

Efficient Shortest Paths in Scale-Free Networks with Underlying Hyperbolic Geometry

THOMAS BLÄSIUS, Karlsruhe Institute of Technology

CEDRIC FREIBERGER, TOBIAS FRIEDRICH, MAXIMILIAN KATZMANN,

FELIX MONTENEGRO-RETANA, and MARIANNE THIEFFRY, Hasso Plattner Institute, University of Potsdam

A standard approach to accelerating shortest path algorithms on networks is the bidirectional search, which explores the graph from the start and the destination, simultaneously. In practice this strategy performs particularly well on scale-free real-world networks. Such networks typically have a heterogeneous degree distribution (e.g., a power-law distribution) and high clustering (i.e., vertices with a common neighbor are likely to be connected themselves). These two properties can be obtained by assuming an underlying hyperbolic geometry.

To explain the observed behavior of the bidirectional search, we analyze its running time on hyperbolic random graphs and prove that it is $\tilde{O}(n^{2-1/\alpha} + n^{1/(2\alpha)} + \delta_{\max})$ with high probability, where $\alpha \in (1/2, 1)$ controls the power-law exponent of the degree distribution, and δ_{\max} is the maximum degree. This bound is sublinear, improving the obvious worst-case linear bound. Although our analysis depends on the underlying geometry, the algorithm itself is oblivious to it.

CCS Concepts: • **Theory of computation** → **Random network models; Shortest paths**; • **Mathematics of computing** → **Random graphs; Paths and connectivity problems; Graph algorithms**;

Additional Key Words and Phrases: Random graphs, hyperbolic geometry, scale-free networks, bidirectional shortest path

ACM Reference format:

Thomas Bläsius, Cedric Freiberger, Tobias Friedrich, Maximilian Katzmann, Felix Montenegro-Retana, and Marianne Thieffry. 2022. Efficient Shortest Paths in Scale-Free Networks with Underlying Hyperbolic Geometry. *ACM Trans. Algorithms* 18, 2, Article 19 (March 2022), 32 pages.

<https://doi.org/10.1145/3516483>

A preliminary version of this article appeared in [4].

This research was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Grant No. 390859508.

Authors' addresses: T. Bläsius, Karlsruhe Institute of Technology, Am Fasanengarten 5, Karlsruhe, Baden-Württemberg, 76131, Germany; email: thomas.blaesius@kit.edu; C. Freiberger, T. Friedrich, M. Katzmann, F. Montenegro-Retana, and M. Thieffry, Hasso Plattner Institute, Prof.-Dr.-Helmert-Straße 2-3, Potsdam, Brandenburg, 14482, Germany; emails: cedric.freiberger@student.hpi.de, {tobias.friedrich, maximilian.katzmann}@hpi.de, {felix.montenegro-retana, marianne.thieffry}@student.hpi.de.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1549-6325/2022/03-ART19 \$15.00

<https://doi.org/10.1145/3516483>

1 INTRODUCTION

One of the most fundamental graph problems consists of finding a shortest path between two vertices in a network. Besides being of independent interest, many algorithms use shortest path queries as a subroutine. On unweighted graphs, such queries can be answered in linear time using a *breadth-first search (BFS)*. Though this is optimal in the worst case, it is not efficient enough when dealing with large networks or problems involving many shortest path queries.

A way to heuristically improve the run time, is to use a *bidirectional BFS* [26]. It runs two searches, simultaneously exploring the graph from the start and the destination. The shortest path is found once the two search spaces touch. Being one of the standard heuristics, the bidirectional BFS is widely used in practice (e.g., in route planning). On homogeneous networks (where most vertices have similar degrees, like road networks) this typically leads to a speedup factor of about two. However, on heterogeneous networks (having many vertices of low degree and only few vertices of very high degree, like social networks or the internet) experiments indicate that the bidirectional BFS yields an asymptotic running time improvement [10].

Despite being such a fundamental heuristic, theory completely fails its main purpose of predicting and explaining the observed behavior. The theoretical worst-case running time overshoots the observations by a lot. A more promising approach is the average case analysis by Borassi and Natale [10], which considers instances that are drawn from certain probability distributions instead of assuming the worst case. Their results are summarized in the first row of Table 1. The analysis covers a variety of random graph models. On one hand, these include homogeneous networks where the degree distribution has bounded variance, e.g., Erdős–Rényi random graphs. On the other hand, they also consider heterogeneous networks where the variance of the degree distribution is unbounded, e.g., Chung–Lu random graphs with power-law exponent $\beta \in (2, 3)$. However, the results, again, do not match what is observed in practice, as it predicts shorter running times on homogeneous networks than on heterogeneous ones.

The fundamental obstacle that prevents the average case analysis from producing convincing explanations is that the considered random graph models are not realistic. They assume that edges in the graph are independent of each other. However, real-world networks typically exhibit *locality*, i.e., edges in an evolving network tend to form between vertices that are already close in the network.

We resolve this discrepancy by modeling edge dependencies using geometry and extend the comparison in Table 1 by adding the second row. Generally, *geometric random graphs* are obtained by randomly distributing vertices in some metric space (e.g., the Euclidean plane) and connecting any two vertices with a probability that depends on their distance. In this framework, heterogeneous networks (i.e., networks on which the bidirectional BFS has been observed to perform particularly well) can be obtained by using the hyperbolic plane as the underlying geometry.

In this article, we analyze the bidirectional BFS on random graph models with an underlying geometry. We prove that, with high probability, the bidirectional BFS has a sublinear worst-case running time on the heterogeneous networks generated by the hyperbolic random graph model. Additionally, it is not hard to see why there is no asymptotic speedup on the homogeneous networks generated by the Euclidean random graph model. Both results match previous empirical observations. Finally, we interpret these insights and discuss how the heterogeneity of the degree distribution and an underlying geometry affect the running time of the bidirectional breadth first search.

Related Work. The research on scale-free networks has gained a lot of attention for quite some time now. Therefore, it is no surprise that the extensively studied problem of computing shortest paths has also been considered in the context of such graphs [1, 20, 23]. However, the bidirectional

Table 1. Probabilistic Bounds on the Running Time of the Bidirectional BFS Obtained by Analyzing Different Random Graph Models

	Homogeneous	Heterogeneous
Independent Edges	Bounded Variance	Unbounded Variance
	$m^{\frac{1}{2}+o(1)}$ [10]	$m^{\frac{4-\beta}{2}+o(1)}$ [10]
Underlying Geometry	Euclidean Random Graphs	Hyperbolic Random Graphs
	$\Theta(n)$ (Folklore)	$\tilde{O}(n^{2\frac{\beta-2}{\beta-1}} + n^{\frac{1}{\beta-1}})$ (This article)

The considered models (and associated results) are arranged by the heterogeneity of the corresponding degree distributions of the graphs and the (in)dependence of edges. Here, n and m denote the number of vertices and edges in the graph, respectively. The parameter $\beta \in (2, 3)$ denotes the power-law exponent of the degree distribution in the considered heterogeneous networks.

search that was introduced in 1969 [26] and that has since become one of the standard search heuristics, has only recently been examined on scale-free networks. In fact, there are only two theoretical explanations for the performance improvements obtained using this heuristic, both using an average-case analysis that considers one or more random graph models [10, 21].

