

Structure and Independence in Hyperbolic Uniform Disk Graphs

Thomas Bläsius   

Karlsruhe Institute of Technology, Germany

Jean-Pierre von der Heydt   

Karlsruhe Institute of Technology, Germany

Sándor Kisfaludi-Bak  

Aalto University, Espoo, Finland

Marcus Wilhelm   

Karlsruhe Institute of Technology, Germany

Geert van Wordragen  

Aalto University, Espoo, Finland

Abstract

We consider intersection graphs of disks of radius r in the hyperbolic plane. Unlike the Euclidean setting, these graph classes are different for different values of r , where very small r corresponds to an almost-Euclidean setting and $r \in \Omega(\log n)$ corresponds to a firmly hyperbolic setting. We observe that larger values of r create simpler graph classes, at least in terms of separators and the computational complexity of the INDEPENDENT SET problem.

First, we show that intersection graphs of disks of radius r in the hyperbolic plane can be separated with $\mathcal{O}((1 + 1/r) \log n)$ cliques in a balanced manner. Our second structural insight concerns Delaunay complexes in the hyperbolic plane and may be of independent interest. We show that for any set S of n points with pairwise distance at least $2r$ in the hyperbolic plane, the corresponding Delaunay complex has outerplanarity $1 + \mathcal{O}(\frac{\log n}{r})$, which implies a similar bound on the balanced separators and treewidth of such Delaunay complexes.

Using this outerplanarity (and treewidth) bound we prove that INDEPENDENT SET can be solved in $n^{\mathcal{O}(1 + \frac{\log n}{r})}$ time. The algorithm is based on dynamic programming on some unknown sphere cut decomposition that is based on the solution. The resulting algorithm is a far-reaching generalization of a result of Kisfaludi-Bak (SODA 2020), and it is tight under the Exponential Time Hypothesis. In particular, INDEPENDENT SET is polynomial-time solvable in the firmly hyperbolic setting of $r \in \Omega(\log n)$. Finally, in the case when the disks have ply (depth) at most ℓ , we give a PTAS for MAXIMUM INDEPENDENT SET that has only quasi-polynomial dependence on $1/\varepsilon$ and ℓ . Our PTAS is a further generalization of our exact algorithm.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Design and analysis of algorithms; Theory of computation $\rightarrow$ Computational geometry; Theory of computation $\rightarrow$ Graph algorithms analysis

Keywords and phrases hyperbolic geometry, unit disk graphs, independent set, treewidth

Digital Object Identifier 10.4230/LIPIcs.SoCG.2025.21

Related Version Full Version: <https://arxiv.org/abs/2407.09362> [9]

Funding Jean-Pierre von der Heydt: Funded by the pilot program Core-Informatics of the Helmholtz Association (HGF).

Sándor Kisfaludi-Bak: Supported by the Research Council of Finland, Grant 363444.

Marcus Wilhelm: funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 524989715.



© Thomas Bläsius, Jean-Pierre von der Heydt, Sándor Kisfaludi-Bak, Marcus Wilhelm, and Geert van Wordragen;
licensed under Creative Commons License CC-BY 4.0


41st International Symposium on Computational Geometry (SoCG 2025).

Editors: Oswin Aichholzer and Haitao Wang; Article No. 21; pp. 21:1–21:16



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



1 Introduction

Given a set of disks in the plane, one can assign to them a *geometric intersection graph* whose vertices are the disks, and edges are added between pairs of intersecting disks. The study of intersection graphs is usually motivated by physically realized networks: such networks require spatial proximity between vertices for successful connections, and the simplest model allows for connections within a distance 2, which is equivalent to the intersection graph induced by disks of radius 1 centered at the vertices. Such graphs are called *unit disk graphs*.

Unit disk graphs have received a lot of attention in the theoretical computer science literature. Their intriguing structural properties yield profound mathematical insights and facilitate the development of efficient algorithms. Historically, unit disk graphs have been studied in the Euclidean plane, where they are well motivated due to their relevance for, e.g., sensor networks. The relevance of unit disk graphs in the hyperbolic plane is less obvious. However, originating in the network science community, it has been observed [26] that the intersection graph of randomly sampled disks of equal radius r yields graphs that resemble complex real-world networks in regards to important properties¹. They are, e.g., heterogeneous with a degree distribution following a power law [19], have high clustering coefficient [19], and exhibit the small-world property [18, 35].

Besides numerous structural results, these hyperbolic random graphs also allow for the design of more efficient algorithms [7, 8, 4, 5]. However, all these results rely on the fact that the disks are chosen randomly. So far, there is only little research on hyperbolic uniform disk graphs from a deterministic, more graph-theoretic perspective. It is important to note that for the intersection graphs of radius- r disks the resulting graph classes are different for different values of r . Moreover, it is natural to allow for the radius r to be a (monotone) function of the number of vertices. Unlike the Euclidean setting where a simple scaling shows that the choice of “unit” does not matter, it is a crucial parameter in the hyperbolic setting.


For those unfamiliar with hyperbolic geometry there is a way to conceptually understand these graphs in a Euclidean setting as follows. One can think of a hyperbolic disk graph of radius r disks as a Euclidean disk graph where the radius of a disk of center p is set to $f_r(\text{dist}(o, p))$, where f_r is decreasing very quickly at a rate set by r , and $\text{dist}(o, p)$ is the distance between the origin and the disk center p . In particular, as r goes to 0, the rate of decrease is negligible, and we get almost Euclidean unit disk graphs, while large r corresponds to a very different graph class. See Figure 1 for an illustration.

This observation is formalized by stating that hyperbolic disk graphs of uniform (that is, equal) radius $1/n^3$ and uniform radius $\log n$ are incomparable: The former class contains the $k \times k$ grid graph with $n = k^2$ for all k , but does not contain stars of size $n \geq 8$. The latter contains all star graphs but does not contain any $k \times k$ grid of size $k \geq 5$. See Figure 1 for an illustration and the full version of this paper for a formal proof of these claims. We call the regime $r \leq 1/\sqrt{n}$ *almost Euclidean*, while $r \in \Omega(\log n)$ is called *firmly hyperbolic*.

The primary goal of our paper is to explore the following:

How does the structure of hyperbolic uniform disk graphs depend on the radius $r = r(n)$?

¹ For this to work, the choice of the radius r is crucial. It is chosen as $r = \log n + c$ for a constant c that controls the average degree, and the disk centers are all sampled within a disk of radius $2r$.



(a) Smaller r enables larger grids. (b) Only small stars for small r . (c) Larger r enables larger stars.

■ **Figure 1** Realizing a grid is only possible for small radii, while large stars are only possible in the firmly hyperbolic setting of $r \in \Omega(\log n)$.

