\mathbb{H}} - \frac{\varepsilon}{1+\varepsilon}|\varphi_{i+1}\varphi_{i+2}|_{\mathbb{H}}$, since each edge is longer than $\nu \frac{1+\varepsilon}{\varepsilon}$. Repeating, we have

$$\begin{aligned} |\varphi_i\varphi_{i+2}|_{\mathbb{H}} &\geq |\varphi_i\varphi_{i+1}|_{\mathbb{H}} + |\varphi_{i+1}\varphi_{i+2}|_{\mathbb{H}} - \frac{\varepsilon}{1+\varepsilon} \cdot |\varphi_{i+1}\varphi_{i+2}|_{\mathbb{H}} \\ |\varphi_i\varphi_{i+3}|_{\mathbb{H}} &\geq |\varphi_i\varphi_{i+2}|_{\mathbb{H}} + |\varphi_{i+2}\varphi_{i+3}|_{\mathbb{H}} - \frac{\varepsilon}{1+\varepsilon} \cdot |\varphi_{i+2}\varphi_{i+3}|_{\mathbb{H}} \\ |\varphi_i\varphi_{i+4}|_{\mathbb{H}} &\geq |\varphi_i\varphi_{i+3}|_{\mathbb{H}} + |\varphi_{i+3}\varphi_{i+4}|_{\mathbb{H}} - \frac{\varepsilon}{1+\varepsilon} \cdot |\varphi_{i+3}\varphi_{i+4}|_{\mathbb{H}} \\ &\dots \\ |\varphi_i\varphi_{i+p}|_{\mathbb{H}} &\geq |\varphi_i\varphi_{i+p-1}|_{\mathbb{H}} + |\varphi_{i+p-1}\varphi_{i+p}|_{\mathbb{H}} - \frac{\varepsilon}{1+\varepsilon} \cdot |\varphi_{i+p-1}\varphi_{i+p}|_{\mathbb{H}} \end{aligned}$$

Adding:

$$|\varphi_i\varphi_{i+p}|_{\mathbb{H}} \geq \sum_{x=i}^{i+p-1} |\varphi_x\varphi_{x+1}|_{\mathbb{H}} - \sum_{x=i}^{i+p-1} \frac{\varepsilon}{1+\varepsilon} |\varphi_x\varphi_{x+1}|_{\mathbb{H}}$$

$$\text{Therefore, } |\varphi_i\varphi_{i+p}|_{\mathbb{H}} \geq \frac{1}{1+\varepsilon} \cdot \sum_{x=i}^{i+p-1} |\varphi_x\varphi_{x+1}|_{\mathbb{H}}.$$

Since τ is fixed for all edges, $|\varphi_i\varphi_{i+p}|_{\mathbb{H}} \geq \frac{1}{1+\varepsilon} \cdot \tau |v_i v_{i+p}|_T$.

All edges are assumed to be scaled by a factor $\tau \geq \eta_{\max}$ such that they are longer than $\nu \frac{(1+\varepsilon)}{\varepsilon}$. We can therefore assume without loss of generality that $\tau > 1$. Now, for vertices v_i and v_j , $|v_i v_j|_{\mathbb{H}} \leq \tau |v_i v_j|_T$. That is, hyperbolic distance is at most τ times the tree distance.

Thus the expansion factor δ_e is dominated by the edges of the tree satisfying equality, and $\delta_e = \tau$. The contraction factor is dominated by the pair with

maximum distortion: $\delta_c \leq (1 + \varepsilon)/\tau$. Therefore, distortion $\delta = \delta_c \cdot \delta_e \leq (1 + \varepsilon)$. \square

Based on the theorem, the algorithm for $1 + \varepsilon$ distortion embedding becomes simple:

Algorithm: $(1 + \varepsilon)$ *Distortion embedding*. Suppose d is the maximum degree of any node. Then a low distortion embedding algorithm follows in these steps:

1. Compute a cone separation angle $\beta < \pi/d$, and an angle for cones $\alpha = 2\pi/d - 2\beta$. Set $\nu = -2k \ln(\tan \frac{\beta}{2})$.
2. Compute η_{\max} as in the previous section.
3. Select $\tau > \eta_{\max}$ such that all edges are longer than $\nu \frac{1+\varepsilon}{\varepsilon}$.
4. Embed edges as before, but into β separated cones, and edges scaled by a factor τ .