A model that yields a better representation of real-world networks than the ones considered before, is the hyperbolic random graph model introduced by Krioukov et al. [19]. The generated graphs feature a heterogeneous degree distribution, high clustering, and a small diameter; properties that are often observed in real-world networks. These properties emerge naturally from the hyperbolic geometry. Moreover, the model is conceptually simple, which makes it accessible to mathematical analysis. For these reasons it has gained popularity in different research areas and has been studied from different perspectives.

From the network-science perspective, the goal is to gather knowledge about real-world networks. This is, for example, achieved by assuming that a real-world network has a hidden underlying hyperbolic geometry, which can be revealed by embedding it into the hyperbolic plane [2, 9].

From the mathematical perspective, the focus lies on studying structural properties. The degree distribution and clustering [17], diameter [15, 22], component structure [7, 18], clique size [6], and separation properties [5] have been studied successfully.

Additionally, there is the algorithmic perspective, which is the focus of this article. Usually algorithms are analyzed by proving worst-case running times. Though this is the strongest possible performance guarantee, it is rather pessimistic as practical instances rarely resemble worst-case instances. Techniques leading to a more realistic analysis include parameterized or average case complexity. The latter is based on the assumption that instances are drawn from a certain probability distribution. For hyperbolic random graphs, the maximum clique, as well as the minimum vertex cover can be computed in polynomial time [3, 6], and there are several algorithmic results based on the fact that hyperbolic random graphs have sublinear tree width [5]. Moreover, there is a compression algorithm that can store a hyperbolic random graph using $O(n)$ bits in expectation [12, 25]. Finally, a close approximation of the shortest path between two vertices can be found using greedy routing, which visits only $O(\log \log n)$ vertices for most start–destination pairs [13]. The downside of most of these algorithms is that they need to know the underlying geometry, i.e., the coordinates of each vertex, which is a rather unrealistic assumption for real-world networks. In contrast to that, we analyze an algorithm that is oblivious to the underlying geometry.




Fig. 1. Left: Points and several line segments in the native representation of the hyperbolic plane. A disk of radius r is centered at p_2 . Right: Geometric shapes and their intersections. Sector S has an angular width of φ .

Outline. After a brief introduction to hyperbolic random graphs in Section 2, we examine the bidirectional BFS in Section 3. We start by briefly arguing why the bidirectional BFS gives no asymptotic speedup over the standard BFS on Euclidean random graphs in Section 3.2. Afterwards, in Section 3.3, we rigorously analyze the bidirectional BFS on hyperbolic random graphs. Section 4 contains concentration bounds that were left out in Section 3 to improve readability. In Section 5, we conclude by comparing our theoretical results to empirical data and interpret them.

2 PRELIMINARIES

Let $G = (V, E)$ be an undirected and unweighted graph. We denote the number of vertices and edges with n and m , respectively. The *neighborhood* of a vertex $v \in V$ is $N(v) = \{w \in V \mid \{v, w\} \in E\}$. The *degree* of v is $\deg(v) = |N(v)|$. We denote the maximum degree with δ_{\max} . The soft \mathcal{O} -notation $\tilde{\mathcal{O}}$ suppresses poly-logarithmic factors in n .

2.1 The Hyperbolic Plane

The major difference between hyperbolic and Euclidean geometry is the exponential expansion of space. In the hyperbolic plane, a circle of radius r has area $2\pi(\cosh(r) - 1)$ and circumference $2\pi \sinh(r)$, with $\cosh(x) = (e^x + e^{-x})/2$ and $\sinh(x) = (e^x - e^{-x})/2$, both growing as $e^x/2 \pm o(1)$. To identify points, we use polar coordinates with respect to a designated origin O and a ray starting at O . A point p is uniquely determined by its *radius* r , which is the distance to O , and the *angle* (or *angular coordinate*) φ between the reference ray and the line through p and O . In illustrations, we use the *native representation*, obtained by interpreting the hyperbolic coordinates as polar coordinates in the Euclidean plane; see Figure 1 (left). Due to the exponential expansion, line segments bend toward the origin O . Let $p_1 = (r_1, \varphi_1)$ and $p_2 = (r_2, \varphi_2)$ be two points. The *angular distance* between p_1 and p_2 is the angle between the rays from the origin through p_1 and p_2 . Formally, it is $\Delta(\varphi_1, \varphi_2) = \pi - |\pi - |\varphi_1 - \varphi_2||$. The hyperbolic distance $\text{dist}(p_1, p_2)$ is given by

$$\cosh(\text{dist}(p_1, p_2)) = \cosh(r_1) \cosh(r_2) - \sinh(r_1) \sinh(r_2) \cos(\Delta(\varphi_1, \varphi_2)).$$

Note how the angular coordinates make simple definitions cumbersome as angles are considered modulo 2π , leading to a case distinction depending on where the reference ray lies. Whenever possible, we implicitly assume that the reference ray was chosen such that we do not have to compute modulo 2π . Thus, the above angular distance between p_1 and p_2 simplifies to $|\varphi_1 - \varphi_2|$. A third point $p = (r, \varphi)$ lies between p_1 and p_2 if $\Delta(\varphi, \varphi_1) + \Delta(\varphi, \varphi_2) = \Delta(\varphi_1, \varphi_2)$.

Throughout the article, we regularly use different geometric shapes that are mostly based on disks centered at the origin O , as can be seen in Figure 1 (right). With $D_r(p)$ we denote the disk of radius r around a point p , i.e., the set of points that have distance r to p . For disks, that are centered in the origin O , we simplify the notation and set $D_r := D_r(O)$. The restriction of a disk D_r to all points with angular coordinates in a certain interval is called *sector*, which we usually denote with the letter S . Its *angular width* is the length of this interval. For an arbitrary set of points A , we use $A|_{r_1}^{r_2}$ to denote the restriction of A to points with radii in $[r_1, r_2]$, i.e., $A|_{r_1}^{r_2} = A \cap (D_{r_2} \setminus D_{r_1})$.

2.2 Hyperbolic Random Graphs

A *hyperbolic random graph* is generated by drawing n points uniformly at random in a disk of the hyperbolic plane and connecting pairs of points whose distance is below a threshold. More precisely, the model depends on two parameters C and α that are assumed to be constants. The generated graphs have a power-law degree distribution with power-law exponent $\beta = 2\alpha + 1$ and a constant average degree depending on C . The parameter α is assumed to be in the range $(1/2, 1)$, yielding power-law exponents $\beta \in (2, 3)$. Exponents outside of this range are atypical for hyperbolic random graphs. For $\beta < 2$ the average degree of the generated networks diverges, while for $\beta > 3$ the graphs decompose into small components (of size sublinear in n), and the variance of the degree distribution is no longer unbounded. In contrast, it is unbounded for $\beta \in (2, 3)$, resulting in very heterogeneous degree distributions. Moreover, in this range the obtained networks have a giant component of size $\Omega(n)$ [8], and all other components have at most polylogarithmic size with high probability [15, Corollary 13]. Note that a bidirectional BFS could completely explore a non-giant component in $\tilde{O}(1)$ time and either return the shortest path (if both vertices are in the same non-giant component) or conclude that the vertices are in different components. Therefore, we only consider the case when the two considered vertices are both in the giant component in the remainder of the article.

