Karlsruhe Institute of Technology

Hyperbolic Sphericity

Bachelor's Thesis of

Lennart Großkreutz

At the Department of Informatics
Institute of Theoretical Informatics (ITI)

Reviewer:	T.T.-Prof. Dr. Thomas Bläsius
Second reviewer:	Dr. rer. nat. Torsten Ueckerdt
Advisor:	M.Sc. Jean-Pierre von der Heydt

15/05/2024 – 16/09/2024

I declare that I have developed and written the enclosed thesis completely by myself. I have not used any other than the aids that I have mentioned. I have marked all parts of the thesis that I have included from referenced literature, either in their original wording or paraphrasing their contents. I have followed the by-laws to implement scientific integrity at KIT.

Karlsruhe, 16/09/2024

. . . . . .

(Lennart Großkreutz)

Abstract

The Euclidean sphericity of a graph G denotes the smallest integer $n \in \mathbb{N}$ such that G can be represented as an intersection graph of equal-sized balls of radius r in n-dimensional Euclidean space. Modern research strongly suggests that extending this type of intersection graphs to hyperbolic space is of great relevance, with applications in network science and machine learning. However, little is known about the corresponding graph classes so far.

To establish a foundation in this new research area, we introduce the concept of *hyperbolic sphericity*. This graph invariant resembles its Euclidean counterpart, with an important addition: different choices of the parameter r generally result in different graph classes in the hyperbolic context. We prove that hyperbolic sphericity is bounded from above by Euclidean sphericity, if r is chosen small. Conversely, large choices of r might increase hyperbolic sphericity, but by at most one. In other cases (e.g. trees), large choices of r can lead to a particularly low hyperbolic sphericity in comparison to the Euclidean one.

Zusammenfassung

Die euklidische Sphärizität eines Graphen G bezeichnet die kleinste natürliche Zahl $n \in \mathbb{N}$, sodass G als Schnittgraph von Bällen mit einheitlichem Radius r im n-dimensionalen euklidischen Raum dargestellt werden kann. Moderne Forschung deutet darauf hin, dass solche Schnittgraphen im hyperbolischen Raum von großer Relevanz sind, mit Anwendungen in der Netzwerkwissenschaft und dem maschinellen Lernen. Jedoch ist bisher kaum etwas über die entsprechenden Graphklassen bekannt.

Um eine Grundlage in diesem neuen Forschungsbereich zu schaffen, führen wir das Konzept der *hyperbolischen Sphärizität* ein. Diese Graphinvariante ähnelt ihrem euklidischen Gegenstück, mit einem wichtigen Zusatz: Verschiedene Werte des Parameters r führen im Allgemeinen zu verschieden Graphklassen im hyperbolischen Kontext. Wir beweisen, dass die hyperbolische Sphärizität von oben durch die euklidische Sphärizität beschränkt ist, wenn r klein gewählt ist. Im Gegensatz dazu können große Werte für r dazu führen, dass die hyperbolische Sphärizität wächst, aber höchstens um eins. In anderen Fällen (z.B. bei Bäumen) können große Werte für r dazu führen, dass die hyperbolische Sphärizität besonders klein ist im Vergleich zur Euklidischen.

Contents

1 Introduction

Often, real-world networks have a geometric aspect: for example, in a wireless communication network, two nodes are more likely to be able to directly communicate if they are positioned closely. Conversely, we may even be certain that some nodes cannot communicate, for example when they are kilometers apart and may only use a Bluetooth connection. Hierarchical real-world networks like the internet, often called a *network of networks*, might also display geometric patterns: clusters of end-devices and home-routers mostly fall into the above category of a wireless communication network. However, there are also a few nodes ensuring the connectivity of the entire network, like internet service providers. While they have many more communication partners than end-devices, spatial proximity still influences their direct connections, although on a much larger scale: for example countrywide.

These observations lead to an interest in graph classes with an underlying geometry. In this context, Euclidean unit disk graphs play an important role. This well-studied class consists of intersection graphs of equal-sized disks of radius r in the Euclidean plane. Any arrangement of such equal-sized disks is called a *representation* of the corresponding unit disk graph. These graphs are naturally suited to model communication networks in which every node has roughly the same transmission range [CCJ90]. Not only does a unit disk graph representation of a network help to better understand its characteristics, it may also yield faster algorithms for certain problems. For example, finding a maximum-sized clique is NP-complete on general graphs, but can be accomplished in polynomial time on unit disk graphs (assuming a corresponding representation is provided) [CCJ90]. Unfortunately, unit disk graphs quickly reach their limits when it comes to modeling complex real-world networks. In particular, they fail to represent the hierarchical structures that are often inherent in these networks.

The class of unit disk graphs can be generalized or adapted in two main ways: first of all, one may consider n-dimensional Euclidean space for any $n \in \mathbb{N}$ instead of just the Euclidean plane, leading to more flexibility. In this setting, the term *unit ball graphs* (UBGs) naturally arises to describe the resulting graph class depending on n. While higher dimensions allow to capture more graphs, geometric aspects are clearly more pertinent in lower dimensions. Similarly, it can easily be shown that any UBG in n dimensions is also a UBG in higher dimensions. For these reasons, we want to determine the *lowest* dimension $n \in \mathbb{N}$ that still permits a UBG representation of G. This integer n is called the *Euclidean sphericity* of the graph G. But does such an integer even exist for every finite graph? Yes, due to a result of Maehara, showing that each graph $G = (V, E)$ is always a UBG in $|V|$-dimensional Euclidean space [Mae84b].

While these insights are surely interesting from a graph-theoretical point of view, they are still unsatisfying: real-world networks like the internet appear to have geometric properties, but a high-dimensional representation as a Euclidean UBG fails to properly reflect such properties. At this point, it is worth asking whether Euclidean space itself might have some undesired characteristics that simply do not go well with the hierarchical aspects of networks like the internet. Indeed, Euclidean space is known to have *polynomial growth*: the volume of a ball of radius r is polynomial in r, no matter the number of dimensions. However, hierarchical networks often display exponential growth. To see this, consider regular trees as the arguably simplest form of a hierarchical network. We know that the number of tree nodes increases exponentially in terms of the tree depth.

This leads to the second main approach: keeping the concept of unit disk graphs, but transferring it to non-Euclidean spaces. In order not to run into the previous problem, we are looking for a space that shows exponential growth instead of polynomial growth. Metric spaces corresponding to hyperbolic geometry are known to have this property. As a bonus, these spaces fulfill almost all axioms of the Euclidean plane. In the hyperbolic context, we speak of *hyperbolic unit disk graphs*. Research in this area is fairly recent, mostly coming from the network community. Results show that randomly generated hyperbolic unit disk graphs indeed capture the main characteristics of complex real-world networks like the internet, most importantly a heterogeneous degree distribution and strong clustering [Kri+10 | BFM15].

Our contribution is to unify the two approaches by considering unit ball graphs in n-dimensional hyperbolic space, obtained by generalizing hyperbolic plane geometry. It combines the advantages of hyperbolic geometry with the dimensional generality that is needed to grasp the broad spectrum of real-world networks, often not perfectly represented in two dimensions. This is especially important when trying to understand and analyze networks without a strong inherent spatial aspect, like social networks.

We are the first to define and explore the *hyperbolic sphericity* of graphs. This new graph invariant resembles its Euclidean counterpart, with an important difference: in the Euclidean case, we do not have to pay much attention to the choice of the radius r of the equal-sized balls, since every representation can be scaled such that it may use any other radius r'. This does not apply to hyperbolic space, due to its exponential growth. So r is an influential parameter in terms of hyperbolic sphericity. We call it the *threshold radius*. Despite this new parameter, we are able to show that each finite graph has finite hyperbolic sphericity, for every choice of the threshold radius r. Following the intuition that polynomial growth and exponential growth are hard to tell apart when zooming in closely, we also prove that a UBG representation in n-dimensional Euclidean space can be reconstructed in n-dimensional hyperbolic space, if r is chosen small enough. This is promising because it means that hyperbolic sphericity is bounded from above by Euclidean sphericity for small r. How do things behave when we choose great threshold radii r, making full use of hyperbolic space's exponential growth? Not surprisingly, not every graph benefits from this exponential growth. We provide examples of non-hierarchical graphs whose hyperbolic sphericity increases for great r. More surprisingly, this potential increase is very small: we show that a graph having hyperbolic sphericity of $n \in \mathbb{N}$ for some threshold radius $r > 0$ is guaranteed to have hyperbolic sphericity of at most $(n+1)$ for every $r' > r$. In particular, hyperbolic sphericity never exceeds Euclidean sphericity by more than one. Conversely, for the right graphs, great threshold radii r are where hyperbolic space displays its true potential: we prove that every tree T is a hyperbolic unit disk graph, i.e. has hyperbolic sphericity of one or two, if $r = r(T)$ is chosen great enough. Meanwhile, Euclidean sphericity of trees becomes arbitrarily great by increasing tree depth and vertex degrees. To paint the big picture, hyperbolic space can be seen as an enhancement of Euclidean space in terms of sphericity. Due to the flexibility introduced by the choice of the threshold radius r, it supports both Euclidean and hierarchical structures.

Some other contributions are worth mentioning: we show that each triangle-free Euclidean or hyperbolic unit disk graph is planar. Concerning complexity theory, we demonstrate that hyperbolic sphericity is computable using a polynomial amount of space. We also investigate sphericity in spherical space, since the latter is closely related to hyperbolic space. However, in contrast to hyperbolic space, spherical space does not bring major advantages over Euclidean space in terms of sphericity.

1.1 Related Work

This thesis is situated within the research on intersection graphs and graph embeddings. Concerning intersection graphs, specific geometric graph classes like unit disk graphs or intersection graphs of axis-aligned rectangles have been studied extensively. For the most part, research on such graph classes can be divided into two domains: the first focuses on structural properties of the graph class, evaluating how the class is related to other common graph classes etc. In terms of unit disk graphs, results are somewhat limited by the fact that their corresponding recognition problem is NP-hard [BK98]. Similar structural research deals with the comparison of unit ball graphs and unit cube graphs in the Euclidean context [Fis83]. The second research domain builds upon these structural properties by constructing efficient algorithms for common graph problems, tailored towards the specific class. For example, the problem of finding a maximum-sized clique is NP-complete for general graphs, but can be solved in polynomial time for unit disk graphs as long as a geometric representation is provided [CCJ90].

Recently, similar research gains traction in the context of hyperbolic unit disk graphs. Concerning structural properties, Bläsius et al. proved that their recognition is $\exists\mathbb{R}$-complete, implying NP-hardness [BBDJ23]. They also examined bounds on balanced separators of hyperbolic unit disk graphs, leading to an efficient algorithm on the Independent Set problem [Blä+24]. More algorithmic results on hyperbolic unit disk graphs come from Kisfaludi-Bak. He presented polynomial time algorithms on the problems Hamiltonian Cycle and 3-Coloring, both NP-complete in the general setting [Kis20]. So far, Kisfaludi-Bak is also the only author to use the term *unit ball graph* in a hyperbolic context, defining the class HUBG(n, s) as the intersection graphs of equal-sized balls with radius s in n-dimensional hyperbolic space [Kis20]. To avoid confusion, note that we usually use s to refer to hyperbolic radii and r to refer to Euclidean radii.

The goal of our work is to deepen the structural understanding of this family of graph classes, by studying the hyperbolic sphericity of graphs G, i.e. the lowest $n \in \mathbb{N}$ such that $G \in$ HUBG(n, s) for some fixed threshold radius $s > 0$. Doing so, we follow in the footsteps of Maehara, who contributed immensely to the understanding of unit ball graphs in n-dimensional Euclidean space and Euclidean sphericity [Mae84b | Mae84a].

All of this is closely related to graph embeddings: asking whether some graph $G = (V, E)$ is in the class HUBG(n, s) is equivalent to searching an *embedding* (also called a representation) in the form of a function

$$\tau : V \to H^n \quad \text{such that } vw \in E \iff d(\tau(v), \tau(w)) < 2s$$

where H^n denotes (some model of) n-dimensional hyperbolic space with the metric d. Interest in such embeddings comes primarily from the machine learning community: neural networks on graph data profit remarkably from embeddings by using them as a pre-processing step to metrically vectorize the graph input. While Euclidean embeddings have already proven useful, hyperbolic embeddings are even more promising when it comes to graphs with a hierarchical structure [Yan+22].

The mathematical fundament of this thesis is provided by research on hyperbolic space. Anderson's and Ratcliffe's books are worth mentioning in this context [And05 | Rat06]. In the context of metric spaces and isometries, some other research branches appeared promising at first glance, but turned out to be unapplicable for this thesis. A mismatch often originates from the question whether a subset of a metric space is simply equipped with the restricted metric or considered as a surface, thereby obtaining a new metric by considering shortest paths on this surface. For this reasons, certain statements relating Riemannian manifolds like hyperbolic space to Euclidean space, known as Nash embedding theorems, do not contradict the results of this thesis [Nas54]. In general, meanings of the term *embedding* may

vary from context to context. For example, the research on so-called low-distortion embeddings is conceptually relevant for this thesis. However, the results are centered around l_p ($p \in \{1, \ldots, \infty\}$) metric spaces so far [IM04].

1.2 Outline

The main part of this thesis is structured into six chapters. In Chapter 2 we introduce core notions of graph theory, study important properties of Euclidean space and define models of hyperbolic space. This preparation allows us to define unit ball graph representations as well as Euclidean and hyperbolic sphericity in Chapter 3. In Chapter 4 we compare Euclidean and hyperbolic sphericity, establishing important inequalities and showing that hyperbolic sphericity is well-defined. On this basis, we end Chapter 4 with a proof that hyperbolic sphericity is computable using a polynomial amount of space. In Chapter 5, we focus on trees. We present two proofs that trees are hyperbolic unit disk graphs if the threshold radius is chosen great enough. In contrast, we demonstrate that Euclidean sphericity of trees cannot be bounded by a constant. In Chapter 6, we develop a graph family whose behavior contrasts the one of trees: here, a small threshold radius leads to a minimal hyperbolic sphericity of two, and great threshold radii produce hyperbolic sphericity of three. Similarly, we show that for all graphs, increasing the threshold radius increases the hyperbolic sphericity by at most one. Finally, we take a brief excursion towards unit ball graphs in spherical space in Chapter 7.

2 Preliminaries

In this chapter, we present core notions of graph theory, study important properties of Euclidean space and define models of hyperbolic space. For the latter, we benefit from the theory on geodesics, formalizing shortest paths in metric spaces. Throughout the thesis, we use the notation $x > 0$ to denote any positive real number and $0 := (0, \dots, 0)$ to denote the origin of $\mathbb{R}^n$.

2.1 Graph Theory

Graph theory is a broad topic in computer science. The definition of a graph varies slightly depending on the context. In this thesis, we are only interested in undirected graphs without loops.

Definition 2.1: A *graph* is a tuple $G = (V, E)$ where V is a finite set, called *vertex set*, and E is a finite set, called *edge set*. The following properties must hold:

(1) Each element $e \in E$ has the form $e = uv$ where $u, v \in V$.

(2) $uv = vu$, i.e. the edges are undirected.

(3) Edges of the form vv (so-called *loops*) are not allowed.

In some cases, we drop the assumption that V and E are finite. We define *infinite graphs* as tuples $G = (V, E)$ where V is infinite while E may or may not be finite, that satisfy all other properties stated in Definition 2.1. Similarly, we sometimes refer to graphs as *finite graphs*. For any finite or infinite graph G, we use the notation $V(G)$ to denote the vertex set of G.

When comparing two graphs G and G', it is easy to see if they are equal or not. A more interesting question is whether they have the same *structure*: is it possible to obtain G' only by renaming the vertices of G or vice versa? This leads to the following definition:

Definition 2.2: Two graphs $G = (V, E)$ and $G' = (V', E')$ are *isomorphic* if and only if there is a bijection $f : V \to V'$ such that

$$vw \in E \iff f(v)f(w) \in E' \quad \text{for each pair of distinct vertices } v, w \in V.$$

A bijection f with this property is called a *graph isomorphism* between G and G'.

When given a family of sets, a graph can be obtained by treating the sets as vertices and connecting distinct vertices by an edge if and only if their intersection is nonempty:

Definition 2.3: Let X be some set and $A \subseteq \mathcal{P}(X)$ where $\mathcal{P}(X)$ denotes the power set of X. Then, the *intersection graph* of the set A is

$$\mathcal{G}(A) := (A, E) \quad \text{where } E := \{ab \mid a, b \in A, a \neq b \text{ and } a \cap b \neq \emptyset\}.$$

If the inducing set A is finite, then $\mathcal{G}(A)$ is a finite graph, otherwise it is an infinite graph.

Definition 2.4: Let $G = (V, E)$ be a graph and $V' \subseteq V$. The *vertex-induced subgraph* of G with respect to V' is

$$G|_{V'} := (V', E') \quad \text{where } E' := \{vw \in E \mid v, w \in V'\}.$$

That is, a vertex-induced subgraph must contain all the edges of G whose endpoints belong to the vertex subset V'. If we only take a subset of these edges, we simply speak of a *subgraph* of G. Clearly, vertex-induced subgraphs and subgraphs are graphs themselves.

2.2 Euclidean Space

The definitions in this section are based on Ratcliffe's work [Rat06]. Let $\langle \cdot, \cdot \rangle$ denote the standard inner product on $\mathbb{R}^n$ given by

$$\langle x, y \rangle := \sum_{i=1}^{n} x_i y_i \quad \text{for each } x, y \in \mathbb{R}^n.$$

Then, the *Euclidean norm* of a vector $x \in \mathbb{R}^n$ is defined as the real number

$$\|x\| := \sqrt{\langle x, x \rangle}.$$

Definition 2.5: A *metric space* is a tuple (X, d) where X is some set and $d : X \times X \to \mathbb{R}$ is a function with the following properties: for all $x, y, z \in X$,

(1) $d(x, y) \geq 0$ (nonnegativity),

(2) $d(x, y) = 0$ if and only if $x = y$ (nondegeneracy),

(3) $d(x, y) = d(y, x)$ (symmetry),

(4) $d(x, z) \leq d(x, y) + d(y, z)$ (triangle inequality).

A function d with the above properties is also called a *metric* of X.

Definition 2.6: The n-dimensional Euclidean space E^n is given by the metric space $(\mathbb{R}^n, d_E)$ where

$$d_E(x, y) := \|x - y\| \quad \text{for each } x, y \in \mathbb{R}^n.$$

For a proof that the tuple $(\mathbb{R}^n, d_E)$ is indeed a metric space, we refer to Ratcliffe [Rat06]. It is worth introducing some notation concerning metric spaces: if a metric d of a given space X is known from context, we do not clearly differentiate between X and (X, d). For example, the notation $x, y \in E^n$ denotes elements of $\mathbb{R}^n$ while highlighting that these elements are subject to the metric d_E. We proceed by proving the following useful property of the Euclidean metric:

Lemma 2.7: Let $x, y \in E^n$ and $\lambda > 0$. Then $d_E(\lambda x, \lambda y) = \lambda \, d_E(x, y)$.

Proof. We calculate, using the fact that the standard inner product $\langle \cdot, \cdot \rangle$ is linear in both components:

$$\begin{aligned} d_E(\lambda x, \lambda y) &= \sqrt{\langle \lambda x - \lambda y, \lambda x - \lambda y \rangle} \\ &= \sqrt{\langle \lambda(x - y), \lambda(x - y) \rangle} \\ &= \sqrt{\lambda^2 \langle x - y, x - y \rangle} \\ &= \lambda \sqrt{\langle x - y, x - y \rangle} \\ &= \lambda \, d_E(x, y). \end{aligned}$$

∎

Definition 2.8: A *line* of E^n is a set of the form $\{x + \lambda y \mid \lambda \in \mathbb{R}\}$ for any $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^n \setminus \{0\}$.

We also refer to lines of E^n as *Euclidean lines*. If we consider distinct points $x, y \in E^n$, one can easily prove that the unique line of E^n is the set

$$\{x + \lambda(y - x) \mid \lambda \in \mathbb{R}\}.$$

The *line segment* in E^n joining x to y is defined as the set

$$\{x + \lambda(y - x) \mid \lambda \in [0, 1]\}.$$

It is worth noting that every line segment in E^n extends to a unique line of E^n.

Definition 2.9: The *Euclidean sphere* of E^n with center a and radius $r > 0$ is the set

$$S_E(a, r) := \{x \in \mathbb{R}^n \mid d_E(a, x) = r\} \subseteq E^n.$$

Furthermore, the *Euclidean ball* of E^n with center a and radius $r > 0$ is the region bounded by the corresponding Euclidean sphere, i.e.

$$B_E(a, r) := \{x \in \mathbb{R}^n \mid d_E(a, x) < r\} \subseteq E^n.$$

This definition clearly depends on the dimension of the surrounding space E^n. This information will always be evident from the context. Spheres of E^2 are also called *circles*.

Lemma 2.10: Let S be a Euclidean sphere of E^n and l a Euclidean line. Then, $|S \cap l| \leq 2$. Furthermore, if l contains a point of the Euclidean ball corresponding to S, then $|S \cap l| = 2$.

Proof. Let $S = S_E(a, r)$ and $l = \{x + \lambda y \mid \lambda \in \mathbb{R}\}$. Let $b = x + \lambda y$ denote some point of l. Then $b \in S$ if and only if

$$\begin{aligned} & d_E(a, x + \lambda y) = r \\ \iff\ & d_E(a, x + \lambda y)^2 - r^2 = 0. \end{aligned}$$

Treating λ as a variable, the left expression is a polynomial of degree 2 which we denote with $p : \mathbb{R} \to \mathbb{R}$. As such, p can have at most two roots, thus $|S \cap l| \leq 2$.

In the following, suppose that there is a point $z = x + \lambda_z y \in l$ such that $d_E(a, z) < r$. By choice of z, we have $p(\lambda_z) < 0$. Due to the fact that

$$\lim_{\lambda \to \infty} p(\lambda) = \lim_{\lambda \to -\infty} p(\lambda) = \infty,$$

p must have at least two zeros due to the intermediate value theorem. Hence p must have precisely two zeros and since the zeros of p correspond to points in $S \cap l$, we have $|S \cap l| = 2$. ∎

Definition 2.11: Let $x, y \in E^n$. The *Euclidean angle* between x and y is the unique real number $\theta(x, y) \in [0, \pi]$ that satisfies

$$\langle x, y \rangle = \|x\| \, \|y\| \cos \theta(x, y).$$

The uniqueness and existence of such a number $\theta(x, y)$ in the interval $[0, \pi]$ is due to Cauchy's inequality and the fact that $\cos : [0, \pi] \to [-1, 1]$ is a bijection. We simply refer to $\theta(x, y)$ as the *angle* between x and y, if the context is clearly Euclidean.

Remark 2.12: The angle between two vectors behaves as one would naturally expect: for example, consider two vectors $x = \exp(i\alpha)$ and $y = \exp(i\beta)$ with $0 \leq \alpha < \beta \leq \pi$. Then the angle between them is indeed $\theta(x, y) = \beta - \alpha$. To show this, recall that $\exp(i\alpha) = \cos\alpha + i\sin\alpha$ and that $|\exp(i\alpha)| = 1$ for any $\alpha \in \mathbb{R}$. Identifying $\mathbb{C}$ and $\mathbb{R}^2$, we calculate

$$\cos\theta(x, y) = \|x\|\,\|y\|\cos\theta(x, y) = \langle x, y\rangle = \cos\alpha\cos\beta + \sin\alpha\sin\beta = \cos(\beta - \alpha)$$

having used a famous trigonometrical identity in the last step. Since cos is injective when restricting the domain to $[0, \pi]$, we conclude that $\theta(x, y) = \beta - \alpha$.

The following equalities concerning the angle between two vectors will become useful in multiple chapters:

Lemma 2.13: Let x, y in E^n. We can express the Euclidean distance d_E using only the Euclidean norm of x, the Euclidean norm of y and the angle between x and y:

(1) $d_E(x, y) = \sqrt{\|x\|^2 + \|y\|^2 - 2\|x\|\|y\|\cos(\theta(x, y))}$.

(2) If $\|x\| = \|y\|$ holds , then $d_E(x, y) = 2\|x\|\sin\left(\frac{\theta(x,y)}{2}\right)$.

Proof. Concerning the first statement, we calculate

$$\begin{aligned} d_E(x, y) = \sqrt{\langle x - y, x - y\rangle} &= \sqrt{\langle x, x\rangle - 2\langle x, y\rangle + \langle y, y\rangle} \\ &= \sqrt{\|x\|^2 + \|y\|^2 - 2\|x\|\|y\|\cos\theta(x, y)} \end{aligned}$$

using Definition 2.11 in the last step. Concerning the second statement, suppose in the following that $\|x\|$ and $\|y\|$ both are equal to some constant $r \geq 0$. Using the above calculation and the trigonometric identity

$$\sin\left(\frac{\alpha}{2}\right) = \sqrt{\frac{1 - \cos\theta(x, y)}{2}}$$

that holds for any $\alpha \in [0, \pi]$, we obtain

$$\begin{aligned} d_E(x, y) = \sqrt{2r^2 - 2r^2\cos\theta(x, y)} = \sqrt{4r^2\left(\frac{1 - \cos\theta(x, y)}{2}\right)} &= \sqrt{4r^2\sin^2\left(\frac{\theta(x, y)}{2}\right)} \\ &= 2r\sin\left(\frac{\theta(x, y)}{2}\right). \end{aligned}$$

■

2.2.1 Circles in the Euclidean Plane

It will be helpful to recall some properties of circles in 2-dimensional Euclidean Space E^2, i.e. in the Euclidean plane. Three distinct points $x, y, z \in E^2$ are called *collinear* if and only if they lie on a common Euclidean line. Otherwise, they are called noncollinear. In this case, they form a triangle, denoted $T(x, y, z)$.

Lemma 2.14: Consider three distinct noncollinear points $x, y, z \in E^2$. Then there is a unique circle C that passes through x, y and z. We call C the *circumcircle* of the triangle $T(x, y, z)$.

For a proof of this lemma, we refer to Koecher's book on plane geometry [KK07]. As a consequence, two distinct circles intersect at most in two points, otherwise they must be identical. We are particularly interested in the case where two circles are *perpendicular* to each other, meaning that they intersect in two points and at each point of intersection, the two tangents are perpendicular. Identifying $\mathbb{R}^2$ with $\mathbb{C}$, the following statement and its proof can be found in Anderson's book on hyperbolic geometry [And05, p. 2].

Lemma 2.15: Let $S^1 = \{z \in \mathbb{C} \mid |z| = 1\}$ be the unit circle in $\mathbb{C}$. Let A be the Euclidean circle in $\mathbb{C}$ with Euclidean center $\lambda \exp(i\theta)$, $\lambda > 1$, and Euclidean radius $r > 1$. Then A is perpendicular to S^1 if and only if $r = \sqrt{\lambda^2 - 1}$.

2.3 Geodesics

In this section, we generalize the concept of a line and a line segment known from E^n to other metric spaces. Doing so will facilitate the understanding of hyperbolic space.

2.3.1 Distance-Preserving Functions and Isometries

The following two definitions are of great relevance throughout the thesis:

Definition 2.16: Let (X, d_X) and (Y, d_Y) denote metric spaces. A function $\varphi : X \to Y$ is *distance-preserving* if and only if

$$d_Y(\varphi(x), \varphi(y)) = d_X(x, y) \quad \text{for all } x, y \in X.$$

As Ratcliffe remarks, a distance-preserving function is a continuous injection [Rat06, p. 15]. Furthermore, the composition of distance-preserving functions is itself distance-preserving.

Definition 2.17: An *isometry* from a metric space X to a metric space Y is a distance-preserving bijection $\phi : X \to Y$.

The inverse of an isometry is clearly an isometry, and the composition of two isometries is an isometry. Two metric spaces X and Y are said to be *isometric* if and only if there is an isometry $\phi : X \to Y$. We observe that being isometric is an equivalence relation among the class of all metric spaces.

Example 2.18: Let $m, n \in \mathbb{N}$ and $m < n$. An example of a distance-preserving function that is not an isometry is the so-called *inclusion*

$$\iota : E^m \to E^n, (x_1, \ldots, x_m) \mapsto (x_1, \ldots, x_m, 0, \ldots, 0).$$

We also consider the identity $\mathrm{id} : E^m \to E^m$ as a trivial inclusion. The concept of inclusions can also be applied to other metric spaces parametrized by dimension.

Example 2.19: We can equip the complex plane $\mathbb{C}$ with a metric $d_{\mathbb{C}}$, called the *complex metric*, by firstly defining the complex norm $|\cdot|$ as

$$|z| := \sqrt{\mathrm{Re}(z)^2 + \mathrm{Im}(z)^2} \quad \text{for each } z \in \mathbb{C}$$

and then defining the complex metric by

$$d_{\mathbb{C}}(z_1, z_2) := |z_1 - z_2| \quad \text{for all } z_1, z_2 \in \mathbb{C}.$$

This strongly resembles the Euclidean metric d_E of E^2. Indeed, E^2 and $(\mathbb{C}, d_{\mathbb{C}})$ are isometric by the isometry

$$\phi : E^2 \to \mathbb{C}, (x, y) \mapsto x + iy.$$

Example 2.20: Let $v \in \mathbb{R}^n$ be any vector and let

$$\phi_v : E^n \to E^n, x \mapsto v + x$$

denote the translation by the vector v. Then, direct calculation shows that ϕ_v is an isometry.

2.3.2 Geodesic Segments and Geodesics

The definitions in this section are due to Ratcliffe [Rat06]. We prove some additional statements that will be useful later on. Throughout this section, we use (X, d) to denote a metric space.

Definition 2.21: A *curve* in a space X is a continuous function $\gamma : [a, b] \to X$ where $[a, b]$ is a closed interval in $\mathbb{R}$ with $a < b$.

Let $\gamma : [a, b] \to X$ be a curve. Then $\gamma(a)$ is called the *initial point* of γ and $\gamma(b)$ is called the *terminal point*. We say that γ is a curve in X from $\gamma(a)$ to $\gamma(b)$.

Definition 2.22: A *geodesic arc* in a metric space X is a distance-preserving function $\alpha : [a, b] \to X$, with $a < b$ in $\mathbb{R}$.

A geodesic arc α is a continuous injection and so is a curve.

Definition 2.23: A *geodesic segment* joining a point x to a point y in a metric space X is the image of a geodesic arc $\alpha : [a, b] \to X$ whose initial point is x and terminal point is y.

For example, it can be shown that the geodesic segments of E^n are its line segments. This intuition can be applied to other metric spaces as well: we can generally think of a geodesic segment as a *shortest path* between two specific points. Similarly, we can think of the corresponding geodesic arcs as walks on this shortest paths with constant, normalized speed.

Remark 2.24: Consider two geodesic arcs $\alpha : [a, b] \to X$ and $\beta : [c, d] \to X$ that have the same image, namely the geodesic segment $[x, y]$. Define the Euclidean isometry $\phi : [a, b] \to [c, d]$ by $\phi(k) := (k - a + c)$ for each $k \in [a, b]$. We claim that α and β only differ by precomposition with the function ϕ, i.e. $\alpha = \beta \circ \phi$. Clearly, $\beta \circ \phi$ is also a geodesic arc that has $[x, y]$ as its image. Fix any $k \in [a, b]$. Since $\beta \circ \phi$ is injective and has the same image as α, there is a unique $k' \in [a, b]$ such that $\alpha(k) = \beta \circ \phi(k')$. We calculate

$$k - a = d_E(k, a) = d(\alpha(k), \alpha(a)) = d(\beta \circ \phi(k'), \beta \circ \phi(a)) = d_E(k', a) = k' - a.$$

Thus $k = k'$. Since k was arbitrary, we conclude $\alpha = \beta \circ \phi$.

Definition 2.25: A metric space X is *geodesically convex* if and only if for each pair of distinct points x, y of X, there is a unique geodesic segment in X joining x to y.

Theorem 2.26: Let $[x, y]$ and $[y, z]$ be geodesic segments joining x to y and y to z, respectively, in a metric space X. Then the set $[x, y] \cup [y, z]$ is a geodesic segment joining x to z in X if and only if the triangle inequality is tight for x, y, z, i.e. $d(x, z) = d(x, y) + d(y, z)$.

For a proof, we refer to Ratcliffe [Rat06, p. 24]. We can also go the other way and pick some point on an existing geodesic segment. Then, the triangle inequality is tight as well:

Corollary 2.27: Let $[x, z]$ be a geodesic segment joining x to z. For each $y \in [x, z]$, we have $d(x, y) + d(y, z) = d(x, z)$.