We note that one of the main challenges for large radii is that unlike the Euclidean plane, the kissing number² of hyperbolic disks of large radius is unbounded, and as a result, there may be arbitrarily many “nearby disks” to any given disks in the independent set. At the same time the larger the hyperbolic disks are, the harder it is to “surround” them with equal radius disks, which intuitively leads to a simpler structure.

The case of constant radius hyperbolic disks (with constant kissing number) has been studied by Kisfaludi-Bak [23], but these methods fail for super-constant kissing number. The case of radius roughly $\log n$ has only been studied in a more specialized setting (inspired by hyperbolic random graph models) in [6]. They call the relevant graph class *strongly hyperbolic uniform disk graphs*. The above-mentioned hyperbolic random graphs are randomly sampled strongly hyperbolic uniform disk graphs. Apart from these particular choices of the radius, we have very limited understanding about the structure of hyperbolic uniform disk graphs, and the results of [6] do not imply anything for our hyperbolic uniform disk graphs.

The radius’ impact on problem complexity was studied in the context of the traveling salesman problem, where one can use the minimum distance between points as a parameter with a similar flavor. In the hyperbolic plane, Kisfaludi-Bak [24] showed that the complexity of TSP decreases as one increases the minimum pairwise distance r among input points. In this paper, we show that the structure of hyperbolic uniform disk graphs behaves similarly: their structure becomes easier for larger values of r , and INDEPENDENT SET can be solved faster for larger r .

Let $\text{HUDG}(r(n))$ denote the set of intersection graphs where each n -vertex graph can be realized as the intersection graph of disks of uniform (equal) radius $r(n)$ in the hyperbolic plane of Gaussian curvature³ -1 . We denote by HUDG the union of these classes for all radii, that is, $\text{HUDG} = \bigcup_{\rho>0} \text{HUDG}(\rho)$. Let UDG denote the class of unit disk graphs in the Euclidean plane. The authors in [6] have shown that $\text{UDG} \subset \text{HUDG}$. When a graph is given without the geometric realization, it is NP-hard (even $\exists\mathbb{R}$ -complete) to decide if the graph is a UDG [33] or HUDG [3]; for this reason, we will assume throughout this article that the input intersection graphs of our algorithms are given by specifying their geometric realization. In the hyperbolic setting, one can specify the disk centers using a so-called *model* of the hyperbolic plane, which is simply an embedding of the hyperbolic plane into some specific part of the Euclidean plane, e.g., inside the Euclidean unit disk. See our preliminaries for a brief introduction to Euclidean models of the hyperbolic plane.

² The kissing number of radius r disks is the maximum number of interior-disjoint disks of radius r that can touch a fixed disk of radius r .

³ Equivalently, one can fix the radius to be 1 and set the Gaussian curvature to be $-r^2$.

From the perspective of graph algorithms, unit disk graphs in the Euclidean plane have been serving an important role as a graph class that is not comparable, but similar to planar graphs. The circle packing theorem [25] states that every planar graph can be realized as an intersection graph of disks (of arbitrary radii); thus, disk graphs serve as a common generalization of planar and unit disk graphs. Using *conformal* hyperbolic models, one can observe that hyperbolic disks appear as Euclidean disks⁴ in the model, which means that all HUDGs are realized as Euclidean disk graphs. Denoting Euclidean disk graphs with DG, we thus have $\text{UDG} \subseteq \text{HUDG} \subseteq \text{DG}$ ⁵.

A key structural tool in the study of both planar and (unit) disk graphs has been *separator theorems*. By a result of Lipton and Tarjan[27], any n -vertex planar graph can be partitioned into three vertex sets, A , B , and the separator S , such that no edges go between A and B , the separator set S has size $\mathcal{O}(\sqrt{n})$, and $\max\{|A|, |B|\} \leq \frac{2}{3}n$. Separator theorems are closely related to the *treewidth*⁶ of graphs [10], for example, the above separator implies a treewidth bound of $\mathcal{O}(\sqrt{n})$ on planar graphs.

For (unit) disk graphs, similar separators and treewidth bounds are not possible, because one can represent cliques of arbitrary size. There have been at least three different approaches to deal with large cliques. The first option is to bound the cliques in some way. One can, for example, obtain separators for a set of disks of bounded *ply*, i.e., assuming that each point of the ambient space is included in at most ℓ disks. There are several separators for disks (and even balls in higher dimensions) that involve ply. The strongest and most general among these is by Miller, Teng, Thurston, and Vavasis [34]. The second way to deal with large cliques is to use so-called clique-based separators [15, 16], where the separator is decomposed into cliques and the cliques are sometimes assigned some small weight depending on their size. In the hyperbolic setting, Kisfaludi-Bak [23] showed that for any constant c , the graph class $\text{HUDG}(c)$ admits a balanced separator consisting of $\mathcal{O}(\log n)$ cliques. The same paper shows that $\text{HUDG}(c)$ also has clique-based separators of weight $\mathcal{O}(\log^2 n)$ and extends these techniques to balls of constant radius in higher-dimensional hyperbolic spaces, along with much of the machinery of [15]. The third option is to use random disk positions to disperse large cliques with high probability. Bläsius, Friedrich, and Krophmer [7] gave high-probability bounds on separators and treewidth for certain radius regimes in hyperbolic random graphs.

The above separators and treewidth can be used to obtain fast divide-and-conquer or dynamic programming algorithms [28, 15]. In the INDEPENDENT SET problem, the goal is to decide if there are k pairwise non-adjacent vertices in the graph. In general graphs, the best known algorithms run in $2^{\mathcal{O}(n)}$ or $n^{\mathcal{O}(k)}$ time, and these running times are optimal under the Exponential Time Hypothesis (ETH) [14] (see [22] for the definition of ETH). In planar and disk graphs however, the separators allow algorithms with running time $2^{\mathcal{O}(\sqrt{n})}$ or $n^{\mathcal{O}(\sqrt{k})}$ [14, 32]. Moreover, the separator implies that planar graphs have treewidth⁶ $\mathcal{O}(\sqrt{n})$ and unit disk graphs have so-called *$\mathcal{P}$ -flattened treewidth* $\mathcal{O}(\sqrt{n})$ for some clique-partition $\mathcal{P}$. As a consequence, one can adapt treewidth-based algorithms [14] to these graph classes. In the hyperbolic setting, the ($\mathcal{P}$ -flattened) treewidth bounds yield an $n^{\mathcal{O}(\log n)}$ quasi-polynomial algorithm for INDEPENDENT SET and many other problems on graphs in $\text{HUDG}(c)$ for any constant c , which is matched by a lower bound under ETH [23] in case of INDEPENDENT SET.