When generating a hyperbolic random graph, the n points are sampled within the disk D_R of radius $R = 2 \log n + C$. For each vertex, the angular coordinate is drawn uniformly from $[0, 2\pi]$. Its radius r is sampled according to the probability density function $f(r)$, which can then be used to define the joint distribution of angles and radii $f(r, \varphi)$. They are given by

$$f(r) = \frac{\alpha \sinh(\alpha r)}{\cosh(\alpha R) - 1} = \Theta(e^{-\alpha(R-r)}) \quad \text{and} \quad f(r, \varphi) = \frac{1}{2\pi} f(r), \quad (1)$$

for $r \in [0, R]$. For $r > R$, $f(r) = f(r, \varphi) = 0$. Two vertices are connected by an edge if and only if their hyperbolic distance is at most R . The above probability distribution is a natural choice as the probability for a vertex ending up in a certain region is proportional to its area (at least for $\alpha = 1$). Note that the exponential growth in r reflects the fact that the area of a disk grows exponentially with the radius. It follows that a hyperbolic random graph has few vertices with high degree close to the center of the disk and many vertices with low degree near its boundary. The following lemma is common knowledge; for the sake of completeness we give a short proof.

LEMMA 2.1. *Let G be a hyperbolic random graph. Furthermore, let v_1, v_2 be two vertices with radii $r_1 \leq r_2 \leq R$, respectively, and with the same angular coordinate. Then, $N(v_2) \subseteq N(v_1)$.*

PROOF. Let $w \in N(v_2)$, i.e., $\text{dist}(v_2, w) \leq R$. Now consider the triangle $v_2 O w$, which is completely contained in the disk of radius R around w (since $\text{dist}(v_2, w) \leq R$ and $r(w) \leq R$). Since disks are convex and v_1 lies on the line from O to v_2 , it is part of the triangle and therefore also contained in this disk. Consequently, $\text{dist}(v_1, w) \leq R$ and thus $w \in N(v_1)$. \square

Given two vertices with fixed radii r_1 and r_2 , their hyperbolic distance grows with increasing angular distance. The maximum angular distance such that they are still adjacent [17, Lemma 3.1] is

$$\begin{aligned}\theta(r_1, r_2) &= \arccos\left(\frac{\cosh(r_1)\cosh(r_2) - \cosh(R)}{\sinh(r_1)\sinh(r_2)}\right) \\ &= 2e^{\frac{R-r_1-r_2}{2}}(1 + \Theta(e^{R-r_1-r_2})),\end{aligned}\quad (2)$$

assuming $r_2 \geq R - r_1$. Otherwise, we have $r_1 + r_2 < R$, meaning two vertices with these radii are adjacent, independent of their angular distance.

The probability that a sampled vertex falls into a given subset $A \subseteq D_R$ of the disk is given by its probability measure $\mu(A) = \iint_A f(r, \varphi) d\varphi dr$, which can be thought of as the area of A . There are two types of regions we encounter regularly: disks D_r with radius r centered at the origin and disks $D_R(r, \varphi)$ of radius R centered at a point (r, φ) . Note that the measure of $D_R(r, \varphi)$ gives the probability that a random vertex lies in the neighborhood of a vertex with position (r, φ) . Gugelmann et al. [17, Lemma 3.2] showed that

$$\mu(D_r) = e^{-\alpha(R-r)}(1 + o(1)), \text{ and} \quad (3)$$

$$\mu(D_R(r, \varphi)) = \Theta(e^{-r/2}). \quad (4)$$

For a given region $A \subseteq D_R$ of the disk, let X_1, \dots, X_n be random variables with $X_i = 1$ if vertex i lies in A and $X_i = 0$ otherwise. Then $X = \sum_{i=1}^n X_i$ is the number of vertices lying in A . By the linearity of expectation, we obtain that the expected number of vertices in A is $\mathbb{E}[X] = \sum_{i=1}^n \mathbb{E}[X_i] = n\mu(A)$.

Often, determining the expected value of a random variable is not sufficient to obtain meaningful statements. Therefore, we additionally classify events depending on how likely they are to occur. We say that an event holds *with high probability*, if it occurs with probability $1 - O(1/n)$. Moreover, we say that an event holds *asymptotically almost surely* if it occurs with probability $1 - o(1)$.

To show that certain random variables are concentrated around their expectation (i.e., with high probability the outcome does not deviate much from the expected value) we regularly use the following Chernoff–Hoeffding bound.

THEOREM 2.2 (CHERNOFF BOUND [14, THEOREM 1.1]). *Let X_1, \dots, X_n be independent random variables with $X_i \in \{0, 1\}$ and let X be their sum. Then,*

$$\begin{aligned}\Pr[X > t] &\leq 2^{-t} && \text{for } t > 2e\mathbb{E}[X], \text{ and} \\ \Pr[X < (1 - \epsilon)\mathbb{E}[X]] &\leq e^{-\epsilon^2/2\mathbb{E}[X]} && \text{for } \epsilon \in (0, 1).\end{aligned}$$

Usually, it suffices to show that a random variable does not exceed a certain upper bound or drop below a lower bound with high probability. The following corollaries show that sufficiently large upper and lower bounds on the expected value suffice to obtain concentration.

COROLLARY 2.3. *Let X_1, \dots, X_n be independent random variables with $X_i \in \{0, 1\}$ and let X be their sum. Furthermore, let $f(n) = \Omega(\log(n))$ be such that $\mathbb{E}[X] \leq f(n)$ and let c be a constant. Then, $X = O(f(n))$ holds with probability $1 - O(n^{-c})$.*

PROOF. We prove the statement by showing that the probability for the complementary event (i.e., X is more than a constant factor larger than $f(n)$) is $O(n^{-c})$ for any c . Since $\mathbb{E}[X] \leq f(n)$, we can choose a constant c_1 sufficiently large such that $c_1 f(n) > 2e\mathbb{E}[X]$. Thus, by Theorem 2.2 it holds that

$$\Pr[X > c_1 f(n)] \leq 2^{-c_1 f(n)}.$$




Fig. 2. Visualization of the proof of Lemma 2.5. When constrained to the disk D_R , the disk D' (bold) with center P' at radius r is completely contained in two consecutive disks D^i and D^{i+1} (red region). Point P_i is between P and P' .