Proof. Let $\alpha : [a, c] \to X$ denote a geodesic arc whose image is $[x, z]$. Let $b \in [a, c]$ such that $\alpha(b) = y$. The restriction $\alpha|_{[a,b]}$ is also distance-preserving, i.e. a geodesic arc as well. The same holds for the restriction $\alpha|_{[b,c]}$. By construction, we have

$$[x, y] \cup [y, z] = \alpha|_{[a,b]}([a, b]) \cup \alpha|_{[b,c]}([b, c]) = \alpha([a, c]) = [x, z].$$

Now the statement follows directly from Theorem 2.26. ∎

The *ball* of X with center $x \in X$ and radius $r > 0$ is defined as the set

$$B(x, r) := \{z \in X \mid d(x, z) < r\} \subseteq X.$$

Testing whether two balls intersect can be done by comparing the distance of their centers with their radii, as we are used to from Euclidean space:

Lemma 2.28: Let (X, d) be a geodesically convex metric space. Let $x, y \in X$ denote distinct points in X and let $r, s > 0$. Then

$$B(x, r) \cap B(y, s) \neq \emptyset \iff d(x, y) < r + s.$$

Proof. The first implication holds due to the triangle inequality: suppose there is $z \in B(x, r) \cap B(y, s)$, then $d(x, y) \leq d(x, z) + d(y, z) < r + s$. Conversely, suppose that $d(x, y) < r + s$. Since X is geodesically convex, there is a geodesic segment joining x to y and thus also a corresponding geodesic arc α. We may assume that α has the form $\alpha : [0, a] \to X$ where $a := d(x, y)$. From $a < r + s$ it follows that

$$\frac{r}{r+s}a < r \quad \text{and} \quad \frac{s}{r+s}a < s.$$

Let $b := \frac{r}{r+s}a < a$ and $z := \alpha(b) \in X$. The fact that α is distance-preserving implies

$$d(x, z) = d(\alpha(0), \alpha(b)) = b - 0 = b < r \quad \text{and} \quad d(z, y) = d(\alpha(b), \alpha(a)) = a - b = a(1 - \frac{r}{r+s}) = a\frac{s}{r+s} < s.$$

We conclude that $z \in B(x, r)$ and $z \in B(y, s)$, so the intersection of the two balls is nonempty. ∎

Similarly, we define the *closed ball* of X with center x and radius $r > 0$ as the set

$$\overline{B}(x, r) := \{z \in X \mid d(x, z) \leq r\} \subseteq X.$$

By slightly adapting the proof of Lemma 2.28, the following equivalence can be shown, again supposing that X is geodesically convex:

$$\overline{B}(x, r) \cap \overline{B}(y, s) \neq \emptyset \iff d(x, y) \leq r + s \quad \text{for each distinct pair of points } x, y \in X \text{ and } r, s > 0.$$

Next, we consider geodesic lines, which one obtains when a geodesic segment is continued indefinitely. A function $\varphi : X \to Y$ between metric spaces is *locally distance-preserving* if and only if for each point $x \in X$ there is an $r > 0$ such that φ preserves the distances between any two points in $B(x, r)$. In that case φ is continuous, because it is continuous at each point of X.

Definition 2.29: A *geodesic line* in a metric space X is a locally distance-preserving function $\lambda : \mathbb{R} \to X$. Furthermore, a *geodesic* in a metric space X is the image of such a geodesic line λ.

Definition 2.30: A metric space X is *geodesically complete* if and only if each geodesic arc $\alpha : [a, b] \to X$ extends to a unique geodesic line $\lambda : \mathbb{R} \to X$.

Example 2.31: The geodesics of E^n are its Euclidean lines. Furthermore, E^n is geodesically convex and geodesically complete [Rat06, p. 24, 25].

In particular, all results presented in this section apply to Euclidean space. Let us explore how the introduced concepts relate to isometries.

Theorem 2.32: Let (X, d_X) and (Y, d_Y) be metric spaces that are isometric by an isometry $\phi : X \to Y$. Then, the following statements hold:

(1) If $\alpha : [a, b] \to X$ is a geodesic arc in X, then $\phi \circ \alpha : [a, b] \to Y$ is a geodesic arc in Y. Furthermore, if s is a geodesic segment in X joining x_1 to x_2, then $\phi(s)$ is a geodesic segment in Y joining $\phi(x_1)$ to $\phi(x_2)$.

(2) X is geodesically convex if and only if Y is geodesically convex.

(3) If $\lambda : \mathbb{R} \to X$ is a geodesic line in X, then $\phi \circ \lambda : \mathbb{R} \to Y$ is a geodesic line in Y. In particular, the image of a geodesic in X under ϕ is a geodesic in Y.

(4) X is geodesically complete if and only if Y is geodesically complete.

Proof. We prove the statements one by one.

(1) Let $\alpha : [a, b]$ denote a geodesic arc in X. As an isometry, ϕ is distance-preserving. The composition of two distance-preserving functions is itself distance-preserving. Hence $\phi \circ \alpha$ is a geodesic arc in Y by definition. Let s be a geodesic segment in X joining x_1 to x_2. In the following, we may assume that s is the image of α. In particular, we have

$$\phi(s) = \phi(\alpha([a, b])) = (\phi \circ \alpha)([a, b])$$

thus $\phi(s)$ is a geodesic segment in Y. Because $\phi(x_1)$ is the initial point of the geodesic arc $\phi \circ \alpha$ and because $\phi(x_2)$ is its terminal point, s joins these two points.

(2) It suffices to show that, if X is geodesically convex, then Y is geodesically convex. The reverse implication can be shown analogously using the isometry $\phi^{-1} : Y \to X$. Hence, suppose that X is geodesically convex. Let y_1, y_2 denote two distinct points in Y. Let $x_1, x_2 \in X$ such that $\phi(x_1) = y_1$ and $\phi(x_2) = y_2$. Let $[x_1, x_2] \subseteq X$ denote the unique geodesic segment joining x_1 to x_2 in X. Due to Theorem 2.32 (1), $s := \phi([x_1, x_2])$ is a geodesic segment in Y that joins y_1 to y_2. It remains to show that it is the unique geodesic segment in Y with this property. Assume another geodesic segment s' in Y also joins y_1 to y_2. By applying Theorem 2.32 (1) in the context of the isometry $\phi^{-1} : Y \to X$, we conclude that $\phi^{-1}(s)$ and $\phi^{-1}(s')$ are both geodesic segments in X joining x_1 to x_2. Since X is geodesically convex, they must be equal. Consequently, s and s' must be equal.

(3) Let $\lambda : \mathbb{R} \to X$ denote a geodesic line in X. Thus λ is locally distance-preserving. Clearly, the composition $\phi \circ \lambda$ is also locally distance-preserving, because ϕ is distance-preserving. Hence $\phi \circ \lambda$ is a geodesic line in Y by definition. Let g be a geodesic in X. In the following, we may assume that g is the image of λ. We have $\phi(g) = (\phi \circ \lambda)(\mathbb{R})$, so $\phi(g)$ is the image of the geodesic line $\phi \circ \lambda$ in Y and thus a geodesic in Y.

(4) Similar to how we proved Theorem 2.32 (2), it suffices to suppose that X is geodesically complete and show that Y then must also be geodesically complete. Let $\alpha : [a, b] \to Y$ denote a geodesic arc in Y. Applying Theorem 2.32 (1) in the context of the isometry $\phi^{-1} : Y \to X$, we get that $\phi^{-1} \circ \alpha$ is a geodesic arc in X. Since X is geodesically complete, $\phi^{-1} \circ \alpha$ extends to a unique geodesic line in X which we denote with $\lambda : \mathbb{R} \to X$. Because of Theorem 2.32 (3), $\phi \circ \lambda$ is a geodesic line in Y and we observe that α extends to this geodesic line since

$$\lambda|_{[a,b]} = \phi^{-1} \circ \alpha \implies (\phi \circ \lambda)|_{[a,b]} = \alpha.$$

It remains to show that $l := \phi \circ \lambda$ is the unique geodesic line in Y to which α extends. Assume another geodesic line l' in Y also has this property. By applying Theorem 2.32 (3) in the context of the isometry $\phi^{-1} : Y \to X$, we conclude that $\phi^{-1} \circ l$ and $\phi^{-1} \circ l'$ are both geodesic lines in X. Furthermore, they are both an extension of the geodesic arc $\phi^{-1} \circ \alpha$ in X since

$$l|_{[a,b]} = \alpha \implies (\phi^{-1} \circ l)\big|_{[a,b]} = \phi^{-1} \circ \alpha$$

and the same is valid for l'. Because X is geodesically complete, we conclude that $\phi^{-1} \circ l = \lambda = \phi^{-1} \circ l'$ and thus $l = l'$. ∎

We finish this section by deriving a rather technical lemma that will be needed later on.

Lemma 2.33: Let (X, d) be a geodesically convex and geodesically complete metric space. Consider two geodesic segments with a common endpoint $[x, y]$ and $[x, z]$ that intersect in a second point $p \neq x$. Then, one geodesic segment is a subset of the other.

Proof. We may assume that $[x, y]$ is the image of some geodesic arc $\alpha : [0, a] \to X$ and that $[x, z]$ is the image of another geodesic arc $\beta : [0, b] \to X$. Considering the point of intersection p, there are values $c \in [0, a]$ and $c' \in [0, b]$ such that $\alpha(c) = p = \beta(c')$. Because α and β are distance-preserving, we obtain $c = d(x, p) = c'$. Note that $\alpha|_{[0,c]}$ and $\beta|_{[0,c]}$ are both geodesic arcs with initial point x and terminal point p. Since X is geodesically convex, we obtain the equality of images $\alpha|_{[0,c]}([0, c]) = \beta|_{[0,c]}([0, c])$. Not only are these two images equal, but even $\alpha|_{[0,c]} = \beta|_{[0,c]}$ holds as a consequence of Remark 2.24. Since X is also geodesically complete, the restriction of α and α itself must extend to the same geodesic line. The same holds for β and its restriction. Now, the fact that the two restrictions are equal implies that α and β must extend to the same, unique geodesic line $\lambda : \mathbb{R} \to X$. That is, we have $\lambda|_{[0,a]} = \alpha$ and $\lambda|_{[0,b]} = \beta$. Without loss of generality, we may assume that $a \leq b$. Consequently, $[x, y] = \lambda([0, a]) \subseteq \lambda([0, b]) = [x, z]$. ∎

2.4 Hyperbolic Space

The following definition of n-dimensional hyperbolic space is often quoted: *the n-dimensional hyperbolic space is a (unique) simply connected complete n-dimensional Riemannian manifold with a constant negative sectional curvature -1* [GN98]. Fully understanding this definition is outside the scope of this thesis and would require at least an entry course on topology and differential geometry. However, it is remarkable that hyperbolic space can be defined *axiomatically*, i.e. via certain properties and not simply as a particular metric space. In fact, there are multiple metric spaces satisfying the above definition. We call them *models* of n-dimensional hyperbolic space and use the variable H^n to denote any such model. Note that we will not use the axiomatic definition of hyperbolic space in this thesis, but it is worth knowing that such a definition exists, to better understand the concept of having multiple models of n-dimensional hyperbolic space.

Similarly, it is also possible to define n-dimensional Euclidean space axiomatically, i.e. by demanding certain properties, and then providing one or multiple models, i.e. metric spaces that fulfill the properties. We did not take this approach because it fully suffices for our purposes to treat $E^n = (\mathbb{R}^n, d_E)$ as (the canonical model of) n-dimensional Euclidean space. However, other models are possible: when focusing on Euclidean plane geometry, one can also base a model on the complex plane $\mathbb{C}$ with the complex metric, as defined in Example 2.19.

Coming back to the axiomatic definition of hyperbolic space, one property is worth highlighting: the constant curvature of -1. To put this in perspective, it is known that the unit sphere $S^n \subseteq \mathbb{R}^{n+1}$ has constant curvature of 1 and that n-dimensional Euclidean space has constant curvature of 0. Informally, we can thus imagine hyperbolic space as the counterpart of spherical space with Euclidean space lying somewhere in between the two. Historically, another approach towards hyperbolic space, in particular 2-dimensional, is common: geometric axiomizations of Euclidean plane geometry include the *parallel postulate*, stating that, given a Euclidean line l and a point p not on l, there is only one Euclidean line through p and parallel to l. *Parallel* simply means that the two lines are disjoint. For a long time, it was uncertain whether the parallel postulate could be derived from the other axioms. The development

of hyperbolic plane geometry settled the debate. Hyperbolic plane geometry satisfies every axiom of Euclidean plane geometry, apart from the parallel postulate: in the hyperbolic case, there are *infinitely many* hyperbolic lines through a point p that are parallel to a given hyperbolic line l, assuming $p \notin l$. For details on this axiomatic approach, we refer to Hartshorne's book [Har00].

Defining an initial model of hyperbolic space in a comprehensive way is a challenging task on its own. In this context, we strongly recommend Ratcliffe's book on hyperbolic manifolds [Rat06]. Ratcliffe primarily uses the *hyperboloid model* to accomplish this task. This model is well suited to define hyperbolic space in a way that does not appear arbitrary, because it highlights relationships between spherical, Euclidean and hyperbolic geometry. After fully defining the hyperboloid model, including measurements like hyperbolic angles and hyperbolic volume with respect to this model, Ratcliffe derives other models by constructing them in a way that they are isometric to the hyperboloid model. Since being isometric is an equivalence relation, this has the following important consequence:

Remark 2.34: Let $n \in \mathbb{N}$. Then, any two models of n-dimensional hyperbolic space are isometric.

In other words, models of n-dimensional hyperbolic space are *unique up to isometry*. This result can also be derived from the axiomatic definition given at the start, justifying the term *unique* used therein. The question remains why it is useful to consider multiple models of hyperbolic geometry. This is because each model has its advantages and disadvantages in terms of geometric intuition and in terms of simplicity of formulas. However, one would expect that certain properties are valid for any model of hyperbolic space, like laws of trigonometry. This is indeed the case. To provide an example, consider *hyperbolic angles*: no matter in which model X of hyperbolic space they are firstly defined, one needs to ensure that they are preserved by isometries between the models. To be precise, suppose that two metric spaces X and Y are given which both model n-dimensional hyperbolic space. As mentioned above, there exists an isometry $\phi : X \to Y$. Then, we expect that the following property holds: if the hyperbolic angle between two geometric objects $o_1, o_2 \subseteq X$ is defined and equals α, then the hyperbolic angle between $\phi(o_1), \phi(o_2) \subseteq Y$ is also defined and equals α as well. The same applies to *hyperbolic volume*: if the hyperbolic volume of $o_1 \subseteq X$ is defined, we expect that the hyperbolic volume of $\phi(o_1) \subseteq Y$ is also defined and that both are equal. This coherency between different models is usually achieved by definition.

Since our focus lies on sphericity and not on studying hyperbolic spaces as an end in itself, our approach differs from Ratcliffe. Instead of the hyperboloid model, we study the *Poincaré disk model* extensively, since the latter has much more to offer in the context of sphericity. We also define hyperbolic angles and hyperbolic volume in this model. In order to understand where the following definitions geometrically originate and why the different models are consistent in terms of angles, volume etc., we refer to Ratcliffe [Rat06].

2.4.1 The Poincaré Disk Model

The name of this model originates from the *Poincaré disk*, a 2-dimensional model of hyperbolic plane geometry, named after Henri Poincaré. Here, we follow the definitions of Ratcliffe [Rat06]. Generalizing the 2-dimensional Poincaré disk yields the following model:

Definition 2.35: The n-dimensional *Poincaré disk model* is given by the metric space $(\mathbb{D}^n, d_D)$ where $\mathbb{D}^n := B_E(0,1) \subseteq \mathbb{R}^n$ and the metric d_D is given by

$$\cosh d_D(x,y) = 1 + \frac{2\|x-y\|^2}{(1-\|x\|^2)(1-\|y\|^2)}.$$

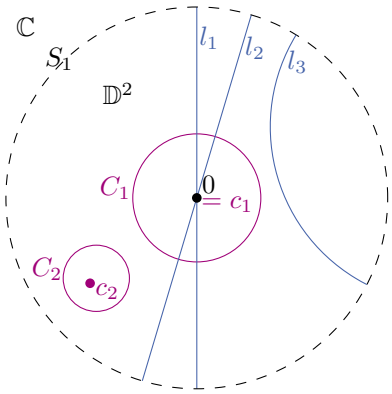

Figure 2.1: Examples of hyperbolic lines l_1, l_2, l_3 of $\mathbb{D}^2$ and two equal-radius hyperbolic circles C_1, C_2 (with center c_1 resp. c_2) of $\mathbb{D}^2$.

The case of $n = 2$ will be most prominent throughout the thesis. In this case, we usually identify $B_E(0, 1) \subseteq \mathbb{R}^2$ with the unit ball of $\mathbb{C}$. This is justified by the isometry between $(\mathbb{C}, d_{\mathbb{C}})$ and E^2 as presented in Example 2.19. This view is helpful because the complex plane $\mathbb{C}$ not only allows the Cartesian-like representation of vectors in the form $z = x + iy \in \mathbb{C}$, but also a representation by angle and vector length, based on the complex exponential function $\exp : \mathbb{C} \to \mathbb{C}$ that satisfies

$$\exp(i\varphi) = \cos\varphi + i \sin\varphi.$$

Consequently, for each $z \in \mathbb{C} \setminus \{0\}$, there is a unique length $r > 0$ and a unique angle $\varphi \in [0, 2\pi]$ such that $z = r \exp(i\varphi)$. In particular, we can capture all points in the Poincaré disk $\mathbb{D}^2$ by requiring $r < 1$. To adapt the above definition for $n = 2$ to this alternative view, one simply has to replace each occurrence of the Euclidean norm $\|\cdot\|$ with the complex norm $|\cdot|$.

Definition 2.36: A subset l of $\mathbb{D}^n$ is a hyperbolic line of $\mathbb{D}^n$ if and only if l is either an open diameter of $\mathbb{D}^n$ or the intersection of $\mathbb{D}^n$ with a circle orthogonal to S^{n-1}.

A hyperbolic line l is an open diameter of $\mathbb{D}^n$ if and only if it contains the origin. In this case, l is both a hyperbolic line and a Euclidean line restricted to $\mathbb{D}^n \subseteq \mathbb{R}^n$. Figure 2.1 depicts the two types of hyperbolic lines of $\mathbb{D}^2$. It also illustrates that the Euclidean parallel postulate does not hold in the context of $\mathbb{D}^2$. The next theorem establishes a connection to the previous section:

Theorem 2.37: The hyperbolic lines of $\mathbb{D}^n$ are its geodesics. Furthermore, $\mathbb{D}^n$ is geodesically convex and geodesically complete.

Proof. Ratcliffe provides proofs that the hyperboloid model $\mathbb{F}^n$ of hyperbolic space has the above properties [Rat06, p. 65, 70]. We have not encountered this model yet and will only do so briefly. This is not an issue since Ratcliffe also presents an isometry $\phi : \mathbb{F}^n \to \mathbb{D}^n$ that bijectively maps hyperbolic lines of $\mathbb{F}^n$ to hyperbolic lines of $\mathbb{D}^n$. Due to Theorem 2.32, we conclude that $\mathbb{D}^n$ inherits the above properties from $\mathbb{F}^n$. ∎

In particular, all results presented in Section 2.3 apply to the Poincaré disk model $\mathbb{D}^n$. With the knowledge of Theorem 2.37, we are also able to describe the geodesic segments of $\mathbb{D}^n$: a geodesic segment s of $\mathbb{D}^n$ is either a Euclidean line segment of an open diameter of $\mathbb{D}^n$ or an arc in $\mathbb{D}^n$ of a circle orthogonal to S^{n-1}. Next, we consider hyperbolic angles. Angles in $\mathbb{D}^n$ are inherited from the surrounding space $\mathbb{R}^n$:

Definition 2.38: The angle between two curves in $\mathbb{D}^n$ is defined to be the angle between the curves when they are considered to be curves in $\mathbb{R}^n$, which in turn is defined to be the angle between their tangent lines.

Hyperbolic spheres and balls play a major role in this thesis. The *hyperbolic sphere* of $\mathbb{D}^n$, with center $a \in \mathbb{D}^n$ and radius $s > 0$, is defined as the set

$$S_D(a, s) := \{x \in \mathbb{D}^n \mid d_D(a, x) = s\} \subseteq \mathbb{D}^n.$$

Furthermore, the *hyperbolic ball* of $\mathbb{D}^n$ with center $a \in \mathbb{D}^n$ and radius $s > 0$ is defined to be the region bounded by the corresponding hyperbolic sphere, i.e.

$$B_D(a, s) := \{x \in \mathbb{D}^n \mid d_D(a, x) < s\} \subseteq \mathbb{D}^n.$$

In the context of $\mathbb{D}^2$, hyperbolic spheres resp. balls are called hyperbolic *circles* resp. *disks*. The Poincaré disk model is particularly well-suited to study hyperbolic spheres and balls, as becomes clear with the following theorem:

Theorem 2.39: A subset S of $\mathbb{D}^n$ is a hyperbolic sphere of $\mathbb{D}^n$ if and only if S is a Euclidean sphere E^n that is contained in $\mathbb{D}^n$.

For a proof, we refer to Ratcliffe [Rat06, p. 124]. As an immediate consequence of this theorem, a subset B of $\mathbb{D}^n$ is a hyperbolic ball of $\mathbb{D}^n$ if and only if B is a Euclidean ball in E^n that is contained in $\mathbb{D}^n$. Two equal-radius hyperbolic circles of $\mathbb{D}^2$ are depicted in Figure 2.1. To better understand the relationship between hyperbolic spheres of $\mathbb{D}^n$ and Euclidean spheres, the following lemma is helpful:

Lemma 2.40: Let $x \in \mathbb{D}^n$ and let $s := d_D(0, x)$ denote the hyperbolic distance from x to the origin. It is related to the Euclidean distance $\|x\|$ from x to the origin by the following two formulas:

$$\|x\| = \tanh\left(\frac{s}{2}\right) \quad \text{and} \quad s = \log\left(\frac{1 + \|x\|}{1 - \|x\|}\right).$$

For a proof, we refer again to Ratcliffe [Rat06, p. 127]. The next two remarks focus on the conversion between hyperbolic spheres and Euclidean spheres in $\mathbb{D}^n$. They can also be applied to balls. Suppose that we know the Euclidean representation of such a sphere, what is its hyperbolic radius? For our purposes, it suffices to consider spheres whose corresponding balls do not contain the origin.

Remark 2.41: Consider a Euclidean sphere $S := S_E(x, r)$ contained in $\mathbb{D}^n$ that satisfies $\|x\| > r$. Our goal is to calculate the hyperbolic radius s of S as a hyperbolic sphere. Let l denote the hyperbolic line through the origin and x. Since l is on open diameter of S^{n-1} and passes through the Euclidean center of S, we know that l must intersect S in exactly two points

$$p := x - r\frac{x}{\|x\|} \quad \text{and} \quad q := x + r\frac{x}{\|x\|}$$

due to Lemma 2.10. Clearly, these two points induce a Euclidean diameter $[p, q]$ of S. As such, $[p, q]$ is orthogonal to S in the Euclidean sense, and thus also in the hyperbolic sense, because the Poincaré disk model is conformal with respect to angles. Since $[p, q]$ is also a geodesic segment in $\mathbb{D}^n$, this means that $[p, q]$ is a hyperbolic diameter of S as well. So the hyperbolic radius must satisfy $s = d_D(p, q)/2$ and it remains to calculate the right-hand side. To do so, we note that p lies between the origin and q on the hyperbolic line l. This is because $\|x\| > r$. Using Theorem 2.26, we obtain

$$d_D(0, p) + d_D(p, q) = d_D(0, q) \implies d_D(p, q) = d_D(0, q) - d_D(0, p).$$

Now Lemma 2.40 yields

$$\begin{aligned} d_D(p,q) &= \log\left(\frac{1+\|q\|}{1-\|q\|}\right) - \log\left(\frac{1+\|p\|}{1-\|p\|}\right) \\ &= \log\left(\frac{(1+\|q\|)(1-\|p\|)}{(1-\|q\|)(1+\|p\|)}\right) \\ &= \log\left(\frac{(1+(\|x\|+r))\cdot(1-(\|x\|-r))}{(1-(\|x\|+r))\cdot(1+(\|x\|-r))}\right) \\ &= \log\left(\frac{(1+r)^2-\|x\|^2}{(1-r)^2-\|x\|^2}\right). \end{aligned} \tag{2.1}$$

The following consequence is worth highlighting: if we have two spheres $S_E(x,r)$ and $S_E(y,r)$ both contained in $\mathbb{D}^n$ and if $\|x\| = \|y\|$ holds, then the two spheres also have the same hyperbolic radius because only the length of the center vector appears in Equation (2.1).

Assuming that we know the hyperbolic representation of a sphere, what is the Euclidean representation?

Remark 2.42: Given a hyperbolic sphere of $\mathbb{D}^n$ with hyperbolic center c_H and hyperbolic radius r_H. Let μ_H denote the hyperbolic distance of c_H from the origin. The Euclidean representation as a sphere with center c_E and radius r_E can be derived via the following steps: let l be the hyperbolic line containing c_H and the origin. It intersects the hyperbolic sphere in two points p_1, p_2 whose hyperbolic distance from the origin is $b_{1,H} := \mu_H - r_H$ resp. $b_{2,H} := \mu_H + r_H$. As argued in Remark 2.41, the geodesic segment $[p_1, p_2]$ in $\mathbb{D}^n$ is both a Euclidean and hyperbolic diameter of the sphere. In particular, c_E must lie on this line segment. Using the first formula from Lemma 2.40, we obtain that the Euclidean distance of p_j from the origin is $b_{j,E} := \tanh(\frac{b_{j,H}}{2})$ for $j \in \{1,2\}$. As a consequence, the Euclidean distance of c_E from the origin must be $\mu_E := \frac{b_{1,E}+b_{2,E}}{2}$ and thus $c_E = \mu_E \frac{c_H}{|c_H|}$. Furthermore, the Euclidean radius must satisfy $r_E = \mu_E - b_{1,E} = \frac{b_{2,E}-b_{1,E}}{2}$.

For sake of completeness, we also provide a formula of hyperbolic volume in the context of the Poincaré disk model. We will never use this definition directly, but always rely on explicit formulas for particular geometric objects.

Definition 2.43: The *hyperbolic volume* of a set X in $\mathbb{D}^n$ is given by the integral

$$\mathrm{Vol}_{\mathrm{H}}(X) := \int_X \frac{2^n dx_1 \dots dx_n}{(1-\|x\|^2)^n},$$

provided that this integral is defined.

As one would naturally expect from a volume notion, the hyperbolic volume of a set $X \subseteq \mathbb{D}^n$ is nonnegative if it is defined. Furthermore, hyperbolic volume is countably additive: if $\{X_i\}_{i\in\mathbb{N}}$ is a sequence of disjoint subsets $X_i \subseteq \mathbb{D}^n$ such that $\mathrm{Vol}_{\mathrm{H}}(X_i)$ is defined for each $i \in \mathbb{N}$, then the hyperbolic volume of $X := \cup_{i\in\mathbb{N}} X_i$ is also defined and equals

$$\mathrm{Vol}_{\mathrm{H}}(X) = \sum_{i\in\mathbb{N}} \mathrm{Vol}_{\mathrm{H}}(X_i).$$

Hyperbolic volume is also invariant under isometries of $\mathbb{D}^n$. In case of $\mathbb{D}^2$, we speak of *hyperbolic area* instead of hyperbolic volume, which we denote with $\mathrm{Area}_{\mathrm{H}}(\cdot)$.

2.4.2 The Upper Half-Space Model

The following definitions and the theorem are taken from Ratcliffe with slight rephrasing [Rat06, p. 128-132].

Definition 2.44: The n-dimensional *upper half-space model* is given by the metric space $(\mathbb{U}^n, d_U)$ where $\mathbb{U}^n := \{x = (x_1, \dots, x_n) \in \mathbb{R}^n \mid x_1 > 0\}$ and the metric d_U is given by

$$\cosh d_U(x, y) = 1 + \frac{\|x - y\|^2}{2x_n y_n}.$$

Analogous to the Poincaré disk model, we define the *hyperbolic sphere* of $\mathbb{U}^n$, with center $a \in \mathbb{U}^n$ and radius $s > 0$, as the set

$$S_U(a, s) := \{x \in \mathbb{U}^n \mid \ d_U(a, x) = s\} \subseteq \mathbb{U}^n.$$

Furthermore, the *hyperbolic ball* of $\mathbb{U}^n$ with center $a \in \mathbb{U}^n$ and radius $s > 0$ is defined to be the region bounded by the corresponding hyperbolic sphere, i.e.

$$B_U(a, s) := \{x \in \mathbb{U}^n \mid \ d_U(a, x) < s\} \subseteq \mathbb{U}^n.$$

The following theorem reminds us of Theorem 2.39 in the way it relates hyperbolic spheres to Euclidean spheres. It even provides a useful formula making the relationship more explicit. This formula is a core reason why we sometimes prefer using $\mathbb{U}^n$ instead of $\mathbb{D}^n$ in the context of sphericity.

Theorem 2.45: A subset S of $\mathbb{U}^n$ is a hyperbolic sphere of $\mathbb{U}^n$ if and only if S is a Euclidean sphere of E^n that is contained in $\mathbb{U}^n$. Furthermore, the following relationship holds:

$$S_U(a, s) = S_E(c(s), a_n \sinh s) \quad \text{where } \ c(s) = (a_1, \dots, a_{n-1}, a_n \cosh s).$$

The same applies to hyperbolic balls of $\mathbb{U}^n$.

2.4.3 The Hyperboloid Model

Here, we model n-dimensional hyperbolic space by embedding it into $\mathbb{R}^{n+1}$. Take the following quadratic form on $\mathbb{R}^{n+1}$:

$$Q_n(x) := -x_0^2 + \sum_{j=1}^{n} x_j^2.$$

Note that in the context of the hyperboloid model, it is useful to use the vector notation $x = (x_0, \dots, x_n) \in \mathbb{R}^{n+1}$, i.e. index the first entry with zero. The quadratic form Q defines a two-sheeted hyperboloid $\{x \in \mathbb{R}^{n+1} \mid Q_n(x) = -1\}$. The hyperbolic space is represented by all points on the forward sheet $\mathbb{F}^n$:

Definition 2.46: The *hyperboloid model* of n-dimensional hyperbolic space is given by the so-called *forward sheet*

$$\mathbb{F}^n := \{x \in \mathbb{R}^{n+1} \mid Q_n(x) = -1 \text{ and } x_0 > 0\}$$

together with the metric d_F given by

$$\cosh d_F(x, y) = B(x, y) \quad \text{with } \ B(x, y) := x_0 y_0 - \left(\sum_{j=1}^{n} x_j y_j \right) \ \text{ for each pair of distinct points } x, y \in \mathbb{F}^n.$$

2.4.4 Model-Independent Properties

As noted in Remark 2.34, for a fixed dimension $n \in \mathbb{N}$, any two models of n-dimensional hyperbolic space are isometric. In the following, we use H^n to denote any model of n-dimensional hyperbolic space.

Remark 2.47: How does hyperbolic space compare to Euclidean space? Ratcliffe shows the following: H^1 is isometric to E^1, but H^n is not isometric to E^n for each $n > 1$ [Rat06, p. 70].

When speaking of *hyperbolic lines* of H^n, we always refer to the geodesics of the particular model H^n. We know that an isometry ϕ maps geodesics onto geodesics, see Theorem 2.32 (4). Taking ϕ^{-1} into account, this induced mapping between the geodesics of the two metric spaces must be bijective. Combining this with the fact that any two models of n-dimensional space are isometric, the hyperbolic lines of one model must correspond bijectively to the hyperbolic lines of any other model. As a consequence, the proof of Theorem 2.37 generalizes to any model H^n of n-dimensional hyperbolic space: each model of n-dimensional space is geodesically convex and geodesically complete. In particular, all results presented in Section 2.3 apply to each model H^n.

Similar to Example 2.18, inclusions are also available for hyperbolic space. We can even use them to convert between models while increasing the dimension:

Lemma 2.48: Let $m, n \in \mathbb{N}$ and $m \leq n$. Let H^m denote some model of m-dimensional hyperbolic space and let H^n denote a potentially different model of n-dimensional hyperbolic space. Then, there is an inclusion, i.e. a distance-preserving function $\iota : H^m \to H^n$.

Proof. If $n = m$, we can define ι as an isometry from H^m to H^n whose existence is provided by Remark 2.34. In the following, suppose that $m < n$. We first provide an inclusion for the special case of $\mathbb{D}^m$ and $\mathbb{D}^n$: clearly, the function

$$\iota_D : \mathbb{D}^m \to \mathbb{D}^n, (x_1, \ldots, x_m) \mapsto (x_1, \ldots, x_m, 0, \ldots, 0)$$

is distance-preserving. Using ι_D, we are able to construct an inclusion for the general case. To do so, consider two isometries $\phi : H^m \to \mathbb{D}^m$ and $\varphi : \mathbb{D}^n \to H^n$ whose existence is provided by Remark 2.34. Then, the function $\iota := \varphi \circ \iota_D \circ \phi : H^m \to H^n$ is distance-preserving as a composition of distance-preserving functions, i.e. an inclusion. ∎

2.4.5 Hyperbolic Plane Geometry

We use H^2 to denote any model of 2-dimensional hyperbolic space. It can be shown that each hyperbolic line l of H^2 divides the set $H^2 \setminus l$ into two components. These components are referred to as *open half-planes* in H^n. The hyperbolic line that induced an open half-plane is called its *bounding line*.