⁴ In the Poincaré disk and half-plane models, all hyperbolic disks are Euclidean disks, but the Euclidean and hyperbolic radii of these disks are different.

⁵ With a little effort, one can show that both containments are strict: $\text{UDG} \subsetneq \text{HUDG} \subsetneq \text{DG}$.

⁶ See the full version of this paper for definitions of treewidth and $\mathcal{P}$ -flattened treewidth.

When approximating INDEPENDENT SET, the main algorithmic tool in planar graphs is Baker's shifting technique [2] that gives a $(1-\varepsilon)$ -approximate independent set in $2^{\mathcal{O}(1/\varepsilon)}n$ time. A conceptually easier grid shifting was discovered for unit disks by Hochbaum and Maas [21], which was later (implicitly) improved by Agarwal, van Kreveld, and Suri [1]. This line of research culminated in the algorithm of Chan [13] with a running time of $n^{\mathcal{O}(1/\varepsilon)}$ for disks, which also can be generalized for balls in higher dimensions. Baker's algorithm and Chan's algorithm (even for unit disks) are optimal under standard complexity assumptions [30]. To the best of our knowledge, there are no published approximation algorithms for $\text{HUDG}(r(n))$ or HUDG other than what is already implied from the algorithms on disk graphs. However, for hyperbolic random graphs, there is an $\mathcal{O}(m \log n)$ algorithm that yields a $1 - o(1)$ -approximation for INDEPENDENT SET asymptotically almost surely [5].

An interesting setting, where the treewidth of planar graphs has been used for intersection graphs, is that of covering and packing problems in the work of Har-Peled and later Marx and Pilipczuk [20, 32]. For the INDEPENDENT SET problem on unit disk graphs, they consider the *Voronoi diagram* of the disk centers of a maximum independent set of size k . For a set of points (often called *sites*) S , the Voronoi diagram is a planar subdivision (a planar graph) that partitions the plane according to the closest site of S , that is, the *cell* of a site $s \in S$ consists of those points whose nearest neighbor from S is s . As the Voronoi diagram is a planar graph on $\mathcal{O}(k)$ vertices, it has a balanced separator of size $\mathcal{O}(\sqrt{k})$. Although this Voronoi diagram is based on the solution, and thus unknown from the algorithmic perspective, the idea of Marx and Pilipczuk is that one can enumerate all possible separators of the Voronoi diagram in $n^{\mathcal{O}(\sqrt{k})}$ time, and one of these separators can then be used to separate the (still unknown) solution in a balanced manner. A similar "separator guessing" technique has been used elsewhere in approximation and parameterized algorithms [31, 12].

Our contribution. Our first structural result is a balanced clique-based separator theorem. We note that all of our results in hyperbolic uniform disk graphs assume that the graph is given via some geometric representation, i.e., with the Euclidean coordinates of its disk centers in some Euclidean model of the hyperbolic plane.


► **Theorem 1.** *Let G be a hyperbolic uniform disk graph with radius r . Then G has a separator S that can be covered with $\mathcal{O}(\log n \cdot (1 + \frac{1}{r}))$ cliques, such that all connected components of $G - S$ have at most $\frac{2}{3}n$ vertices. The separator can be computed in $\mathcal{O}(n \log n)$ time.*

Our separator is a carefully chosen line; the proof bounds the number of cliques intersecting the separator via decomposing some neighborhood of the chosen line into a small collection of small diameter regions. Notice that the theorem guarantees a separator with $\mathcal{O}(\log n)$ cliques in the hyperbolic setting, and deteriorates linearly as a function of $1/r$. In the almost-Euclidean setting of $r = 1/\sqrt{n}$ the separator has $\mathcal{O}(\sqrt{n} \log n)$ cliques, which almost matches the $\mathcal{O}(\sqrt{n})$ clique separator for unit disk graphs [15]. This is also a direct generalization of the 2-dimensional separator of Kisfaludi-Bak [23] from $r \in \Theta(1)$ to general radii.

Moreover, we observe that the techniques of [15, 23] can be utilized as the above separator implies $\mathcal{P}$ -flattened treewidth $\mathcal{O}(\log^2 n (1 + \frac{1}{r}))$ for any natural clique partition $\mathcal{P}$.

► **Corollary 2.** *Let G be a hyperbolic uniform disk graph with radius r . Then a maximum independent set of G can be computed in $n^{\mathcal{O}((1+1/r) \log^2 n)}$ time.*

Thus, a quasi-polynomial $n^{\mathcal{O}(\log^2 n)}$ algorithm is possible for all $r \in \Omega(1)$. If $r \leq 1$ – or even $r \in \mathcal{O}(1)$ – then for any $G \in \text{HUDG}(r)$ the neighborhood of a disk can be covered by a constant number of cliques of G (see also [23] for the case $r \in \Theta(1)$). Together with our separator in Theorem 1, the machinery of [15, 23] now yields the following result.



■ **Figure 2** A Euclidean Delaunay complex of n points with outerplanarity $n/3$.

► **Corollary 3.** *Let G be a hyperbolic uniform disk graph with radius $r \in \mathcal{O}(1)$. Then DOMINATING SET, VERTEX COVER, FEEDBACK VERTEX SET, CONNECTED DOMINATING SET, CONNECTED VERTEX COVER, CONNECTED FEEDBACK VERTEX SET can be solved in $n^{\mathcal{O}((1/r) \log n)}$ time, and q -COLORING for $q \in \mathcal{O}(1)$, HAMILTONIAN PATH and HAMILTONIAN CYCLE can be solved in $n^{\mathcal{O}(1/r)}$ time.*

While the above separator is already powerful and can be used as a basis for quasi-polynomial divide-and-conquer algorithms, the separator size stops improving when increasing the radius r beyond a constant. Nonetheless, a more powerful result is possible for super-constant radius r , which requires a different type of separator. Rather than separating hyperbolic uniform disk graphs, our second structural result is about separating the Delaunay complex, which is the dual of the Voronoi diagram, of a set of sites with pairwise distance at least $2r$. Note that in some graph $G \in \text{HUDG}(r)$, the disk centers corresponding to an independent set will have pairwise distance more than $2r$, which is what motivated us to study such point sets.

A plane graph (a planar graph with a fixed planar embedding) is called *outerplanar* or *1-outerplanar* if all of its vertices are on the unbounded face. A plane graph is called *k -outerplanar* if deleting the vertices of the unbounded face (and all edges incident to them) yields a $(k - 1)$ -outerplanar graph. Our result concerns the outerplanarity of a Delaunay complex. We prove the following tight result on the outerplanarity of the Delaunay complex $\mathcal{D}(S)$ of such a site set S . We believe that this result may be of independent interest.