Moreover, we have $f(n) = \Omega(\log n)$. Consequently, there exists another constant c_2 such that $f(n) \geq c_2 \log n$ for sufficiently large n . We obtain

$$\Pr[X > c_1 f(n)] \leq 2^{-c_1 c_2 \log n} \leq n^{-c_1 c_2},$$

for n sufficiently large. Finally, we can choose c_1 such that $c_1 > c/c_2$, which yields the claim. \square

COROLLARY 2.4. *Let X_1, \dots, X_n be independent random variables with $X_i \in \{0, 1\}$ and let X be their sum. Furthermore, let $f(n) = \omega(\log n)$ be such that $f(n) \leq \mathbb{E}[X]$ and let c be a constant. Then, $X \in \Omega(f(n))$ holds with probability $1 - O(n^{-c})$.*

PROOF. Analogous to the proof of Corollary 2.3 we prove the statement by showing that the probability for the complementary event (i.e., X is more than a constant factor smaller than $f(n)$) is $O(n^{-c})$ for any c . Let ε be a constant with $\varepsilon \in (0, 1)$. The following inequalities are obtained by first using the fact that $f(n) \leq \mathbb{E}[X]$, applying the second statement of Theorem 2.2, again applying $f(n) \leq \mathbb{E}[X]$, and finally using $f(n) \in \omega(\log n)$:

$$\Pr[X < (1-\varepsilon)f(n)] \leq \Pr[X < (1-\varepsilon)\mathbb{E}[X]] \leq e^{-\varepsilon^2/2 \cdot \mathbb{E}[X]} \leq e^{-\varepsilon^2/2 \cdot f(n)} = e^{-\varepsilon^2/2 \cdot \omega(\log n)} = n^{-\omega(1)}. \quad \square$$

Finally, the following lemma shows that statements about the neighborhood of a vertex with fixed angular coordinate can be extended to hold for arbitrary angular coordinates, with a small penalty in certainty.

LEMMA 2.5. *Let G be a hyperbolic random graph, let $X_w \geq 0$ for $w \in V$ be random variables, and let $X(D) = \sum_{w \in D} X_w$ for $D \subseteq D_R$. Furthermore, let $D_R(r)$ be the set of disks of radius R with center at radius r . If for each $D \in D_R(r)$ it holds that $\Pr[X(D) \leq f(n)] \geq 1 - p$, then $\Pr[\forall D \in D_R(r) : X(D) \leq 2f(n)] \geq 1 - O(np)$.*

PROOF. Let $D' \in D_R(r)$ be a disk with radius R centered at radius r and arbitrary angular coordinate. To bound $X(D')$, we cover the disk D_R with a circular sequence of n' disks $D^1, \dots, D^{n'}$, such that D' is completely contained in two consecutive disks (when constrained to the whole disk D_R). That is, there exists an $i \in \{1, \dots, n'\}$ such that $D' \subseteq D^i \cup D^{i+1}$. Since $X_w \geq 0$ for all $w \in V$, it then holds that

$$X(D') = \sum_{w \in D'} X_w \leq \sum_{D^i \cup D^{i+1}} X_w \leq \sum_{w \in D^i} X_w + \sum_{w \in D^{i+1}} X_w = X(D^i) + X(D^{i+1}).$$

Since $\Pr[X(D) \leq f(n)] \geq 1 - p$ holds for each $D \in D_R(r)$, we can apply the union bound to conclude that $X(D^i) \leq f(n)$ holds for all $i \in \{0, \dots, n'\}$ with probability $1 - n'p$. Consequently, $X(D') \leq 2f(n)$ with probability $1 - n'p$.

To complete the proof, it remains to show that there exists such a sequence $D^1, \dots, D^{n'}$ with $n' \in O(n)$. See Figure 2 for an illustration of how the sequence is constructed. All disks D^i for $i \in \{1, \dots, n'\}$ have their center at radius r . The center of the first disk is placed at angular coordinate 0, and each subsequent disk is placed at an angular distance of $2\theta(r, R)$ (see Equation (2)) to its predecessor until the whole disk is covered. Note that, as a consequence, the boundaries of two consecutive disks intersect at the boundary of the whole disk D_R .

Let P' be the center of D' . To see that D' is contained in two consecutive disks D^i and D^{i+1} (when constrained to the whole disk D_R), first note that there exists an $i \in \{1, \dots, n'\}$ such that P' is between the centers P_i and P_{i+1} of two consecutive disks D^i and D^{i+1} . We show that any point $P \in D'$ is contained in $D^i \cup D^{i+1}$. Clearly, $D^i \cup D^{i+1}$ contains all points between P_i and P_{i+1} (blue region in Figure 2). For the case where P does not lie between P_i and P_{i+1} , assume without loss of generality, that P_i is between P and P' , as depicted in Figure 2. Since $\text{dist}(P, P') \leq R$ and since P' and P_i have the same radius but P_i is between P and P' , it follows that $\text{dist}(P, P_i) \leq R$, and thus $P \in D^i$. Finally, it remains to show that $n' = O(n)$ disks are sufficient to cover the whole disk D_R . Since two consecutive disks are placed at an angular distance of $2\theta(r, R)$, we need $n' = 2\pi / (2\theta(r, R)) = O(1/\theta(r, R))$ disks. Since $\theta(r, R) \geq \theta(R, R)$, it follows that $n' = O(1/\theta(R, R)) = O(e^{R/2})$ due to Equation (2). Substituting $R = 2 \log(n) + C$ then yields the claim. \square

3 BIDIRECTIONAL BREADTH-FIRST SEARCH

In this section, we analyze the running time of the bidirectional BFS and obtain an upper bound on the maximum running time over all possible start–destination pairs. Our results are summarized in the following main theorem.

THEOREM 3.1. *Let G be a hyperbolic random graph. With high probability the shortest path between any two vertices in G can be computed in $\tilde{O}(n^{2-1/\alpha} + n^{1/(2\alpha)} + \delta_{\max})$ time.*

We note that this bound on the running time also holds in expectation. Our bound fails with probability $O(1/n)$, in which case the worst-case running time is still bounded by the size of the hyperbolic random graph, which is $O(n)$. Consequently, this case contributes $O(1)$ to the expectation, which is dominated by the above bound.

To prove Theorem 3.1, we make use of the hyperbolic geometry in the following way; see Figure 3. As long as the two searches visit only low-degree vertices, all explored vertices lie within a small region, i.e., the searches operate locally. Once the searches visit high-degree vertices closer to the center of the hyperbolic disk (green area in Figure 3), it takes only few steps to complete the search, as hyperbolic random graphs have a densely connected core. Thus, we split our analysis in two phases: a first phase in which both searches advance toward the center and a second phase in which both searches meet in the center. Note that this strategy assumes that we know the coordinates of the vertices as we would like to stop a search once it reached the center. To resolve this




Fig. 3. Visualization of the two phases of each BFS in a hyperbolic random graph. Vertices that are visited during the first phase are red. The red edges denote the first encounter of a vertex in the inner disk $D_R|_0^\rho$ (green region). This corresponds to the first step in the second phase. The last step then leads to a common neighbor via the blue edges.

issue, we first show in Section 3.1 that there exists an alternation strategy that is oblivious to the geometry but performs not much worse than any other alternation strategy. We note that this result is independent of hyperbolic random graphs and thus interesting in its own right. Afterwards, in Section 3.2, we examine the performance of the bidirectional BFS on Euclidean random graphs, before focusing on hyperbolic random graphs in Section 3.3.