It will be useful to explore the isometries taking H^2 to itself. Note that a characterization of these isometries via elementary functions can be accomplished, at least in the 2-dimensional case and for the models we considered [And05]. For our purposes, it is much more relevant how these isometries of the form $\phi : H^2 \to H^2$ relate to certain geometric objects of the model H^2. The following definition formalizes such relationships:

Definition 2.49: Let X denote a metric space and let Y denote a set of certain geometric objects with respect to X, e.g. its geodesics. Then the isometries of X *act transitively* on Y if and only if for each pair of distinct elements $y_1, y_2 \in Y$, there is an isometry $\phi : X \to X$ such that $\phi(y_1) = \phi(y_2)$.

Focusing on $H^2 = \mathbb{D}^2$, the following useful transitivity properties hold:

Theorem 2.50: (1) The isometries of $\mathbb{D}^2$ act transitively on $\mathbb{D}^2$.

(2) The isometries of $\mathbb{D}^2$ act transitively on the set of hyperbolic lines, i.e. geodesics, of $\mathbb{D}^2$.

(3) The isometries of $\mathbb{D}^2$ act transitively on the set of triples (l, r, z) where l is a hyperbolic line, z is a point on l and r is one of the two rays in l determined by z.

(4) The isometries of $\mathbb{D}^2$ act transitively on the set of open half-planes of $\mathbb{D}^2$.

(5) The isometries of $\mathbb{D}^2$ act transitively on the set of tuples (H, p) where H is an open half-plane and p is a point on the bounding line of H.

Anderson provides proofs for the first four statements [And05, p. 62-65]. His proof of Theorem 2.50 (4) easily generalizes to show the statement of Theorem 2.50 (5). Note that Anderson refers to the isometries of $\mathbb{D}^2$ as *Möbius transformations*, a special type of functions that can be shown to characterize the isometries of the form $\phi : \mathbb{D}^2 \to \mathbb{D}^2$. We will only use this properties in the context of $\mathbb{D}^2$, but it can be shown that they hold for any model H^2 of 2-dimensional hyperbolic space. This is essentially a consequence of Remark 2.34.

2.4.5.1 Hyperbolic Trigonometry

Three points $x, y, z \in H^2$ are said to be *hyperbolically collinear* if and only if there exists a hyperbolic line of H^2 that contains all three points. Otherwise, they are called hyperbolically noncollinear.

Definition 2.51: Let $x, y, z \in H^2$ be three hyperbolically noncollinear points. The *hyperbolic triangle* $T(x, y, z)$ with vertices x, y, z is the region of H^2 bounded by the geodesic segments $[x, y]$, $[y, z]$ and $[z, x]$.

We refer to the geodesic segments $[x, y]$, $[y, z]$ and $[z, x]$ as the *sides* of the hyperbolic triangle $T(x, y, z)$. Each side is contained in the triangle. The context will always clarify whether the notation $T(\dots)$ refers to a Euclidean or hyperbolic triangle. For a proof of the following four well-known theorems on hyperbolic triangles, we refer to Ratcliffe [Rat06, p. 79, 80].

Theorem 2.52: If α, β, γ are the angles of a hyperbolic triangle, then

$$\alpha + \beta + \gamma < \pi.$$

Informally speaking, hyperbolic triangles are *slimmer* than Euclidean triangles. This also reflects in their area, as we will soon observe.

Theorem 2.53 (Law of Sines)**:** If α, β, γ are the angles of a hyperbolic triangle and a, b, c are the lengths of the opposite sides, then

$$\frac{\sinh a}{\sin \alpha} = \frac{\sinh b}{\sin \beta} = \frac{\sinh c}{\sin \gamma}.$$

Theorem 2.54 (The First Law of Cosines)**:** If α, β, γ are the angles of a hyperbolic triangle and a, b, c are the lengths of the opposite sides, then

$$\cos \gamma = \frac{\cosh a \cosh b - \cosh c}{\sinh a \sinh b}.$$

Theorem 2.55 (The Second Law of Cosines)**:** If α, β, γ are the angles of a hyperbolic triangle and a, b, c are the lengths of the opposite sides, then

$$\cosh c = \frac{\cos \alpha \cos \beta + \cos \gamma}{\sin \alpha \sin \beta}.$$

Analogous to the definition of congruent Euclidean triangles, we define two hyperbolic triangles T_1 and T_2 in H^2 to be *congruent* if and only if there exists an isometry of H^2 that maps T_1 to T_2. So-called *congruence theorems* provide sufficient conditions for congruence of two hyperbolic triangles. For example, the congruence theorem *side-angle-side*, commonly abbreviated with *SAS*, has the following meaning: if two sides and the included angle of one hyperbolic triangle T_1 are equal to two sides and the included angle of another hyperbolic triangle T_2, then T_1 and T_2 must be congruent. Note that the order is significant: the angle is the one enclosed by the pair of sides. The following theorem lists some congruence theorems for hyperbolic triangles. For a proof, we refer to the website of the university of Glasgow [Sto].

Theorem 2.56 (Hyperbolic Congruence Theorems)**:** The following conditions are sufficient for two hyperbolic triangles to be congruent:

(1) side-angle-side (SAS),

(2) angle-side-angle (ASA),

(3) side-angle-angle (SAA),

(4) side-side-side (SSS),

(5) angle-angle-angle (AAA).

Note that all of the above congruence theorems also have a Euclidean counterpart, apart from the AAA congruence theorem: clearly, scaling a Euclidean triangle preserves its angles, but does not yield a congruent triangle. The next theorem characterizes isosceles hyperbolic triangles, i.e. hyperbolic triangles that have two sides of equal length. The same characterization is known from Euclidean isosceles triangles.

Theorem 2.57: Let T denote a hyperbolic triangle with vertices x, y, z. Then the sides $[x, y]$ and $[y, z]$ are of equal length if and only if the angle at x and the angle at z are equal.

For a proof, we refer again to the website of the university of Glasgow [Sto].

2.4.5.2 Hyperbolic Area

Since the formula of hyperbolic area as noted in Definition 2.43 is model-dependent and unhandy, it is worth to collect explicit formulas for important geometric objects. A formula for the hyperbolic area of a hyperbolic disk with radius s can be found in Ratcliffe's book [Rat06, p. 77]:

Theorem 2.58: Let $B(x, s)$ be the hyperbolic disk centered at a point x of H^2 of radius s. Then,

$$\mathrm{Area}_{\mathrm{H}}(B(x, s)) = 2\pi(\cosh(s) - 1).$$

We define, as in the script of Walkden [Wal19]:

Definition 2.59: Let $z_1, \ldots, z_n \in H^2$. Then the hyperbolic n-gon P with vertices at $z_1, \ldots, z_n$ is the region of H^2 bounded by the geodesic segments $[z_1, z_2], \ldots, [z_{n-1}, z_n], [z_n, z_1]$.

Similar to triangles, the geodesic segments occurring in the above definition are called the *sides* of P and they are subsets of P. One particularity of hyperbolic geometry is that the area of a polygon can be derived solely by the sum of its internal angles, as stated in the following famous theorem, which can also be found in Walkden's script [Wal19]:

Theorem 2.60 (Gauss-Bonnet Theorem for a hyperbolic polygon)**:** Let P be an n-sided hyperbolic polygon with vertices $v_1, \ldots, v_n$ and internal angles $\alpha_1, \ldots, \alpha_n$. Then

$$\mathrm{Area}_{\mathrm{H}}(P) = (n-2)\pi - (\alpha_1 + \cdots + \alpha_n).$$

In particular, the area of a hyperbolic triangle T with internal angles α, β, γ is given by

$$\mathrm{Area}_{\mathrm{H}}(T) = \pi - (\alpha + \beta + \gamma).$$

3 Unit Ball Graphs and Sphericity

We are finally ready to define the core concepts subject to this thesis.

3.1 Unit Ball Graph Representations

For the rest of the section, let $n \in \mathbb{N}$ and use $(X, d) \in \{E^n, H^n\}$ to denote n-dimensional Euclidean space or any model of n-dimensional hyperbolic space.

Definition 3.1 (Unit Ball Graph Representation)**:** Let $G = (V, E)$ be any graph. Let $\tau : V \to X$ be an injective function from the vertices of G to X and let $r > 0$. Then the tuple (τ, r) is a *unit ball graph representation* of G in X with threshold radius r if and only if the following holds:

$$vw \in E \iff d(\tau(v), \tau(w)) < 2r \quad \text{for each pair of distinct vertices } v, w \in V.$$

To understand the idea behind the term *unit ball graph* and why we multiply r with the factor 2 in the above definition, it is worth recalling Lemma 2.28: when considering two balls of the same radius r, this lemma states that

$$d(x, y) < 2r \iff B(x, r) \cap B(y, r) \neq \emptyset.$$

This means that we may use balls to characterize a unit ball graph representation. By defining $x := \tau(v) \in X$ and $y := \tau(w) \in X$ for any pair of distinct vertices $v, w \in V$, the following corollary follows immediately:

Corollary 3.2: Let $G = (V, E)$ be a graph, let $\tau : V \to X$ be an injective function and let $r > 0$. Then the tuple (τ, r) is a unit ball graph representation of G in X with threshold radius r if and only if

$$vw \in E \iff B(\tau(v), r) \cap B(\tau(w), r) \neq \emptyset \quad \text{for each pair of distinct vertices } v, w \in V.$$

To avoid confusion, note that the term *unit* refers to a constant threshold radius r and not necessarily to $r = 1$. The stated characterization of Definition 3.1 is often very useful to construct unit ball graph representations. In order to verify whether a function $\tau : V(G) \to X$ is a unit ball graph representation of some graph G, the initial definition is usually handier.

If a graph $G = (V, E)$ has a unit ball graph representation $(\tau : V \to X, r)$, we also say that G is a *unit ball graph* (UBG) in X with threshold radius r. Throughout the thesis, the case $X \in \{E^2, H^2\}$ will be of particular interest. Since balls in the plane are usually called disks, we speak of Euclidean, respectively hyperbolic, *unit disk graph representations* and *unit disk graphs* (UDGs) in this context.

Remark 3.3: Corollary 3.2 might remind us of intersection graphs, see Definition 2.3. Indeed, there is a strong connection: consider the set A consisting precisely of the balls with center $\tau(v)$ and radius r for each $v \in V$. In particular, we have $A \subseteq \mathcal{P}(X)$. Let $\mathcal{G}(A)$ denote the intersection graph of A. Then, we observe that (τ, r) is a UBG representation of G if and only if the function that identifies vertices with their corresponding balls

$$f : V(G) \to A, v \mapsto B(\tau(v), r)$$

is a graph isomorphism between G and G'.

This also illustrates why we require that a UBG representation $\tau : V \to X$ is injective: if we boil a UBG representation down to some balls positioned in a space X, we still want to be able to recover the represented graph G, at least up to graph isomorphism. If multiple vertices correspond to the same ball, this is not possible because we cannot even count the number of vertices.

Remark 3.4: Another viewpoint is occasionally useful: start with a tuple $(\tau : V \to X, r)$ where V is some finite set, τ is injective and $r > 0$. We then say that (τ, r) *induces* the graph $G = (V, E)$ whose edge set is defined by $vw \in E :\iff d(\tau(v), \tau(w)) < 2r$ for each pair of distinct vertices $v, w \in V$. Clearly, (τ, r) then is a UBG representation of the induced graph G.

How relevant is the threshold radius when dealing with UBG representations? In the Euclidean case, it can be neglected, as the next theorem justifies:

Theorem 3.5: Let G be any graph, $n \in \mathbb{N}$ and $r > 0$. Suppose that G is a UBG in E^n with threshold radius r. Then, for each $r' > 0$, the graph G is also a UBG in E^n with threshold radius r'.

Proof. Let V denote the vertices of G. Choose a UBG representation $(\tau : V \to E^n, r)$ of G. Set $\lambda := \frac{r'}{r}$. Using the fact that $d_E(\lambda x, \lambda y) = \lambda\, d_E(x, y)$ as observed in Lemma 2.7, we obtain

$$d_E(\lambda x, \lambda y) < r' \iff \lambda\, d_E(x, y) < r' \iff d_E(x, y) < r \quad \text{for all } x, y \in E^n.$$

Thus, define the injection $\tau' : V \to E^n$ by $\tau'(v) := \lambda\tau(v)$ for each $v \in V$. Due to the above observation, the tuple (τ', r') is a UBG representation of G in E^n with threshold radius r'. ∎

In this thesis, the following situation will often occur: we start with a UBG representation $(\tau : V \to X, r)$ of G in some metric space X and want to convert it to a UBG representation $(\rho : V \to Y, s)$ of G in another metric space Y. Theorem 3.5 addresses the special case of $X = Y = E^n$ and focuses on the threshold radius. What if X and Y are distinct metric spaces? If a distance-preserving function $\phi : X \to Y$ is available, then $(\phi \circ \tau : V \to Y, r)$ is clearly a UBG representation of G in Y. As an important consequence, the existence of inclusions $\iota : H^m \to H^n$ for $m \leq n$, see Lemma 2.48, implies that a UBG representation of G in some model H^m of m-dimensional hyperbolic space can be converted to a UBG representation of G in some model H^n of n-dimensional hyperbolic space, when keeping the same threshold radius. We can use this for both increasing the dimension of a model and for converting between different models. Similarly, this increase of dimension also applies to UBG representations in Euclidean space, due to Euclidean inclusions, see Example 2.18.

What if X and Y are not related by an isometry or distance-preserving function? This is usually the most interesting case. Similar to above, the idea of composing $\tau : V \to X$ with some special function $f : X \to Y$ is worth further investigation. Note that in this scenario, we are willing to tolerate a change of threshold radius, i.e. do not require $r = s$. Then, the following type of function arises naturally:

Definition 3.6 (Threshold-Preserving Functions)**:** Let (X, d_X) and (Y, d_Y) be metric spaces. A function $f : X \to Y$ is *threshold-preserving* with respect to the thresholds $r, s > 0$ if and only if f is injective and satisfies

$$d_X(x, y) < r \iff d_Y(f(x), f(y)) < s \quad \text{for all } x, y \in X.$$

Note that the order of the thresholds $r, s > 0$ clearly matters. Clearly, any distance-preserving function $f : X \to Y$ is also threshold-preserving when choosing any equal thresholds $r = s > 0$. The scalar multiplication function used in the proof of Theorem 3.5 is an example of a threshold-preserving function with respect to distinct thresholds.

Remark 3.7: If $X \in \{E^n, H^n\}$, we can drop the requirement of $f : X \to Y$ being injective in Definition 3.6, because it follows from the rest of the definition. To see this, consider a function $f : X \to Y$ that satisfies the rest of Definition 3.6 with respect to some fixed thresholds $r, s > 0$. Assume that f is not injective, i.e. there are two distinct points $x, y \in X$ such that $f(x) = f(y) \in Y$. If $X \in \{E^n, H^n\}$, there is a point $z \in X$ such that $d_X(z, x) < r$ and $r < d_X(z, y)$. For example, such a point can be found by considering the unique geodesic of X containing x and y. Then, we obtain $s < d_Y(f(z), f(y)) = d_Y(f(z), f(x)) < s$, a contradiction. So if f satisfies the rest of Definition 3.6, it is guaranteed to be injective.

Note that this argument does not work for any metric space, for example if the distance of any two points is bounded by some constant. Spherical space has this property, we will encounter it in Chapter 7.

The natural application of threshold-preserving functions is shown in the following theorem:

Theorem 3.8: Let $f : X \to Y$ be a threshold-preserving function with respect to the thresholds $2r, 2s > 0$. Let $G = (V, E)$ be a graph with a UBG representation $(\tau : V \to X, r)$ of G in X. Then, $(f \circ \tau : V \to Y, s)$ is a UBG representation of G in Y.

Proof. The function $f \circ \tau$ is injective as a composition of two injective functions. By definition of f, we have

$$d_X(\tau(v), \tau(w)) < 2r \iff d_Y(f \circ \tau(v), f \circ \tau(w)) < 2s \quad \text{for each } v, w \in V.$$

Following the terminology introduced in Remark 3.4, we conclude that the tuples (τ, r) and $(f \circ \tau, s)$ must induce the same graph, i.e. they are both UBG representations of G. ∎

It is worth highlighting that this type of conversion is *graph independent*: it can be applied to any graph, as long as it has a UBG representation in X. Another important property of UBG representations is that they can be restricted to vertex-induced subgraphs:

Remark 3.9: Let $G = (V, E)$ be any graph with a UBG representation $(\tau : V \to X, r)$ in a space X. Consider any vertex subset $V' \subseteq V$. By restriction of τ to the domain V', we obtain the tuple $(\tau|_{V'}, r)$. Let $G' = (V', E')$ denote the induced graph of this tuple. By construction, two vertices $v, w \in V' \subseteq V$ are adjacent with respect to E' if and only if they are adjacent with respect to E. Thus G' is the vertex-induced subgraph of G by the set V' and the tuple $(\tau|_{V'}, r)$ is a UBG representation of G' in the same space X.

3.1.1 Strict UBG Representations

From the perspective of Remark 3.3, we only considered intersection graphs of open balls of a fixed radius r so far. In this thesis, we will not treat the setting where different balls can have different radii, but it is worth considering what happens if one uses only *closed* balls instead of open ones. Due to

$$\overline{B}(x, r) \cap \overline{B}(y, r) \neq \emptyset \iff d(x, y) \leq 2r$$

as mentioned in the context of Lemma 2.28, using closed balls is equivalent to replacing the comparator $<$ with $\leq$ in Definition 3.1. In this case, we speak of *unit closed ball graph representations* and *unit closed ball graphs* (UCBGs). At first glance, UCBG representations do not differ much from UBG representations. For example, they can also be restricted to vertex-induced subgraphs. Indeed, both types of representations are equivalent, as we show in the remainder of this section. However, in practice, it is often easier to construct UCDG representations of certain graphs. As a first result, a UBG representation can be viewed as a UCBG representation when slightly adapting the threshold radius, and vice versa:

Lemma 3.10: Let G be any graph. The following two implications hold:

(1) If G is a UBG in X with threshold radius $r > 0$, then there is $r_0 < r$ such that for any $r' \in [r_0, r)$: G is a UCBG in X with threshold radius r'.

(2) If G is a UCBG in X with threshold radius $r > 0$, then there is $r_0 > r$ such that for any $r' \in (r, r_0]$: G is a UBG in X with threshold radius r.

Proof. To prove the first statement, choose a UBG representation $(\tau : V \to X, r)$ of G. Then,

$$r_0 := \frac{1}{2} \max\{d(\tau(v), \tau(w)) \mid v, w \in V \text{ and } vw \in E\} < r$$

has the desired property: if $r' \in [r_0, r)$, then each tuple (τ, r') is a UCBG representation of G.

Concerning the second statement, choose a UCBG representation $(\tau : V \to X, r)$ of G. Similarly,

$$r_0 := \frac{1}{2} \min\{d(\tau(v), \tau(w)) \mid v, w \in V \text{ and } vw \notin E\} > r$$

has the desired property: if $r' \in (r, r_0]$, then each tuple (τ, r') is a UBG representation of G. ∎

For our purposes, this way of converting between the two types of representations is impractical. For example, suppose that G is a UCBG in X for each threshold radius $r > 0$. Then we would expect that G is also a UBG in X for each threshold radius $r > 0$. Unfortunately, we cannot derive such a statement from Lemma 3.10. To see this, let $r_0 = r_0(r)$ denote the value provided by Theorem 3.10 (2) for each threshold radius $r > 0$. This means that we can be certain that G is a UBG in X for each threshold radius r' contained in the union

$$A := \bigcup_{r>0} (r, r_0(r)].$$

But this set does not necessarily contain every positive real number. For example, $1 > r_0(r) = \frac{1+r}{2} > r$ might hold for each $r < 1$. Then the set A does not contain $r' = 1$.

So we must take another approach. Instead of modifying the threshold radius of a UBG or UCBG representation $(\tau : V \to G, r)$, we must modify the function τ. If done correctly, we can even convert a given UBG or UCBG representation to a *strict* UBG representation, meaning tuples (τ, r) that are simultaneously UBG and UCBG representations.

Some terminology is helpful: suppose a graph $G = (V, E)$ together with a UBG or UCBG representation $(\tau : V \to X, r)$ of G is known from context. If $vw \in E$, we refer to $d(\tau(v), \tau(w))$ as the *length of the edge* vw. Similarly, if $vw \notin E$, we refer to $d(\tau(v), \tau(w))$ as the *length of the non-edge* vw. Inspired by scalar multiplication functions in E^n, we define:

Definition 3.11: Let $a < 0 < b$ and $X_0 \subseteq X$. A function $f_\varepsilon : X_0 \to X$ is *distance-scaling* if and only if for each pair of distinct points $x, y \in X_0$, the following holds:

(1) $d(x, y) = d(f_0(x), f_0(y))$,

(2) $d(x, y) > d(f_\varepsilon(x), f_\varepsilon(y)) > 0$ for each $\varepsilon \in (a, 0)$,

(3) $d(x, y) < d(f_\varepsilon(x), f_\varepsilon(y))$ for each $\varepsilon \in (0, b)$,

(4) the function $g_{x,y} : (a, b) \to \mathbb{R}, \varepsilon \mapsto d(f_\varepsilon(x), f_\varepsilon(y))$ is continuous.

In the above definition, the values $a = -\infty$ and $b = \infty$ are allowed as well. A distance-scaling function f_ε is injective for each ε in the provided interval (a, b). With the help of distance-scaling functions, we can convert UBG representations to strict UBG representations while keeping the same threshold radius:

Lemma 3.12: Let $(\tau : V \to X, r)$ be a UBG representation of $G = (V, E)$ in X. Set $X_0 = \tau(V)$. Suppose that $f_\varepsilon : X_0 \to X$ is distance-scaling function where $\varepsilon \in (a, b)$. Then there is an $\varepsilon_0 \in (0, b)$ such that $R := (f_{\varepsilon_0} \circ \tau, r)$ is a strict UBG representation of G.

Proof. We restrict ourselves to values $\varepsilon_0 > 0$, implicating that every non-edge has length strictly greater than $2r$ with respect to the representation R, due to Lemma 3.11 (3). So it suffices to ensure that we do not lose any edge of G. To do so, define the function

$$g : (a, b) \to \mathbb{R}, \varepsilon \mapsto \max_{vw \in E} g_{\tau(v),\tau(w)}(\varepsilon)$$

where functions of the form $g_{\tau(v),\tau(w)}$ are defined as in Lemma 3.11 (4) and are thereby continuous. As a consequence, the function g is continuous as the maximum of a finite number of continuous functions. Since (τ, r) is a UBG representation, we have

$$g(0) = \max_{vw \in E} d(f_0(\tau(v)), f_0(\tau(w))) = \max_{vw \in E} d(\tau(v), \tau(w)) < 2r.$$

Due to the continuity of g, we may choose $\varepsilon_0 \in (0, b)$ such that $g(\varepsilon_0) < 2r$ also holds. This means that every edge of G has length strictly smaller than $2r$ with respect to the representation R. In summary, R is a strict UBG representation of G with the above choice of $\varepsilon_0 > 0$. ∎

Using a similar argument, we can also convert UCBG representations to strict UBG representations while keeping the same threshold radius:

Lemma 3.13: Let $(\tau : V \to X, r)$ be a UCBG representation of $G = (V, E)$ in X. Set $X_0 = \tau(V)$. Suppose that $f_\varepsilon : X_0 \to X$ is distance-scaling function where $\varepsilon \in (a, b)$. Then there is an $\varepsilon_0 \in (a, 0)$ such that $R := (f_{\varepsilon_0} \circ \tau, r)$ is a strict UBG representation of G.

Proof. We restrict ourselves to values $\varepsilon_0 < 0$, implicating that every edge of G has length strictly less than $2r$ with respect to the representation R, due to Lemma 3.11 (3). So it suffices to ensure that every non-edge of G stays a non-edge with respect to R. To do so, define the function

$$h : (a, b) \to \mathbb{R}, \varepsilon \mapsto \min_{vw \notin E} g_{\tau(v),\tau(w)}(\varepsilon)$$

where functions of the form $g_{\tau(v),\tau(w)}$ are defined as in Lemma 3.11 (4) and are thereby continuous. As a consequence, the function h is continuous as the minimum of a finite number of continuous functions. Since (τ, r) is a UCBG representation, we have

$$g(0) = \min_{vw \notin E} d(f_0(\tau(v)), f_0(\tau(w))) = \min_{vw \notin E} d(\tau(v), \tau(w)) > 2r.$$

Due to the continuity of G, we may choose $\varepsilon \in (a, 0)$ such that $g(\varepsilon_0) > 2r$ also holds. This means that every non-edge of G has length strictly greater than $2r$ with respect to the representation R. In summary, R is a strict UBG representation of G with the above choice of $\varepsilon_0 < 0$. ∎

It remains to provide concrete distance-scaling functions. For $X = E^n$, define

$$f_\varepsilon : E^n \to E^n, x \mapsto (1 + \varepsilon) \cdot x \quad \text{for each } \varepsilon \in (-1, \infty).$$

Due to Lemma 2.7, we conclude that f satisfies the first three properties of a distance-scaling function. It also satisfies the fourth property since the function

$$g_{x,y} : (-1, \infty) \to \mathbb{R}, \varepsilon \mapsto d_E(f_\varepsilon(x), f_\varepsilon(y)) = (1 + \varepsilon)\, d_E(x, y)$$

is clearly continuous for each $x, y \in E^n$. It is also evident that f remains distance-scaling when restricted to any subset $X_0 \subseteq E^n$.

Finding a distance-scaling function in the hyperbolic context is a little trickier. Consider the upper half-space model H^n of n-dimensional hyperbolic space. Definition 2.44 suggests that distances in $\mathbb{U}^n$ become smaller when points are translated upwards. This impression is indeed correct: given any finite set $X_0 \subseteq H^n$, define $b := \min\{x_n \mid x \in X_0\} > 0$. For any $\varepsilon \in (-\infty, b)$, define the function

$$F_\varepsilon : X_0 \to \mathbb{U}^n, (x_1, \ldots, x_n) \mapsto (x_1, \ldots, x_{n-1}, x_n - \varepsilon)$$

that subtracts ε from the last component of the argument $x \in X_0$. By choice of b, we have $F_\varepsilon(x) \in \mathbb{U}^n$ since $(F_\varepsilon(x))_n > 0$. How does F_ε influence the hyperbolic distance d_U? Since F_ε keeps the Euclidean distance of two points invariant, we have

$$d_U(F_\varepsilon(x), F_\varepsilon(y)) = \operatorname{arcosh}\left(1 + \frac{\|x - y\|^2}{2(x_n - \varepsilon)(y_n - \varepsilon)}\right) \quad \text{for each } x, y \in X_0 \subseteq \mathbb{U}^n. \tag{3.1}$$

Treating this as a function of ε for fixed and distinct points $x, y \in \mathbb{U}^n$, we get

$$G_{x,y} : (-\infty, b) \to \mathbb{R}, \varepsilon \mapsto d_U(F_\varepsilon(x), F_\varepsilon(y)).$$

By evaluating Equation (3.1), we conclude that $G_{x,y}$ is continuous as a composition of continuous functions. Furthermore, $G_{x,y}$ is strictly increasing since arcosh is strictly increasing. Combining this with the fact that $G_{x,y}(0) = d_U(x, y)$, we conclude that $F_\varepsilon : X_0 \to \mathbb{U}^n$ where $\varepsilon \in (-\infty, b)$ is indeed a distance-scaling function. The fact that we used the upper half-space model H^n is not a restriction since UBG and UCBG representations in any model of n-dimensional hyperbolic space may be converted to $\mathbb{U}^n$ and back.

The existence of the functions f_ε and F_ε together with Lemma 3.12 and Lemma 3.13 yield the following corollary:

Corollary 3.14: Let $X \in \{E^n, H^n\}$. If a graph G is a UBG or UCBG in X with threshold radius $r > 0$, then G is also a strict UBG in X with the same threshold radius r.

In particular, each UBG is a UCBG with the same threshold radius, and vice versa.

3.2 Sphericity

Given a graph G, studying in which spaces G has a UBG representation provides interesting insights concerning the graph's structure. When dealing with spaces of the type $X \in \{E^n, H^n\}$, the fact that we can always convert a UBG representation in X to a UBG representation in a higher-dimensional version of X suggests that finding UBG representations in low-dimensional spaces provides more structural insights about the represented graphs. This leads to the concept of *sphericity*. We start with the case of $X = E^n$.

Definition 3.15 (Euclidean Sphericity)**:** Let G be any graph. The *Euclidean sphericity* of G, denoted $\mathrm{ESph}(G)$, is the smallest integer $n \in \mathbb{N}$ such that there is a UBG representation $(\tau : V \to E^n, 1)$ of G in n-dimensional Euclidean space.

Note that the threshold radius of 1 in the above definition can be replaced with any other threshold radius $r > 0$ without altering the Euclidean sphericity, because these UBG representations can all be converted to each other while keeping the number of dimensions, as seen in Theorem 3.5. Immediately, the question arises whether Euclidean sphericity is well-defined. Could there be some finite graph G that does not have a UBG representation in each finite-dimensional Euclidean space E^n? Such graphs do not exist, as the following theorem by Maehara shows, by providing a first upper bound of Euclidean sphericity [Mae84b]:

Theorem 3.16: For any graph $G = (V, E)$, we have $\mathrm{ESph}(G) \leq |V|$.

This upper bound is of theoretical relevance, but it is not tight: in the same paper, Maehara proves that

$$\mathrm{ESph}(G) \leq |V| - \omega(G),$$

where $\omega(G) \geq 1$ denotes the clique number of G, i.e. the size of the largest complete subgraph of G.

When defining hyperbolic sphericity, we need to pay attention to the threshold radius. In the Euclidean case, it could be neglected because of Theorem 3.5. As we do not have a hyperbolic version of this theorem, we need to make hyperbolic sphericity depend on the threshold radius used. We will encounter examples which demonstrate that this is indeed necessary, e.g. trees, see Chapter 5.

Definition 3.17 (Hyperbolic Sphericity)**:** Let G be any graph, $s > 0$ and let H^n denote any model of n-dimensional hyperbolic space. The *hyperbolic sphericity* of G with threshold radius s, denoted $\mathrm{HSph}_s(G)$, is the smallest integer $n \in \mathbb{N}$ such that there is a UBG representation $(\rho : V \to H^n, s)$ of G in H^n.

To prove that this definition is well-defined, we need to show that it does not depend on the choice of the model H^n and that $\mathrm{HSph}_s(G)$ is finite for any finite graph G and any threshold radius $s > 0$. The independence of the model H^n is something we already noted in the previous section: if one finds a UBG representation in some model, it can be converted to a UBG representation in any other model, assuming a fixed dimension n and a fixed threshold radius s. This is an important insight because it permits to pick a model that favors specific constructions or geometric intuitions without worrying about the other models. Concerning the second aspect of $\mathrm{HSph}_s(G) \in \mathbb{N}$, we postpone its proof to Section 4.1 where we compare Euclidean and hyperbolic sphericity.

It is worth noting that Euclidean sphericity and hyperbolic sphericity are both graph invariants, i.e. they are invariant under graph isomorphisms. This is because we can freely convert between UBG representation of isomorphic graphs. To see this, let $G = (V, E)$ and $G' = (V', E')$ be isomorphic graphs via a graph isomorphism $f : V' \to V$. If $(\tau : V \to X, r)$ is a UBG representation of G, then $(\tau \circ f : V' \to X, r)$ clearly is a UBG representation of G' in the same space. This even holds for any metric space X.

Remark 3.18: If we defined Euclidean or hyperbolic sphericity not via UBG representations, but via UCBG representations or via strict UBG representations, this would lead to equivalent definitions due to Corollary 3.14.

Concerning vertex-induced subgraphs, an immediate consequence of Remark 3.9 is the following corollary:

Corollary 3.19: Let $G = (V, E)$ be any graph and let G' be a vertex-induced subgraph of G. Then, for each $s > 0$, we have

$$\mathrm{ESph}(G') \leq \mathrm{ESph}(G) \quad \text{and} \quad \mathrm{HSph}_s(G') \leq \mathrm{HSph}_s(G).$$

Such a statement does not hold for subgraphs in general: each graph $G = (V, E)$ is isomorphic to a subgraph of $K_{|V|}$, the complete graph with $|V|$ vertices. But complete graphs have Euclidean and hyperbolic sphericity of 1, as the next example shows.

Example 3.20: If G is a complete graph, then for any $s > 0$ we have that $\mathrm{HSph}_s(G) = \mathrm{ESph}(G) = 1$.

Proof. Let $G := K_n$ and let $V := \{v_0, \dots, v_{n-1}\}$ denote its vertices. We start by providing a UBG representation $(\tau : V \to E^1, s)$ of G in E^1 with threshold radius s. Define $\tau(v_i) := \frac{i}{n}s \in \mathbb{R}$ for each $i \in \{0, \dots, n-1\}$. Then

$$d_E(\tau(v_i), \tau(v_j)) = \frac{|i-j|}{n}s < s < 2s \quad \text{for each pair of distinct vertices } v_i, v_j \in V.$$

Hence (τ, s) has the desired property. $\mathrm{ESph}(G) = 1$ follows because the threshold radius does not matter in the Euclidean case.