The bounds of embedding arbitrary metrics into trees with low distortion extend to the hyperbolic plane via this result. From [9] we have that for any metric, there is a distribution over embeddings in the hyperbolic plane with expected distortion $O(\log n)$. Similarly [7] implies that for every graph there is an embedding into hyperbolic plane such that the average distortion of edges is bounded by $O(\log^2 n \log \log n)$. As before, the algorithm applies to infinite trees with known degree and length bounds.

Embedding into Euclidean Models. Our presentation is in abstract hyperbolic geometry, and local for each vertex. Computationally, sometimes it is useful to have an embedding in one of the Euclidean models of hyperbolic metric, such as the Poincare Disk or the Half Plane model. Translating our algorithms to one of these models is fairly simple, based on Möbius transformations applied to the individual edges. We omit the detailed description for lack of space. See [12] for a distributed algorithm with such transforms.

6 Conclusion

We presented a method to embed trees into hyperbolic plane that simultaneously has several desirable properties. It is a Delaunay realization of the tree, preserves edge lengths exactly, and distance between non-neighbors is distorted by at most $1 + \varepsilon$. This suggests hyperbolic plane as a useful general target to investigate embedding questions. It remains to be seen what additional properties can be obtained in \mathbb{H} , for the more general embedding questions.

Acknowledgement. Thanks to Jie Gao and the anonymous reviewers for many helpful comments on the original draft of the article. This work is funded by the German Research Foundation (DFG) through the research training group Methods for Discrete Structures (GRK 1408).

References

1. I. Abraham, Y. Bartal, and O. Neimany. Advances in metric embedding theory. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, STOC '06, pages 271–286, 2006.

2. Y. Aumann and Y. Rabani. An $o(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM J. Comput.*, 27:291–301, 1998.
3. J. Bourgain. On lipschitz embedding of finite metric spaces in hilbert space. *Israel J. Math.*, 52:46–52, 1985.
4. V. Chepoi, F. Dragan, B. Estellon, M. Habib, Y. Vaxès, and Y. Xiang. Additive spanners and distance and routing labeling schemes for hyperbolic graphs. *Algorithmica*, pages 1–20, 2010. 10.1007/s00453-010-9478-x.
5. A. Cvetkovski and M. Crovella. Hyperbolic embedding and routing for dynamic graphs. In *Proceedings of Infocom 2009*, Apr. 2009.
6. K. Dhamdhere, A. Gupta, and H. Räcke. Improved embeddings of graph metrics into random trees. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, SODA '06, pages 61–69, 2006.
7. M. Elkin, Y. Emek, D. A. Spielman, and S.-H. Teng. Lower-stretch spanning trees. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, STOC '05, pages 494–503, 2005.
8. D. Eppstein and M. T. Goodrich. Graph drawing. chapter Succinct Greedy Graph Drawing in the Hyperbolic Plane, pages 14–25. Springer-Verlag, 2009.
9. J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 448–455, 2003.
10. M. J. Greenberg. *Euclidean and Non-Euclidean Geometries*. W.H. Freeman, 1993.
11. M. Gromov. Essays in group theory. chapter Hyperbolic groups, pages 75–263. Springer, New York, 1987.
12. R. Kleinberg. Geographic routing using hyperbolic space. In *Proceedings of the 26th Conference of the IEEE Communications Society (INFOCOM'07)*, pages 1902–1909, 2007.
13. J. Lee, A. Naor, and Y. Peres. Trees and markov convexity. *Geometric And Functional Analysis*, 18:1609–1659, 2009. 10.1007/s00039-008-0689-0.
14. T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J. ACM*, 46:787–832, 1999.
15. C. Monma and S. Suri. Transitions in geometric minimum spanning trees (extended abstract). In *Proceedings of the seventh annual symposium on Computational geometry*, pages 239–249, 1991.
16. F. Papadopoulos, D. Krioukov, M. Boguñá, and A. Vahdat. Greedy forwarding in dynamic scale-free networks embedded in hyperbolic metric spaces. In *Proceedings of the 29th conference on Information communications*, INFOCOM'10, pages 2973–2981, 2010.
17. S. Rao. Small distortion and volume preserving embeddings for planar and euclidean metrics. In *Proceedings of the fifteenth annual symposium on Computational geometry*, SCG '99, pages 300–306, 1999.
18. R. Sarkar. Low distortion delaunay embedding of trees in hyperbolic plane. <http://page.inf.fu-berlin.de/sarkar/papers/HyperbolicDelaunayFull.pdf>.
19. T. Tanuma, H. Imai, and S. Moriyama. Revisiting hyperbolic voronoi diagrams from theoretical, applied and generalized viewpoints. *International Symposium on Voronoi Diagrams in Science and Engineering*, 0:23–32, 2010.
20. W. Zeng, R. Sarkar, F. Luo, X. D. Gu, and J. Gao. Resilient routing for sensor networks using hyperbolic embedding of universal covering space. In *Proc. of the 29th Annual IEEE Conference on Computer Communications (INFOCOM'10)*, April 2010.