3.1 Bidirectional Search and Alternation Strategies

In an unweighted and undirected graph $G = (V, E)$, a BFS finds the shortest path between two vertices $s, t \in V$ by starting at s and exploring the graph in levels, where the i th level L_i^s contains the vertices with distance i to s . More formally, the BFS starts with the set $L_0^s = \{s\}$ on level 0. Assuming the levels L_0^s, \dots, L_i^s have been computed already, one obtains the next level L_{i+1}^s as the set of neighbors of vertices in level L_i^s that are not contained in earlier levels. Computing L_{i+1}^s from L_i^s is called an *exploration step*, obtained by *exploring the edges* between vertices in L_i^s and L_{i+1}^s .

The bidirectional BFS runs two BFSs simultaneously. The *forward search* starts at s and the *backward search* starts at t . The shortest path between the two vertices can then be obtained, once the search spaces of the forward and backward search touch. Since the two searches cannot actually be run simultaneously, they alternate depending on their progress. When exactly the two searches alternate is determined by the *alternation strategy*. Note that we only swap after full exploration steps, i.e., we never explore only half of level i of one search before continuing with the other. This has the advantage that we can be certain to know the shortest path once a vertex is found by both searches.

In the following we define the *greedy alternation strategy* as introduced by Borassi and Natale [10] and show that it is not much worse than any other alternation strategy. Assume the latest levels of the forward and backward searches are L_i^s and L_j^t , respectively. Then the next exploration step of the forward search would cost time proportional to $c_i^s := \sum_{v \in L_i^s} \deg(v)$, while the cost for the backward search is $c_j^t := \sum_{v \in L_j^t} \deg(v)$. The *greedy alternation strategy* then greedily continues with the search that causes the fewer cost in the next exploration step, i.e., it continues with the forward search if $c_i^s \leq c_j^t$ and with the backward search otherwise.

THEOREM 3.2. *Let G be a graph with diameter d . If there exists an alternation strategy such that the bidirectional BFS explores $f(n)$ edges, then the bidirectional BFS with greedy alternation strategy explores at most $d \cdot f(n)$ edges.*

PROOF. Let A be the alternation strategy that explores only $f(n)$ edges. First note that the number of explored edges only depends on the number of levels explored by the two different searches and not on the actual order in which they are explored. Thus, if the greedy alternation strategy is different from A , we can assume without loss of generality that the greedy strategy performed more exploration steps in the forward search and fewer in the backward search compared to A . Let c^s and c^t be the number of edges explored by the forward and backward search, respectively, when using the greedy strategy. Moreover, let j be the last level of the backward search (which is actually not explored) and, accordingly, let c_j^t be the number of edges the next step in the backward search would have explored. Then, $c^t + c_j^t \leq f(n)$ as, when using A , the backward search still explores level j . Moreover, the forward search with the greedy strategy explores at most $c^t + c_j^t$ (and therefore at most $f(n)$) edges in each step, as exploring the backward search would be cheaper otherwise. Consequently, each step in the forward and backward search costs at most $f(n)$. As there are at most d steps in total, we obtain the claimed bound. \square

3.2 Bidirectional Search in Euclidean Random Graphs

Euclidean random graphs, commonly known as *random geometric graphs*, are generated by distributing n vertices uniformly at random in the unit square $[0, 1]^2$ and connecting any two vertices if the Euclidean distance between them is at most some threshold $R \in \mathbb{R}$ [24]. One can imagine, that each vertex is equipped with a disk of radius R and an edge is added to all other vertices that lie in this disk. The threshold R affects the properties of the generated network and in order to obtain graphs with a giant component of linear size (as is the case for hyperbolic random graphs), R has to be chosen from the so called *supercritical regime* [24]. In contrast to hyperbolic random graphs, the uniform sampling of the vertices in the Euclidean space leads to a distribution where the number of vertices falling into each disk is roughly the same, which, in turn, leads to a homogeneous degree distribution.

We examine how a BFS explores such a graph, by considering the region of the plane containing the vertices visited after several exploration steps. For Euclidean random graphs with R chosen from the supercritical regime, it was shown that for two vertices at graph theoretic distance d , it holds that $R \cdot d$ is at most a constant factor larger than the Euclidean distance between them, if d is super-logarithmic [16]. Additionally, it is easy to see that the Euclidean distance between them can be at most $R \cdot d$. Therefore, we can assume that after k (sufficiently many) steps the region in the plane that contains the visited vertices resembles a disk of radius proportional to k . Since the area of a disk with radius r grows as πr^2 , the expected number of explored vertices is in $\Theta(nk^2)$ (since the vertices are distributed uniformly).

In this scenario, it is easy to see that the performance of a bidirectional BFS improves by a constant factor, compared to a standard BFS. Let s and t be two vertices with (sufficiently large)

graph theoretic distance d from each other. Then, the expected number of vertices explored by a standard BFS from s to t is $\Theta(nd^2)$. If we run two searches instead (one starting at s , the other at t), then the expected explored search space is minimized when the two BFSs touch after half as many steps, exploring two disks of half the radius. (Note that this holds independent of the chosen alternation strategy.) In that case the expected number of explored vertices is proportional to $2n\pi(d/2)^2$, which is again $\Theta(nd^2)$, indicating that the bidirectional variant yields no asymptotic speedup over the standard BFS.

In the remainder of this article we focus on the performance of the bidirectional BFS on hyperbolic random graphs. In contrast to Euclidean random graphs, they feature a heterogeneous degree distribution, leading to significant differences in the performance of the bidirectional BFS.

3.3 Bidirectional Search in Hyperbolic Random Graphs

To analyze the size of the search space of the bidirectional BFS in hyperbolic random graphs, we separate the whole disk D_R into two parts. One is the *inner disk* $D_R|_0^\rho$ centered at the origin. Its radius ρ is chosen in such a way that any two vertices in $D_R|_0^\rho$ have a common neighbor with high probability. The second part is the *outer band* $D_R|_\rho^R$, the remainder of the whole disk. A single BFS now explores the graph in two phases. In the first phase, the BFS explores vertices in the outer band. The phase ends, when the next vertex to be encountered lies in the inner disk. Once both BFSs completed the first phase, they only need at most two more steps for their search spaces to share a vertex. One step to encounter the vertex¹ in the inner disk and another step to meet at their common neighbor that any two vertices in the inner disk have with high probability; see Figure 3.

Note that this scenario describes the worst case. Depending on the positions of the two considered vertices the two searches may touch earlier, e.g., when both vertices are close to each other in the outer band or when at least one of them is already contained in the inner disk. However, since we want to determine an upper bound on the running time, we consider the case where both vertices lie in the outer band and the two searches touch in the inner disk. In the remainder of the article we only consider how one of the two searches explores the graph. The obtained bounds also hold for the other search, meaning the total search space increases only by a constant factor when considering both searches instead of only one.

For our analysis we assume an alternation strategy in which each search stops once it explored one additional level after finding the first vertex in the inner disk $D_R|_0^\rho$. Of course, this cannot be implemented without knowing the underlying geometry of the network. However, by Theorem 3.2 the search space explored using the greedy alternation strategy is only a poly-logarithmic factor larger, as the diameter of hyperbolic random graphs is poly-logarithmic with high probability [15].² The following lemma shows for which choice of ρ the above sketched strategy works.