To see $\mathrm{HSph}_s(G) = 1$, let H^1 denote any model of n-dimensional hyperbolic space and let $\phi : E^1 \to H^1$ denote an isometry according to Remark 2.47. Then $(\phi \circ \tau : V \to H^1, s)$ must be a UBG representation of G in H^1 with threshold radius s. ∎

Remark 3.21: The above technique easily generalizes to showing that each UBG representation $(\tau : V \to E^1, 1)$ of some graph $G = (V, E)$, not necessarily complete, can be converted to a UBG representation $(\rho : V \to H^1, s)$ of G with any threshold radius $s > 0$. Similarly, with the help of the isometry ϕ^{-1}, we can convert in the other direction. This means that the class of graphs G with $\mathrm{ESph}(G) = 1$ corresponds to the class of graphs G with $\mathrm{HSph}_s(G) = 1$ and that the threshold radius can be neglected in the hyperbolic 1-dimensional case. This graph class is commonly known as *unit interval graphs*. It is already studied extensively, so we may focus on graphs with Euclidean and hyperbolic sphericity greater than one in this thesis.

We proceed with a nontrivial example of such graphs:

Example 3.22: Let $m \in \mathbb{N}, m \geq 2$ and let $G = K_{m(2)}$ be the complete m-partite graph on m sets of size two, i.e. there are $2m$ vertices and every vertex in G is adjacent to every other vertex except for one. Then, for any $s > 0$, we have that $\mathrm{HSph}_s(G) = \mathrm{ESph}(G) = 2$.

Proof. Let $V := \{v_1, \dots, v_m, w_1, \dots, w_m\}$ denote the vertices of G such that v_i is adjacent to each other vertex except w_i and vice versa, for each $i \in \{1, \dots, m\}$. Consider the vertex subset $V' := \{v_1, v_2, w_1, w_2\} \subseteq V$. The vertex-induced subgraph $G|_{V'}$ corresponds to the cycle graph with four vertices C_4. Unit interval graphs are known to be C_4-free, which implies that $\mathrm{ESph}(G) > 1$ and $\mathrm{HSph}_s(G) > 1$ for each $s > 0$.

It remains to show that two dimensions are sufficient. Starting with Euclidean sphericity, we may define an injective function $\tau : V \to S_E(0, 1) \subseteq E^2$ which places the vertices on the unit circle, such that v_i and w_i are placed on antipodal points, i.e. $\tau(v_i) = -\tau(w_i)$ for each $i \in \{1, \dots, m\}$. This means that $d_E(\tau(v_i), \tau(w_i)) = 2$ and $d_E(\tau(v_i), \tau(v')) < 2$ for each $v' \in V \setminus \{w_i\}$. Since the distance must be strictly smaller than 2 for two vertices to be adjacent when using threshold radius $r = 1$, the tuple $(\tau, 1)$ is a UDG representation of G in E^2.

This idea also works in the context of hyperbolic sphericity. We choose to work with the Poincaré disk model $\mathbb{D}^2$. Let $s > 0$ denote the threshold radius that shall be used. Then there is an injective function $\rho : V \to S_D(0, s) \subseteq \mathbb{D}^2$ such that v_i and w_i are placed on antipodal points of the sphere $S_D(0, s)$, i.e. $\rho(v_i) = -\rho(w_i)$ for each $i \in \{1, \dots, m\}$. Similarly to the Euclidean case, $d_D(\rho(v_i), \rho(w_i)) = 2s$ holds for each i because the geodesic segment $[\rho(v_i), \rho(w_i)]$ contains the origin, making the triangle inequality tight, see Corollary 2.27. Let $v' \in V \setminus \{v_i, w_i\}$. We claim that the geodesic segment $a := [\tau(v_i), \tau(v')]$ cannot contain the origin: otherwise, a would be a Euclidean line segment containing the origin, see Section 2.4.1. As such, it could intersect the sphere $S_D(0, s)$ in at most two points, due to Theorem 2.39 and Lemma 2.10. This would imply $\tau(v') = \tau(w_i)$ and thus $v' = w_i$ since we required that τ is injective. However we also required $v' \neq w_i$. This proves $0 \notin a$. With this knowledge, Theorem 2.26 yields that

$$d_D(\tau(v_i), \tau(v')) < d_D(\tau(v_i), 0) + d_D(\tau(v'), 0) = s + s = 2s.$$

We conclude that (ρ, s) is a UDG representation of G in $\mathbb{D}^2$. ∎

4 Comparing Euclidean and Hyperbolic Sphericity

In this chapter, we compare Euclidean and hyperbolic sphericity. We show that hyperbolic sphericity is bounded from above by Euclidean sphericity, if the threshold radius is chosen small enough. If the threshold radius may not depend on the graph, a similar bound holds: in this case, hyperbolic sphericity may exceed Euclidean sphericity, but by at most one. With the help of these results, we show in the last section that hyperbolic sphericity is computable.

4.1 Hyperbolic Sphericity Is Well-Defined

In this section, we focus on the scenario where the threshold may not depend on the graph.

Theorem 4.1: For any graph G and any $s > 0$, we have

$$\mathrm{HSph}_s(G) \leq \mathrm{ESph}(G) + 1.$$

Proof. Let $G = (V, E)$ and $n := \mathrm{ESph}(G) \in \mathbb{N}$. Fix any threshold radius $s > 0$. We will construct a UBG representation $(\rho : V \to \mathbb{U}^{n+1}, s)$ of G in the $(n+1)$-dimensional upper half-space model $\mathbb{U}^{n+1}$ with threshold radius s. Then, the statement of the theorem follows immediately.

Start by setting $s' := 2s$. Using Theorem 3.5, we can choose a UBG representation $(\tau : V \to E^n, r)$ of G in E^n with threshold radius r such that

$$(2r)^2 = 2(\cosh(s') - 1) \iff r = \frac{1}{2}\sqrt{2(\cosh(s') - 1)}$$

since the last term is always greater than zero, due to $\cosh(s') > 1$. We define the function $\rho : V \to \mathbb{U}^{n+1}$ by setting

$$\rho(v)_i := \tau(v)_i \text{ for each } i \in \{1, \ldots, n\} \text{ and } \rho(v)_{n+1} := 1$$

for each vertex $v \in V$. Here, $\rho(v)_i$ denotes the i-th entry of the vector $\rho(v)$. In particular, $\rho(v) \in \mathbb{U}^{n+1}$ holds for each vertex $v \in V$ because the last coordinate is positive. Clearly, ρ inherits injectivity from τ. By construction, the following holds: $d_E(\tau(v), \tau(w)) = \|\tau(v) - \tau(w)\| = \|\rho(v) - \rho(w)\|$. Now, let $v, w \in V$ be any pair of distinct vertices. Then the following equivalence transformation shows that the tuple (ρ, s) has the desired property, see Definition 3.1:

$$\begin{aligned}
vw \in V &\iff d_E(\tau(v), \tau(w)) < 2r \\
&\iff d_E(\tau(v), \tau(w))^2 < (2r)^2 \\
&\iff \|\rho(v) - \rho(w)\|^2 < (2r)^2 \\
&\iff \frac{\|\rho(v) - \rho(w)\|^2}{\rho(v)_{n+1} \cdot \rho(w)_{n+1}} < 2(\cosh(s') - 1) \\
&\iff 1 + \frac{\|\rho(v) - \rho(w)\|^2}{2 \cdot \rho(v)_{n+1} \cdot \rho(w)_{n+1}} < \cosh(s') \\
&\iff \cosh(d_U(\rho(v), \rho(w))) < \cosh(s') \\
&\iff d_U(\rho(v), \rho(w)) < s' = 2s.
\end{aligned}$$

In the last two steps, we used the definition of d_U (see Definition 2.44) and the fact that cosh is a strictly increasing function when restricted to positive real numbers. ∎

Using Theorem 3.16, the following corollary is an immediate consequence of the preceding theorem:

Corollary 4.2: For any graph $G = (V, E)$ and any $s > 0$, we have $\mathrm{HSph}_s(G) \leq |V| + 1$.

This justifies that hyperbolic sphericity is well-defined, in the sense that it is finite if the given graph is finite (no matter the threshold radius used).

4.2 Euclidean UBGs as a Special Case of Hyperbolic UBGs

It is a well-known fact that models of hyperbolic n-dimensional space look *locally* almost like Euclidean n-dimensional space. This leads to the idea that one could copy a Euclidean UBG representation into some model of hyperbolic space and hope that with enough downscaling, the hyperbolic representation induces the same graph as the Euclidean representation. Bläsius et al. show that this works in two dimensions [BFKS23]. In the following, we generalize this statement to any dimension $n \in \mathbb{N}$. To do so, we need an auxiliary lemma. Informally speaking, the lemma states that, given a UBG representation in E^n, the radius of each vertex's ball can be reduced by a small amount without loss of graph edges.

Lemma 4.3: Let $G = (V, E)$ be any graph with a UBG representation $(\tau : V \to E^n, r)$. Then, there is an $r_0 < r$ such that the following holds: if we attribute to each vertex $v \in V$ some arbitrary radius $r'(v) \in (r_0, r]$, then the intersection graph $\mathcal{G}(A)$ of the set of Euclidean balls

$$A := \{B_E(\tau(v), r'(v)) \mid v \in V\}$$

is isomorphic to G by identifying each vertex $v \in V$ with the ball centered at $\tau(v)$ in A. Furthermore, the value r_0 still has the above property for UBG representations $(\phi \circ \tau, r)$ of G that originate by composing τ with an isometry $\phi : E^n \to E^n$.

Proof. Because $G = (V, E)$ is finite, we may define the value

$$\alpha := \max\{d_E(\tau(v), \tau(w)) \mid vw \in E\}. \tag{4.1}$$

Note that we set $\alpha := 0$ in case of $E = \emptyset$. In particular, $\alpha < 2r$ holds. Define $r_0 := \alpha/2 < r$. Clearly, the values α and r_0 are the same if we replace τ with $\phi \circ \tau$ in Equation (4.1), where $\phi : E^n \to E^n$ is any isometry. We proceed by proving that r_0 has the desired property. Let $r' : V \to (r_0, r]$ be any function. We need to verify that the function

$$f : V \to A, v \mapsto B_E(\tau(v), r'(v))$$

is a graph isomorphism from G to the intersection graph $\mathcal{G}(A)$, as claimed above. Clearly, f is a bijection. Let $vw \in E$ be any edge. Using Lemma 2.28, we get

$$vw \in E \implies d_E(\tau(v), \tau(w)) \leq \alpha = 2r_0 < r'(v) + r'(w) \implies f(v) \cap f(w) \neq \emptyset$$

i.e. the intersection graph $\mathcal{G}(A)$ has an edge $f(v)f(w)$.

Conversely, consider an edge $f(v)f(w)$ of $\mathcal{G}(A)$. We need to show that $vw \in E$. Now using the other implication of Lemma 2.28, we get

$$\begin{aligned} f(v) \cap f(w) \neq \emptyset &\implies B_E(\tau(v), r'(v)) \cap B_E(\tau(w), r'(w)) \neq \emptyset \\ &\implies d_E(\tau(v), \tau(w)) < r'(v) + r'(w) \leq 2r. \end{aligned}$$

Hence, $vw \in E$ follows from the fact that (τ, r) is a UBG representation. ∎

Theorem 4.4: For any graph G, there is a value $s_0 > 0$ such that for every $s \in (0, s_0]$, we have

$$\mathrm{HSph}_s(G) \leq \mathrm{ESph}(G).$$

Proof. Let $G = (V, E)$ and $n := \mathrm{ESph}(G) \in \mathbb{N}$. As in the previous section, we work with the upper half-space model $\mathbb{U}^n$ of n-dimensional hyperbolic space. To prove the statement, we provide suitable UBG representations of G in $\mathbb{U}^n$. Start by choosing some UBG representation $(\tau : V \to E^n, 1)$ of G in E^n with threshold radius 1. Define the function

$$c_\Delta : V \to E^n, v \mapsto (\tau(v)_1, \ldots, \tau(v)_{n-1}, \tau(v)_n + \Delta)$$

which adds a fixed value $\Delta \in \mathbb{R}$ to the last component of $\tau(v)$. In particular, c_Δ is the composition of τ with a Euclidean isometry, namely a translation. From Section 3.1 we know that the tuple $(c_\Delta, 1)$ is a UBG representation of G in E^n as well, for each $\Delta \in \mathbb{R}$. In the following, two vertices are of particular interest: let $u \in V$ denote some vertex who is positioned *the highest*, i.e. $\tau(u)_n = \max\{\tau(v)_n \mid v \in V\}$. Similarly, let $w \in V$ denote some vertex who is positioned *the lowest*, i.e. $\tau(w)_n = \min\{\tau(v)_n \mid v \in V\}$. Since G is finite, u and w exist, but are not necessarily unique. Note that u and w are still positioned the highest resp. the lowest with respect to the function c_Δ. In the following, we may assume that $\tau(u)_n \geq \tau(w)_n > 1$, otherwise replace τ with c_Δ for a value of Δ great enough. Under this circumstances, we may restrict ourselves to values of $\Delta \geq 0$. This ensures that for each $v \in V$, the Euclidean ball $B_E(c_\Delta(v), 1)$ is completely contained in $\mathbb{U}^n$.

Recall that hyperbolic balls of the upper half-space model $\mathbb{U}^n$ correspond to Euclidean balls contained in $\mathbb{U}^n$ and vice-versa, see Theorem 2.45. Our goal is to construct a function $\rho_\Delta : V \to \mathbb{U}^n$ and a value $s(\Delta) > 0$ such that the following holds: for each vertex $v \in V$, the hyperbolic ball $B_U(\rho_\Delta(v), s(\Delta))$ shall have the Euclidean center $c_\Delta(v)$ and shall be contained in the Euclidean ball $B_E(c_\Delta(v), 1)$ (that is, its Euclidean radius shall be ≤ 1). Furthermore, in the case of $v = u$, we ask for equality, i.e. $B_U(\rho_\Delta(u), s(\Delta)) = B_E(c_\Delta(u), 1)$. In this context, let $\mathcal{B}_U(\Delta)$ denote the set of hyperbolic balls of the form $B_U(\rho_\Delta(v), s(\Delta))$. Everything we need to fulfill these requirements is provided by the formula of Theorem 2.45: fix any vertex $v \in V$. By calculation, we see that the ball with hyperbolic radius s and Euclidean center $c := c_\Delta(v)$ must have the hyperbolic center

$$a := a_\Delta(v) := \left(c_1, \ldots, c_{n-1}, \frac{c_n}{\cosh s(\Delta)}\right)$$

and the Euclidean radius

$$r := r_\Delta(v) := c_n \cdot \frac{\sinh s(\Delta)}{\cosh s(\Delta)} = c_n \tanh s(\Delta).$$

Both are well-defined since $\cosh x > 0$ for each $x \in \mathbb{R}$. Set $\rho_\Delta(v) := a_\Delta(v)$ for each $v \in V$. We have not yet decided on the value of $s(\Delta)$. Here, the last requirement comes into play:

$$\begin{aligned} &B_U(\rho_\Delta(u), s(\Delta)) = B_E(c_\Delta(u), 1) \\ \iff &r_\Delta(u) = 1 \iff (c_\Delta(u))_n \cdot \tanh s(\Delta) = 1 \iff s(\Delta) = \tanh^{-1}\left(\frac{1}{(c_\Delta(u))_n}\right). \end{aligned}$$

We define $s(\Delta)$ accordingly. Plugging this into the definition of $r_\Delta(v)$, we obtain the simplified formula

$$r_\Delta(v) = \frac{(c_\Delta(v))_n}{(c_\Delta(u))_n} = \frac{\tau(v)_n + \Delta}{\tau(u)_n + \Delta} \quad \text{for each } v \in V. \tag{4.2}$$

Division by zero cannot occur, since we required that $\tau(u)_n > 0$ and $\Delta \geq 0$. We observe that the Euclidean radius $r_\Delta(v)$ is the greatest for $v = u$ because u maximizes c_n by definition. So each vertex $v \neq u$ has $r_\Delta(v) \leq 1$. Hence, by construction,

$$B_U(\rho_\Delta(v), s(\Delta)) \subseteq B_E(c_\Delta(v), 1) \quad \text{for each } v \in V$$

and the two balls have the same Euclidean center, as desired. So all requirements are fulfilled.

Finally, Lemma 4.3 comes into play: we have a UBG representation $(c_\Delta, 1)$ of G in E^n. Let r_0 denote the threshold given by Lemma 4.3. Since c_Δ is the composition of τ with a Euclidean isometry for each Δ, the value r_0 does not depend on Δ. If we can guarantee that every ball in the set $\mathcal{B}_U(\Delta)$ has Euclidean radius greater than r_0, then the intersection graph of $\mathcal{B}_U(\Delta)$ must be isomorphic to G according to the lemma. Indeed, we can guarantee this by choosing Δ great enough, as the following argument shows: by definition of $w \in V$, the smallest occurring Euclidean radius of balls in $\mathcal{B}_U(\Delta)$ is $r_\Delta(w) \in (0, 1]$. Considering Equation (4.2), we may treat $r_\Delta(w)$ as a function of $\Delta \geq 0$. With means of calculus, it can easily be checked that this function is increasing, bounded from above by 1 and satisfies $\lim_{\Delta\to\infty} r_\Delta(w) = 1$ since the values $\tau(u)_n \geq \tau(w)_n > 1$ are constant. We conclude that by choosing Δ great enough, $r_\Delta(w)$ grows arbitrarily close to 1 and will eventually surpass $r_0 < 1$ for some finite value Δ_0. In particular, $r_{\Delta_0}(v) \in (r_0, 1]$ holds for each vertex $v \in V$.

Hence, the intersection graph of $\mathcal{B}_U(\Delta_0)$ is isomorphic to G by identifying each vertex $v \in V$ with the ball having the Euclidean center $c_{\Delta_0}(v)$ resp. the hyperbolic center $\rho_{\Delta_0}(v)$. By construction, this intersection graph corresponds to the tuple $(\rho_{\Delta_0}, s(\Delta_0))$, so we conclude from Remark 3.3 that this tuple is a UBG representation of G in $\mathbb{U}^n$.

By setting $s_0 := s(\Delta_0)$, the last thing we need to show is that every value $s' \in (0, s_0]$ can be used as a hyperbolic threshold radius as well. Note that the function $s(\Delta) = \tanh^{-1}\left(\frac{1}{(c_\Delta(u))_n}\right)$ is continuous and strictly decreasing with respect to $\Delta > 0$. Furthermore, we have $\lim_{\Delta\to\infty} s(\Delta) = 0$. Using the intermediate value theorem, we obtain that $s([\Delta_0, \infty)) = (0, s_0]$. In particular, for every $s' \in (0, s_0]$, there is a $\Delta' \geq \Delta_0$ such that $s' = s(\Delta')$. Since $r_\Delta(w)$ is increasing with respect to Δ, we conclude that $r_{\Delta'}(w) \geq r_{\Delta_0}(w) > r_0$. Thus the tuple $(\rho_{\Delta'}, s')$ is a UBG representation of G as well. ∎

Again using Theorem 3.16, the next corollary follows immediately:

Corollary 4.5: For any graph $G = (V, E)$, there is an $s_0 > 0$ such that for every $s \in (0, s_0]$, we have $\mathrm{HSph}_s(G) \leq |V|$.

In the remainder of this section, we focus on deriving a graph class $\mathcal{G} = \{G_1, G_2, \dots\}$ with the following property: for each $m \in \mathbb{N}$,

$$\mathrm{HSph}_s(G_m) \geq \mathrm{ESph}(G_m) = m \quad \text{for each } s > 0.$$

In particular, such a result demonstrates that Theorem 4.4 cannot be improved in general and that hyperbolic sphericity can get arbitrarily great.

The graph class can be found in one of Maehara's papers on Euclidean sphericity [Mae84b]. For $m = 1$, set $G_1 := P_4$, the path graph of order 4. Suppose now that $m \geq 2$. Let $k := 2^{m+1}$. Observe that the set $\{1, 2, \dots, m+1\} \subseteq \mathbb{N}$ has exactly k distinct subsets, label them $A_1, \dots, A_k$. Define G_m to be the following graph: start with the complete graph with vertex set $\{u_1, \dots, u_k\}$. Each vertex u_i represents the set $A_i \subseteq \{1, \dots, m+1\}$. Next, add $(m+1)$ vertices $v_1, \dots, v_{m+1}$ and add an edge $e = v_j u_i$ if and only if $j \in A_i$ where $j \in \{1, \dots, m+1\}, i \in \{1, \dots, k\}$. Figure 4.1 depicts the graph G_2.

For a proof that $\mathrm{ESph}(G_m) \leq m$, we refer to the original paper [Mae84b]. Here, we will reproduce Maehara's proof that $\mathrm{ESph}(G) \geq m$ and explain how it can be applied to hyperbolic sphericity as well. The idea is based on a set-theoretical observation:

Remark 4.6: Given $P_1, \dots, P_m$ subsets of some set P. Let $P_j^c := P \setminus P_j$ denote the complement of P_j. Assume that, for each $j \in \{1, \dots, m\}$, we have to pick either P_j or its complement and then intersect all the m sets we picked. It is clear that there are 2^m possible ways of constructing such intersections. If each intersection obtained in this way is nonempty, we say that $P_1, \dots, P_m$ are *independent subsets* of P. Formally, this condition can be expressed as

$$|\{T \mid T = T_1 \cap \dots \cap T_m \text{ where each } T_j \text{ is either } P_j \text{ or } P_j^c \text{ and } T \neq \emptyset\}| = 2^m.$$

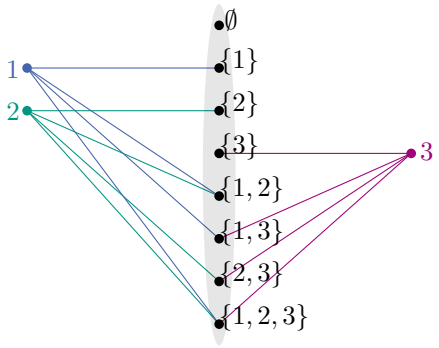

Figure 4.1: The graph G_2. The sets in the middle correspond to the vertices $u_1, \dots, u_8$. For a better overview, we have omitted the edges making these vertices a complete subgraph. The other vertices correspond to v_1, v_2 and v_3.

This concept should not be confused with independent sets in graph theory. Now, consider the particular case of some Euclidean balls $P_j := B_E(x_j, r_j)$ as subsets of Euclidean space E^n. By a result of Rényi et al., there may be at most $(n+1)$ independent balls in E^n [RRS50]. Note that it is not required that the balls have a common radius.

How does this concept relate to the graph family $\mathcal{G}$? Fix any graph $(V, E) := G_m \in \mathcal{G}$. The case $m = 1$ is trivial, so let $m \geq 2$ in the following. By construction, any UBG representation $(\tau : V \to E^n, r)$ of G_m in n-dimensional Euclidean space must contain $(m+1)$ independent balls. To see this, consider the balls $P_j := B_E(\tau(v_j), 2r) \subseteq E^n$ for $j \in \{1, \dots, m+1\}$. Then, for any fixed set A_i like above, the following implications hold:

$$j \in A_i \implies v_j u_i \in E \implies d_E(\tau(v_j), \tau(u_i)) < 2r \implies \tau(u_i) \in P_j$$

and similarly,

$$j \notin A_i \implies v_j u_i \notin E \implies d_E(\tau(v_j), \tau(u_i)) \geq 2r \implies \tau(u_i) \in P_j^c.$$

Combining this, we conclude that the following intersection is nonempty because it contains $\tau(u_i)$:

$$\tau(u_i) \in \left(\bigcap_{j \in A_i} P_j \right) \cap \left(\bigcap_{j \notin A_i} P_j^c \right) \neq \emptyset \quad \text{for each } i \in \{1, \dots, k\}.$$

Since A_i can be any subset of $\{1, \dots, m+1\}$, it follows that $P_1, \dots, P_{m+1}$ are independent subsets of E^n. Since E^n can contain at most $(n+1)$ independent balls, $n+1 \geq m+1$ must hold, i.e. $n \geq m$. In particular, $\mathrm{ESph}(G) \geq m$ by the definition of Euclidean sphericity.

The same argument applies when one assumes a UBG representation (ρ, s) of G_m in the upper half-space model $\mathbb{U}^n$ of n-dimensional hyperbolic space. Recall that assuming a certain model for a hyperbolic UBG representation is not a restriction, since all models are isometric for a fixed dimension. Similarly to above, define $P_j := B_U(\rho(v_j), 2s) \subseteq \mathbb{U}^n$. For the same reason as previously, $P_1, \dots, P_{m+1}$ are independent subsets of $\mathbb{U}^n$. However, it is well known that each P_j can also be interpreted as a Euclidean ball in E^n, see Theorem 2.45. As Euclidean balls, $P_1, \dots, P_{m+1}$ potentially have different radii, but as noted in remark Remark 4.6, this case is covered as well. We summarize:

Theorem 4.7: Let $\mathcal{G} = \{G_1, G_2, \dots\}$ denote the previous graph family. Then, for each $m \in \mathbb{N}$ and each $s > 0$:

$$\mathrm{ESph}(G_m) = m \quad \text{and} \quad \mathrm{HSph}_s(G_m) \geq m.$$

4.3 Computability of Hyperbolic Sphericity

A natural question that arises when considering graph invariants like Euclidean and hyperbolic sphericity is the computational complexity of calculating them algorithmically. Before searching for efficient algorithms, one should ensure that the invariant is computable for every finite graph. For common invariants like *size of largest clique*, this is usually the case and can be accomplished with simple exponential-time brute-force algorithms.

When it comes to geometric graph invariants like sphericity, such brute-force algorithms usually are not evident: if we want to know whether some graph G satisfies $\mathrm{ESph}(G) \leq n$, we cannot simply test every tuple of the form $(\tau : V(G) \to E^n, 1)$ to see if it is a UBG representation of G since there are uncountably many of these tuples. The challenge lies in the nature of the real numbers, so it is a good idea to study the following complexity class that deals with true sentences over $\mathbb{R}$.

Definition 4.8 ($\exists\mathbb{R}$)**:** The *existential theory of the reals* is the set of all true sentences of the form $\exists X \in \mathbb{R}^n : \phi(X)$, where $\phi(X)$ is a quantifier-free formula consisting of polynomial equations and inequalities. We denote the decision problem whether such a sentence is true by ETR and define the complexity class $\exists\mathbb{R}$ to contain all decision problems that polynomial-time reduce to ETR.

The wording of the above definition is taken from Bläsius et al., since their paper will be our guideline in this section [BBDJ23]. For a more technical definition and deeper analysis of this complexity class, we refer to the work of Schaefer and Štefankovič [SS17]. In particular, they state the following result about the class $\exists\mathbb{R}$:

$$\mathsf{NP} \subseteq \exists\mathbb{R} \subseteq \mathsf{PSPACE}.$$

The second part of this statement is due to the dissertation of Canny and can be read in the following way: every decision problem in $\exists\mathbb{R}$ is decidable by a deterministic Turing machine using a polynomial amount of space [Can88].

We focus on hyperbolic sphericity. One can easily verify that all results presented in this section also apply to Euclidean sphericity, using the the same or even simpler arguments. Bläsius et al. show that the problem of recognizing hyperbolic unit disk graphs is $\exists\mathbb{R}$-complete [BBDJ23]. This recognition problem corresponds to the question whether, for a given graph G, there is a threshold radius $s = s(G) > 0$ such that $\mathrm{HSph}_s(G) \leq 2$. The natural extension of this is to define the graph class $\mathsf{HUBG}(n)$ as the set of graphs G for whom there is a threshold radius $s = s(G) > 0$ such that $\mathrm{HSph}_s(G) \leq n$. Accordingly, we define the generalized recognition problem $\textsc{Recog}(\mathsf{HUBG}(n))$ as the problem of testing whether a given graph is part of $\mathsf{HUBG}(n)$. We suspect that $\textsc{Recog}(\mathsf{HUBG}(n))$ is $\exists\mathbb{R}$-complete for any dimension $n \geq 2$, but for our purposes it suffices to show membership of $\exists\mathbb{R}$:

Lemma 4.9: The recognition problem $\textsc{Recog}(\mathsf{HUBG}(n))$ is in $\exists\mathbb{R}$.

Proof. Bläsius et al. have provided a proof for the 2-dimensional case [BBDJ23]. Their proof is based on the hyperboloid model and generalizes well to any dimension n, as shown in the following:

by a result from Erickson, van der Hoog and Miltzow we can prove $\exists\mathbb{R}$-membership by describing a polynomial-time verification algorithm on a real RAM machine [EvM20]. Given a graph $G = (V, E)$ and a function $\rho : V \to \mathbb{F}^n$ where $\rho(v)$ represents the center of an equal-radius ball in the hyperboloid model for each vertex $v \in V$. We may assume that $n < |V|$, otherwise the recognition problem is trivial due to Corollary 4.5. Compute

$$d_{\text{adj}} := \max_{vw \in E} B(\rho(v), \rho(w)) \quad \text{and} \quad d_{\text{non-adj}} := \min_{vw \notin E} B(\rho(v), \rho(w))$$

where B is the bilinear form as in Definition 2.46. Note that $B(x, y)$ can be evaluated using $(n + 1)$ multiplication operations and n summation operations of real numbers, so we can evaluate $B(x, y)$ in time $\mathcal{O}(n)$ on a real RAM machine, for any arguments $x, y \in \mathbb{R}^{n+1}$. We can think of d_{adj} and $d_{\text{non-adj}}$ as distances in the hyperboloid model of hyperbolic space (actually they are the hyperbolic cosine of a distance), but since arcosh = $\cosh^{-1}$ is a strictly increasing function, this view is justified. Now if and only if $d_{\text{adj}} < d_{\text{non-adj}}$, there is a threshold radius s such that the tuple (ρ, s) is a UBG representation of G (choose s such that $d_{\text{adj}} < \cosh(s)/2 \leq d_{\text{non-adj}}$). The algorithm takes $\mathcal{O}(n \cdot |V|^2)$ time. Since we have assumed that $n < |V|$, we can bound the time by $\mathcal{O}(|V|^3)$ which is polynomial in the input time, proving $\exists\mathbb{R}$-membership. ∎

One can also define the graph class $\mathsf{HUBG}(n, s)$ as the set of graphs G that satisfy $\mathrm{HSph}_s(G) \leq n$. In contrast to $\mathsf{HUBG}(\mathrm{n})$, we do not assume an optimal threshold radius in this case, but have to work with a given threshold radius s. For this new graph class, we can also define the corresponding recognition problem $\textsc{Recog}(\mathsf{HUBG}(n, s))$ as the problem of testing whether a given graph is part of $\mathsf{HUBG}(n, s)$. By slight modification of the proof of Lemma 4.9, we derive the similar lemma:

Lemma 4.10: For each $n \in \mathbb{N}$ and $s > 0$, the recognition problem $\textsc{Recog}(\mathsf{HUBG}(n, s))$ is in $\exists\mathbb{R}$.

Proof. We can use the same real RAM verification algorithm as in the proof of Lemma 4.9, with a slight difference: the given function ρ now depends not only on n, but also on s. Also, we do not choose a threshold radius anymore, but verify whether $d_{\text{adj}} < \cosh(s)/2 \leq d_{\text{non-adj}}$ holds. The runtime is still bounded by $\mathcal{O}(|V|^3)$ because each case where $n \geq |V| + 1$ is trivial due to Corollary 4.2, i.e. we only need to run the algorithm if $n \leq |V|$. ∎

Theorem 4.11: Let $\mathcal{G}$ denote the set containing all finite graphs. Consider the two functions

$$f : \mathcal{G} \to \mathbb{N},\ G \mapsto \min\{\mathrm{HSph}_s(G) \mid s > 0\},$$
$$g : \mathcal{G} \times \mathbb{R}_{>0} \to \mathbb{N},\ (G, s) \mapsto \mathrm{HSph}_s(G).$$

They are both computable by a deterministic Turing machine using a polynomial amount of space.

Proof. A simple algorithm $\mathcal{A}$ that calculates $f(G)$ for any graph $G = (V, E)$ is given by the following instructions: start with $n = 1$. Run a PSPACE subroutine A_n that decides whether $G \in \mathsf{HUBG}(\mathrm{n})$. Such a subroutine exists for every $n \in \mathbb{N}$ because the corresponding decision problem is in $\exists\mathbb{R} \subseteq \mathsf{PSPACE}$, see Lemma 4.9. If A_n outputs that G is a member of $\mathsf{HUBG}(\mathrm{n})$, return n. Else, increment n by one and proceed as before. There will be at most $|V|$ iterations because we know that $f(G) \leq |V|$ due to Corollary 4.5. In particular, $\mathcal{A}$ always terminates. Furthermore, $\mathcal{A}$ only uses a polynomial amount of space since this is true for each iteration and after each iteration, we can clear everything from storage expect the variable n.