► **Theorem 4.** *Let S be a set of n sites (points) in $\mathbb{H}^2$ with pairwise distance at least $2r$. Then the Delaunay complex $\mathcal{D}(S)$ is $1 + \mathcal{O}(\frac{\log n}{r})$ -outerplanar.*

In plane graphs, the treewidth, the outerplanarity, and the dual graph's treewidth and outerplanarity are all within a constant multiplicative factor of each other [10, 11]. Consequently, we get the bound $\mathcal{O}(1 + \frac{\log n}{r})$ for the treewidth and outerplanarity of the Voronoi diagram and Delaunay complex of such point sets. Note that this implies constant outerplanarity for the firmly hyperbolic setting of $r \in \Omega(\log n)$.

In particular, the theorem shows that for some constant c and $r > c \log n$ the Delaunay complex is outerplanar. We recover the outerplanarity bound of [23] for $r = \Theta(1)$, and get a sublinear outerplanarity bound ($\mathcal{O}(\sqrt{n} \log n)$) even in the almost-Euclidean setting of $r = 1/\sqrt{n}$. This is surprising in light of the fact that Delaunay-triangulations can have outerplanarity $\Omega(n)$ in the Euclidean setting, as demonstrated by the set of sites in Figure 2. Notice however that the construction requires a point set where the ratio of the maximum and minimum distance among the points is $\Omega(n)$; mimicking this construction in $\mathbb{H}^2$ with a point set of minimum distance $r = 1/\sqrt{n}$ is not possible, as at distance $n \cdot r = \sqrt{n}$ the hyperbolic distortion is very significant compared to the Euclidean setting.

Using the outerplanarity bound, we are able to give the following algorithm for INDEPENDENT SET.

► **Theorem 5.** *Let G be a hyperbolic uniform disk graph with radius r and let $k \geq 0$. Then we can decide if there is an independent set of size k in G in $\min \{n^{\mathcal{O}(1+\frac{1}{r} \log k)}, 2^{\mathcal{O}(\sqrt{n})}, n^{\mathcal{O}(\sqrt{k})}\}$ time.*

Our algorithm yields the first component of the running time $(n^{\mathcal{O}(1+\frac{1}{r} \log k)})$, but for very small r we can switch to the general disk graph algorithms with running times $2^{\mathcal{O}(\sqrt{n})}$ by the authors in [15] and $n^{\mathcal{O}(\sqrt{k})}$ by Marx and Pilipczuk [32]. Our algorithm is best possible under ETH for $r = \Theta(1)$ as [23] showed a lower bound of $n^{\Omega(\log n)}$ (for large k). Notice moreover that the almost-Euclidean case of $r = 1/\sqrt{n}$ gives a running time of $n^{\mathcal{O}(\sqrt{n} \log k)}$, which almost recovers the Euclidean running time of $2^{\mathcal{O}(\sqrt{n})}$. Recall that the Euclidean running time is also ETH-tight for Euclidean unit disks [15]. While this Euclidean lower bound cannot be directly applied to HUDG($1/\sqrt{n}$), it can be adapted to the this setting, so the running time lower bound $2^{\Omega(\sqrt{n})}$ holds also in the hyperbolic plane (assuming ETH). See [24] for a similar adaptation of a Euclidean lower bound to the setting of $r = 1/\sqrt{n}$.


Finally, but most surprisingly, we get a polynomial running time for INDEPENDENT SET in the firmly hyperbolic setting of $r \in \Omega(\log n)$; in fact, our algorithm is polynomial already for $r \in \Omega(\log k)$. In particular, this provides a polynomial exact algorithm for hyperbolic random graphs. It thereby directly generalizes the algorithm of [4] which gives a polynomial running time in hyperbolic random graphs with high probability: our algorithm has no assumptions on the input distribution and provides a polynomial worst-case running time.

The underlying idea of our exact algorithm can be regarded as a dynamic programming algorithm along the (unknown) tree decomposition of the Voronoi diagram of the solution disks, inspired by Marx and Pilipczuk [32]. More precisely, both [32] and the present paper use so-called *sphere cut decompositions*, a variant of branch decompositions for plane graphs [17] to guide the algorithm. However, due to our setting, a simple adaptation of the divide-and-conquer approach of [32] does not give the desired running time for us (it loses another logarithmic factor in the exponent). To get around this problem, we extend Marx and Pilipczuk's technique of noose-based separators to a noose-based dynamic programming algorithm. Additionally, we introduce *well-spaced* and *valid* sphere cut decompositions, and prove that they are in one-to-one correspondence with independent sets of G .

Finally, we consider approximation algorithms for MAXIMUM INDEPENDENT SET when the graph in question has low ply, that is, each point of $\mathbb{H}^2$ is covered by at most ℓ disks. We show the following.

► **Theorem 6.** *Let $\varepsilon \in (0, 1)$ and let G be a hyperbolic uniform disk graph of radius r and ply ℓ . Then a $(1 - \varepsilon)$ -approximate maximum independent set of G can be found in $\mathcal{O}(n^4 \log n) + n \cdot (\frac{\ell}{\varepsilon})^{\mathcal{O}(1+\frac{1}{r} \log \frac{\ell}{\varepsilon})}$ time.*

An important component of the algorithm's correctness proof is bounding the so-called *degeneracy* of hyperbolic uniform disk graphs in terms of their clique size and in terms of their ply. This generalizes the Euclidean results of [29]. When compared to the $n^{\mathcal{O}(1/\varepsilon)}$ algorithm for disk graphs by Chan [13] or the $2^{\mathcal{O}(1/\varepsilon)}n$ algorithm in planar graphs by Baker [2], both of which are conditionally optimal [30], our algorithm has a surprising quasi-polynomial dependence on $1/\varepsilon$ rather than exponential. The dependence on r exhibits the classic Euclidean-to-hyperbolic scaling seen in our earlier structural results. For $\varepsilon = 1/n$, we are guaranteed to get the optimal solution, and setting $\ell = n$ covers the general case. The resulting running time for $1/\varepsilon = \ell = n$ is $n^{\mathcal{O}(1+\frac{1}{r} \log n)}$, which matches our exact algorithm. Consequently, our approximation scheme is a generalization of our exact algorithm.



■ **Figure 3** Part (a) shows a visualization of the Poincaré disk including (parallel) lines, a circle with its center, a triangle, and polygons that include ideal points. Part (b) shows a set of sites $S \subseteq \mathbb{H}^2$ in the Poincaré disk model, $\mathcal{D}(S)$ (black), and $\mathcal{V}(S)$ (gray); Voronoi vertices are filled squares, ideal Voronoi vertices are empty squares.