LEMMA 3.3. *Let G be a hyperbolic random graph. With high probability, G contains a vertex that is adjacent to every other vertex in $D_R|_0^\rho$, for $\rho = \frac{1}{\alpha}(\log n - \log \log n)$.*

PROOF. Assume v is a vertex with radius at most $R - \rho$. Note that the distance between two points is upper bounded by the sum of their radii. Thus, every vertex in $D_R|_0^\rho$ has distance at most R to v , and is therefore adjacent to v . Hence, to prove the claim, it suffices to show the existence of this vertex v with radius at most $R - \rho$. As described in Section 2, the probability for a single vertex

¹Note that this vertex has a degree of $\tilde{\Omega}(n^{1-1/(2\alpha)})$ with high probability. Consequently, a non-giant component of size $\tilde{O}(1)$ is detected (at the latest) before exploring this vertex (see Section 2).

²We note that there is a tighter bound of $\mathcal{O}(\log(n))$ on the diameter of hyperbolic random graphs, which holds asymptotically almost surely [22].

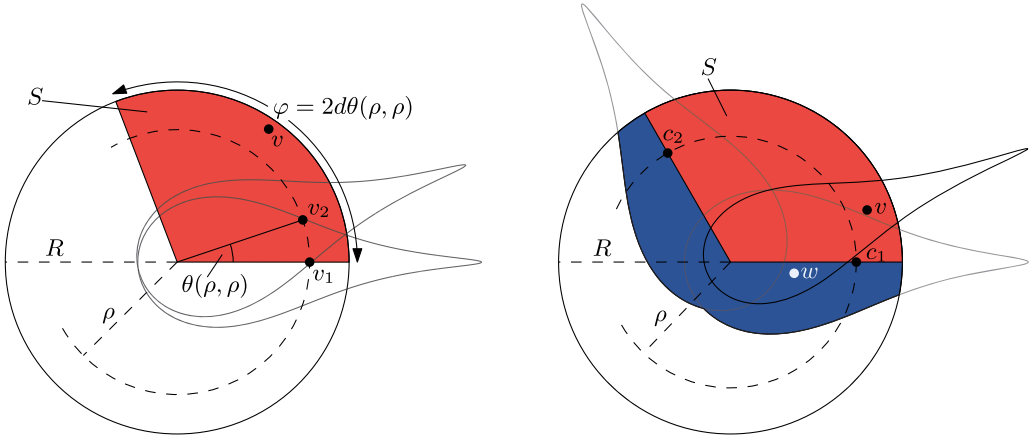


Fig. 4. Left: The sector S (red) of angular width φ contains the search space of a BFS starting at v , in the outer band $D_R|_\rho^R$. The vertices v_1 and v_2 are at maximum angular distance to still be adjacent. Right: Neighbor w of vertex v is in S (red) or a neighbor of c_1 or c_2 (blue).

to have radius at most $R - \rho$ is given by the measure $\mu(D_R|_0^{R-\rho})$. Using Equation (3), we obtain

$$\mu(D_R|_0^{R-\rho}) = e^{-\alpha\rho}(1 + o(1)) = \frac{\log n}{n}(1 + o(1)).$$

Thus, the probability that none of the n vertices lies in $D_R|_0^{R-\rho}$ is given by $(1 - \mu(D_R|_0^{R-\rho}))^n$. That is,

$$\Pr[\{v \in D_R|_0^{R-\rho}\} = \emptyset] = \left(1 - \frac{\log n}{n}(1 + o(1))\right)^n.$$

Since $(1 - x) \leq e^{-x}$ for all $x \in \mathbb{R}$, this term can be bounded by

$$\Pr[\{v \in D_R|_0^{R-\rho}\} = \emptyset] \leq e^{-\frac{\log(n)}{n}(1+o(1)) \cdot n} = e^{-\log(n)(1+o(1))} = n^{-(1+o(1))} = O(1/n).$$

Hence, there is at least one vertex with radius at most $R - \rho$ with high probability. \square

In the following, we first bound the search space explored in the first phase, i.e., before we enter the inner disk $D_R|_0^\rho$. Afterwards, we bound the search space explored in the second phase, which consists of two exploration steps. The first one to enter $D_R|_0^\rho$ and the second one to find a common neighbor, which exists due to Lemma 3.3.

3.3.1 Search Space in the First Phase. To bound the size of the search space in the outer band, we make use of the geometry in the following way. For two vertices in the outer band to be adjacent, their angular distance has to be small. Moreover, the number of exploration steps is bounded by the diameter of the graph. Thus, the maximum angular distance between vertices visited in the first phase cannot be too large. Note that the following lemma restricts the search to a sublinear portion of the disk, which we later use to show that also the number of explored edges is sublinear.

LEMMA 3.4. *With high probability, all vertices that a BFS on a hyperbolic random graph explores before finding a vertex with radius at most $\rho = \frac{1}{\alpha}(\log n - \log \log n)$ lie within a sector of angular width $\tilde{O}(n^{-(1/\alpha-1)})$.*

PROOF. For an illustration of the proof see Figure 4 (left). Recall from Section 2 that $\theta(r_1, r_2)$ denotes the maximum angular distance between two vertices of radii r_1 and r_2 such that they are

still adjacent. Since r_1 and r_2 only appear as negative exponents in the expression for $\theta(r_1, r_2)$ (see Equation (2)), this angle increases with decreasing radii. Thus, $\theta(r_1, r_2) \leq \theta(\rho, \rho)$ holds for all vertices in the outer band $D_R|_\rho^R$.

Now assume we start a BFS at a vertex $v \in D_R|_\rho^R$ and perform d exploration steps without leaving the outer band $D_R|_\rho^R$. Then, no explored vertex has angular distance more than $d\theta(\rho, \rho)$ from v . Thus, the whole search space lies within a disk sector of angular width $2d\theta(\rho, \rho)$. The number of steps d is at most poly-logarithmic as the diameter of a hyperbolic random graph is poly-logarithmic with high probability [15]. Using Equation (2) for $\theta(\rho, \rho)$, we obtain

$$\begin{aligned} \theta(\rho, \rho) &= 2e^{\frac{R-2\rho}{2}} (1 + \Theta(e^{R-2\rho})) \\ &= 2e^{C/2} n^{1-1/\alpha} \log(n)^{1/\alpha} (1 + \Theta((\log n/n^{1-\alpha})^{2/\alpha})) \\ &= O(n^{-(1/\alpha-1)} \log(n)^{1/\alpha}), \end{aligned}$$

which proves the claimed bound. \square

Note that the expected number of vertices in a sector S of angular width φ is linear in $n\varphi$ due to the fact that the angular coordinate of each vertex is chosen uniformly at random. Thus, Lemma 3.4 already shows that the expected number of vertices visited in the first phase of the BFS is $\tilde{O}(n^{2-1/\alpha})$, which is sublinear in n . It is also not hard to see that this bound holds with high probability (see Corollary 2.3). To also bound the number of explored edges, we sum the degrees of vertices in S . It is not surprising that this yields the same asymptotic bound in expectation, as the expected average degree in a hyperbolic random graph is constant. However, to obtain meaningful results, we need a bound that holds with high probability. Though we can use techniques similar to those that have been used to show that the average degree of the whole graph is constant with high probability [11, 17], the situation is complicated by the restriction to a sublinear portion of the disk. Nonetheless, we obtain the following theorem.