The same principle can be used to construct an algorithm $\mathcal{B}$ that calculates $g(G, s)$ for any graph $G = (V, E)$ and any $s > 0$. In this case, in the n-th iteration, call a PSPACE subroutine B_n that decides whether $G \in \mathsf{HUBG}(\mathrm{n}, \mathrm{s})$. Such a subroutine exists due to Lemma 4.10 and the fact that $\exists\mathbb{R} \subseteq \mathsf{PSPACE}$. Now, there will be at most $|V| + 1$ iterations because we know that $g(G, s) \leq |V| + 1$ for every $s \in S$ due to Corollary 4.2. The algorithm $\mathcal{B}$ only uses a polynomial amount of space for the same reason that $\mathcal{A}$ does. $\mathcal{A}$ and $\mathcal{B}$ are both deterministic since we can require each subroutine to be deterministic by definition of PSPACE. ∎

5 Trees

In this chapter, we prove that every finite tree is a hyperbolic UDG, if the threshold radius is chosen great enough. In contrast, Euclidean sphericity of trees is unbounded. Let T_m^k denote the k-ary tree of depth m, i.e. every vertex in T_m^k is either a leaf or has k children. Note that a depth of $m = 0$ corresponds to the tree consisting of a single vertex.

5.1 Euclidean Sphericity of Trees

Using a simple volume argument, we can show that the Euclidean sphericity of a tree T_m^k is unbounded in terms of k and m.

Theorem 5.1: If $m, k \in \mathbb{N}$ and $k \geq 2$, then the following lower bound holds:

$$\mathrm{ESph}(T_m^k) \geq \frac{m \log(k)}{\log(2m+1)}.$$

Proof. Let $(V, E) := T_m^k$. Suppose that $(\tau : V \to E^n, 1)$ is a UBG representation of T_m^k in n-dimensional Euclidean space. Let $v_0 \in V$ denote the root vertex. Let $W \subseteq V$ denote the leaves. Note that $|W| = k^m$. For any leaf $w \in W$, there is a unique path $P = (p_0 = v_0, p_1, \ldots, p_{m-1}, p_m = w)$ from the root v_0 to the leaf w. We use this path to approximate $d_E(\tau(v_0), \tau(w))$ by the triangle inequality:

$$d_E(\tau(v_0), \tau(w)) \leq \sum_{i=1}^{m} d_E(\tau(p_{i-1}), \tau(p_i)) < \sum_{i=1}^{m} 2 = 2m$$

We conclude that each ball of the form $B_E(\tau(w), 1)$ for some leaf w is completely contained in the ball $B_E(\tau(v_0), 2m+1)$. Since the leafs W form an independent set of T_m^k, balls of the form $B_E(\tau(w), 1)$ must be pairwise disjoint. The Euclidean volume of a ball of E^n with radius $R > 0$ is given by the formula

$$V_n(R) := \frac{\pi^{\frac{n}{2}}}{\Gamma(\frac{n}{2}+1)} R^n$$

where Γ denotes the gamma-function [Gip14]. Our previous observations now yield the desired inequality, since Euclidean volume is countably additive:

$$k^m \, V_n(1) \leq V_n(2m+1) \implies k^m 1^n \leq (2m+1)^n \implies \log_{2m+1}(k^m) \leq n.$$

Using the change of base formula and the power rule for logarithms, we may rewrite $\log_{2m+1}(k^m)$ as $\frac{m \log(k)}{\log(2m+1)}$ to obtain the desired inequality. ∎

This shows that the Euclidean sphericity of regular trees increases logarithmically with k and almost linearly with the depth m, as long as $k \geq 2$. Note that the inequality can be further refined: we just used the fact that the leafs form an independent set of the tree. This independent set can be increased by adding every second layer of vertices from the bottom to the top. However, the above inequality suffices for our purposes.

5.2 Trees Are Hyperbolic UDGs

In this section, we develop a recursive construction of the k-ary tree T_m^k with an arbitrary depth $m \in \mathbb{N}_0$ as a hyperbolic unit closed disk graph (UCDG) for each $k \geq 1$. The threshold radius depends only on k and not on the number of layers m, i.e. an arbitrary depth is possible for an appropriate, fixed threshold radius. In the following, we let $\mathbb{D}^+ := \mathbb{D}^2 \cap \{z \in \mathbb{C} \mid \mathrm{Im}(z) > 0\}$ denote the upper half of the 2-dimensional Poincaré disk model.

Lemma 5.2: Let g be a hyperbolic line of $\mathbb{D}^2$ that induces the two open half-planes H_1 and H_2. Then, the following implication holds for any two points $x_1, x_2 \in \mathbb{D}^2$:

$$\overline{B}_D(x_i, s/2) \subseteq H_i \text{ for } i \in \{1,2\} \implies d_D(x_1, x_2) > s. \tag{5.1}$$

Proof. Suppose that $x_1, x_2 \in \mathbb{D}^2$ satisfy the condition on the left side of Equation (5.1). Because g separates $\mathbb{D}^2 \setminus g$ into two sides and x_1 and x_2 lie on distinct sides, the geodesic segment $[x_1, x_2]$ joining x_1 to x_2 must intersect g in some point y. Formally, this is known as *plane separation* and it holds for both Euclidean and hyperbolic plane geometry. Since y lies on g, it cannot be contained in $\overline{B}_D(x_i, s/2) \subseteq H_i$ for $i \in \{1,2\}$. As a consequence, we have $d_D(x_i, y) > s/2$ for $i \in \{1,2\}$. Combining this inequality with Corollary 2.27 yields $d_D(x_1, x_2) = d_D(x_1, y) + d_D(x_2, y) > s$. ∎

Lemma 5.3: For any fixed $k \in \mathbb{N}$ and $t \in (0,1)$, consider the points $y_j = t \cdot \exp\left(i\frac{j\pi}{k+1}\right)$ for $j \in \{1, \ldots, k\}$ on the upper half of the Poincaré disk $\mathbb{D}^+$. Due to Lemma 2.40, their hyperbolic distance to the origin is $s > 0$ such that $t = \tanh(s/2)$. Then there is a threshold $s_0(k) > 0$ such that any choice of $s > s_0(k)$ leads to the following constraints being fulfilled simultaneously:

(1) For all distinct $j, l \in \{1, \ldots, k\}$: $d_D(y_j, y_l) > s$.

(2) For each $j \in \{1, \ldots, k\}$, the closed hyperbolic disk of radius $s/2$ centered at y_j is a subset of $\mathbb{D}^+$.

(3) For each $j \in \{1, \ldots, k\}$, there is a hyperbolic line $g_j \subseteq \mathbb{D}^+$ such that $y_j \in g_j$. Each hyperbolic line g_j is induced by a Euclidean circle C_j orthogonal to S^1. Let p_j denote the Euclidean center of C_j. Each hyperbolic line g_j induces an open half-plane $H_j \subseteq \mathbb{D}^+$ as the side of g_j that does not contain the origin.

(4) The closed hyperbolic disk of radius $s/2$ centered at the origin is disjoint from each g_j and thus from each H_j.

(5) The hyperbolic lines g_j are pairwise disjoint and so are their corresponding open half-planes H_j.

(6) For all distinct $j, l \in \{1, \ldots, k\}$, the closed hyperbolic disk of radius $s/2$ centered at y_j is disjoint from g_l and thus from H_l.

Proof. The proof has the following structure: we show for each constraint individually that it is equivalent to some inequality of the form $f(s) < c(k)$ or $f(s) > c(k)$ where $c(k)$ is constant for fixed k. In case of the first form, we then show that the function f satisfies $\lim_{s\to\infty} f(s) < c(k)$. In case of the second form, we show that f satisfies $\lim_{s\to\infty} f(s) > c(k)$. This means we can fulfill all constraints simultaneously by choosing s great enough, i.e. there is a $s_0(k)$ such that each choice $s > s_0(k)$ is valid. We do not calculate $s_0(k)$ explicitly, because there is not much insight in doing so. For the rest of the proof, fix any distinct $j, l \in \{1, \ldots, k\}$. We may assume that $j < l$.

(1) The origin, y_j and y_l form an isosceles hyperbolic triangle. Let α denote the angle at the origin of this triangle. Then $\alpha = \theta(y_j, y_l) = \frac{(l-j)\pi}{k+1}$ holds because hyperbolic angles correspond to Euclidean angles in the Poincaré disk model. Let $x := d_D(y_j, y_l)$ denote the side length of the triangle's side connecting y_j and y_l. We need to ensure that $x > s$ if s is chosen great enough. Using the fact that cosh is strictly increasing for positive arguments and the first law of cosines, see Theorem 2.54, we derive the following equivalence:

$$\begin{aligned} & x > s \\ \iff & \cosh(x) > \cosh(s) \\ \iff & \frac{\cosh^2(s) - \cosh(s)}{\sinh^2(s)} > \frac{\cosh^2(s) - \cosh(x)}{\sinh^2(s)} \\ \iff & f_1(s) := \frac{\cosh^2(s) - \cosh(s)}{\sinh^2(s)} > \cos\alpha. \end{aligned}$$

Since $\alpha \in (0, \pi)$, we have $\cos\alpha < 1$. Note that α does not depend on s. With means of calculus one can verify that $\lim_{s\to\infty} f_1(s) = 1 > \cos\alpha$. Due to the presented equivalence, this implies that the constraint $x > s$ can always be fulfilled by choosing s great enough.

(2) Let $\beta := \frac{j\pi}{k+1}$ denote the angle of y_j with respect to the real line. It is helpful to find a representation of $\overline{B}_D(y_j, s/2)$ as a Euclidean circle with Euclidean center c_E and Euclidean radius r_E. To fulfill the constraint, we then need to ensure that $\mathrm{Im}(c_E) - r_E > 0$. Following the instructions of Remark 2.42, we set

$$\begin{aligned} & \mu_H = d_D(0, y_j) = s,\ r_H = s/2, \\ & b_{1,H} = \frac{1}{2}s,\ b_{2,H} = \frac{3}{2}s,\ b_{1,E} = \tanh\left(\frac{s}{4}\right),\ b_{2,E} = \tanh\left(\frac{3s}{4}\right), \\ & \mu_E = \frac{1}{2}\left(\tanh\left(\frac{s}{4}\right) + \tanh\left(\frac{3s}{4}\right)\right),\ c_E = \mu_E \exp(i\beta), \\ & r_E = \frac{1}{2}\left(\tanh\left(\frac{3s}{4}\right) - \tanh\left(\frac{s}{4}\right)\right). \end{aligned} \tag{5.2}$$

Since $\lim_{x\to\infty} \tanh(x) = 1$, the following two limits hold:

$$\lim_{s\to\infty} \mu_E(s) = 1,\ \lim_{s\to\infty} r_E(s) = 0. \tag{5.3}$$

Now we see that we can always fulfill the constraint by choosing s great enough because

$$\mathrm{Im}(c_E(s)) - r_E(s) > 0 \iff \mu_E(s)\sin(\beta) - r_E(s) > 0 \iff \sin(\beta) > \frac{r_E(s)}{\mu_E(s)} =: f_2(s)$$

and $\lim_{s\to\infty} f_2(s) = 0/1 = 0$.

(3) For $\beta = \frac{j\pi}{k+1}$ like above and some value $\lambda > 1$, define the point p_j and the Euclidean circle C_j via

$$p_j := \lambda \exp(i\beta) \quad \text{and} \quad C_j := S_E(p_j, \lambda - t).$$

The choice of the Euclidean radius ensures that $y_j \in C_j$. For $g_j := C_j \cap \mathbb{D}^2$ to be a hyperbolic line of $\mathbb{D}^2$, the circle C_j must be orthogonal to S_1. This translates to the condition

$$\lambda - t = \sqrt{\lambda^2 - 1} \iff t = \lambda - \sqrt{\lambda^2 - 1} \iff \lambda = \frac{t^2 + 1}{2t}$$

according to Lemma 2.15. Set $\lambda = \lambda(t)$ accordingly in the following. Because of the relationship $t = \tanh(s/2)$ we can treat t and λ as a function of s. This leads to the following two limits:

$$\lim_{s\to\infty} t(s) = 1 = \lim_{s\to\infty} \lambda(s). \tag{5.4}$$

In order to show $g_j \in \mathbb{D}^+$, it suffices that show that $C_j \in \mathbb{D}^+$, i.e. $\operatorname{Im}(p_j) - (\lambda - t) > 0$. This is equivalent to

$$\lambda(s)\sin\beta = \operatorname{Im}(p_j) > \lambda(s) - t(s) \iff \sin\beta > 1 - \frac{t(s)}{\lambda(s)} =: f_3(s).$$

This corresponds to the desired form because β does not depend on s and $\lim_{s\to\infty} f_3(s) = 1 - \frac{1}{1} = 0 < \sin\beta$. Define the open half-plane H_j as the side of g_j that does not contain the origin. Then H_j is the region bounded by C_j, in particular $H_j \subseteq B_E\big(p_j, \lambda - t\big) \subseteq \mathbb{D}^+$.

(4) This result is always fulfilled, no matter the choice of s: by the previous construction, $\overline{B}_D(0,s)$ and g_j intersect in a single point, namely y_j, that has hyperbolic distance s from the origin. Since $\overline{B}_D(0,s/2) \subseteq \overline{B}_D(0,s) \setminus \{y_j\}$, it follows that $\overline{B}_D(0,s/2)$ and g_j must be disjoint.

(5) In order to show that g_j and g_l are disjoint, it suffices to show that their corresponding circles C_j and C_l are disjoint, i.e. we need to show that

$$d_E(p_j, p_l) > 2(\lambda(s) - t(s)). \tag{5.5}$$

By construction, $\alpha = \frac{l-j}{k+1}\pi$ is the angle between p_j and p_l as well. Due to Lemma 2.13 (2), Inequality (5.5) is then equivalent to

$$2\lambda(s)\sin(\alpha/2) > 2(\lambda(s) - t(s)) \iff \sin(\alpha/2) > 1 - \frac{t(s)}{\lambda(s)} = f_3(s).$$

Since we already observed that $\lim_{s\to\infty} f_3(s) = 0$ and since $\sin(\alpha/2) > 0$, we are done.

(6) To check this constraint, it is again useful to take the Euclidean perspective. Like before, use $c := c_E$ and $r := r_E$ to denote the Euclidean center and radius of $\overline{B}_D\big(y_j, s/2\big)$ viewed as a Euclidean closed disk, see Equation (5.2). Similarly, let $\mu := \mu_E = \|c\|$. Recall that C_l refers to the Euclidean circle inducing the hyperbolic line g_l. We know that C_l has the Euclidean center $p := p_l$ and Euclidean radius $r' := \lambda - t$. Furthermore, $\lambda = \|p\|$ is the norm of the center. The angle between the vectors p and c is given by $\alpha = \frac{(l-j)\pi}{k+1}$. With these variables defined, the constraint translates to ensuring that $d_E(c(s), p(s)) > r(s) + r'(s)$ for s great enough. Because of Lemma 2.13 (1), this inequality is equivalent to

$$0 > r(s) + r'(s) - \sqrt{\mu^2(s) + \lambda^2(s) - 2\mu(s)\lambda(s)\cos\alpha} =: f_4(s) \tag{5.6}$$

One could isolate $\cos\alpha$ to obtain the usual form, but it is easier to take limits directly: using the limits we have already stated in Equation (5.3) and Equation (5.4), we obtain that

$$\lim_{s\to\infty} f_4(s) = 0 + 0 - \sqrt{1 + 1 - 2\cos\alpha} = -\sqrt{2(1 - \cos\alpha)} < 0$$

since $\alpha \in (0, \pi)$. In particular, Inequality (5.6) can always be fulfilled by choosing s great enough. ∎

The existence of such a configuration for the case $k = 3$ and an appropriate value of s is illustrated in Figure 5.1. We now benefit from the lemma by deriving the following theorem:

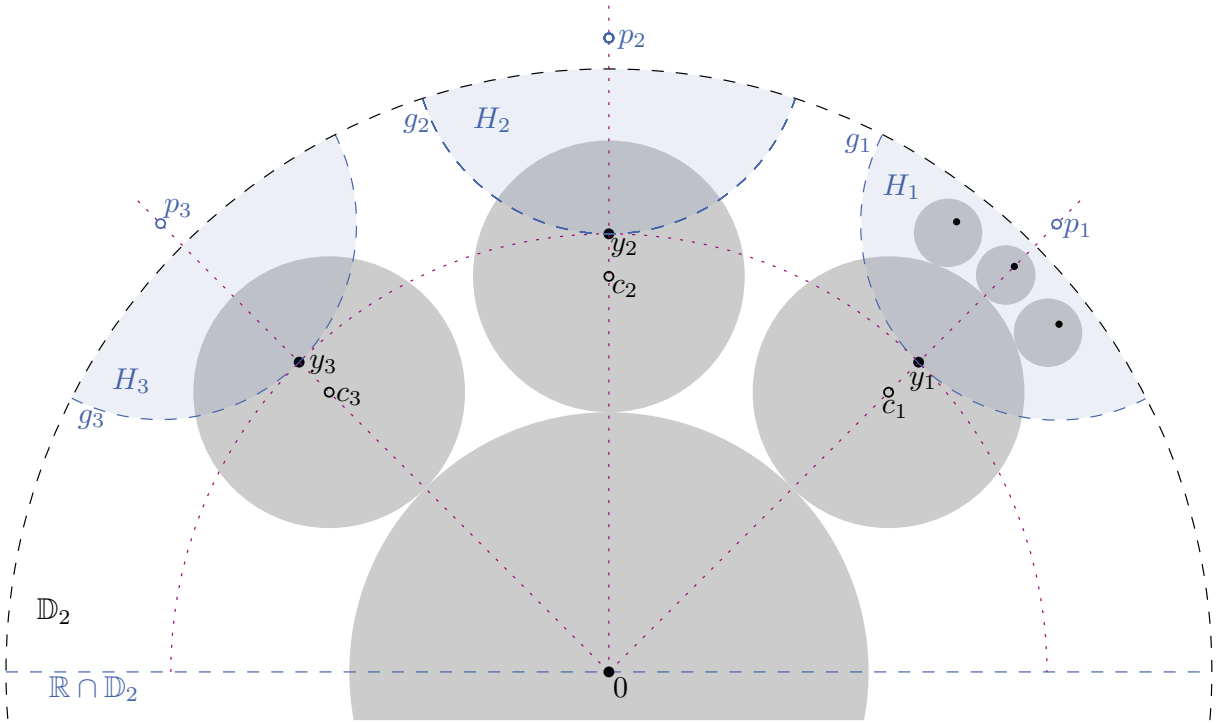

Figure 5.1: Points $y_1, y_2, y_3 \in \mathbb{D}^+$ that fulfill the constraints of Lemma 5.3 for $k = 3$ and serve as a baseline for constructing the 3-ary tree as a hyperbolic UCDG recursively.

Theorem 5.4: For each $k \in \mathbb{N}$, there is a threshold $s_0(k) > 0$ such that the following holds:

$$\mathrm{HSph}_{s/2}(T_m^k) \leq 2 \quad \text{for each } s > s_0(k) \text{ and each } m \in \mathbb{N}_0.$$

Proof. Let s_0 be the function from Lemma 5.3. Fix any $k \in \mathbb{N}$ and any $s > s_0(k)$. Due to Remark 3.18, we may prove the theorem by providing a UCDG representation of T_m^k in $\mathbb{D}^2$ for each $m \in \mathbb{N}_0$. We do so by induction over the tree depth m. The corse idea of this induction is visualized in Figure 5.1: the star-like configuration of closed disks provided by Lemma 5.3 may be copy-pasted into each open half-plane H_j, leading to a tree of depth two. This process can be repeated since new open half-planes appear in every step that are nested inside the other open half-planes. Formally, the statement of induction in terms of $m \in \mathbb{N}_0$ is as follows:

The tree T_m^k has a UCDG representation $(\rho : V(T_m^k) \to \mathbb{D}^2, s/2)$ (where ρ depends on m) with the following two properties:

(1) The root vertex v_0 is placed on the origin, i.e. $\rho(v_0) = 0 \in \mathbb{D}^2$.

(2) For every other vertex $v \in V(T_m^k) \setminus \{v_0\}$, the closed disk of radius $s/2$ centered at $\rho(v)$ is completely contained in $\mathbb{D}^+$.

The start of induction $m = 0$ is trivial since there are no vertices other than the root. We proceed with the induction step from m to $(m+1)$. Let $(V, E) := T_{m+1}^k$. Note that the tree T_{m+1}^k can be seen as a root vertex v_0 with k children $c_1, \ldots, c_k$ such that the subtree rooted at c_j is a copy of T_m^k for each $j \in \{1, \ldots, k\}$. Label each such copy $T_j := (V_j, E_j)$. In conclusion, the vertex set and edge set of T_{m+1}^k are given by

$$V = \{v_0\} \cup \bigcup_{j=1}^{k} V_j \quad \text{and} \quad E = \bigcup_{j=1}^{k} \{v_0 c_j\} \cup \bigcup_{j=1}^{k} E_j. \tag{5.7}$$

Furthermore, let $(\rho_j : V_j \to \mathbb{D}^2, s/2)$ be a UCDG representation of each T_j according to the induction hypothesis. Let y_j and H_j be the points resp. open half-planes from Lemma 5.3 and let g_j denote the hyperbolic line bounding H_j. Observe that $\mathbb{D}^+$ is also an open half-plane. Because of Theorem 2.50 (5) there is an isometry $\varphi_j : \mathbb{D}^2 \to \mathbb{D}^2$ of the Poincaré disk that satisfies $\varphi_j(\mathbb{D}^+) = H_j$ and $\varphi_j(0) = y_j$. Define the function $\rho : V \to \mathbb{D}^2$ by

$$\rho(v_0) := 0 \quad \text{and} \quad \rho(v) := \varphi_j \circ \rho_j(v) \quad \text{for each } j \in \{1, \dots, k\} \text{ and each } v \in V_j \subseteq V.$$

The function ρ is clearly injective because the open half-planes $H_1, \dots, H_k \subseteq \mathbb{D}^+$ are pairwise disjoint, see Lemma 5.3 (5). It remains to verify that $(\rho, s/2)$ is a UCDG representation of T_m^k. Start by showing that each edge $e = vw \in E$ satisfies $d_D(\rho(v), \rho(w)) \leq s$. To do so, distinguish two cases based on Equation (5.7). If $e \in E_j$ for some j, then the desired inequality is satisfied because

$$d_D(\rho(v), \rho(w)) = d_D(\rho_j(v), \rho_j(w)) \leq s$$

since φ_j is an isometry. The remaining edges have the form $e = v_0 c_j$ for some j. In this case, we get

$$d_D(\rho(v_0), \rho(c_j)) = d_D(0, \phi_j \circ \rho_j(c_j)) = d_D(0, \phi_j(0)) = d_D(0, y_j) = s \leq s,$$

having used that $\rho_j(c_j) = 0$ due to the induction hypothesis. In summary, all edges are preserved.

It remains to show that non-edges are preserved as well, i.e. that each pair of distinct vertices $u, u' \in V$ with $uu' \notin E$ satisfies $d_D(\rho(u), \rho(u')) > s$. There are also two types of non-edges: the first one corresponds to $u \in V_j$ and $u' \in V_l$ for some distinct $j, l \in \{1, \dots, k\}$. The second type corresponds to $u = v_0$ and $u' \notin \{c_1, \dots, c_k\}$. Suppose that e' is of the first type. If $e' = uu' = c_j c_l$, we are done since

$$d_D(\rho(c_j), \rho(c_l)) = d_D(\varphi_j(0), \varphi_l(0)) = d_D(y_j, y_l) > s$$

as noted in Lemma 5.3 (1). Otherwise, we may assume that $u \in V_j \setminus \{c_j\}$. Then

$$\overline{B}_D\big(\rho_j(u), s/2\big) \subseteq \mathbb{D}^+ \implies \overline{B}_D(\rho(u), s/2) \subseteq H_j$$

by choice of φ_j. The left side holds due to the second part of the induction hypothesis. Either $u' = c_l$, then

$$\overline{B}_D(\rho(u'), s/2) = \overline{B}_D(y_l, s/2) \subseteq \mathbb{D}^2 \setminus (H_j \cup g_j)$$

due to Lemma 5.3 (6). Or $u' \in V_l \setminus \{c_l\}$, then we similarly have

$$\overline{B}_D(\rho(u'), s/2) \subseteq H_l \subseteq \mathbb{D}^2 \setminus (H_j \cup g_j)$$

due to Lemma 5.3 (5). Either way, we obtain from Lemma 5.2 that $d_D(\rho(u), \rho(u')) > s$ as desired.

Now, suppose that e' is of the second type, i.e. $e' = v_0 u'$ for some $u' \in V_j \setminus \{c_j\}$ and $j \in \{1, \dots, k\}$. As argued above, we have $\overline{B}_D(\rho(u'), s/2) \subseteq H_j$. Due to Lemma 5.3 (4), we have $\overline{B}_D(0, s/2) \subseteq \mathbb{D}^2 \setminus (H_j \cup g_j)$. Now Lemma 5.2 implies that

$$d_D(\rho(v_0), \rho(u')) = d_D(0, \rho(u')) > s$$

as desired.

We conclude that the tuple $(\rho, s/2)$ is indeed a UCBG representation of T_{m+1}^k. Clearly, it satisfies the first property stated in the induction hypothesis. It also satisfies the second property: let $v \in V \setminus \{v_0\}$, i.e. $v \in V_j$ for some j. If $v = c_j$, then the property follows from Lemma 5.3 (2). Otherwise, we have $\overline{B}_D(\rho(v), s/2) \subseteq H_j \subseteq \mathbb{D}^+$ where the first inclusion was argued above and the second inclusion follows from Lemma 5.3 (3). ∎

With this knowledge about k-ary trees, we are able to conclude the following statement about arbitrary trees:

Corollary 5.5: For each finite tree T, there is a threshold $\tilde{s}_0 = \tilde{s}_0(\text{maxdeg}(T))$ only dependent of the maximum degree of T such that the following holds:

$$\text{HSph}_{s/2}(T) \leq 2 \quad \text{for each } s > \tilde{s}_0.$$

Proof. Let $k := \text{maxdeg}(T) - 1$. Then, since T is finite, there is a depth $m \in \mathbb{N}_0$ such that T is a vertex-induced subgraph of T_m^k. As such, T inherits a UDG representation with threshold radius $s/2 > 0$ if one exists for T_m^k with the same threshold radius, see Remark 3.9. The statement now follows from Theorem 5.4 and the therein provided function s_0: define the function $\tilde{s}_0$ by

$$\tilde{s}_0(\text{maxdeg}(T)) := s_0(\text{maxdeg}(T) - 1) = s_0(k). \qquad \blacksquare$$

So all trees are hyperbolic UDGs if the threshold radius is chosen great enough, but which trees are even unit interval graphs? Clearly, trees of the form T_m^k where $m = 0$ or $k = 1$ are path graphs, and therefore unit interval graphs. In turns out that this description captures every tree that is a unit interval graph: every tree T which is not of the above form must contain the star $K_{1,3}$ as a vertex-induced subgraph, but unit interval graphs are known to be $K_{1,3}$-free.

Note that the techniques presented in this section are not necessarily restricted to trees. They could also be helpful in the context of other graphs with a recursive structure.

5.2.1 Hyperbolic Tilings

In this section, we explore hyperbolic tilings and present a way to find each k-ary tree as a vertex-induced subgraph of an appropriate hyperbolic tiling. This can be used to provide a second proof that each k-ary tree is a hyperbolic UDG by showing that tiling graphs are hyperbolic UDGs. Initially, we hoped that this approach would be less technical and more concise than the previous approach by direct construction. This did not turn out to be the case. Nevertheless, tilings appear to be a strong tool in the context of hyperbolic UDGs which could also be applied to other graph families, not just trees.

In the remainder of this chapter, we only need the Poincaré disk $\mathbb{D}^2$ as a model of the hyperbolic plane, i.e. 2-dimensional hyperbolic space, so we will use the two terms synonymously in the context of tilings.

Definition 5.6: Let $p, q \in \mathbb{N}$. A $\{p, q\}$ *regular tiling* of the hyperbolic plane $\mathbb{D}^2$ is a collection $\mathcal{A}$ of convex regular polygons of p vertices, called *tiles*, such that the following holds:

(1) $\mathcal{A}$ is a covering, i.e. $\mathbb{D}^2 = \bigcup_{A \in \mathcal{A}} A$.

(2) $\mathcal{A}$ is a packing, i.e. the interiors of two distinct polygons $A, A' \in \mathcal{A}$ are disjoint.

(3) If two distinct polygons $A, A' \in \mathcal{A}$ touch, they either share one edge and two vertices or just one vertex and nothing else.

(4) At each vertex, exactly q polygons touch.

The notation $\{p, q\}$ is called a *Schläfli symbol* and must not be confused with notation of sets because the order of p and q clearly matters. As an example, the $\{4, 5\}$ regular tiling of $\mathbb{D}^2$ is depicted in Figure 5.2. Regular tilings have a high degree of symmetry, which might not be immediately obvious from Definition 5.6. In particular, they are known to be *vertex-transitive*: for any two vertices of the

regular tiling, there exists a symmetry of the tiling mapping the first vertex isometrically onto the second. For a proof, we refer to *The Tiling Book* by Adams which characterizes regular tilings via the existence of such (and even stronger) symmetries [Ada22, p. 55]. Informally speaking, if we have no coordinate system, the regular tiling looks the same at every vertex.

Concerning the existence of regular hyperbolic tilings, the following theorem states an easy-to-check necessary and sufficient condition. For a proof, we refer again to Adams [Ada22, p. 227].

Theorem 5.7: A $\{p, q\}$ regular tiling of the hyperbolic plane $\mathbb{D}^2$ exists if and only if $\frac{1}{p} + \frac{1}{q} < \frac{1}{2}$.

Suppose fixed values p and q admitting a $\{p, q\}$ regular tiling of the hyperbolic plane $\mathbb{D}^2$. As Adams mentions, each interior angle of each tile must be $\varphi := \frac{2\pi}{q}$ and the side length of the tiles is thereby uniquely determined [Ada22, p. 228]. So any two tiles are congruent.

Remark 5.8: Similarly to regular tilings of the hyperbolic plane, one can define regular tilings of the Euclidean plane E^2 and of the spherical plane (i.e. the surface of the sphere S^2). Then Theorem 5.7 has the following extension: a $\{p, q\}$ regular tiling of the Euclidean plane exists if and only if $\frac{1}{p} + \frac{1}{q} = \frac{1}{2}$. Furthermore, a $\{p, q\}$ regular tiling of the spherical plane exists if and only if $\frac{1}{p} + \frac{1}{q} > \frac{1}{2}$. In particular, for each combination of natural numbers p and q, there exists a $\{p, q\}$ regular tiling in exactly one of the three planes.

Each $\{p, q\}$ regular tiling $\mathcal{A}$ of the hyperbolic plane $\mathbb{D}^2$ induces an infinite graph $G_{p,q} = (V, E)$ by defining

$$V := \{x \in \mathbb{D}^2 \mid x \text{ is a vertex of some tile of } \mathcal{A}\},$$
$$E := \{xy \mid \text{there is a side of a tile of } \mathcal{A} \text{ that has } x \text{ and } y \text{ as its two vertices}\}.$$

The main idea is the find the k-ary tree T_m^k as a vertex-induced subgraph of an appropriate $G_{p,q}$. Similarly to the previous approach, we allow for p and q to depend on k, but not on the depth m of the tree. We may exclude the case $k = 1$, since this makes T_m^k degenerate to the path graph P_{m+1}. This graph clearly is a hyperbolic UDG (even a unit interval graph).

Theorem 5.9: Let $k \in \mathbb{N}, k \geq 2$. The infinite graph of the $\{4, k + 3\}$ regular tiling of $\mathbb{D}^2$ contains T_m^k as a vertex-induced subgraph for any depth $m \in \mathbb{N}_0$.

Proof. Fix any k and m. Let $q = k + 3 \geq 5$. Then, there is a $\{4, q\}$ regular tiling of $\mathbb{D}^2$ because $\frac{1}{4} + \frac{1}{q} < \frac{1}{2}$ (see Theorem 5.7). Following the notation introduced above, we set $G := (V, E) := G_{4,q}$. The $\{4, q\}$ regular tiling of $\mathbb{D}^2$ that induced G also induces a planar drawing of G. In particular, for each vertex v of G, the neighbors of v possess a cyclic order with respect to the planar drawing. Based on this observation, we define the following algorithm to choose a vertex subset $V' \subseteq V$:

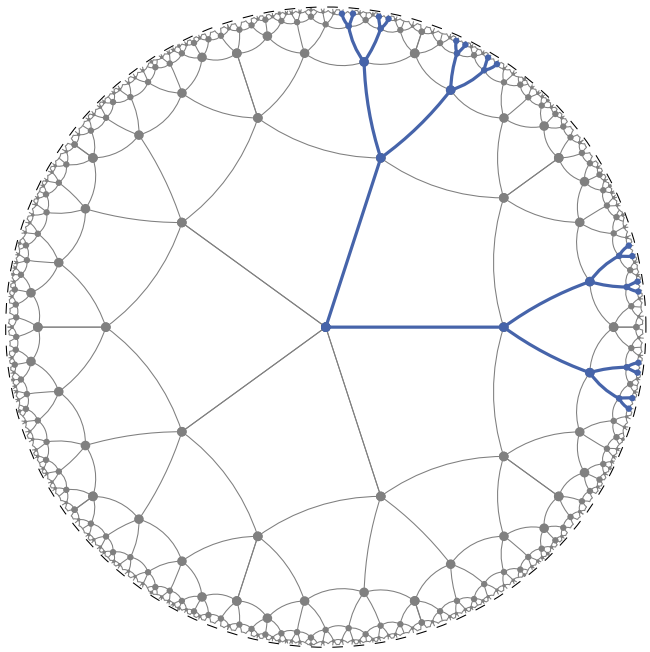

Figure 5.2: Constructing the binary tree of arbitrary depth on the polygons' vertices and sides of the $\{4, 5\}$ regular tiling of $\mathbb{D}^2$.