2 Preliminaries

We introduce fundamental definitions and notation used throughout the paper. Additional concepts and notation are introduced in their relevant sections as needed.

Problem definition. An *independent set* in a graph is a set of pairwise non-adjacent vertices. In the INDEPENDENT SET problem, we are given⁷ a graph G and a number k , and the goal is to decide if there is an independent set of size k . In MAXIMUM INDEPENDENT SET, the input is a graph G , and the goal is to output the maximum value k such that G has an independent set of size k .

Let $G = (V, E)$ be a graph. A vertex set $S \subseteq V$ is a *separator* if V can be partitioned into non-empty subsets S , V_1 , and V_2 such that there is no edge between V_1 and V_2 , i.e., for every $v_1 \in V_1$ and $v_2 \in V_2$, it holds that $\{v_1, v_2\} \notin E$. The separator is β -balanced if $|V_1|, |V_2| \leq \beta n$ for some $\beta < 1$.

Hyperbolic geometry. We write $\mathbb{H}^2$ to refer to the hyperbolic plane and $\mathbb{R}^2$ to refer to the Euclidean plane. The *Poincaré disk* is a model of $\mathbb{H}^2$ that maps the whole hyperbolic plane into the interior of a Euclidean unit disk; see Figure 3a. We refer to the center of the Poincaré disk as *origin*. Straight lines in $\mathbb{H}^2$ are represented as circular arcs perpendicular to the boundary of the Poincaré disk or as chords through the origin. Hyperbolic circles are represented as Euclidean circles. The center of the hyperbolic circle is, however, farther from the origin than the Euclidean center of its representation. The Poincaré disk is *conformal*, i.e., angle preserving. Points on the boundary of the Poincaré disk are called *ideal points*. These are not part of the hyperbolic plane, but form a boundary of infinitely far points.

Let a be the sum of interior angles of a hyperbolic triangle. Then $a < \pi$ and the area of the triangle is $\pi - a$. Note that this implies that the area of a triangle is upper bounded by π . More generally, the area of a k -gon is bounded by $(k - 2)\pi$. The area of a disk with hyperbolic radius r is $4\pi \sinh^2(r/2)$, which is in $\Theta(e^r)$ for $r \rightarrow \infty$ and in $\Theta(r^2)$ for $r \rightarrow 0$.

Delaunay complexes and Voronoi diagrams. We consider Voronoi diagrams with, for the most part, the usual definition: given a set $S \subset \mathbb{H}^2$ of sites, the Voronoi diagram $\mathcal{V}(S)$ is a plane graph where every inner face (*Voronoi cell*) corresponds to a site $s \in S$ and consists

⁷ In this article, the graphs are always intersection graphs given through their geometric representation.

of the points in $\mathbb{H}^2$ for which s is the closest site. However, we deviate from the Euclidean definition by setting the outer face of $\mathcal{V}(S)$ to be the boundary of the Poincaré disk, which means that some of the Voronoi vertices are ideal points. This is shown in Figure 3b.

The weak dual of the Voronoi diagram is called the *Delaunay complex* and denoted by $\mathcal{D}(S)$; it is illustrated in black in Figure 3b. Here, sites $s_1 \in S$ and $s_2 \in S$ are connected in $\mathcal{D}(S)$ if and only if their Voronoi cells share a boundary.

3 Overview of our techniques

3.1 Balanced separators in HUDGs

Let G be a hyperbolic uniform disk graph with n vertices and radius r . We show that G has a balanced separator that can be covered by $\mathcal{O}((1 + 1/r) \log n)$ cliques; see Theorem 1. The overall argument is as follows. We find a double wedge bounded by two lines ℓ_1 and ℓ_2 that contains no vertex in its interior and separates the other vertices of G in a balanced fashion, which is ensured by selecting the center of the double wedge to be the hyperbolic *centerpoint* [24]. In Figure 4a, the double wedge with apex p is shaded gray and contains no vertices and the regions above and below the wedge each contain a constant fraction of the vertices.

Let m be the angular bisector of the wedge. As separator, we use the set of all vertices that have distance at most r from m . The set of points with distance exactly r from m forms two curves that are called *hypercycles* with axis m . Thus, vertices belong to the separator if they lie on or between the two hypercycles. The hypercycles are shown in Figure 4a and the region where separator vertices can lie is shaded blue. In the Poincaré disk, hypercycles are arcs of Euclidean circles that meet the boundary of the disk at the same ideal points as their axis but at a non-right angle. Observe that any vertex from above the top hypercycle has distance greater than $2r$ to any vertex below the bottom hypercycle. Thus, the vertices in the blue region indeed form a balanced separator.

It remains to show that the graph induced by vertices in the blue region can be covered with few cliques. For this, we cover the blue region with boxes as shown in Figure 4b. We show that each of these boxes has diameter at most $2r$, implying that the vertices within one box induce a clique. Moreover, we show that we only need $\mathcal{O}((1 + 1/r) \log n)$ such boxes. For the latter, we need a lower bound on the opening angle of the empty double wedge. Observe that a larger opening angle, φ , in Figure 4c, shrinks the blue region, which reduces the number of boxes required to cover it.

In the full version of this article, we first give a bound on the number of cliques required to cover the separator, parameterized by the opening angle of the empty double wedge. Afterwards, we show that there is a wedge with not too small opening angle that yields a balanced separator. The proof is constructive, i.e., given the graph with vertex positions, we can efficiently find the wedge and thereby the balanced separator of Theorem 1.

3.2 Outerplanarity of hyperbolic Delaunay complexes

Let S be a set of n sites in $\mathbb{H}^2$ with pairwise distance at least $2r$. Note that S interpreted as a hyperbolic uniform disk graph with radius r forms an independent set. To prove Theorem 4, giving an upper bound on the outerplanarity of the Delaunay complex $\mathcal{D}(S)$, we use two different types of arguments. The first argument is based on the observation that sufficiently large radius implies that any inner vertex of $\mathcal{D}(S)$ requires many other sites around it to shield it from being an outer vertex. This is similar to the observation in the introduction




Figure 4 Illustration of our separator using the Poincaré disk model. (a) The axis m is chosen such that it separates the vertices in a balanced fashion. The double wedge between ℓ_1 and ℓ_2 contains no vertices. The separator consists of the blue set of all vertices of distance at most r to the axis m . (b) We cover the separator with boxes that have diameter $2r$ and thus form cliques. (c) If the opening angle φ is sufficiently large, then x is sufficiently small and we only need few boxes to cover the whole separator.