THEOREM 3.5. *Let G be a hyperbolic random graph. The degrees of vertices in every sector of angular width φ sum to $\tilde{O}(\varphi n + n^{1/(2\alpha)} + \delta_{\max})$ with high probability if $\varphi = \Omega(\log(n)^2/n^{1/2})$.*

We note that δ_{\max} has to be included here, as the theorem states a bound for every sector, and thus in particular for sectors containing the vertex of maximum degree. Recall, that $\delta_{\max} = \tilde{O}(n^{1/(2\alpha)})$ holds almost surely [17]. Moreover, we note that the condition $\varphi = \Omega(\log(n)^2/n^{1/2})$ is crucial for our proof, i.e., the angular width of the sector has to be sufficiently large for the concentration bound to hold. We note that, depending on α , the angular width determined in Lemma 3.4 may be smaller than this lower bound. However, if this is the case, we can choose $\varphi = \tilde{O}(n^{-1/2})$ as an upper bound for the angular width of the sector and obtain $\tilde{O}(\varphi n) = \tilde{O}(n^{1/2}) = \tilde{O}(n^{1/(2\alpha)})$ for $\alpha \in (1/2, 1)$. Consequently, the bound holds for the previously determined angular width $\tilde{O}(n^{-(1/\alpha-1)})$ for all $\alpha \in (1/2, 1)$.

As the proof for Theorem 3.5 is rather technical, we defer it to Section 4. Together with Lemma 3.4, we obtain the following corollary. Note that since $\alpha \in (1/2, 1)$, this shows that the running time spend in the first phase (not accounting for the maximum degree) is sublinear in n with high probability.

COROLLARY 3.6. *On a hyperbolic random graph, the first phase of the bidirectional search explores with high probability only $\tilde{O}(n^{2-1/\alpha} + n^{1/(2\alpha)} + \delta_{\max})$ many edges.*

3.3.2 Search Space in the Second Phase. The first phase of the BFS is completed when the next vertex to be encountered lies in the inner disk. Thus, the second phase consists of only two exploration steps. One step to encounter the vertex in the inner disk and another step to meet the

other search. Thus, to bound the running time of the second phase, we have to bound the number of edges explored in these two exploration steps. To do this, let V_1 be the set of vertices encountered in the first phase. Recall that all these vertices lie within a sector S of angular width $\varphi = \tilde{O}(n^{-(1/\alpha-1)})$ (Lemma 3.4). The number of explored edges in the second phase is then bounded by the sum of degrees of all neighbors $N(V_1)$ of vertices in V_1 . To bound this sum, we divide the neighbors of V_1 into two categories: $N(V_1) \cap S$ and $N(V_1) \setminus S$. Note that we already bounded the sum of degrees of vertices in S for the first phase (see Theorem 3.5), which clearly also bounds this sum for $N(V_1) \cap S$. Thus, it remains to bound the sum of degrees of vertices in $N(V_1) \setminus S$.

To bound this sum, we introduce two *hypothetical vertices* (i.e., vertices with specific positions that are not actually part of the graph) c_1 and c_2 such that every vertex in $N(V_1) \setminus S$ is a neighbor of c_1 or c_2 . Then, it remains to bound the sum of degrees of neighbors of these two vertices. To define c_1 and c_2 , recall that the first phase was restricted to points in the sector S that have a radius greater than ρ , i.e., all vertices in V_1 lie within $S|_\rho^R$. The hypothetical vertices c_1 and c_2 are basically positioned at the corners of this region, i.e., they both have radius ρ , and they assume the maximum and minimum angular coordinate within S , respectively. Figure 4 (right) shows these positions. We obtain the following.

LEMMA 3.7. *Let G be a hyperbolic random graph, let S be a sector, and let $v \in S|_\rho^R$ be a vertex. Then, every neighbor of v lies in S or is a neighbor of one of the hypothetical vertices c_1 or c_2 .*

PROOF. Let $v = (r, \varphi) \in S|_\rho^R$ and $w \in N(v) \setminus S$. Without loss of generality, assume that c_1 lies between v and w , as is depicted in Figure 4 (right). Now consider the point $v' = (\rho, \varphi)$ obtained by moving v to the same radius as c_1 . According to Lemma 2.1, we have $N(v) \subseteq N(v')$. In particular, it holds that $w \in N(v')$ and therefore $\text{dist}(v', w) \leq R$. Since v' and c_1 have the same radial coordinate and c_1 is between v' and w , it follows that $\text{dist}(c_1, w) \leq R$. \square

By the above argument, it remains to sum the degrees of neighbors of c_1 and c_2 . In the following, we show that the degrees of the neighbors of a vertex with radius r sum to $\Theta(ne^{-(\alpha-1/2)r})$ in expectation. We note that, for large values of r , i.e., for a vertex lying close to the boundary of the disk, this term is surprisingly large. This is due to the fact that, although vertices near the center of the disk are rather unlikely to exist in the first place, their degree would be sufficiently large such that they dominate the expected degree sum.

LEMMA 3.8. *Let G be a hyperbolic random graph. The degrees of the neighbors of a vertex v sum to $\Theta(ne^{-(\alpha-1/2)r(v)})$ in expectation.*

PROOF. Let Z_v be the sum of the degrees of the neighbors of v , which is a random variable that depends on the positions of all vertices in the graph. Formally, we can express Z_v by assigning each vertex $w \in V \setminus \{v\}$ two random variables X_w and Y_w . The first is an indicator random variable with $X_w = 1$ if w is a neighbor of v and $X_w = 0$ otherwise. Additionally, the random variable Y_w denotes the degree of w . The sum of the degrees of the neighbors of v can then be written as

$$Z_v = \sum_{w \in V \setminus \{v\}} X_w \cdot Y_w.$$

The expected value of Z_v is given by

$$\mathbb{E}[Z_v] = \mathbb{E} \left[\sum_{w \in V \setminus \{v\}} X_w \cdot Y_w \right] = \sum_{w \in V \setminus \{v\}} \mathbb{E}[X_w \cdot Y_w],$$

where the second equality holds due to the linearity of expectation. To compute the expected value of $X_w \cdot Y_w$ we can apply the law of total expectation and obtain

$$\mathbb{E}[Z_v] = \sum_{w \in V \setminus \{v\}} \sum_{x \in \{0,1\}} \mathbb{E}[X_w \cdot Y_w \mid X_w = x] \cdot \Pr[X_w = x].$$

Clearly, the case where $X_w = 0$ does not contribute anything to the sum, which can thus be simplified as

$$\mathbb{E}[Z_v] = \sum_{w \in V \setminus \{v\}} \mathbb{E}[Y_w \mid X_w = 1] \cdot \Pr[X_w = 1].$$