Algorithm 5.1: Construction of T_m^k via regular tiling.

Input: Tuple of integers (k, m) with $k \geq 2$.
Output: Subset V' of the vertices of $\{4, k+3\}$ regular hyperbolic tiling.

$q \longleftarrow k + 3$
$G = (V, E) \longleftarrow G_{4,q}$
`// pick any root vertex, init parent and layer function`
$V' \longleftarrow \{v_0\}$ where $v_0 \in V$ is chosen arbitrarily
parent : $V' \to V'$, parent$(v_0) \longleftarrow v_0$
layer : $V' \to \mathbb{N}_0$, layer$(v_0) \longleftarrow 0$
for $i = 0$ to $m - 1$ **do** `// expand each vertex in the current layer, similar to BFS`
  forall $v \in V'$ with layer$(v) = i$ **do**
    let $w_1, w_2, \ldots, w_p$ be the neighbors of v in cyclic order where $w_1 = \text{parent}(v)$
    (if $v = v_0$, let w_1 be any neighbor of v_0)
    `// expand the edges` $vw_3, \ldots, vw_{p-1}$ `(skip the vertices` w_2 `and` w_p`)`
    $V' \longleftarrow V' \cup \{w_3, \ldots, w_{p-1}\}$
    parent$(w_j) \longleftarrow v$ for each $j \in \{3, \ldots, p-1\}$
    layer$(w_j) \longleftarrow (i+1)$ for each $j \in \{3, \ldots, p-1\}$
return V'

Figure 5.2 illustrates the algorithm's behavior for $k = 2$. In the following, let G' denote the vertex-induced subgraph $G|_{V'}$. Our goal is to show that G' corresponds to T_m^k. We start by showing that G' is a tree. Recall that a tree is a connected acyclic graph. Also, recall that cycle detection can be done with breath-first-search (BFS): a graph contains a cycle if and only if BFS visits a vertex more than once. It is clear that G' is connected, because the following invariant holds at any step of the algorithm: $G|_{V'}$ is connected. We prove by contradiction that G' is acyclic.

Assume that G' contains a cycle C. The cycle C may contain edges along which we expanded (see line 10 of the algorithm) and edges $uv \in E$ along which we did not expand. This second type of edges $uv \in E$ is characterized by the fact that u and v are not in a parent-relationship (i.e. u is neither the parent of v nor vice-versa). As soon as V' contains two vertices u, v with this property, the vertex-induced subgraph $G|_{V'}$ contains a cycle: this a because a BFS search on this graph encounters the vertex u twice, once via its parent and once via v. This has the following important consequence: by assuming that C is the first cycle that occurs during the algorithm's execution, we are ensured that C contains at most one edge uv along which we did not expand. By identifying the vertex u_1 of C in the lowest layer, we may thus decompose C into the following three components:

(1) a path $u_1, \dots, u_k$ of length k where $\text{parent}(u_{i+1}) = u_i$ for each $i \in \{1, \dots, k-1\}$,

(2) a path $v_1, \dots, v_l$ of length l with $v_1 = u_1$ where $\text{parent}(v_{j+1}) = v_j$ for each $j \in \{1, \dots, l-1\}$,

(3) an edge $u_k v_l$ closing the cycle (potentially an edge along which we did not expand).

Let $g := k+l-1$ denote the length of C. The fact that C is the first cycle that occurs during the algorithm's execution implies that the vertex-induced subgraph

$$G|_W \quad \text{where} \quad W := \{u_1, \dots, u_{k-1}\} \cup \{v_1, \dots, v_{l-1}\} \subseteq V(C)$$

is acyclic, thus simply a path graph. In the $\{4, q\}$ regular tiling, C induces a hyperbolic g-gon P. We continue by studying the internal angles and the hyperbolic area of P. Let α_i denote the internal angle at u_i for each $i \in \{1, \dots, k\}$. Let β_j denote the internal angle at v_j for each $j \in \{1, \dots, l\}$. As the edges of P align with the tiling, each internal angle α_i and β_j must be a positive integer multiple of the constant internal angle of each tile, namely $\varphi := \frac{2\pi}{q}$. Due to the parent-relationship and line 10 of the algorithm, we know that

$$\alpha_i \geq 2\varphi \quad \text{for each } i \in \{2, \dots, k-1\} \quad \text{and} \quad \beta_j \geq 2\varphi \quad \text{for each } i \in \{2, \dots, l-1\}.$$

Only the three angles $\alpha_1 = \beta_1$, α_k and β_l are potentially equal to φ. It follows that the sum of internal angles Γ of the polygon P satisfies

$$\Gamma = \alpha_1 + \cdots + \alpha_k + \beta_2 + \cdots + \beta_l \geq 3\varphi + (g-3)2\varphi = (2g-3)\varphi. \tag{5.8}$$

Next, we want to estimate the hyperbolic area of P. Because the edges of P align with the tiling, it can be decomposed into distinct tiles, each tile being a copy of the regular quadrilateral Q whose every interior angle is $\varphi = \frac{2\pi}{q}$. We argue that P must contain at least $\frac{g-1}{2}$ distinct tiles: we may assume that $k \leq l$. Consider the $(l-1)$ edges of the form $e_j := v_{j-1}v_j$ of C, for $j \in \{2, \dots, l\}$. Two neighbored edges e_j and e_{j+1} cannot be edges of the same tile since each tile's internal angle is φ, but the corresponding angle formed by the two edges is $\beta_j \geq 2\varphi$. Furthermore, any two non-neighbored edges cannot be edges of the same tile since this would contradict the fact that $G|_W$ is a path graph, as noted above. Consequently, each edge e_j corresponds to a unique tile contained in P. Since there are $l-1$ of these edges, P must contain at least $l-1$ distinct tiles. Furthermore, we have

$$l \geq k \implies 2(l-1) \geq k-1+l-1 = g-1 \implies l-1 \geq \frac{g-1}{2},$$

concluding the statement about the number of tiles P contains. In particular, the area of P must be at least $\frac{g-1}{2}$ times the area of each tile. Using Theorem 2.60 we can express the area of P and of each tile, i.e. of the regular quadrilateral Q, using the sum of their internal angles:

$$\text{Area}_{\text{H}}(P) = (g-2)\pi - \Gamma \quad \text{and} \quad \text{Area}_{\text{H}}(Q) = 2\pi - 4\varphi.$$

Linking this with the above observation and Inequality (5.8), we obtain

$$\begin{aligned}
&(g-2)\pi-(2g-3)\varphi \geq (g-2)\pi-\Gamma = \mathrm{Area}_{\mathrm{H}}(P) \geq \frac{g-1}{2}\,\mathrm{Area}_{\mathrm{H}}(Q) = (g-1)(\pi-2\varphi)\\
\Longrightarrow\ &(g-2)\pi-(2g-3)\varphi-(g-1)(\pi-2\varphi) \geq 0\\
\Longrightarrow\ &(g-2)\pi-(2g-3)\frac{2\pi}{q}-(g-1)\pi\left(1-\frac{4}{q}\right) \geq 0\\
\Longrightarrow\ &(g-2)-\frac{2(2g-3)}{q}-(g-1)\left(1-\frac{4}{q}\right) \geq 0\\
\Longrightarrow\ &g-2-\frac{4g}{q}+\frac{6}{q}-(g-1)+\frac{4(g-1)}{q} \geq 0\\
\Longrightarrow\ &-1+\frac{2}{q} \geq 0\\
\Longrightarrow\ &2 \geq q.
\end{aligned}$$

This contradicts our choice of $q = k+3 \geq 5$ made in the beginning, so G' must be acyclic and thus a tree.

Now, it follows easily that G' corresponds to T_m^k: the fact that G' is a tree implies that the algorithm constructs a BFS layering of G'. In particular, the layer of each vertex of G' is set exactly once (if BFS encounters a vertex more than once, this would imply a cycle). As a consequence, each vertex in layer $i < m$ is parent of exactly $q-3 = k$ other vertices in the layer $i+1$ and each vertex in layer m is parent of no other vertex, i.e. a leaf. This is precisely the structure of T_m^k. ∎

It remains to show that the obtained representation of T_m^k is a hyperbolic UCDG representation with some threshold radius. This will follow from Corollary 3.19 if we can show that the corresponding tiling graph $G_{4,q}$ with $q = k+3$ is an infinite hyperbolic UCDG. Doing so is rather technical and will take up the remainder of this chapter. Some preparation is necessary. To start with, we explore properties of individual tiles, i.e. of convex regular quadrilaterals in $\mathbb{D}^2$.

Lemma 5.10: Each convex regular hyperbolic quadrilateral P is circumscribed by a unique hyperbolic circle C whose center is contained in P. Let $x > 0$ denote the length of each side of P and let φ denote the angle at each vertex of P. Then, the hyperbolic radius of C is given by $y > 0$ such that

$$\sinh(y) = \sinh(x)\sin(\varphi/2).$$

Proof. We will make frequent use of the congruency conditions for hyperbolic triangles in this proof, see Theorem 2.56. Refer to Figure 5.3 for a visual aid. Let v_1, v_2, v_3, v_4 denote the vertices of P in circular order. Because of the side-angle-side congruency theorem, the two triangles $T(v_1, v_2, v_3)$ and $T(v_1, v_3, v_4)$ are congruent. Consequently, the geodesic segment $[v_1, v_3]$ partitions the angle φ at v_1 into two equal-sized angles. We conclude that $[v_1, v_3]$ is the angle bisector at v_1. It is also the angle bisector at v_3 for reasons of symmetry. Similarly, $[v_2, v_4]$ must be the angle bisector at v_2 and at v_4. Since v_2 and v_4 lie on distinct sides of $[v_1, v_3]$, we know that $[v_1, v_3]$ and $[v_2, v_4]$ must intersect in a unique point which we call c. Since P is convex, $c \in P$ holds. We observe that c partitions P into four triangles, each having c as one vertex and two adjacent v_i, v_j as the other two. Due to the angle-side-angle congruency theorem, these triangles are congruent. In particular, c has the same distance to every vertex v_i. Let y denote this distance. To prove the statement concerning $\sinh(y)$, consider the triangle $T(v_1, v_2, c)$. Clearly, this triangle has a right angle at c. Taking this into consideration when using the law of sines (Theorem 2.53), we get

$$\frac{\sinh(y)}{\sin(\varphi/2)} = \sinh(x).$$

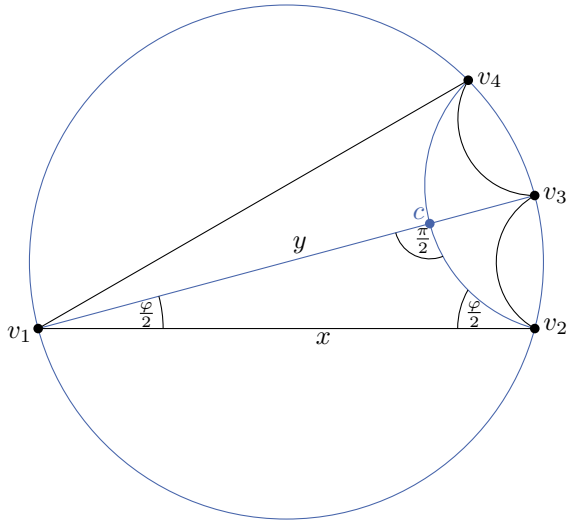

Figure 5.3: Circumcircle of a regular hyperbolic quadrilateral.

To show uniqueness, assume there is another center point c' of P. In particular, c' must also be equidistant from all vertices. By the side-side-side congruency theorem, all four triangles of the form $T(v_i, v_j, c')$ where v_i and v_j are adjacent must be congruent. Consequently, c' must also lie on each angle bisector of P. Since the angle bisectors of P intersect in the unique point c as shown above, we obtain $c' = c$. ∎

In the following, let $\overrightarrow{ab}$ denote the ray starting at point a through point b. We use this notation both in Euclidean and hyperbolic context. Given three noncollinear points in E^2 or three hyperbolically noncollinear points in $\mathbb{D}^2$, we use the notation $\measuredangle(b, a, c)$ to describe the region bounded by the rays $\overrightarrow{ab}$ and $\overrightarrow{ac}$, both rays included. We call this region an *angular sector*.

Lemma 5.11: Consider an isosceles Euclidean triangle with vertices u, v, w such that $d_E(u, v) = d_E(u, w)$. Let r denote this common side length. Let K denote the Euclidean circle centered at u with radius r. Let C denote the unique circumcircle of the triangle provided by Lemma 2.14. Then, every point in $C \cap \measuredangle(v, u, w) \setminus \{v, w\}$ lies outside K.

Proof. Figure 5.4 depicts the situation. We start by naming some variables: let $s = d_E(v, w)$ denote the remaining side length of the triangle. Let c denote the center of C and let R denote the radius of C. Let u' denote the antipodal point of u with respect to C. We observe that c must lie inside the triangle (more precisely, on the angle bisector at u due to the Euclidean side-side-side congruency theorem applied to the Euclidean triangles $T(u, v, c)$ and $T(u, c, w)$). Since u' lies on the ray $\overrightarrow{uc}$, if follows that u' is contained in the angular sector $\measuredangle(v, u, w)$. Wolfram MathWorld provides a formula to calculate the circumradius R [Wei]:

$$R = \frac{r^2 s}{\sqrt{(2r+s)(2r-s)s^2}} = \frac{r^2}{\sqrt{(2r+s)(2r-s)}} = \frac{r^2}{\sqrt{(2r)^2 - s^2}} > \frac{r^2}{2r} = \frac{r}{2}.$$

Consequently, $d_E(u, u') = 2R > r$, i.e. the point $u' \in C \cap \measuredangle(v, u, w)$ lies outside the circle $K = S_E(u, r)$. Because C and K are distinct circles, they may intersect at most in two points due to Lemma 2.14, namely v and w. We know that $C' := C \cap \measuredangle(v, u, w)$ is a (topologically) connected subset of C. Consider the continuous function

$$f : C \to \mathbb{R}, x \mapsto d_E(u, x) - r.$$

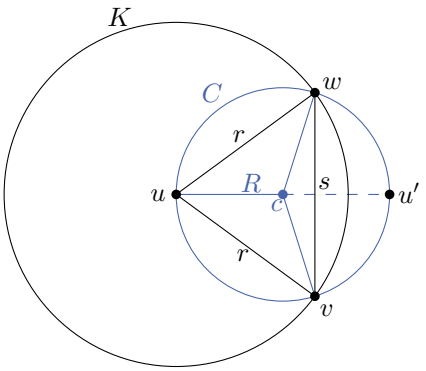

Figure 5.4: Comparing the circumcircle C of an isosceles Euclidean Triangle $T(u, v, w)$ with apex u to the circle K centered at u through v and w.

Clearly, zeros of f correspond to intersection points of C and K, so v and w are the only zeros of f. Furthermore, we know that $f(u') > 0$. Let x denote any point in $C' \setminus \{v, w\}$. We need to show that $f(x) > 0$ to prove the statement. Clearly, $f(x) = 0$ is impossible since $x \notin \{v, w\}$. Since f is continuous and C' is connected, $f(C') \subseteq \mathbb{R}$ is also connected. Similar to the intermediate value theorem, this implies that a point $x \in C'$ with $f(x) < 0$ would lead to a third zero of f since $f(w) = 0$, $f(u') > 0$ and $f(v) = 0$. But we know that f has only two zeros. This means that every point in $C' \setminus \{v, w\}$ must lie outside K. ∎

Lemma 5.12: Let P denote a convex regular hyperbolic quadrilateral with side length x and internal angle φ at each vertex of P. Let v_1, v_2, v_3, v_4 denote the vertices of P in circular order. Then,

$$B_D(v_1, z) \cap \measuredangle(v_2, v_1, v_4) \subseteq P \quad \text{where} \quad \sinh(z) = \sinh(x)\sin(\varphi).$$

Proof. Figure 5.5 depicts the situation. As argued before, P may be decomposed into the two congruent hyperbolic triangles $T_1 := T(v_1, v_2, v_3)$ and $T_2 := T(v_1, v_3, v_4)$. Similarly, the angular sector $A := \measuredangle(v_2, v_1, v_4)$ may be decomposed into the two sub-sectors $A_1 := \measuredangle(v_2, v_1, v_3)$ and $A_2 := \measuredangle(v_3, v_1, v_4)$. If we can show that

$$B_D(v_1, z) \cap A_i \subseteq T_i \quad \text{for each } i \in \{1, 2\},$$

the statement of the theorem follows from the distributive law of set theory. In order to show this inclusion, we may focus on $i = 1$ since the triangles T_1 and T_2 are congruent. The following observation is helpful: let w denote some point on the geodesic segment $[v_2, v_3]$. Then, the points v_1, v_2 and w induce a hyperbolic triangle. Concerning this triangle, the internal angle at v_2 is φ and we know that $d_D(v_1, v_2) = x$. Let α denote the internal angle at w. Using the law of sines (Theorem 2.53), we obtain

$$\sinh(d_D(v_1, w)) = \frac{\sinh(x)\sin(\varphi)}{\sin(\alpha)} \geq \sinh(x)\sin(\varphi) = \sinh(z),$$

thus $d_D(v_1, w) \geq z$ since sinh is strictly increasing.

Now consider any point $p \in B_D(v_1, z) \cap A_1$. Due to the crossbar theorem known from axiomatic geometry, we know that the hyperbolic ray $\overrightarrow{v_1 p}$ must intersect $[v_2, v_3]$ in some point w. Since $d_D(v_1, p) < z$ and $d_D(v_1, w) \geq z$, we may conclude that p lies between v_1 and w on the ray $\overrightarrow{v_1 p}$. Since hyperbolic triangles are convex, $p \in [v_1, w]$ implies that $p \in T_1$, since $v_1, w \in T_1$. This concludes the proof. ∎

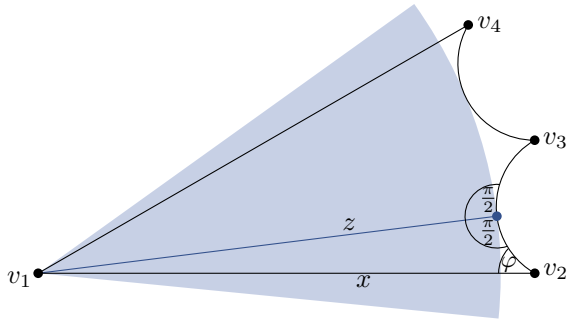

Figure 5.5: Illustration of Lemma 5.12.

Theorem 5.13: Let $q \in \mathbb{N}, q \geq 5$. Let x denote the side length of the tiles used in the regular $\{4, q\}$ tiling of the hyperbolic plane $\mathbb{D}^2$. Then, $G_{4,q}$ is a hyperbolic UCDG with threshold radius $x/2$.

Proof. Let $(V, E) := G_{4,q}$. Let φ denote the internal angle at each vertex of each tile. The vertices V are already a subset of $\mathbb{D}^2$, so it suffices to show that $(\mathrm{id} : V \to \mathbb{D}^2, x/2)$ is a UCDG representation. Since the tiling is vertex-transitive, it suffices to focus on any vertex v_1. Our goal is to show that each vertex $w \in V \setminus \{v_1\}$ such that $d_D(v_1, w) \leq x$ also satisfies $v_1 w \in E$. Consider a tile P in the tiling having v_1 as a vertex. Let v_1, v_2, v_3, v_4 denote the vertices of this tile P in circular order.

Due to the transitivity properties of hyperbolic isometries, in particular Theorem 2.50 (3), we may assume without loss of generality that v_1 is placed on the origin and that v_2 is placed on the positive real axis. Now, the tiling is arranged like the tiling depicted in Figure 5.2. Under this circumstances, the tiling's rotational symmetries allow us to restrict our task to showing that the only vertices $w \in V \setminus \{v_1\}$ contained in the angular sector

$$A := \measuredangle(v_2, v_1, v_4) = \{z \in \mathbb{D}^2 \mid \arg z \in [0, \varphi]\}$$

with $d_D(v_1, w) \leq x$ are precisely v_2 and v_4. In one formula, we equivalently need to show that

$$\overline{B}_D(v_1, x) \cap A \cap (V \setminus \{v_1\}) = \{v_2, v_4\}.$$

Clearly, v_2 and v_4 are contained in the set on the left side, so we may restrict ourselves to showing that

$$\overline{B}_D(v_1, x) \cap A \cap (V \setminus \{v_1, v_2, v_4\}) = \emptyset. \tag{5.9}$$

Let $c \in P \subseteq \mathbb{D}^2$ denote the center of P whose existence is provided by Lemma 5.10. Let $y := d_D(v_1, c)$. We claim that there is no vertex $w \in V$ such that $d_D(w, c) < y$. To prove this claim by contradiction, assume that such a vertex w exists. Then, there must also exist a center c' of some other tile such that $d_D(v_1, c') = d_D(w, c) < y$ because the tiling is vertex-transitive. This means that the ball $B_D(v_1, y)$ must contain the center of some tile. Let z denote the hyperbolic length defined in Lemma 5.12. From

$$\sinh(y) = \sinh(x)\sin(\varphi/2) < \sinh(x)\sin(\varphi) = \sinh(z)$$

we conclude that $y < z$. It is helpful to consider all tiles $P_1, P_2, \ldots, P_q$ of the tiling having v_1 as a vertex. Assuming that $P = P_1$, all of these tiles originate from P by rotation around the origin, i.e. around v_1. Applying Lemma 5.12 to all of these tiles, we get

$$c' \in B_D(v_1, y) \subseteq B_D(v_1, z) \subseteq P_1 \cup P_2 \cup \cdots \cup P_q. \tag{5.10}$$

Because a tile's center is contained in the tile, c' must be the center of some P_j, $j \in \{1, \dots, q\}$. However, since we can obtain P_j by rotating P around v_1, we know that the center of c_j of P_j satisfies $d_D(v_1, c_j) = d_D(v_1, c) = y$. Thus $c' \neq c_j$ for each j, which brings us to a contradiction. We sum up that $B_D(c, y) \cap V = \emptyset$.

Now, we can apply Lemma 5.11 to the vertices v_1, v_2, v_4 that induce an isosceles Euclidean triangle. Since each hyperbolic circle corresponds to a Euclidean circle inside $\mathbb{D}^2$, see Theorem 2.39, the circle K from the lemma must be the boundary of $\overline{B}_D(v_1, x)$ and the circle C from the lemma must be the boundary of $B_D(c, y)$. Hence, the lemma implies that

$$\overline{B}_D(v_1, x) \cap A \setminus \{v_1, v_2, v_4\} \subseteq B_D(c, y) \cap A$$

Intersecting both sides with V, we obtain the following relationship:

$$\overline{B}_D(v_1, x) \cap (A \setminus \{v_1, v_2, v_4\}) \cap V \subseteq B_D(c, y) \cap V \cap A = \emptyset$$

which directly implies that Equation (5.9) holds. This concludes the proof. ∎

We strongly suspect that each graph $G_{p,q}$ induced by a regular hyperbolic tiling is an infinite hyperbolic UCDG. Many elements of the above proof can be generalized. For example, one can easily show that any regular convex hyperbolic p-gon can be circumscribed by a unique circle, similar to Lemma 5.10. The only part where the arguments fail to generalize is in Equation (5.10). By visual inspection of $\{p, q\}$ regular hyperbolic tilings with large p, one observes that the stated subset relationship does not generally hold. Here, a more sophisticated argument must be found to show that, given the center c of some tile P, the vertices closest to c are precisely the vertices of P.

6 Hyperbolic UDGs: Influence of Threshold Radius

In the last chapter, we observed that each tree is a hyperbolic unit disk graph (HUDG) if the threshold radius is chosen great enough. In this chapter, we will develop on this by examining other graphs, paying particular attention to how their hyperbolic sphericity changes with respect to the threshold radius used.

6.1 Star Graphs

Clearly, every star graph $K_{1,n}$ can be viewed as the n-ary tree of depth 1. In particular, every star graph is a HUDG, if the threshold radius is chosen great enough in terms of n. A well-known fact about the Euclidean sphericity of stars is [Mae84b]:

$$\mathrm{ESph}(K_{1,m}) = 1 \text{ for } m \in \{1,2\}, \quad \mathrm{ESph}(K_{1,n}) = 2 \text{ for } n \in \{3,4,5\} \quad \text{and} \quad \mathrm{ESph}(K_{1,6}) = 3. \tag{6.1}$$

Since models of the hyperbolic plane look locally almost Euclidean, one could suspect that $\mathrm{HSph}_s(K_{1,6}) = 2$ if and only if s is great enough, in the sense that $s \approx 0$ would lead to a higher hyperbolic sphericity. A little surprisingly, this is not the case, due to the word *almost* in the above statement. In fact, we have

Lemma 6.1: $\mathrm{HSph}_s(K_{1,6}) = 2$ for each $s > 0$.

Proof. It is clear that $K_{1,6}$ is not a unit interval graph (see Remark 3.21), since $\mathrm{ESph}(K_{1,6}) = 3 > 1$. To show that $\mathrm{HSph}_s(G) \leq 2$, let $v \in V(K_{1,6})$ denote the central vertex of $K_{1,6}$ and let $w_0, \ldots, w_5$ denote the other vertices. Working with the Poincaré disk model $\mathbb{D}^2$, we claim the following: given any threshold radius $s > 0$, the tuple $(\rho : V(K_{1,6}) \to \mathbb{D}^2, s)$ defined by

$$\rho(v) := 0 \quad \text{and} \quad \rho(w_k) := \tanh(s) \exp\left(ik\frac{\pi}{3}\right) \quad \text{for each } \; k \in \{0, \ldots, 5\}$$

is a UCDG representation of $K_{1,6}$ in $\mathbb{D}^2$. To verify this, first note that $d_D(\rho(v), \rho(w_k)) = 2s$ holds for each k because of $\|\rho(w_k)\| = \tanh(s)$, see Lemma 2.40. So all edges are respected by (ρ, s). It remains to show that $d_D(\rho(w), \rho(w')) > 2s$ for each pair of distinct vertices $w, w' \in \{w_0, \ldots, w_k\}$. If the hyperbolic angle between $\rho(w)$ and $\rho(w')$ is π, then the geodesic segment $[w, w']$ must contain the origin. Hence, we have $d_D(\rho(w), \rho(w')) = 4s > 2s$ in this case due to Corollary 2.27. If the angle between w and w' is less then π, then $0 = \rho(v), \rho(w)$ and $\rho(w')$ must form an isosceles hyperbolic triangle. Let $\alpha < \pi$ denote the angle at the origin. By construction, $\alpha \in \{\frac{\pi}{3}, \frac{2\pi}{3}\}$ holds. Because the triangle is isosceles, the two other angles must be equal as noted in Theorem 2.57. Denote them with β. From the fact that the angle sum in a hyperbolic triangle is always less than π, see Theorem 2.52, we conclude that $\beta < \frac{\pi}{3}$. As a consequence, $\sin\alpha > \sin\beta$ due to

$$\sin\alpha = \sin\left(\frac{2\pi}{3}\right) = \sin\left(\frac{\pi}{3}\right) > \sin\beta > 0$$

where we used the fact that sin is strictly increasing on $[0, \frac{\pi}{3}]$. Finally, let $x := d_D(w, w')$ denote the length of the triangle's side opposite the angle α. Using the hyperbolic law of sines (Theorem 2.53) and the fact that $\frac{\sin\alpha}{\sin\beta} > 1$, we get

$$\begin{aligned} & \frac{\sinh(2s)}{\sin(\beta)} = \frac{\sinh(x)}{\sin(\alpha)} \\ \implies & \frac{\sin(\alpha)}{\sin(\beta)} \sinh(2s) = \sinh(x) \\ \implies & \sinh(2s) < \sinh(x) \\ \implies & 2s < x \end{aligned}$$

because sinh is strictly increasing. In summary, the tuple (ρ, s) is a UCDG representation of $K_{1,6}$. Due to Remark 3.18, this concludes the proof. ∎

Equation (6.1) suggests that at least $K_{1,7}$ does not have a HUDG representation in $\mathbb{D}^2$ with small threshold radii. Proving this formally is a little tedious. We will use the Mathematica software as a support for technical calculus statements, like evaluation of limits and checking function properties [Inc].

Lemma 6.2: There is a threshold $s_0 > 0$ such that for each $s \in [0, s_0)$,

$$\mathrm{HSph}_s(K_{1,7}) > 2.$$

Proof. Without loss of generality, we may again work with the Poincaré disk model. Assume a UDG representation $(\rho : V(K_{1,7}) \to \mathbb{D}^2, s)$ of $K_{1,7}$. We show that this implies certain conditions concerning the threshold radius s. Let $v \in V(K_{1,7})$ denote the central vertex and let $w_0, \ldots, w_6$ denote the other vertices. As usually, we may assume that $\rho(v) = 0 \in \mathbb{D}^2$, otherwise compose ρ with an isometry of $\mathbb{D}^2$ that maps $\rho(v)$ to the origin (such an isometry exists due to Theorem 2.50 (1)). Let $q_k = \rho(w_k) \in \mathbb{D}^2$ denote the image of each vertex v_k for $k \in \{0, \ldots, 6\}$. By writing each point q_k in polar coordinates, we may assume that $q_0, \ldots, q_6$ are sorted in circular angular order. With respect to this order, let $\alpha_0, \ldots, \alpha_6$ denote the hyperbolic angles between neighbored vectors q_k, q_l. Clearly, we have

$$\sum_{k=0}^{6} \alpha_k = 2\pi.$$

It follows that at least one angle α_k satisfies $\alpha_k \leq \frac{2\pi}{7}$. Because the order is circular, we may assume that this angle corresponds to the vectors q_0 and q_1. Together with the origin, q_0 and q_1 induce a hyperbolic triangle. Let $\alpha := \alpha_k$ be the angle at the origin. In particular,

$$\cos\alpha \geq \cos\left(\frac{2\pi}{7}\right)$$

since cos is strictly decreasing on $[0, \pi]$. Because ρ is a UDG representation of $K_{1,7}$ in $\mathbb{D}^2$, the side lengths of the triangle must satisfy

$$x := d_D(0, q_0) < 2s, \quad y := d_D(0, q_1) < 2s \quad \text{and} \quad z := d_D(q_0, q_1) \geq 2s.$$

We may assume that $x \geq y$. We also know that $\cosh x \leq \cosh z$ because cosh is strictly increasing for positive real numbers. Furthermore, the following identity will be useful:

$$\tanh\left(\frac{y}{2}\right) = \frac{\cosh(y) - 1}{\sinh y} \quad \text{for each } y > 0.$$

It can be obtained from basic identities about hyperbolic functions or via Mathematica. Combining the above facts with the first law of cosines, see Theorem 2.54, we derive

$$\begin{aligned}\cos\alpha &= \frac{\cosh x\cosh y-\cosh z}{\sinh x\sinh y}\\ &\leq \frac{\cosh x\cosh y-\cosh x}{\sinh x\sinh y}\\ &= \frac{\cosh x}{\sinh x}\cdot\frac{\cosh(y)-1}{\sinh y}\\ &= \coth x\cdot\tanh\left(\frac{y}{2}\right)\\ &\leq \coth x\cdot\tanh\left(\frac{x}{2}\right) =: f(x)\end{aligned}$$

having used that tanh is strictly increasing in the last step. Clearly, the function f is defined for positive real numbers and is continuous. With the help of Mathematica, we verify that $f:\mathbb{R}_{>0}\to\mathbb{R}$ is strictly increasing. This implies

$$\cos\left(\frac{2\pi}{7}\right)\leq\cos\alpha\leq f(x)<f(2s)=\coth(2s)\tanh(s)=:g(s) \tag{6.2}$$

giving us a necessary condition on s for the existence of a HUDG representation (ρ,s) of $K_{1,7}$. Cleary, the function $g:\mathbb{R}_{>0}\to\mathbb{R}$ is strictly increasing since f is. Thus g has an inverse g^{-1} that is also strictly increasing. While $g(0)$ is undefined, we again use Mathematica to see that

$$\lim_{s\to 0}g(s)=\frac{1}{2}=\cos\left(\frac{2\pi}{6}\right)<\cos\left(\frac{2\pi}{7}\right)$$

We also have $\lim_{s\to\infty}g(s)=1$, so the domain of g^{-1} is given by the open interval $\left(\frac{1}{2},1\right)$. Consequently, the necessary condition on s as stated in Inequality (6.2) is nontrivial because it is violated by every choice of s below the threshold

$$s_0:=g^{-1}\left(\cos\left(\frac{2\pi}{7}\right)\right).$$

■

Clearly, the above proof can be applied to star graphs $K_{1,n}$ with $n\geq 7$ as well.