A Appendix for Section 2: Basics

Lemma 5. *The Voronoi edge between vertices at P and R lies in the perpendicular bisector of PR .*

Proof. Consider point W on the Voronoi edge between P and R , and Q is the mid point of PR . Then by the *SSS* congruency of $\overline{PQ} = \overline{RQ}$, $\overline{QW} = \overline{QW}$, $\overline{WP} = \overline{WR}$ we have $\angle PQW = \angle RQW = \pi/2$. Thus any point on the Voronoi edge is on the perpendicular bisector. \square

Let us denote by $D_P(r)$ the disk centered at P and of radius r .

Lemma 6. *If Q is the mid point of segment PR , and interior of $D_Q(|QP|_{\mathbb{H}})$ is empty, then PR is a Delaunay edge.*

Proof. Since $D_Q(|QP|_{\mathbb{H}})$ is empty, distance of Q from any vertex is at least $|QP|_{\mathbb{H}}$. Both P and R are at a distance exactly $|QP|_{\mathbb{H}}$ from Q , therefore the closure of their Voronoi cells intersect at Q . Thus, PR is a Delaunay edge \square

Lemma 7. *If PR is an edge in the minimum spanning tree T of a set of vertices in \mathbb{H} , then the open disk with the segment PR as a diameter is empty of vertices.*

Proof. Suppose to the contrary that vertex Z is in the open disk. Since PR is the diameter of the disk, we have $|PZ|_{\mathbb{H}} < |PR|_{\mathbb{H}}$ and $|RZ|_{\mathbb{H}} < |PR|_{\mathbb{H}}$. Now consider T after removing PR . Let us say T_P and T_R are the two components that contain P and R respectively. Without loss of generality we can assume that Z belongs to T_P . Then inserting the edge ZR creates a spanning tree whose weight is smaller than that of T by $|PR|_{\mathbb{H}} - |RZ|_{\mathbb{H}}$, contradicting that T is a minimum spanning tree. \square

Proof (of Theorem 2: Delaunay contains MST). Follows from the two lemmas above. Given this, it follows that if the Delaunay graph is a tree, it must be the MST. \square

B Appendix for Section 3: Delaunay Embedding

A line l separates two open half-spaces. If a half space contains a point P , let us denote that by $l|_P$, and the half space not containing P by $l|_{\overline{P}}$.