(recall Figure 1) that for high radius r , large stars can be realized. The argument then goes roughly as follows: If all inner vertices have degree more than 6, then we need a linear fraction of vertices on the outer face to account for the fact that the average degree in planar graphs is bounded by 6. This already implies that iteratively deleting the vertices on the outer face takes at most $\log n$ steps as each step deletes a constant fraction of the vertices. A larger radius r leads to higher degrees of inner vertices, which in turn yields stronger bounds for the outerplanarity.

This type of argument can only work if the radius r is sufficiently large. For smaller radii, the second type of argument considers layers of bounded Voronoi cells (which correspond to inner vertices of $\mathcal{D}(S)$) around one fixed center cell. As the union of these layers is bounded by a polygon with a linear number of vertices, its area is linear. In contrast to that, we will see that the area of the layers grows exponentially with the layer (with the base depending on r), showing that there cannot be too many layers.

In the following, we give a sketch of how the second argument based on area works for small radii. Although this is the argument of choice for radii that are small constants or even decreasing with n , we first give a simple area-based argument for why $r \geq \log n$ implies that the Delaunay complex has no inner vertices. This gives a good intuition how the area behaves in the hyperbolic plane and how this can help us in proving outerplanarity, even though our argument for small radii is more involved. Let $s \in S$ be a site. If s is an inner vertex of $\mathcal{D}(S)$, then its Voronoi cell is bounded and its boundary is a polygon of at most $n - 1$ vertices. Thus, the area of this Voronoi cell is less than $(n - 3)\pi$. Simultaneously, the disk of radius r around s is included in the Voronoi cell of s as all sites have pairwise distance at least $2r$. The area of this disk is $4\pi \sinh^2(\frac{r}{2})$ which is larger than e^r for sufficiently large r . From this it follows that bounded cells can only exist when $r < \log n$.




Figure 5 Illustration of the small-radius case. (a) shows two consecutive layer polygons $P_{\ell-1}$ and P_{ℓ} in the Delaunay complex around s . The Voronoi diagram is shown in blue. We show an upper bound of the area inside P_{ℓ} by covering it with the red triangles (b). We show a lower bound for the area between $P_{\ell-1}$ to P_{ℓ} via the red triangles in (c). Relating these two bounds yields an exponential area growth with a basis depending on r .

To make an argument that works for smaller r , let $s \in S$ be an arbitrary but fixed vertex of the Delaunay complex $\mathcal{D}(S)$. We partition the vertices of $\mathcal{D}(S)$ into *layers* by hop distance from s in $\mathcal{D}(S)$, i.e., V_{ℓ} is the set of vertices with distance ℓ from s . Let L be the largest integer such that for all $\ell \leq L$ the layer V_{ℓ} contains only inner vertices. Note that our goal is to prove an upper bound on L , as this bounds the distance of s to the outer face.

As the Delaunay complex is an internally triangulated plane graph, each layer V_{ℓ} induces a cycle. We denote the vertices in V_{ℓ} with v_{ℓ}^i and number the vertices modulo $|V_{\ell}|$, such that vertex v_{ℓ}^i is adjacent to vertices v_{ℓ}^{i-1} and v_{ℓ}^{i+1} ; see Figure 5a. Moreover, we assume the Delaunay graph to be drawn as follows. For every edge $\{v_{\ell}^i, v_{\ell}^{i+1}\}$ we choose a point on its dual edge, i.e., the boundary between the two corresponding Voronoi cells, denote it with $v_{\ell}^{i+0.5}$, and connect v_{ℓ}^i and v_{ℓ}^{i+1} with two line segments via $v_{\ell}^{i+0.5}$. We denote the resulting polygon with P_{ℓ} and call it a *layer polygon* of s . Note that this yields a sequence of nested polygons where $P_{\ell-1}$ lies in the interior of P_{ℓ} .

Our goal then is to show that the area inside P_{ℓ} , denoted by $\text{Area}(P_{\ell})$, grows exponentially in ℓ with a basis $b(r)$ depending on r . As $\text{Area}(P_L)$ for the outermost layer L is upper bounded by something linear in n , there cannot be too many layers. The interesting part is proving the exponential growth. For this, we show that the area gain $\text{Area}(P_{\ell}) - \text{Area}(P_{\ell-1})$ in layer ℓ makes up at least some sufficiently large fraction of the area $\text{Area}(P_{\ell})$. For this, we give an upper bound for $\text{Area}(P_{\ell})$ and relate it to a lower bound for $\text{Area}(P_{\ell}) - \text{Area}(P_{\ell-1})$.

How we derive these bounds is illustrated in Figure 5b and Figure 5c, respectively. For the upper bound on $\text{Area}(P_{\ell})$, we cover P_{ℓ} with triangles connecting every edge of P_{ℓ} with the vertex s ; see the two red triangles in Figure 5b for the two edges $\{v_{\ell}^{0.5}, v_{\ell}^1\}$ and $\{v_{\ell}^1, v_{\ell}^{1.5}\}$ of P_{ℓ} . For the lower bound, we find a set of disjoint triangles that lie between P_{ℓ} and $P_{\ell-1}$. For this, observe that for every vertex $v_{\ell}^i \in V_{\ell}$ in layer ℓ , the Voronoi cell of v_{ℓ}^i completely contains the disk of radius r around v_{ℓ}^i as the sites have pairwise distance at least $2r$. Thus, the two triangles illustrated in red for v_{ℓ}^1 in Figure 5c satisfy the property of lying between P_{ℓ} and $P_{\ell-1}$. As the two triangles can in principle intersect, we choose for each vertex in layer ℓ the larger of the two triangles. Note that this gives a collection of $|V_{\ell}|$ disjoint triangles, as each chosen triangle lies in a different Voronoi cell. Thus, the total area of these triangles gives a lower bound for $\text{Area}(P_{\ell}) - \text{Area}(P_{\ell-1})$.

It then remains to relate the upper bound for $\text{Area}(P_\ell)$, i.e., the area of the red triangles in Figure 5b, with the lower bound for $\text{Area}(P_\ell) - \text{Area}(P_{\ell-1})$, i.e., the area of larger red triangle in Figure 5c. Intuitively, this does not seem too far fetched for the following reasons. First, the triangles in Figure 5b share a side with the triangles in Figure 5c. Secondly, the other sides of the triangles in Figure 5c have length at least r . Thus, the triangles in Figure 5c cannot be too much smaller than those in Figure 5b.

The details of this area-based argument for small radii can be found in the full version of this paper. Together with the degree-based bound for large radii we obtain Theorem 4.