6.2 Relation Between Unit Disk Graphs and Planarity

In this section, we study the relationship between planar graphs and unit ball graphs in metric spaces with certain properties. We focus mainly on E^2 and $\mathbb{D}^2$, but the results are not restricted to these two spaces. Nevertheless, spaces in this section can generally be imagined as *plain-like*. For this reason, we only speak of unit disk graphs (UDGs) in this section. At first glance, a strong relationship between UDGs and planar graphs does not exist. For example, complete graphs K_n are Euclidean and hyperbolic UDGs (even unit interval graphs), but they are famously nonplanar if $n\geq 5$. On the other hand, star graphs $K_{1,n}$ are always planar, but they are Euclidean UDGs if and only if $n\leq 5$. However, as the main result of this section, a relationship can be established if we take the property of having triangles, i.e. a K_3 subgraph, into account.

To pursuit, a formal definition of planar graphs is necessary, based on graph drawings in the plane $\mathbb{R}^2$.

Definition 6.3: Let $G = (V, E)$ be any graph. A tuple (f, Γ) is an $\mathbb{R}^2$-*drawing* of G if and only of the following holds:

(1) $f : V \to \mathbb{R}^2$ is an injective function.

(2) There is a bijection from E to Γ. Moreover, the element γ_e corresponding to the edge $e = vw$ is an injective, continuous curve $\gamma_e : [a_e, b_e] \to \mathbb{R}^2$ that satisfies

$$\{\gamma_e(a_e), \gamma_e(b_e)\} = \{f(v), f(w)\} \quad \text{and} \quad \gamma_e(t) \notin f(V) \ \text{ for each } a_e < t < b_e.$$

Note that *continuity* in this context is defined via the standard topology on closed real intervals and the standard topology on $\mathbb{R}^2$. Furthermore, we call an $\mathbb{R}^2$-drawing (f, Γ) *planar* if and only if for each pair of distinct edges $e, e' \in E$

$$\gamma_e(t) \neq \gamma_{e'}(t') \quad \text{for all } t \in (a_e, b_e),\ t' \in (a_{e'}, b_{e'}).$$

A *planar graph* is a graph that has at least one planar $\mathbb{R}^2$-drawing. Note that the particular wording of our definition of planarity might alter slightly from usual definitions, but it is equivalent.

For the remainder of this section, let (X, d) be a geodesically convex metric space, see Definition 2.25. Recall that this means that for each pair of distinct points $x, y \in X$, there is a unique geodesic segment $[x, y]$ joining x to y. Call x and y the *endpoints* of the geodesic segment $[x, y]$. Furthermore, this geodesic segment is the image of some geodesic arc $\alpha : [a, b] \to X$. We also require (X, d) to be geodesically complete, as this allows us to use Lemma 2.33 later on. We come to the core definition of this section:

Definition 6.4: Let $G = (V, E)$ be a graph with a UDG representation $R := (\tau : V \to X, r)$ of G in X with threshold radius $r > 0$. Then R induces the *UDG-drawing* $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ of G defined by the following two components:

(1) The *drawn vertices* are given by the set $\mathcal{V} := \{\tau(v) \mid v \in V\} \subseteq X$.

(2) The *drawn edges* $\mathcal{E}$ are given by the set $\mathcal{E} := \{[\tau(v), \tau(w)] \mid vw \in E\}$.

For a fixed representation R, such a UDG-drawing $\mathcal{D}$ exists and is unique because we asked for X to be geodesically convex. Furthermore, we call the UDG-drawing $\mathcal{D}$ *crossing-free* if and only if the following holds: any two distinct geodesic segments $[x_1, x_2], [y_1, y_2] \in \mathcal{E}$ are either disjoint or they intersect in exactly one point p that is a common endpoint of both geodesic segments.

Remark 6.5: How do UDG-drawings and $\mathbb{R}^2$-drawings relate? This question arises if the considered space X is a subset of $\mathbb{R}^2$, e.g. $X = \mathbb{D}^2$. As a technical side note, we should require that an open set $A \subseteq X$ is also an open set in $\mathbb{R}^2$ and that $B \cap X$ is an open set in X for each open set $B \subseteq \mathbb{R}^2$. This then leads to the nice property that a function $\alpha : [a, b] \to X$ is continuous if and only if it is continuous when treating it as a function from $[a, b]$ to $\mathbb{R}^2$. For example, E^2 trivially has this property. Another example of a space with this property is $\mathbb{D}^2$, since $\mathbb{D}^2 \subseteq \mathbb{R}^2$ is an open set in $\mathbb{R}^2$ and since the metric topology of $\mathbb{D}^2$ is the same as the subspace-topology of $\mathbb{D}^2 \subseteq \mathbb{R}^2$ due to Theorem 2.39.

In the following, suppose that the space X satisfies the said property. Given a UDG-drawing $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ corresponding to a UDG representation $(\tau : V \to X, r)$ of some graph G, we attempt to convert it to an $\mathbb{R}^2$-drawing of G: to do so, define the set $\mathcal{E}'$ by converting each geodesic segment $[x, y]$ in $\mathcal{E}$ to some corresponding geodesic arc $\alpha : [a, b] \to X \subseteq \mathbb{R}^2$. Then, one could hope that $(\tau, \mathcal{E}')$ is an $\mathbb{R}^2$-drawing of G because each geodesic arc in $\mathcal{E}'$ can also be viewed as a continuous injection from a closed interval to $\mathbb{R}^2$. There is only one problem: this construction might not satisfy the last part of Definition 6.3 (2). Informally speaking, a vertex might be drawn on an edge connecting two other vertices. However, if we additionally suppose that the UDG-drawing is crossing-free, this issue is impossible. In this case, we obtain a valid $\mathbb{R}^2$-drawing using the above conversion. Furthermore, it follows immediately that this $\mathbb{R}^2$-drawing is planar.

Now let us consider some fixed graph $G = (V, E)$ with a UDG representation $(\tau : V \to X, r)$ and an induced UDG-drawing $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ like in Definition 6.4. We suppose that $\mathcal{D}$ is not crossing-free, so there are at least two distinct elements $[x_1, x_2], [y_1, y_2] \in \mathcal{E}$ that violate the condition of being crossing-free. It is helpful to differentiate two cases.

First of all, suppose that $[x_1, x_2]$ and $[y_1, y_2]$ have a common endpoint, say $x_1 = y_1$. Then $x_2 \neq y_2$, otherwise the two geodesic segments would be identical. So there are three vertices involved here: label them $u, v, w \in V$ such that $x_1 = y_1 = \tau(u)$, $x_2 = \tau(v)$ and $y_2 = \tau(w)$. Because the two geodesic segments cause a crossing, they must intersect in some point $z \neq x_1$. Here, Lemma 2.33 comes into play: it implies that $[x_1, x_2] \subseteq [x_1, y_2]$ or $[x_1, y_2] \subseteq [x_1, x_2]$. In particular, $x_2 \in [x_1, y_2]$ or $y_2 \in [x_1, x_2]$. Either way, all three vertices u, v and w are mapped onto a geodesic segment of length less than $2r$. Corollary 2.27 yields that each pairwise distance of the points $\tau(u), \tau(v), \tau(w) \in X$ must be less than $2r$. Consequently, all three vertices are pairwise adjacent, that is we have found a triangle. The existence of a triangle will be the core observation of the second case as well.

Concerning the second case, suppose that $[x_1, x_2]$ and $[y_1, y_2]$ do not have a common endpoint. Because τ is injective, this means that there are four distinct vertices involved: label them $u_1, u_2, v_1, v_2 \in V$ such that $e := u_1u_2 \in E$, $e' := v_1v_2 \in E$ and $x_i = \tau(u_i), y_i = \tau(v_i)$ for $i \in \{1, 2\}$. Since $\mathcal{D}$ is not crossing-free, there is $z \in [x_1, x_2] \cap [y_1, y_2]$. Due to Corollary 2.27, we obtain $d(x_1, z) + d(x_2, z) = d(x_1, x_2) < 2r$. Because d is a metric, each term is nonnegative and we conclude that at least one of the terms $d(x_1, z)$ and $d(x_2, z)$ is strictly less than r. The same argument applies to the geodesic segment $[y_1, y_2]$. Without loss of generality, we may assume that $d(x_1, z) < r$ and $d(y_1, z) < r$. Using the triangle inequality which holds for any metric space, we obtain $d(x_1, y_1) \leq d(x_1, z) + d(z, y_1) < 2r$. Because the UDG-drawing $\mathcal{D}$ was induced by a unit disk graph representation R, there must be an edge $e'' = u_1v_1 \in E$.

We just observed that crossings in a UDG-drawing enforce the existence of certain graph edges. This leads to the question if a single crossing in a UDG-drawing enforces even more edges or if one is already able to provide a concrete UDG-drawing with a single crossing of the graph consisting of four nodes v_1, v_2, u_1, u_2 and edges v_1v_2, u_1u_2, v_1u_1, thus of the path graph P_4. To answer this question, we derive the system of conditions concerning distances that needs to be satisfiable if P_4 has an UDG-drawing with one crossing:

$$d(x_1, x_2) < 2r,\ d(y_1, y_2) < 2r,\ d(x_1, y_1) < 2r \tag{6.3}$$

$$d(x_1, y_2) \geq 2r,\ d(x_2, y_1) \geq 2r,\ d(x_2, y_2) \geq 2r \tag{6.4}$$

$$d(x_1, x_2) = d(x_1, z) + d(x_2, z),\ d(y_1, y_2) = d(y_1, z) + d(y_2, z) \tag{6.5}$$

Inequality (6.3) decodes the existing edges in P_4. Inequality (6.4) decodes the forbidden edges in P_4. Inequality (6.5) is a consequence of the fact that z has to lie on the geodesic segment $[x_1, x_2]$ and on the geodesic segment $[y_1, y_2]$, see above. Now that we have convinced ourselves that all conditions are necessary, we try to find some contradiction in the system allowing us to conclude that there is no UDG-drawing of P_4 like described. In fact, by assuming without loss of generality that $d(x_1, z) \leq d(y_1, z)$, we get

$$2r > d(y_1, y_2) = d(y_1, z) + d(y_2, z) \geq d(x_1, z) + d(y_2, z) \geq d(x_1, y_2) \geq 2r \tag{6.6}$$

Thus $r > r$, a contradiction no matter the choice of r.

So the next question is what other edges we have to allow in order to obtain a UDG-drawing with at least one crossing of P_4 plus some other edges. If we simply allow all edges, we obtain the complete graph K_4 and a UDG-drawing with one or more crossings is clearly possible, but this might be more than strictly necessary.

Continuing to assume that $d(x_1, z) \leq d(y_1, z)$, we first test what happens if we ask for an edge u_2v_1, thus flip the second condition in Inequality (6.4). Flipping denotes change of the $\geq$-comparator to $<$-comparator and vice versa. We see that adding this edge is still not enough, because we did not use

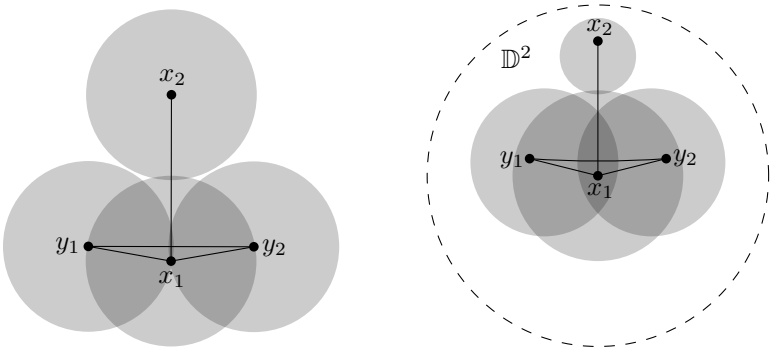

Figure 6.1: Euclidean (left) and hyperbolic (right) UDG-drawing of P_4' with one crossing.

x_2 in the contradiction found in Inequality (6.6). The same applies if we ask for an additional edge u_2v_2 instead, thus flip the third condition in Inequality (6.4). However, the situation is different if we try to add the edge u_1v_2 to P_4. Now, $d(x_1, y_2) \geq 2r$ does not hold anymore, thus the contradiction is no longer valid for the new system of inequalities. Label the obtained graph $P_4' := (V, E)$ where $V := \{u_1, u_2, v_1, v_2\}$, $E := \{u_1u_2, v_1v_2, u_1v_1, u_1v_2\}$. We see immediately that this graph contains a triangle, consisting of the vertices $\{u_1, v_1, v_2\}$. In fact, we are now able to provide a UDG-drawing of P_4' with one crossing, at least for the two cases that interest us the most: $X \in \{E^2, \mathbb{D}^2\}$, see Figure 6.1. The new system might not be satisfiable for *any* metric space X, but we are certain that there are no longer any problems due to the triangle inequality, because it is satisfiable for $X \in \{E^2, \mathbb{D}^2\}$. The key aspect here is that no matter if we add only the edge u_1v_2 or this edge and one or both of the two others, the resulting graph will always contain a triangle. In other words, we have seen that transforming P_4 into the triangle-free cycle graph C_4 is not enough to fix the contradiction noted in Inequality (6.6).

In summary, we observed that in both cases, a graph G that has a UDG-drawing in X with at least one crossing must also have a triangle as a subgraph. Consequently, if a graph G is a UDG in X and also triangle-free, every UDG-drawing of G is crossing-free. If $X \in \{E^2, \mathbb{D}^2\}$, we can choose any such crossing-free UDG-drawing $\mathcal{D}$ (at least one exists) and convert it to a planar $\mathbb{R}^2$-drawing of G, as noted in Remark 6.5. As this result will be helpful in the next section, we capture it in the following theorem:

Theorem 6.6: Let $X \in \{E^2, \mathbb{D}^2\}$ and let G be a graph without triangles that has a UDG representation in X. Then, each UDG-drawing of G is crossing-free and its corresponding $\mathbb{R}^2$-drawing is planar. In particular, G is planar.

Note that *any* UDG representation of G is sufficient: we just have to find a single threshold radius that works. In terms of G being planar, we note that the condition of having a UDG representation is necessary: the complete bipartite graph $K_{3,3}$ is an example of a graph without triangles that is famously nonplanar.

6.3 Great Threshold Radius Potentially Suboptimal

Recall that Euclidean UDGs are HUDGs, at least if the threshold radius s is chosen small enough, see Theorem 4.4. In this section, we construct graphs that are Euclidean UDGs but ceise to be HUDGs when the threshold radius s is chosen too great. In particular, this shows that the inequality of Theorem 4.1 can be tight.

The idea of constructing such graphs is based on the following informal observation with respect to the Poincaré disk model $\mathbb{D}^2$: when representing star graphs or trees as HUDGs, we usually have one large disk centered at the origin. The other disks are small in comparison, so we can arrange many of

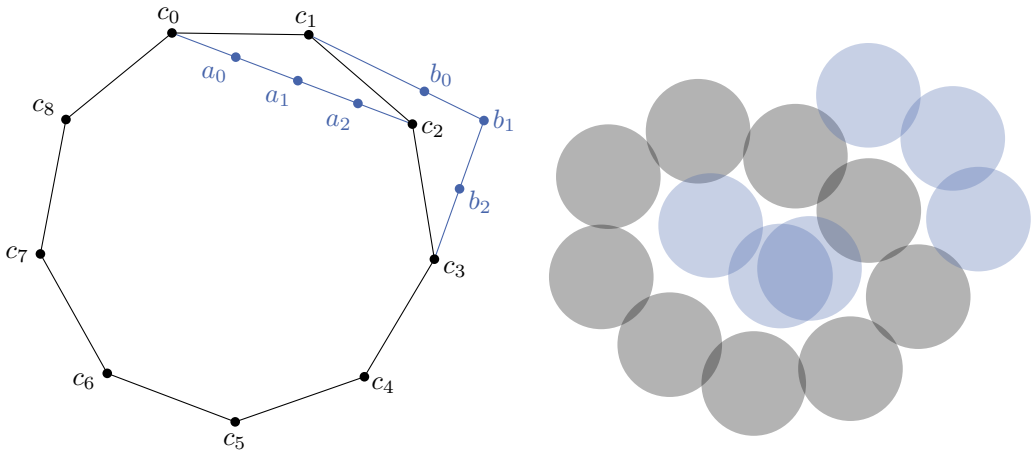

Figure 6.2: The graph $\tilde{C}_9$ (left) with a Euclidean UDG representation (right).

them to touch the large disk but to not touch each other. Note that the terms *large* and *small* refer to the Euclidean perspective. What happens if we want each small disk to touch two other small disks in a circular manner, but not the large disk? In such a situation, the small disks would surround the large disk, and this can only work if there are sufficiently many of them.

We proceed by formalizing the described idea. Consider the graph family obtained by the following construction: start with the cycle graph C_n of nodes $c_0, \dots, c_{n-1}$. Augment it in the following way: add nodes a_0, a_1, a_2 and edges $c_0a_0, a_0a_1, a_1a_2, a_2c_2$. Add further nodes b_0, b_1, b_2 and edges $c_1b_0, b_0b_1, b_1b_2, b_2c_3$. Call the resulting graph $\tilde{C}_n$. Ensure that $\tilde{C}_n$ is a Euclidean UDG, otherwise increase n. Figure 6.2 depicts the graph $\tilde{C}_9$ and a Euclidean UDG representation of $\tilde{C}_9$. By visual inspection, we suspect that $n = 9$ is the smallest n that still allows a Euclidean UDG representation. More importantly, we observe visually that choosing $n \geq 9$ ensures that a Euclidean UDG representation exists. We continue by showing that this graph family has the desired property:

Theorem 6.7: Let $n \in \mathbb{N}, n \geq 9$ and $s > \operatorname{arcosh}\left(\frac{n-2}{2} + 1\right)$. Then, $\mathrm{HSph}_s(\tilde{C}_n) > 2 = \mathrm{ESph}(\tilde{C}_n)$.

Proof. Fix any graph $(V, E) := \tilde{C}_n$ with $n \geq 9$. The fact that $\mathrm{ESph}(\tilde{C}_n) = 2$ was argued above. We prove the rest of the statement by contraposition: let $s > 0$ such that $\mathrm{HSph}_s(\tilde{C}_n) \leq 2$ holds. In this case, there is a UDG representation $(\rho : V \to \mathbb{D}^2, s)$ of G in the 2-dimensional Poincaré disk model with threshold radius s. We will show that $s \leq \operatorname{arcosh}\left(\frac{n-2}{2} + 1\right)$ must necessarily hold.

Clearly, the graph $\tilde{C}_n$ contains no triangles. Due to Theorem 6.6, the UDG-drawing $\mathcal{D} = (\mathcal{V}, \mathcal{E})$ induced by (ρ, s) must be crossing-free. As argued in Remark 6.5, we can convert $\mathcal{D}$ to a planar $\mathbb{R}^2$-drawing $(\rho, \mathcal{E}')$ of $\tilde{C}_n$ by converting each geodesic segment in $\mathcal{E}$ to a geodesic arc in $\mathcal{E}'$. In this planar $\mathbb{R}^2$-drawing, three curves are of particular interest: let α denote the curve corresponding to the path c_0, a_0, a_1, a_2, c_2 in $\tilde{C}_n$. Refer to Figure 6.2 for a visual aid. Let β denote the curve corresponding to the path c_1, b_0, b_1, b_2, c_3 in $\tilde{C}_n$. Finally, let γ denote the closed curve corresponding to the circle $c_0, \dots, c_{n-1}$ in $\tilde{C}_n$. We know that the region of $\mathbb{D}^2$ bounded by γ is a hyperbolic n-gon as defined in Definition 2.59, because the curve γ is composed of the geodesic segments $[\rho(c_0), \rho(c_1)], \dots, [\rho(c_{n-2}), \rho(c_{n-1})], [\rho(c_{n-1}), \rho(c_0)] \in \mathcal{E}$. Label this n-gon P. In particular, γ has a well-defined inner and outer region. From planar graph theory we know that one of the points $\rho(a_1)$ and $\rho(b_1)$ must be in the inner region of γ and the other one in the outer region. Otherwise, the curves α and β would intersect, but the $\mathbb{R}^2$-drawing is planar. We may assume that $\rho(a_1)$ is in the inner region, i.e. contained in the n-gon P. We proceed by studying the hyperbolic disk $B_D(\rho(a_1), s)$. We claim that it is contained in P as well. If it was not, there would be

some point $x \in B_D(\rho(a_1), s)$ that would lie on the boundary of P. Since the boundary of P consists of the geodesic segments induced by the circle $c_0, \dots, c_{n-1}$, there would be a geodesic segment $[\rho(c_i), \rho(c_j)]$ where $c_i c_j \in E$ which contains x. We obtain

$$d_D(\rho(c_i), x) + d_D(\rho(c_j), x) = d_D(\rho(c_i), \rho(c_j)) < 2s$$

due to Corollary 2.27. So at least one of the terms $d_D(\rho(c_i), x)$ and $d_D(\rho(c_j), x)$ must be smaller than s. We may assume that the first one is. Then, the triangle inequation would imply

$$d_D(\rho(c_i), \rho(a_1)) \leq d_D(\rho(c_i), x) + d_D(x, \rho(a_1)) < s + s = 2s.$$

Because (ρ, s) is a UDG representation of $\tilde{C}_n$, an edge $a_1 c_i \in E$ would have to exist, but this is not the case. Consequently, such a point x cannot exist and $B_D(\rho(a_1), s) \subseteq P$ must hold. Now, we obtain an upper bound on s by calculating the hyperbolic area of this ball and of the n-gon P: as an immediate consequence of the Gauss-Bonnet-Theorem (Theorem 2.60) we have $\mathrm{Area}_{\mathrm{H}}(P) \leq (n-2)\pi$. Combining this with the formula of the area of a hyperbolic disk with radius s provided in Theorem 2.58, we get

$$\begin{aligned} & 2\pi(\cosh(s) - 1) = \mathrm{Area}_{\mathrm{H}}(B_D(\rho(a_1), s)) \leq \mathrm{Area}_{\mathrm{H}}(P) \leq (n-2)\pi \\ \implies & s \leq \operatorname{arcosh}\left(\frac{n-2}{2} + 1\right) \end{aligned}$$

because arcosh = $\cosh^{-1}$ is strictly increasing. ∎

It is worth noting that this bound on s could be further improved, due to the rough estimate concerning the area of a hyperbolic n-gon made in the proof.

So far, we observed graphs that allow a HUDG representation if one chooses the threshold radius either *great enough* (star graphs) or *small enough* (graphs of the family $\tilde{C}_n,\ n \geq 9$). Without much effort, this allows us to construct examples of graphs that are HUDGs if and only if the threshold radius is chosen *well-balanced*, i.e. neither too great nor too small. The main idea is to create a graph having a vertex-induced subgraph that requires a small threshold radius and having another vertex-induced subgraph that requires a great threshold radius. This is based on the fact that, given a fixed threshold radius, a UDG representation of a graph induces a UDG representation of each of its vertex-induced subgraphs, as proven in Corollary 3.19. For example, we can consider the disjoint union of some $\tilde{C}_n$ graph and the star graph $K_{1,7}$. If we choose n too small, the restriction on the threshold radius as noted in Theorem 6.7 will prevent a simultaneous representation of $K_{1,7}$. By trial and error, we notice that $n \geq 16$ is sufficient, see Figure 6.3.

6.3.1 Upper Bound for Great Threshold Radii

Previously, we observed graphs that are HUDGs for small threshold radii, but have hyperbolic sphericity of at least three when the threshold radius is increased. However, visually we get the impression that models of 3-dimensional hyperbolic space are always sufficient to find UBG representations of these particular graphs. In the following, we generalize and prove this impression, by comparing n-dimensional hyperbolic space to $(n+1)$-dimensional hyperbolic space in terms of sphericity.

Theorem 6.8: Let $G = (V, E)$ be any graph. Let $s_0 > 0$ such that $\mathrm{HSph}_{s_0}(G) = n \in \mathbb{N}$. Then, for each $s > s_0$, $\mathrm{HSph}_s(G) \leq n + 1$.

Proof. We may assume a UBG representation $(\rho : V \to \mathbb{U}^n, s_0)$ of G in the upper half-space model $\mathbb{U}^n$ of n-dimensional hyperbolic space. For each $t > 0$, define the function

$$f_t : \mathbb{U}^n \to \mathbb{U}^{n+1}, (x_1, \dots, x_n) \mapsto \left(\sqrt{1+t}\, x_1, \dots, \sqrt{1+t}\, x_{n-1}, \sqrt{t}\, x_n, x_n\right).$$

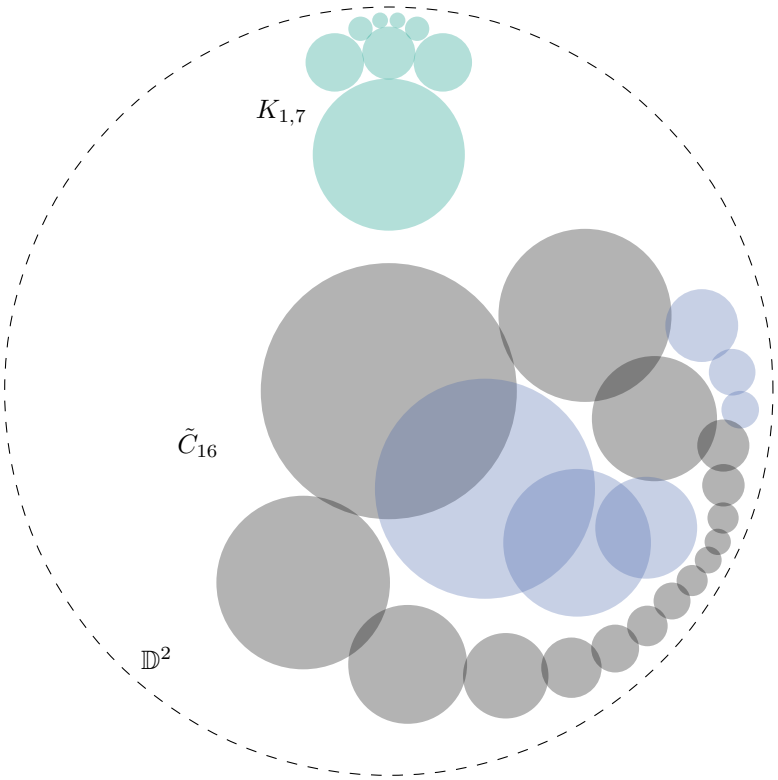

Figure 6.3: The disjoint union of $K_{1,7}$ and $\tilde{C}_{16}$ has a HUDG representation, but only if the threshold radius is chosen well-balanced.

This is well-defined because $x_n > 0$ implies $f_t(x)_{n+1} = x_n > 0$. Clearly, f_t is injective for each $t > 0$. From Definition 2.44, recall that the metric on $\mathbb{U}^n$ is given by

$$d_U(x, y) = \operatorname{arcosh}\left(1 + \frac{\|x - y\|^2}{2x_n y_n}\right) \quad \text{for all } x, y \in \mathbb{U}^n$$

where $\operatorname{arcosh} = \cosh^{-1}$ denotes the inverse hyperbolic cosine function. Some properties of this function are worth highlighting: the domain of arcosh is given by $[1, \infty)$ and $\operatorname{arcosh} : [1, \infty) \to \mathbb{R}_{\geq 0}$ is a strictly increasing bijection.

For any $x, y \in \mathbb{U}^n$, using the definition of d_U and the fact that arcosh is strictly increasing, we calculate for any value $\alpha > 0$

$$\begin{aligned}
& d_U(f_t(x), f_t(y)) < \operatorname{arcosh}(1+\alpha) \\
\iff & \frac{\|f_t(x) - f_t(y)\|^2}{2f_t(x)_{n+1} f_t(y)_{n+1}} < \alpha \\
\iff & \frac{\left(\sum_{i=1}^{n-1}\left(\sqrt{1+t}\,(x_i - y_i)\right)^2\right) + \left(\sqrt{t}(x_n - y_n)\right)^2 + (x_n - y_n)^2}{2x_n y_n} < \alpha \\
\iff & \frac{(1+t)\left(\sum_{i=1}^{n-1}(x_i - y_i)^2\right) + (1+t)(x_n - y_n)^2}{2x_n y_n} < \alpha \\
\iff & (1+t)\frac{\|x-y\|^2}{2x_n y_n} < \alpha \\
\iff & \frac{\|x-y\|^2}{2x_n y_n} < \frac{\alpha}{1+t} \\
\iff & d_U(x,y) < \operatorname{arcosh}\left(1 + \frac{\alpha}{1+t}\right).
\end{aligned}$$

Note that the metric d_U in the first line refers to the space $\mathbb{U}^{n+1}$ while in the last line, d_U refers to $\mathbb{U}^n$. By setting $\beta := \frac{\alpha}{1+t} > 0$, we conclude that the injection $f_t : \mathbb{U}^n \to \mathbb{U}^{n+1}$ is a threshold-preserving function with respect to the thresholds $\operatorname{arcosh}(1+\beta)$ and $\operatorname{arcosh}(1+(1+t)\beta)$ for any values $\beta, t > 0$, see Definition 3.6.

Now, back to the graph G and its UBG representation $(\rho : V \to \mathbb{U}^n, s_0)$. Fix any value $s > s_0$. We prove the statement by providing a UBG representation of G in $\mathbb{U}^{n+1}$ with threshold radius s. Because $\operatorname{arcosh} : [1, \infty) \to \mathbb{R}_{\geq 0}$ is a strictly increasing bijection, there is a unique value $\beta > 0$ such that $2s_0 = \operatorname{arcosh}(1+\beta)$. For the same reason, there is a unique value $\beta'(s) > \beta$ such that $2s = \operatorname{arcosh}(1+\beta'(s))$. By setting $t(s) := \frac{\beta'(s)}{\beta} - 1 > 0$, we have $\beta'(s) = (1+t(s))\beta$. Summarizing this construction, we get

$$2s_0 = \operatorname{arcosh}(1+\beta) < \operatorname{arcosh}(1+(1+t(s))\beta) = 2s.$$

That is, we have chosen the parameter $t(s) > 0$ in a way such that the two thresholds of the threshold-preserving function $f_{t(s)}$ match precisely $2s_0$ and $2s$. Consequently, Theorem 3.8 yields that

$$(f_{t(s)} \circ \rho : V \to \mathbb{U}^{n+1}, s)$$

is a UBG representation of G in $\mathbb{U}^{n+1}$ with threshold radius s. Since $s > s_0$ was chosen arbitrarily, this concludes the proof. ∎

The presented theorem reminds us of the inequalities obtained in Chapter 4 where we compared Euclidean and hyperbolic space in terms of sphericity. Here, we compared hyperbolic space with itself. It is worth mentioning that combining the above theorem with Theorem 4.4 directly implies the statement of Theorem 4.1.

6.4 Open Question: Medium Threshold Radius Potentially Suboptimal

So far, we observed examples of graphs that are HUDGs if and only if the threshold radius is either chosen great enough (large trees), small enough ($\tilde{C}_n$ with $n \geq 9$) or well-balanced (disjoint union of appropriate $\tilde{C}_n$ and $K_{1,7}$). Apart from graphs that are HUDGs for every threshold radius, the question arises whether the influence of the threshold radius is fully classified by the above three categories.

We conjecture that this is not the case. The three categories above have in common that the range of threshold radii permitting a HUDG representation can be described by a real interval. In contrast, we are looking for a graph G where the said range is not connected, i.e. there are threshold radii $0 < s_1 < s_2 < s_3$ such that s_1 and s_3 permit HUDG representations of G, but s_2 does not.

What could such a graph G look like? Our idea resembles the construction of $\tilde{C}_n$ from above: start with a cycle graph C_n. Fix any vertex c_0 of C_n. Add two other graphs G' and G'' and connect both of them to c_0 by adding an edge each. Let G denote the obtained graph from this construction. For now, we treat the subgraphs G' and G'' as black boxes, but we assume that both of them are connected and triangle-free. Again, this allows us to use Theorem 6.6: suppose a HUDG representation of G with some threshold radii s exists. We may assume a representation in the Poincaré disk $\mathbb{D}^2$. Since this representation induces a planar $\mathbb{R}^2$-drawing of G, the subgraph G' must be drawn either completely inside or completely outside the region $P \subseteq \mathbb{D}^2$ bounded by the closed curve corresponding to the cycle C_n. The same applies to G''. We know that this region P is a hyperbolic polygon and that its area is bounded by a constant, no matter how great the threshold radius is chosen.

Now it comes down to finding graphs G' and G'' such that the following holds: for a small threshold radius s_1, the subgraph G' can be drawn as a HUDG inside the cycle C_n since the corresponding region provides enough area. Meanwhile, G'' can be drawn as a HUDG outside the cycle. For a medium threshold radius s_2, the area necessary to draw G' as a HUDG shall exceed the available area inside the cycle C_n (and the same shall hold for G''), so both subgraphs G' and G'' must be drawn outside C_n. However, since they are both attached to C_n, drawing G' and G'' as HUDGs outside C_n shall also not work, unless the threshold radius is further increased to an adequate value s_3.