Lemma 8. *If PR is Delaunay for angle α , any point outside $C(\overrightarrow{PR}, \alpha)$ is nearer to $P = \varphi_0$ than to $R = \varphi_1$, and any point outside $C(\overrightarrow{RP}, \alpha)$ is nearer to R than to P .*

Proof. Say Q is the mid point of PR . Then we know that $l \perp PR$ at Q is contained in $C(\overrightarrow{PR}, \alpha)$. The Voronoi cell of φ_0 is the half space $l|_P$. The region outside $C(\overrightarrow{PR}, \alpha)$ is contained in this half space, therefore any point outside $C(m, \alpha)$ is nearer to φ_0 . The other statement follows symmetrically. \square

Corollary 1. *The Voronoi cell of φ_1 is contained in $C(\overrightarrow{\varphi_0\varphi_1}, \alpha)$, and the Voronoi cell of φ_0 is contained in $C(\overrightarrow{\varphi_1\varphi_0}, \alpha)$.*

Proof. Follows directly from lemma 8. \square

Corollary 2. *Any point outside the cone $C(\overrightarrow{\varphi_1\varphi_0}, \alpha)$ is contained in the cone $C(\overrightarrow{\varphi_0\varphi_1}, \alpha)$, and vice versa.*

Lemma 9. *The Voronoi cell of φ_j is contained in a cone $C(\overrightarrow{\varphi_i\varphi_j}, \beta) \subset C(\overrightarrow{\varphi_h\varphi_i}, \alpha)$ where v_h is the parent of v_i , and v_i is the parent of v_j , in the rooted tree T .*

Proof. By the construction of Φ , the point φ_j is in a cone $C(\overrightarrow{\varphi_i\varphi_j}, \beta)$. We know by corollary 1 that the Voronoi cell $\mathcal{V}(\varphi_j) \subset C(\overrightarrow{\varphi_i\varphi_j}, \beta)$.

By construction, $C(\overrightarrow{\varphi_i\varphi_j}, \beta)$ is outside the cone $C(\overrightarrow{\varphi_i\varphi_h}, \alpha)$, and therefore, by corollary 2, $C(\overrightarrow{\varphi_i\varphi_j}, \beta) \subset C(\overrightarrow{\varphi_h\varphi_i}, \alpha)$. Therefore, $\mathcal{V}(\varphi_j) \subset C(\overrightarrow{\varphi_i\varphi_j}, \beta) \subset C(\overrightarrow{\varphi_h\varphi_i}, \alpha)$. \square

Proof (of lemma 2).

Claim 1. Suppose $v_i, v_{i+1}, v_{i+2} \dots$ is a descending path. By lemma 9, $\mathcal{V}(\varphi_{i+2}) \subset C(\overrightarrow{\varphi_{i+1}\varphi_{i+2}}, \alpha_2) \subset C(\overrightarrow{\varphi_i\varphi_{i+1}}, \alpha_1)$. Similarly, $\mathcal{V}(\varphi_{i+3}) \subset C(\overrightarrow{\varphi_{i+2}\varphi_{i+3}}, \alpha_3) \subset C(\overrightarrow{\varphi_{i+1}\varphi_{i+2}}, \alpha_2)$. For any v_{i+w} on the descending path, $\mathcal{V}(\varphi_{i+w}) \subset C(\overrightarrow{\varphi_{i+w-1}\varphi_{i+w}}, \alpha_w) \subset C(\overrightarrow{\varphi_{i+w-2}\varphi_{i+w-1}}, \alpha_{w-1}) \subset \dots \subset C(\overrightarrow{\varphi_i\varphi_{i+1}}, \alpha_1)$.

Claim 2. Suppose v_g is an ancestor of v_i and $v_a, v_b, v_c \dots$ are the children of v_g . Then the Voronoi cells of nodes in subtrees rooted at $v_a, v_b, v_c \dots$ are contained in disjoint cones at v_g . Suppose without loss of generality that v_a is on the path from v_g to v_i and the cone C_a at v_g contains $\mathcal{V}(\varphi_i)$. Then the Voronoi cells of nodes in subtree v_b, v_c, \dots are contained in disjoint cones $C_b, C_c \dots$ and are disjoint from C_a .