3.3 Exact algorithm


We give a brief description of the exact algorithm for Theorem 5 and refer to the full version of this paper for the remaining details. Instead of computing the independent set directly, we define some additional structure for a given independent set. The algorithm then finds this additional structure such that the corresponding independent set is maximized. The additional structure is illustrated for two example independent sets in Figure 6. In the following, we explain this from top to bottom.

Let G be a hyperbolic uniform disk graph of radius r together with an independent set S (top row in Figure 6). As there are no edges between vertices in S , they have pairwise distance at least $2r$. Thus, the Delaunay complex $\mathcal{D}(S)$ (second row in Figure 6) has low outerplanarity due to Theorem 4. Low outerplanarity implies small treewidth, which implies small branchwidth [36]. Moreover, as $\mathcal{D}(S)$ is a planar graph, we can wish for a so-called *sphere cut decomposition* of low width, which is a branch decomposition with some additional properties [37, 17, 32]. In a nutshell, we get a hierarchy of *nooses*, where each noose is a closed curve that intersects the graph only at vertices and functions as a vertex separator, splitting interior edges from exterior ones. The hierarchy is a binary tree in the sense that every noose containing more than one edge in its interior has two child nooses partitioning these edges further. Row three in Figure 6 shows three nooses of the Delaunay complex $\mathcal{D}(S)$. The solid blue noose is the parent of the dashed red and green noose.

The width of a sphere cut decomposition is the maximum number of vertices on a noose. As $\mathcal{D}(S)$ has low outerplanarity, we have a low-width sphere cut decomposition, i.e., each noose goes through only few vertices. As a last step, we normalize the nooses by specifying their exact geometry, keeping their combinatorial structure intact; see the last row of Figure 6. At its core, the normalized nooses are polygons that alternate between vertices in S and Voronoi vertices, with some special treatment for the outer face of $\mathcal{D}(S)$.

Observe that the two examples (left and right column) illustrate different independent sets S , but share the green normalized noose. This is not a coincidence, but actually happens a lot. In fact, despite the potentially exponential number of independent sets, we can bound the total number of normalized nooses that we can obtain in this way by $n^{\mathcal{O}(1 + \frac{\log n}{r})}$. This follows from the fact that each noose visits only $\mathcal{O}(1 + \frac{\log n}{r})$ vertices due to Theorem 4. Thus, the normalized nooses are essentially polygons with $\mathcal{O}(1 + \frac{\log n}{r})$ corners and the set of potential corners is polynomial, as each corner is a vertex of G or a Voronoi vertex (which is defined by three vertices of G).

Our independent set algorithm works as follows. We start by listing all candidates for normalized nooses. We then use a dynamic program to combine nooses into a hierarchy of polygons. We make sure that the resulting hierarchy is valid in the sense that it corresponds to a sphere cut decomposition. Additionally, we enforce that it is *well-spaced*, which lets us show that all vertices of G visited by the polygons in the hierarchy form an independent set. Then, finding a maximum independent set is equivalent to maximizing the number of vertices visited by polygons in this polygon hierarchy, which can be done with the dynamic program.



■ **Figure 6** Illustration of our exact algorithm in Theorem 5.

3.4 Approximation algorithm

The approximation algorithm of Theorem 6 (see also the full version of this paper) works similarly to Lipton and Tarjan's version for planar graphs [28]: we repeatedly apply a separator until we get small patches and then solve each patch individually using the exact algorithm of Theorem 5. Vertices in the separators get ignored, so we get an $(1 - \varepsilon)$ -approximation when these make up only an ε fraction of the maximum independent set. We guarantee this using a lower bound of $\Omega(n/\ell)$ on the size of a maximum independent set, which follows from a proof that hyperbolic uniform disk graphs of ply ℓ always have a vertex of degree at most $4\ell - 1$, making them $(4\ell - 1)$ -degenerate.

4 Future directions

There are several intriguing future directions to explore in this space. First, is there a fully polynomial approximation scheme for INDEPENDENT SET for $r \geq 1$? If not, is there at least an approximation scheme that is polynomial in n , quasi-polynomial in $1/\varepsilon$, but independent of the ply?

Second, it would be interesting to explore separators in higher-dimensional hyperbolic spaces. The only case studied there is $r = \Theta(1)$ [23]. One may expect that separator size should also decrease with growing r . Is there a constant c_d such that INDEPENDENT SET can be solved in polynomial time in $\mathbb{H}^d$ when $r \geq c_d \log n$?

Third, we should explore uniform disk graphs in surfaces of constant curvature, i.e., on the sphere, flat torus, and especially hyperbolic surfaces. What is the size of the best clique-based separator for a uniform disk graph of radius r on a hyperbolic surface of genus g ?

Fourth, it would be interesting to investigate the complexity of problems other than INDEPENDENT SET, and study how their complexity scales with the radius of the disks, or equivalently, with the curvature of the underlying space.

Finally, it is unclear whether our bound on balanced separators is tight for all values of r . For constant r , regular hyperbolic tilings are hyperbolic uniform disk graphs with constant clique number and logarithmic treewidth, making our bound asymptotically tight. However, for larger radii, it could be possible to reduce the logarithmic factor.

References

- 1 Pankaj K. Agarwal, Marc J. van Kreveld, and Subhash Suri. Label placement by maximum independent set in rectangles. *Comput. Geom.*, 11(3-4):209–218, 1998. doi:10.1016/S0925-7721(98)00028-5.
- 2 Brenda S. Baker. Approximation algorithms for np-complete problems on planar graphs. *J. ACM*, 41(1):153–180, 1994. doi:10.1145/174644.174650.
- 3 Nicholas Bieker, Thomas Bläsius, Emil Dohse, and Paul Jungeblut. Recognizing unit disk graphs in hyperbolic geometry is $\exists\mathbb{R}$ -complete. *CoRR*, abs/2301.05550, 2023. doi:10.48550/arXiv.2301.05550.
- 4 Thomas Bläsius, Philipp Fischbeck, Tobias Friedrich, and Maximilian Katzmann. Solving vertex cover in polynomial time on hyperbolic random graphs. *Theory Comput. Syst.*, 67(1):28–51, 2023. doi:10.1007/S00224-021-10062-9.
- 5 Thomas Bläsius, Tobias Friedrich, and Maximilian Katzmann. Efficiently approximating vertex cover on scale-free networks with underlying hyperbolic geometry. *Algorithmica*, 85(12):3487–3520, 2023. doi:10.1007/S00453-023-01143-X.
- 6 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*, volume 254 of *LIPICs*, pages 13:1–13:17, 2023. doi:10.4230/LIPICs.STACS.2023.13.
- 7 Thomas Bläsius, Tobias Friedrich, and Anton Krohmer. Hyperbolic random graphs: Separators and treewidth. In *24th Annual European Symposium on Algorithms, ESA*, volume 57 of *LIPICs*, pages 15:1–15:16, 2016. doi:10.4230/LIPICs.ESA.2016.15.
- 8 Thomas Bläsius, Cedric Freiberger, Tobias Friedrich, Maximilian Katzmann, Felix Montenegro-Retana, and Marianne Thieffry. Efficient shortest paths in scale-free networks with underlying hyperbolic geometry. *ACM Transactions on Algorithms (TALG)*, 18(2):19:1–19:32, 2022. doi:10.1145/3516483.
- 9 Thomas Bläsius, Jean-Pierre von der Heydt, Sándor Kisfaludi-Bak, Marcus Wilhelm, and Geert van Wordragen. Structure and independence in hyperbolic uniform disk graphs. *Computing Research Repository (CoRR)*, abs/2407.09362, 2024. doi:10.48550/arXiv.2407.09362.
- 10 Hans L. Bodlaender. A partial k -arboretum of graphs with bounded treewidth. *Theor. Comput. Sci.*, 209(1-2):1–45, 1998. doi:10.1016/S0304-3975(97)00228-4.
- 11 Vincent Bouchitté, Frédéric Mazoit, and Ioan Todinca. Chordal embeddings of planar graphs. *Discret. Math.*, 273(1-3):85–102, 2003. doi:10.1016/S0012-365X(03)00230-9.