After some experimentation, the following choices appear promising: let G' and G'' be copies of the tree T whose root has four children that each have two children. Furthermore, we choose C_n as C_{19} since this seems to be the smallest cycle that provides enough area to contain G' for small threshold radii. Figure 6.4 depicts two successful HUDG representations of the resulting graph G, one with a small threshold radius s_1 where G' lies inside C_{19} and one with a great threshold radius s_3 where G' and G'' both lie outside C_{19}. Moreover, Figure 6.4 depicts two failed attempts to construct a HUDG representation of G with a specific threshold radius s_2 that satisfies $s_1 < s_2 < s_3$. Again, in one attempt, G' shall be placed inside C_{19} and in the other attempt, G' and G'' shall both be placed outside C_{19}. Both attempts seem to make nearly optimal use of the available space. Unfortunately, the tools we have at hand are not precise enough to formally prove that the graph G has no HUDG representation with the threshold radius s_2 as specified in Figure 6.4.

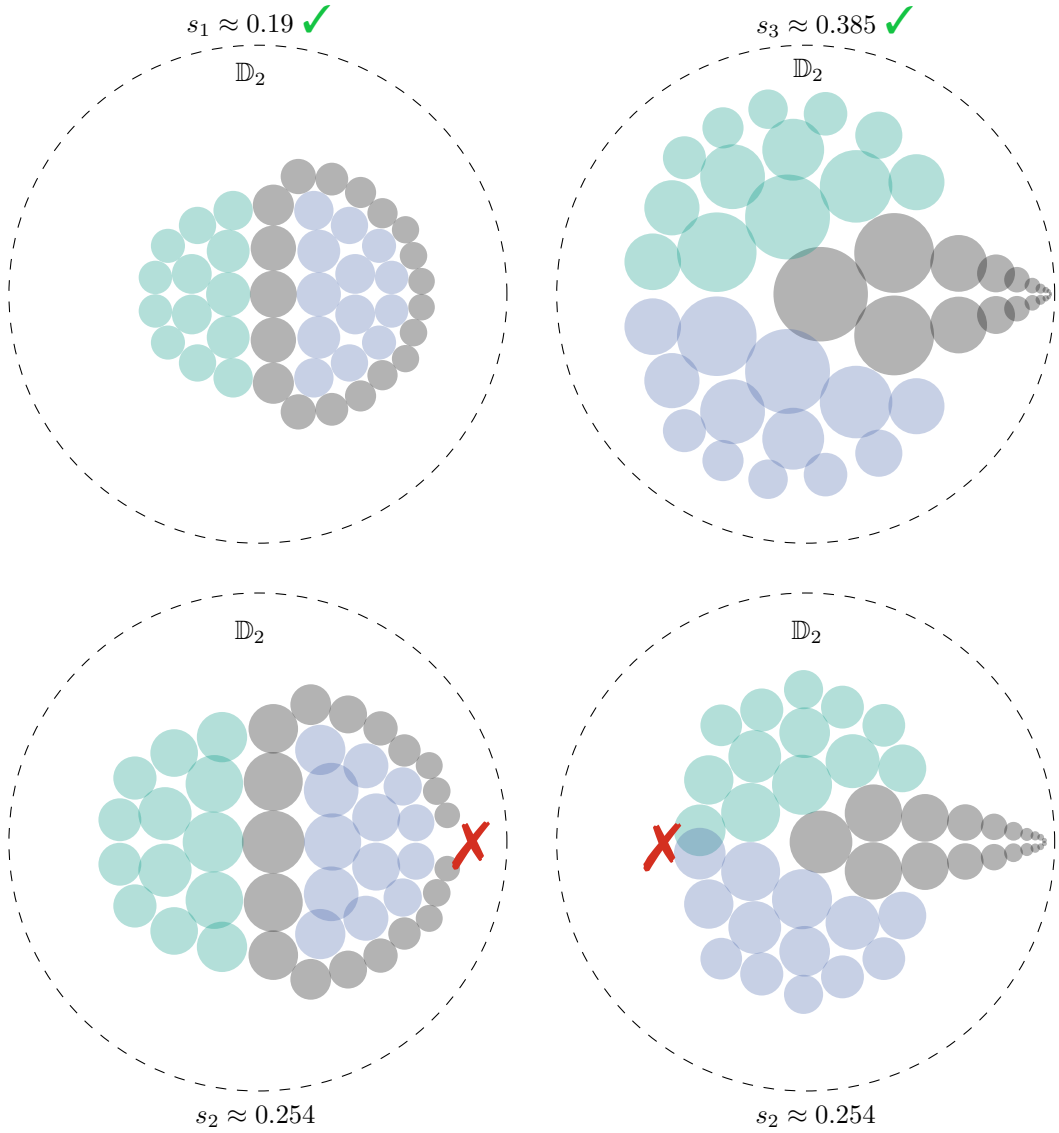

Figure 6.4: Graph G that has a HUDG representation with certain threshold radii s_1 and s_3, but potentially not with every threshold radius s_2 in between.

7 Unit Ball Graphs in Spherical Space

Spherical space is obtained when considering the length of shortest paths between points on the unit sphere S^n embedded in $\mathbb{R}^{n+1}$. There are many analogies between spherical and hyperbolic space. In certain aspects, the two spaces even appear to be counterparts of each other. For example, the angle sum of a hyperbolic triangle is always less than π, see Theorem 2.52. In contrast, the angle sum of a spherical triangle is always greater than π, as Ratcliffe proves [Rat06, p. 47]. For these reasons, we are also interested in exploring unit ball graphs in spherical space. While we initially defined UBGs in the context of Euclidean and hyperbolic space, Definition 3.1 can clearly be applied to any metric space (X, d).

Definition 7.1: The n-dimensional spherical space is given by the metric space (S^n, d_S) where

$$S^n = \{x \in \mathbb{R}^{n+1} \mid \|x\| = 1\} \quad \text{and} \quad d_S(x, y) = \theta(x, y) \text{ for each } x, y \in S^n.$$

Here, $\theta(x, y) \in [0, \pi]$ denotes the angle between the vectors x and y, as defined in Definition 2.11. For a proof that the spherical distance function d_S is indeed a metric on S^n, we refer to Ratcliffe [Rat06, p. 36]. Ratcliffe also shows that the geodesics of S^n are its *great circles*, i.e. intersections of S^n with a 2-dimensional vector subspace of $\mathbb{R}^{n+1}$. As with Euclidean space, a single model of spherical space is sufficient for our purposes. This is why we do not need to make a clear distinction between the abstract spherical geometry (that could also be defined axiomatically) and the metric space as given in the previous definition.

We proceed by studying unit ball graphs in spherical space and sphericity in this context. As mentioned above, Definition 3.1 can also be applied to spherical space. When trying to define the *spherical sphericity* of graphs, we soon encounter some particularities of spherical space. Similar to hyperbolic sphericity, the threshold radius generally has an influence on the the spherical sphericity of a graph. But unlike the hyperbolic case, not every threshold radius leads to a finite sphericity in spherical space: to see this, consider some threshold radius $\alpha > \pi/2$. Because of Definition 3.1, two distinct vertices v, w are adjacent if and only if they are mapped to distinct points on the sphere $x, y \in S^n$ that satisfy $d_S(x, y) < 2\alpha$. But because $2\alpha > \pi$ and $d_S(x, y) \leq \pi$ in any case, this condition is always fulfilled, i.e. two distinct vertices are always adjacent. As a consequence, with such a threshold radius $\alpha > \pi/2$, we are only able to represent complete graphs. These thoughts lead to the following definition:

Definition 7.2 (Spherical Sphericity)**:** Let G be any graph and $\alpha > 0$. The *spherical sphericity* of G with threshold radius α, denoted $\mathrm{SSph}_\alpha(G)$, is the smallest integer $n \in \mathbb{N}$ such that there is a UBG representation $(\psi : V \to S^n, \alpha)$ of G in n-dimensional spherical space. If such an integer $n \in \mathbb{N}$ does not exist, we leave $\mathrm{SSph}_\alpha(G)$ undefined.

In the following theorem, we compare spherical UBG representations to Euclidean ones.

Theorem 7.3: Let $G = (V, E)$ be a graph with a UBG representation $(\psi : V \to S^n, \alpha)$ of G in n-dimensional spherical space with threshold radius $\alpha \in (0, \pi/2)$. Then, there exists a UBG representation $(\tau : V \to E^{n+1}, r)$ of G in $(n+1)$-dimensional Euclidean space. Furthermore, this UBG representation satisfies $\|\tau(v)\| = c$ for each $v \in V$ where $c > r$ is some constant.

Proof. Define the function

$$f : S^n \to E^{n+1}, x \mapsto \frac{x}{\cos\alpha}.$$

This function is well-defined because $\cos\alpha > 0$ holds. Let $x, y \in S^n$. Then

$$\begin{aligned}
& d_E(f(x), f(y)) < 2\tan\alpha \\
\iff\ & d_E\left(\frac{x}{\cos\alpha}, \frac{y}{\cos\alpha}\right) < 2\tan\alpha && \text{| using Lemma 2.7} \\
\iff\ & \frac{1}{\cos\alpha}\, d_E(x, y) < 2\tan\alpha \\
\iff\ & d_E(x, y) < 2\sin\alpha && \text{| using Lemma 2.13 (2)} \\
\iff\ & 2\sin\left(\frac{\theta(x,y)}{2}\right) < 2\sin\alpha && \text{| sin strictly increasing on domain } (0, \pi/2) \\
\iff\ & \frac{\theta(x,y)}{2} < \alpha \\
\iff\ & d_S(x, y) < 2\alpha.
\end{aligned}$$

Since f is clearly injective, we conclude that f is a threshold-preserving function with respect to the thresholds 2α and $2\tan\alpha$ for each $\alpha \in (0, \pi/2)$, see Definition 3.6. Consequently, Theorem 3.8 implies that the tuple $(\tau, r) := (f \circ \psi, \tan\alpha)$ is also a UBG representation of G, but in E^{n+1}.

It remains to verify that all vectors $\tau(v)$ have the same length in $\mathbb{R}^{n+1}$. This follows from

$$\|\tau(v)\| = \left\|\frac{1}{\cos\alpha}\psi(v)\right\| = \frac{1}{\cos\alpha}\|\psi(v)\| = \frac{1}{\cos\alpha} =: c$$

because the Euclidean norm $\|\cdot\|$ is absolute homogeneous and because $\psi(v) \in S^n$. This concludes the proof since $\alpha \in (0, \pi/2)$ implies the inequality

$$c = \frac{1}{\cos\alpha} > \frac{\sin\alpha}{\cos\alpha} = \tan\alpha = r.$$

∎

Remark 7.4: Where do the values in the above proof come from? This is best seen from a geometric approach: we want to "simulate" the spherical ball $B_S(x, \alpha) \subseteq S^n \subseteq \mathbb{R}^{n+1}$ with some Euclidean ball $B_E(\lambda x, r) \subseteq \mathbb{R}^{n+1}$. By considering the simplest case, namely S^1, we get the impression that it would be a good idea to make $B_E(\lambda x, r)$ and S^1 intersect precisely in $B_S(x, \alpha)$. This ensures that an edge of the spherical representation is kept in the Euclidean representation. To avoid new, undesired edges in the Euclidean representation, we need to ensure that two Euclidean balls intersect only if the corresponding spherical balls intersect. This can be generally achieved by requiring that $B_E(\lambda x, r)$ is contained in the infinite cone $C(x, \alpha) := \{y \in \mathbb{R}^{n+1} \mid \theta(x, y) < \alpha\}$. We refer to Figure 7.1 for a visualization. Next, we notice that we can combine the two requirements if we make $S_E(\lambda x, r)$ (i.e. the boundary of $B_E(\lambda x, r)$) and S^1 intersect *orthogonally*. In this case, we obtain a triangle induced by the origin 0, the point λx and one of the two points where $S_E(\lambda x, r)$ and S^1 intersect. Call the last point p. We know that the angle at p is $\pi/2$ and that the angle at the origin is α. Because $d_E(0, p) = 1$, we obtain (see Figure 7.1)

$$\begin{aligned}
\cos\alpha = \frac{1}{d_E(0, \lambda x)} &\implies \lambda = d_E(0, \lambda x) = \frac{1}{\cos\alpha}, \\
\tan\alpha = \frac{d_E(\lambda x, p)}{1} &\implies r = d_E(\lambda x, p) = \tan\alpha.
\end{aligned}$$

Next, we focus on Euclidean UBG representations with the special property provided by Theorem 7.3, i.e. all vertices are placed on points equidistant to the origin. This special property allows a conversion to the Poincaré disk model $\mathbb{D}^n$ for each threshold radius, without increase of dimension:

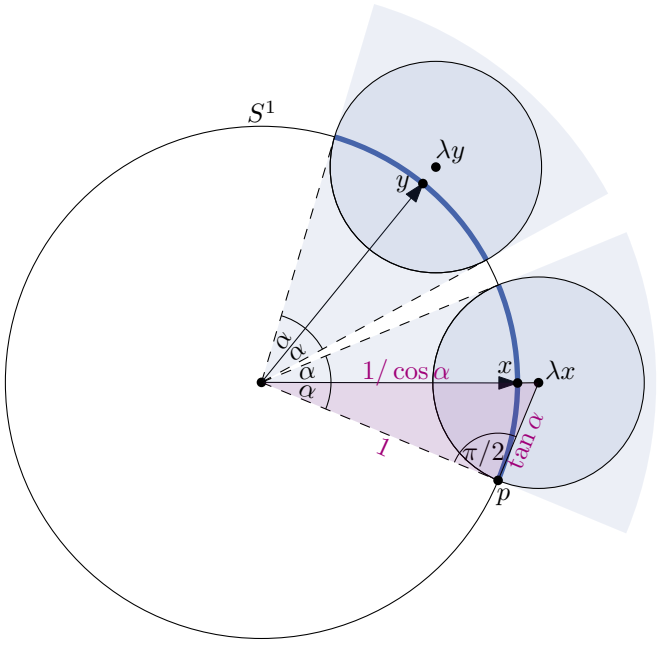

Figure 7.1: Illustration of Theorem 7.3 for S^1.

Lemma 7.5: Let $G = (V, E)$ be a graph with a UBG representation $(\tau : V \to E^n, 1)$ of G. Suppose that this UBG representation satisfies $\|\tau(v)\| = c$ for each $v \in V$ where $c > 1$ is some constant. Then, for any $s > 0$, there is a UBG representation $(\rho_s : V \to \mathbb{D}^n, s)$ of G in n-dimensional hyperbolic space with threshold radius s.

Proof. First of all, we may scale the UBG representation by a factor $\lambda > 0$. Define the function $\tau_\lambda : V \to E^n, v \mapsto \lambda\,\tau(v)$. Then, the tuple (τ_λ, λ) is a UBG representation of G as well because of Theorem 3.5.

For a fixed value of λ, consider the sphere $S_v := S_E(\tau_\lambda(v), \lambda)$ for each $v \in V$. Note that $\|\tau_\lambda(v)\| = \lambda c$ holds. In order to achieve that all these spheres are contained in $\mathbb{D}^n$, we require that the point on each sphere that is furthest from the origin has Euclidean norm less than one. For each sphere, this point has the Euclidean norm $\lambda(c + 1)$, so all choices $0 < \lambda < \frac{1}{c+1}$ lead to $S_v \subseteq \mathbb{D}^n$ for each $v \in V$. Given such a value of λ in the above range, we know from Theorem 2.39 that each S_v is a hyperbolic sphere in $\mathbb{D}^n$ as well. Furthermore, we observed in Remark 2.41 that these spheres all have the same hyperbolic radius $s > 0$ since their Euclidean radius is the same, namely λ, and their Euclidean centers all have the same Euclidean norm, namely λc. This implies that for a fixed value of $\lambda \in (0, \frac{1}{c+1})$, we can convert (τ_λ, λ) to a hyperbolic UBG representation $(\rho : V \to \mathbb{D}^n, s)$ by setting $\rho(v)$ as the hyperbolic center of the sphere S_v and by letting s denote the common hyperbolic radius of these spheres like above. By construction, the induced graphs of (τ_λ, λ) and (ρ, s) both correspond the the intersection graph of the set $\{S_v \mid v \in V\}$, so both tuples must be UBG representations of G.

It remains to show that *every* threshold radius $s > 0$ can result from this construction. Again, Remark 2.41 is helpful: for a fixed $\lambda \in (0, \frac{1}{c+1})$, we deduce that the corresponding value $s = s(\lambda)$ is given by

$$s(\lambda) = \frac{1}{2} \cdot \log\left(\frac{(1+\lambda)^2 - \lambda^2 c^2}{(1-\lambda)^2 - \lambda^2 c^2}\right) =: \frac{1}{2} \cdot \log\left(\frac{f(\lambda)}{g(\lambda)}\right).$$

It will be helpful to allow the case $\lambda = 0$ as well. By calculation, we see that the denominator $g(\lambda)$ in the log is greater than zero if $\lambda \in [0, \frac{1}{c+1})$. This means that s as a function of λ is well-defined for the domain $\lambda \in [0, \frac{1}{c+1})$. Furthermore, s is continuous as a composition of continuous functions. We calculate

$$f\left(\frac{1}{c+1}\right) = \left(1 + \frac{1}{c+1}\right)^2 - \frac{c^2}{(c+1)^2} = \frac{4}{c+1},$$

i.e. some positive constant. More interestingly, we have $g\left(\frac{1}{c+1}\right) = 0$. In summary, we can get the denominator g arbitrarily close to zero by approaching $\lambda \to \frac{1}{c+1}$ from below, while the enumerator f converges to some positive constant. Since log is strictly increasing, we conclude

$$\lim_{\lambda \to \frac{1}{c+1}} s(\lambda) = \infty.$$

Furthermore, we have $s(0) = 0$. When combining the above observations, the intermediate value theorem implies that the function s can take any positive value for arguments $\lambda \in (0, \frac{1}{c+1})$. This concludes the proof. ∎

Next, we treat the case $\alpha = \pi/2$:

Lemma 7.6: Let $G = (V, E)$ be a graph with a UBG representation $(\psi : V \to S^n, \pi/2)$ of G in n-dimensional spherical space with threshold radius $\pi/2$. Then, G is a vertex-induced subgraph of some graph $K_{m(2)}$, i.e. of a complete m-partite graph on m sets of size 2.

Proof. Following Definition 3.1, two distinct vertices $v, w \in V$ are adjacent if and only if $d_S(\psi(v), \psi(w)) < \pi$. Recall that the greatest possible spherical distance between two points $x, y \in S^n$ is π which is attained if and only if $x = -y$. So for every vertex $v \in V$, the following holds: either there is some other vertex $w \in V$ such that $\psi(w) = -\psi(v)$, then v is not adjacent to w but to every other vertex (expect itself), because ψ is an injection. Or there is no such w, then v is adjacent to every other vertex (expect itself). In the second case, we can add a new vertex v_0 to V and define $\psi(v_0) := -\psi(v)$. Again, v_0 is adjacent to every vertex expect v and itself. Repeat this process for every vertex to obtain a new vertex set $V' \supseteq V$. Let G' denote the induced graph of the tuple $(\psi : V' \to S^n, \pi/2)$. By construction, G' is of the type $K_{m(2)}$. In each step of the process of creating G' from G, we either did nothing or added one new vertex and added edges between this new vertex and existing vertices. In particular, we have neither added nor deleted edges between vertices in the initial vertex set V. This implies that G is a vertex-induced subgraph of G'. ∎

Finally, we are ready to combine all results of this chapter in the following corollary.

Corollary 7.7: Let $G = (V, E)$ denote any graph. For each $\alpha > 0$ such that $\mathrm{SSph}_\alpha(G)$ is defined in the sense of Definition 7.2, the following holds:

$$\mathrm{ESph}(G) \leq \mathrm{SSph}_\alpha(G) + 1 \quad \text{and} \quad \mathrm{HSph}_s(G) \leq \mathrm{SSph}_\alpha(G) + 1 \ \text{ for each } s > 0.$$

Proof. We differentiate between three cases. First of all, if $\alpha > \pi/2$, then G is necessarily a complete graph as noted in the introduction of this chapter. From Example 3.20 we know that $\mathrm{ESph}(G) = \mathrm{HSph}_s(G) = 1$ in this case, which implies the statement. Next, consider the case $\alpha = \pi/2$. As observed in Lemma 7.6, G must be vertex-induced subgraph of some graph $K_{m(2)}$. As such, we have

$$\mathrm{HSph}_s(G) \leq \mathrm{HSph}_s(K_{m(2)}) \in \{1, 2\} \quad \text{and} \quad \mathrm{ESph}(G) \leq \mathrm{ESph}(K_{m(2)}) \in \{1, 2\}$$

because of Corollary 3.19 and Example 3.22. This concludes this case since $2 \leq \mathrm{SSph}_{\alpha}(G) + 1$. The final case corresponds to $\alpha \in (0, \pi/2)$. The statement concerning Euclidean sphericity follows immediately from Theorem 7.3. For the statement concerning hyperbolic sphericity, Theorem 7.3 is helpful as well: take the therein provided UBG representation of G in E^{n+1}. Then, scale it to a UBG representation of G in E^{n+1} with threshold 1. This is possible due to Theorem 3.5. Finally, use Lemma 7.5 to convert the obtained UBG representation to a UBG representation of G in H^{n+1} with an arbitrary threshold radius $s > 0$. ∎

To sum up, spherical space is not particular advantageous when it comes to a low sphericity of graphs. Furthermore, it illustrates that sphericity with respect to a metric space and some threshold radius might not behave as nicely as in the Euclidean and hyperbolic case.

8 Conclusion

We have established the notion of hyperbolic sphericity $\mathrm{HSph}_s(\cdot)$ as a graph invariant in dependence of the threshold radius $s > 0$ and introduced related concepts like threshold-preserving functions. We made several contributions towards a better understanding of this graph invariant and the corresponding graph classes, mainly by proving the following upper bounds: for each graph G,

$$\mathrm{HSph}_s(G) \leq \mathrm{ESph}(G) + 1 \quad \text{for each } s > 0, \tag{8.1}$$

$$\mathrm{HSph}_s(G) \leq \mathrm{ESph}(G) \quad \text{for all } s \text{ in a certain interval } (0, s_0(G)], \tag{8.2}$$

$$\mathrm{HSph}_{s_0}(G) = n \implies \mathrm{HSph}_s(G) \leq n + 1 \quad \text{for each } s \geq s_0. \tag{8.3}$$

While Inequality (8.1) can be derived immediately from the other two inequalities, it is still worth mentioning, as this inequality shows that hyperbolic sphericity is always finite for finite graphs. Concerning tightness of the three inequalities, we obtained the following insights: we identified a graph family showing that Inequality (8.2) cannot be refined for any dimension, see Theorem 4.7. We also identified a graph G that satisfies $\mathrm{HSph}_s(G) = \mathrm{ESph}(G) = 2$ for small threshold radii s, but $\mathrm{HSph}_s(G) = 3$ when s is too great, see Theorem 6.7. This demonstrates that Inequality (8.1) cannot be refined in general. In the process, we observed that every triangle-free graph G having a UDG representation in $X \in \{E^2, \mathbb{D}^2\}$ is planar.

Confirming the empiric evidence that hyperbolic UDGs are well suited to model hierarchical networks [Kri+10], we showed that every tree is a hyperbolic UDG, as long as the threshold radius is chosen great enough. In contrast, Euclidean sphericity of trees is unbounded since the volume of Euclidean balls only grows polynomially in terms of the radius. Moreover, we took a brief excursion into complexity theory, proving that hyperbolic sphericity is computable using a polynomial amount of space. In the last chapter, we took another excursion towards spherical space, since it is often mentioned in the same context as hyperbolic space. However, we observed that, unlike hyperbolic space, spherical space does not produce any graphs with particular low sphericity in comparison to Euclidean space.

8.1 Future Work

We end this thesis by calling attention to three interesting directions for future research:

(1) *Deepening the understanding of hyperbolic sphericity.* As explained in Section 6.4, the influence of the threshold radius might not be fully understood yet. Furthermore, concerning Inequalities (8.1) to (8.3), we suspect that all of them cannot be refined for any dimension $n \geq 2$ (i.e. they can be tight with respect to every dimension $n \geq 2$). We were only able to prove so for Inequality (8.2). Similarly, it would be interesting to attempt characterizing the graphs that are particular *non-hyperbolic* in the sense that they satisfy $\mathrm{ESph}(G) \leq \mathrm{HSph}_s(G)$ for every $s > 0$. Conversely, another challenge is to identify more graphs with a great Euclidean sphericity, but a low hyperbolic sphericity for adequate threshold radii, similar to trees. Our approach based on hyperbolic tilings has framework character and can potentially be applied to other graph classes with a recursive structure.

(2) *Research on efficient algorithms on hyperbolic unit ball graphs.* In the multi-dimensional setting, this topic was only addressed by Kisfaludi-Bak until now [Kis20]. Can our insights on hyperbolic unit ball graphs be exploited algorithmically?

(3) *Research on other graph embeddings.* The concepts of unit ball graphs and sphericity can be applied to other metric spaces, for example $\mathbb{R}^n$ with the l_p norm for $p \in \{1, \dots, \infty\}$ (note that $p = 2$ corresponds to Euclidean sphericity). Which types of graphs display a particular low sphericity with respect to these spaces? It may even be interesting to consider non-metric spaces: for example, Fiduccia et al. define the *dot product dimension* of a graph $G = (V, E)$ as the smallest $n \in \mathbb{N}$ such that there is a function with the following property [FSTZ98]:

$$f : V \to \mathbb{R}^n \quad \text{such that } vw \in E \iff \langle f(v), f(w) \rangle \geq 1.$$

This can be seen as a generalization of unit ball graphs in spherical space due to Definition 2.11. So far, this graph invariant has received little attention. Its beauty lies in the simplicity of the representation, also making it attractive for practical applications like machine learning. In contrast, dealing with hyperbolic distances is computationally expensive.

Bibliography

[Ada22] C. Adams. *The Tiling Book*. 1st edition. AMS, 2022. ISBN: 9781470468972.

[And05] J. Anderson. *Hyperbolic Geometry*. 2nd edition. Springer London, 2005. ISBN: 9781852339340.

[BBDJ23] Nicholas Bieker, Thomas Bläsius, Emil Dohse, and Paul Jungeblut. "Recognizing Unit Disk Graphs in Hyperbolic Geometry is ∃ℝ-Complete". In: *CoRR* Volume abs/2301.05550 (2023). arXiv: *2301.05550*.

[BFKS23] Thomas Bläsius, Tobias Friedrich, Maximilian Katzmann, and Daniel Stephan. "Strongly Hyperbolic Unit Disk Graphs". In: *40th International Symposium on Theoretical Aspects of Computer Science, STACS 2023, March 7-9, 2023, Hamburg, Germany*. Edited by Petra Berenbrink, Patricia Bouyer, Anuj Dawar, and Mamadou Moustapha Kanté. Vol. 254. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, 13:1–13:17. DOI: *10.4230/LIPICS.STACS.2023.13*.

[BFM15] Michel Bode, Nikolaos Fountoulakis, and Tobias Müller. "On the Largest Component of a Hyperbolic Model of Complex Networks". In: *Electron. J. Comb.* Volume 22 (2015), p. 3. DOI: *10.37236/4958*.

[BK98] Heinz Breu and David G. Kirkpatrick. "Unit disk graph recognition is NP-hard". In: *Comput. Geom.* Volume 9 (1998), pp. 3–24. DOI: *10.1016/S0925-7721(97)00014-X*.

[Blä+24] Thomas Bläsius, Jean-Pierre von der Heydt, Sándor Kisfaludi-Bak, Marcus Wilhelm, and Geert van Wordragen. *Structure and Independence in Hyperbolic Uniform Disk Graphs*. 2024. arXiv: *2407.09362*.

[Can88] John F. Canny. "Some Algebraic and Geometric Computations in PSPACE". In: *Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA*. Edited by Janos Simon. ACM, 1988, pp. 460–467. DOI: *10.1145/62212.62257*.

[CCJ90] Brent N. Clark, Charles J. Colbourn, and David S. Johnson. "Unit disk graphs". In: *Discret. Math.* Volume 86 (1990), pp. 165–177. DOI: *10.1016/0012-365X(90)90358-O*.

[EvM20] Jeff Erickson, Ivor van der Hoog, and Tillmann Miltzow. "Smoothing the gap between NP and ER". In: *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*. Edited by Sandy Irani. IEEE, 2020, pp. 1022–1033. DOI: *10.1109/FOCS46700.2020.00099*.

[Fis83] Peter C. Fishburn. "On the sphericity and cubicity of graphs". In: *J. Comb. Theory B* Volume 35 (1983), pp. 309–318. DOI: *10.1016/0095-8956(83)90057-6*.

[FSTZ98] Charles M. Fiduccia, Edward R. Scheinerman, Ann N. Trenk, and Jennifer S. Zito. "Dot product representations of graphs". In: *Discret. Math.* Volume 181 (1998), pp. 113–138. DOI: *10.1016/S0012-365X(97)00049-6*.

[Gip14] Jake Gipple. "The Volume of n-balls". In: *Rose-Hulman Undergraduate Mathematics Journal* Volume 15 (2014), Article 14.

[GN98] Alexander Grigor'yan and Masakazu Noguchi. "The Heat Kernel on Hyperbolic Space". In: *Bulletin of the London Mathematical Society* Volume 30 (1998), pp. 643–650. eprint: *https://londmathsoc.onlinelibrary.wiley.com/doi/pdf/10.1112/S0024609398004780*.

[Har00] Robin Hartshorne. "Non-Euclidean Geometry". In: *Geometry: Euclid and Beyond*. New York, NY: Springer New York, 2000, pp. 295–433. ISBN: 978-0-387-22676-7. DOI: *10.1007/978-0-387-22676-7_8*.

[IM04] Piotr Indyk and Jiri Matousek. "Low-Distortion Embeddings of Finite Metric Spaces". In: *Handbook of Discrete and Computational Geometry, Second Edition*. Edited by Jacob E. Goodman and Joseph O'Rourke. Chapman and Hall/CRC, 2004, pp. 177–196. DOI: *10.1201/9781420035315.CH8*.

[Inc] Wolfram Research Inc. "Mathematica, Version 14.0". URL: *https://www.wolfram.com/mathematica*.

[Kis20] Sándor Kisfaludi-Bak. "Hyperbolic intersection graphs and (quasi)-polynomial time". In: *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*. Edited by Shuchi Chawla. SIAM, 2020, pp. 1621–1638. DOI: *10.1137/1.9781611975994.100*.

[KK07] K. Koecher and A. Krieg. "Das Dreieck und seine Kreise". In: *Ebene Geometrie*. 3rd edition. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 143–196. ISBN: 978-3-540-49328-0. DOI: *10.1007/978-3-540-49328-0_5*.

[Kri+10] Dmitri V. Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá. "Hyperbolic Geometry of Complex Networks". In: *CoRR* Volume abs/1006.5169 (2010). arXiv: *1006.5169*.

[Mae84a] Hiroshi Maehara. "On the sphericity for the join of many graphs". In: *Discret. Math.* Volume 49 (1984), pp. 311–313. DOI: *10.1016/0012-365X(84)90169-9*.

[Mae84b] Hiroshi Maehara. "Space graphs and sphericity". In: *Discret. Appl. Math.* Volume 7 (1984), pp. 55–64. DOI: *10.1016/0166-218X(84)90113-6*.

[Nas54] John Nash. "C1 Isometric Imbeddings". In: *Annals of Mathematics* Volume 60 (1954), pp. 383–396. ISSN: 0003486X, 19398980.

[Rat06] J. Ratcliffe. *Foundations of Hyperbolic Manifolds*. 3rd edition. Springer New York, 2006. ISBN: 9780387331973.

[RRS50] A. Rényi, C. Rényi, and J. Surányi. "Sur l'indépendance de domaines simples dans l'espace euclidien à n dimensions". In: *Colloquium Mathematicum* Volume 2 (1950), pp. 130–135.

[SS17] Marcus Schaefer and Daniel Stefankovic. "Fixed Points, Nash Equilibria, and the Existential Theory of the Reals". In: *Theory Comput. Syst.* Volume 60 (2017), pp. 172–193. DOI: *10.1007/S00224-015-9662-0*.

[Sto] Wilson Stothers. "Congruence of hyperbolic triangles". URL: *https://www.maths.gla.ac.uk/wws/cabripages/hyperbolic/congruence.html* (visited on July 31, 2024).

[Wal19] Charles Walkden. "MATH32502 Hyperbolic Geometry". 2019. URL: *https://personalpages.manchester.ac.uk/staff/charles.walkden/hyperbolic-geometry/hyperbolic_geometry_1920.pdf* (visited on July 31, 2024).

[Wei] Eric W. Weisstein. "Circumradius. From MathWorld—A Wolfram Web Resource". URL: *https://mathworld.wolfram.com/Circumradius.html* (visited on July 31, 2024).

[Yan+22] Menglin Yang, Min Zhou, Zhihao Li, Jiahong Liu, Lujia Pan, Hui Xiong, and Irwin King. "Hyperbolic Graph Neural Networks: A Review of Methods and Applications". In: *CoRR* Volume abs/2202.13852 (2022). arXiv: *2202.13852*.