Now suppose v_0, v_1, \dots, v_{i+1} is the path from the root to v_i and v_d is a node not on this path and not in the subtree rooted at v_{i+1} , and v_h is the least common ancestor of v_{i+1} and v_d . Then there are disjoint cones C_{i+1} and C_d at φ_h containing Voronoi cells $\mathcal{V}(\varphi_{i+1})$ and $\mathcal{V}(\varphi_d)$. In particular implying that $\mathcal{V}(\varphi_d)$ lies outside the cone $C(\overrightarrow{\varphi_h\varphi_{h+1}}, \alpha_{h+1})$ containing the subtree rooted at v_{h+1} .

This implies that $\mathcal{V}(\varphi_d)$ lies outside the cone $C(\overrightarrow{\varphi_i\varphi_{i+1}}, \alpha)$ containing the subtree rooted at φ_{i+1} . Therefore, $\mathcal{V}(\varphi_d)$ lies in the cone $C(\overrightarrow{\varphi_{i+1}\varphi_i}, \alpha)$, by corollary 2. \square

Proof (of theorem 3: Φ creates a Delaunay embedding). We need to show that the Delaunay graph D_Φ under embedding Φ consists only of edges of T or equivalently that $\mathcal{V}(\varphi_i)$ and $\mathcal{V}(\varphi_j)$ intersect if and only if $v_i v_j$ is an edge in T .

Necessity. Suppose to the contrary that $\mathcal{V}(\varphi_i)$ and $\mathcal{V}(\varphi_j)$ intersect but $v_i v_j$ is not an edge in T . Consider the unique path p_{ij} from v_i to v_j in T . This path must contain a third node, say v_h , that is a parent of either v_i or v_j , but not

both. Without loss of generality, let us suppose that v_h is the parent of v_i . This means that v_j is not in the subtree rooted at v_i . Now we consider two possibilities:

Case 1: v_j is not in the subtree rooted at v_h . Then, if v_{h-1} is the parent of v_h , there are cones $C(\overrightarrow{\varphi_h \varphi_{h-1}}, \beta)$ and $C(\overrightarrow{\varphi_h \varphi_i}, \alpha)$ containing $\mathcal{V}(\varphi_j)$ and $\mathcal{V}(\varphi_i)$ respectively (by lemma 2). these cones are disjoint by construction, therefore $\mathcal{V}(\varphi_j)$ and $\mathcal{V}(\varphi_i)$ are also disjoint – contradicting the hypothesis.

Case 2: v_j is in the subtree rooted at v_h . Then there are disjoint cones C_i and C_j at φ_h containing the two vornoi cells and leading to contradiction as before.

Sufficiency. We have shown above that edges not in T are also not in D_Φ . For D_Φ to be connected, all edge of T must be in D_Φ , since T is tree on the same set of vertices.

Now suppose for contradiction that $v_i v_j$ is an edge in T , but $\mathcal{V}(\varphi_i)$ and $\mathcal{V}(\varphi_j)$ do not intersect. This means that $v_i v_j$ is not present in D_Φ and D_Φ is not connected. This contradicts the fact that the Delaunay graph is connected. \square

Theorem 7. *The function Φ embeds T in \mathbb{H} as the minimum spanning tree of its vertices.*

Proof. Follows directly from Theorem 2 and Theorem 3. \square

Proof (of Theorem 4: Delaunay graph guarantees greedy routing).

Suppose we wish to route from embedded vertex s to t . Let m be the line segment from s to t . If st is a Delaunay edge, this immediately solves the problem. If not, then consider the Voronoi diagram of the points, and say l' is the first (nearest to s) Voronoi edge intersected by m , and sv is the Delaunay edge dual to it. Suppose $l \supset l'$ is the perpendicular bisector of segment sv . Then $t \in l|_v$, that is, t is on the same side of l as v . Therefore, $|vt|_{\mathbb{H}} < |st|_{\mathbb{H}}$.

Thus, s has a neighbor v nearer to t , and can forward the message to v . Repeating this argument inductively, we have guaranteed greedy routing. \square