- 12 Karl Bringmann, Sándor Kisfaludi-Bak, Michal Pilipczuk, and Erik Jan van Leeuwen. On geometric set cover for orthants. In *27th Annual European Symposium on Algorithms, ESA*, volume 144 of *LIPIcs*, pages 26:1–26:18, 2019. doi:10.4230/LIPICS.ESA.2019.26.
- 13 Timothy M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *J. Algorithms*, 46(2):178–189, 2003. doi:10.1016/S0196-6774(02)00294-8.
- 14 Marek Cygan, Fedor V Fomin, Łukasz Kowalik, Daniel Lokshantov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.
- 15 Mark de Berg, Hans L. Bodlaender, Sándor Kisfaludi-Bak, Dániel Marx, and Tom C. van der Zanden. A framework for exponential-time-hypothesis-tight algorithms and lower bounds in geometric intersection graphs. *SIAM J. Comput.*, 49(6):1291–1331, 2020. doi:10.1137/20M1320870.
- 16 Mark de Berg, Sándor Kisfaludi-Bak, Morteza Monemizadeh, and Leonidas Theodorou. Clique-based separators for geometric intersection graphs. *Algorithmica*, 85(6):1652–1678, 2023. doi:10.1007/S00453-022-01041-8.
- 17 Frederic Dorn, Eelko Penninkx, Hans L. Bodlaender, and Fedor V. Fomin. Efficient exact algorithms on planar graphs: Exploiting sphere cut decompositions. *Algorithmica*, 58(3):790–810, 2010. doi:10.1007/s00453-009-9296-1.
- 18 Tobias Friedrich and Anton Krophmer. On the diameter of hyperbolic random graphs. *SIAM Journal on Discrete Mathematics*, 32(2):1314–1334, 2018. doi:10.1137/17M1123961.
- 19 Luca Gugelmann, Konstantinos Panagiotou, and Ueli Peter. Random hyperbolic graphs: Degree sequence and clustering. In *39th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 573–585, 2012. doi:10.1007/978-3-642-31585-5_51.
- 20 Sarel Har-Peled. Quasi-polynomial time approximation scheme for sparse subsets of polygons. In *30th Annual Symposium on Computational Geometry, SoCG*, page 120. ACM, 2014. doi:10.1145/2582112.2582157.
- 21 Dorit S. Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM*, 32(1):130–136, 1985. doi:10.1145/2455.214106.
- 22 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.
- 23 Sándor Kisfaludi-Bak. Hyperbolic intersection graphs and (quasi)-polynomial time. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 1621–1638. SIAM, 2020. doi:10.1137/1.9781611975994.100.
- 24 Sándor Kisfaludi-Bak. A quasi-polynomial algorithm for well-spaced hyperbolic TSP. *J. Comput. Geom.*, 12(2):25–54, 2021. doi:10.20382/JOCG.V12I2A3.
- 25 Paul Koebe. *Kontaktprobleme der konformen Abbildung*. Hirzel, 1936.
- 26 Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá. Hyperbolic geometry of complex networks. *Phys. Rev. E*, 82:036106, 2010. doi:10.1103/PhysRevE.82.036106.
- 27 Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.
- 28 Richard J. Lipton and Robert Endre Tarjan. Applications of a planar separator theorem. *SIAM Journal on Computing*, 9(3):615–627, 1980. doi:10.1137/0209046.
- 29 Madhav V. Marathe, H. Breu, Harry B. Hunt III, S. S. Ravi, and Daniel J. Rosenkrantz. Simple heuristics for unit disk graphs. *Networks*, 25(2):59–68, 1995. doi:10.1002/NET.3230250205.
- 30 Dániel Marx. On the optimality of planar and geometric approximation schemes. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007)*, pages 338–348. IEEE Computer Society, 2007. doi:10.1109/FOCS.2007.26.
- 31 Dániel Marx, Marcin Pilipczuk, and Michał Pilipczuk. On subexponential parameterized algorithms for steiner tree and directed subset TSP on planar graphs. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 474–484. IEEE Computer Society, 2018. doi:10.1109/FOCS.2018.00052.

- 32 Dániel Marx and Michal Pilipczuk. Optimal parameterized algorithms for planar facility location problems using Voronoi diagrams. *ACM Trans. Algorithms*, 18(2):13:1–13:64, 2022. doi:10.1145/3483425.
- 33 Colin McDiarmid and Tobias Müller. Integer realizations of disk and segment graphs. *J. Comb. Theory B*, 103(1):114–143, 2013. doi:10.1016/J.JCTB.2012.09.004.
- 34 Gary L. Miller, Shang-Hua Teng, William P. Thurston, and Stephen A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *Journal of the ACM*, 44(1):1–29, 1997. doi:10.1145/256292.256294.
- 35 Tobias Müller and Merlijn Staps. The diameter of KPKVB random graphs. *Advances in Applied Probability*, 51(2):358–377, 2019. doi:10.1017/apr.2019.23.
- 36 Neil Robertson and P. D Seymour. Graph minors. x. obstructions to tree-decomposition. *Journal of Combinatorial Theory, Series B*, 52(2):153–190, 1991. doi:10.1016/0095-8956(91)90061-N.
- 37 Paul D. Seymour and Robin Thomas. Call routing and the ratcatcher. *Comb.*, 14(2):217–241, 1994. doi:10.1007/BF01215352.