Recall that $X_w = 1$ denotes the event where w is a neighbor of v . That is,

$$\Pr[X_w = 1] = \Pr[w \in N(v)] = \Pr[w \in D_R(v)] = \mu(D_R(v)).$$

Moreover, recall that Y_w denotes the random variable representing the degree of w . Consequently, we can now write $\mathbb{E}[Z_v]$ as

$$\mathbb{E}[Z_v] = \sum_{w \in V \setminus \{v\}} \mu(D_R(v)) \cdot \mathbb{E}[Y_w \mid w \in D_R(v)] = (n-1) \cdot \mu(D_R(v)) \cdot \mathbb{E}[\deg(w) \mid w \in D_R(v)]. \quad (5)$$

We continue by computing the expected degree of a vertex w *conditioned* on the fact that it is contained in $D_R(v)$. To this end, we first consider the expected value without the condition, analogous to how it was done previously [17] (see the proof of Theorem 2.3), and afterwards explain how to incorporate the condition. The expected degree of a vertex w with fixed radius r is given by

$$\mathbb{E}[\deg(w) \mid r(w) = r] = (n-1)\mu(D_R(w)).$$

To obtain the expected degree of w without fixing its radius (or angle for that matter) we then integrate $\mathbb{E}[\deg(w) \mid r(w) = r \wedge \varphi(w) = \varphi] \cdot f(r, \varphi)$ (note the joint distribution) over the whole disk. That is,

$$\begin{aligned} \mathbb{E}[\deg(w)] &= \iint_{D_R} \mathbb{E}[\deg(w) \mid r(w) = r \wedge \varphi(w) = \varphi] \cdot f(r, \varphi) \, d\varphi \, dr \\ &= \iint_{D_R} \mathbb{E}[\deg(w) \mid r(w) = r] \cdot f(r, \varphi) \, d\varphi \, dr, \end{aligned}$$

where the second step follows from the fact that the expected degree of a vertex is independent of its angular coordinate.

It remains to include the condition on the fact that w cannot be anywhere in the whole disk but lies in $D_R(v)$ instead. First, we have to accommodate for the fact that if w is a neighbor of v , then conversely v is also a neighbor of w . Consequently, we know that w has at least one neighbor, which we reflect in the expected value by introducing the condition on the position $p(v)$ of v . Moreover, in general the conditional expectation of a random variable X conditioned on an event A (with $\Pr[A] > 0$) is given by $\mathbb{E}[X \mid A] = \int_{-\infty}^{\infty} x f_{X|A}(x) \, dx$, where $f_{X|A}$ is defined as

$$f_{X|A}(x) = \begin{cases} \frac{f_X(x)}{\Pr[A]}, & x \in A, \\ 0, & x \notin A. \end{cases}$$

Therefore, the above expression for the expected degree of w can be adjusted to include the condition as

$$\mathbb{E}[\deg(w) \mid w \in D_R(v)] = \iint_{D_R(v)} \mathbb{E}[\deg(w) \mid r(w) = r \wedge p(v) = (r, \varphi)] \frac{f(r, \varphi)}{\Pr[w \in D_R(v)]} \, d\varphi \, dr.$$




Fig. 5. Situation in the proof of Lemma 3.8. Vertex w is a neighbor of v . To integrate $D_R(v) \cap D_R(w)$, we split the region into two parts: $D_R(v)|_0^{R-r(v)} = D_{R-r}(v)$ (blue) and $D_R(v)|_{R-r(v)}^R$ (red).

Note that the probability $\Pr[w \in D_R(v)]$ in the denominator is, again, the measure of the disk of radius R centered at v . Substituting this expression in Equation (5) for the expected sum $\mathbb{E}[Z_v]$ of the degrees of the neighbors of v , we get

$$\begin{aligned} \mathbb{E}[Z_v] &= (n-1) \cdot \mu(D_R(v)) \cdot \mathbb{E}[\deg(w) \mid w \in D_R(v)] \\ &= (n-1) \cdot \mu(D_R(v)) \cdot \iint_{D_R(v)} \mathbb{E}[\deg(w) \mid r(w) = r \wedge p(v) = (r, \varphi)] \frac{f(r, \varphi)}{\mu(D_R(v))} d\varphi dr \\ &= (n-1) \cdot \iint_{D_R(v)} \mathbb{E}[\deg(w) \mid r(w) = r \wedge p(v) = (r, \varphi)] \cdot f(r, \varphi) d\varphi dr. \end{aligned}$$

To compute the integral, we determine the expected degree of w conditioned on the fact that $r(w) = r$ and on the position of v , which is a neighbor of w deterministically. Therefore, we obtain the expected degree by adding 1 (for v) to the expected number of vertices among the remaining $V \setminus \{v, w\}$ that are sampled into $D_R(w)$ and obtain

$$\mathbb{E}[\deg(v) \mid r(w) = r \wedge p(v) = (r, \varphi)] = 1 + (n-2)\mu(D_R(w)),$$

which is $1 + \Theta(ne^{-r/2})$ due to Equation (4). Note that $\Theta(ne^{-r/2})$ is $\Omega(1)$ for all $r \in [0, R]$, allowing us to further simplify the expected value to

$$\mathbb{E}[\deg(v) \mid r(w) = r \wedge p(v) = (r, \varphi)] = \Theta(ne^{-r/2}).$$

Moreover, recall that $f(r, \varphi) = 0$ for $r > R$ and that it can otherwise be bounded by $f(r, \varphi) = \Theta(e^{-\alpha(R-r)})$ (see Equation (1)). We obtain

$$\begin{aligned} \mathbb{E}[Z_v] &= \Theta \left((n-1) \cdot \iint_{D_R(v) \cap D_R} ne^{-r/2} \cdot e^{-\alpha(R-r)} d\varphi dr \right) \\ &= \Theta \left(n^2 e^{-\alpha R} \cdot \iint_{D_R(v) \cap D_R} e^{(\alpha-1/2)r} d\varphi dr \right). \end{aligned}$$

We can now split the integral into two parts: one containing the disk $D_R(v)|_0^{R-r(v)} = D_{R-r}(v)$ and the other containing the remainder of $D_R(v) \cap D_R$, which is given by $D_R(v)|_{R-r(v)}^R$ (see Figure 5).

For the second part, we can use Equation (2) to bound the angle $\theta(r(v), r)$ up to which we need to integrate depending on r . As a result, we get

$$\mathbb{E}[Z_v] = \Theta \left(n^2 e^{-\alpha R} \cdot \left(\int_0^{R-r(v)} \int_0^{2\pi} e^{(\alpha-1/2)r} d\varphi dr + \int_{R-r(v)}^R \int_0^{\theta(r(v), r)} e^{(\alpha-1/2)r} d\varphi dr \right) \right).$$

Regarding the first part of the sum, note that evaluating the inner integral only contributes a constant factor that can be dropped due the Θ -notation. Computing the outer integral then yields $\Theta(e^{(\alpha-1/2)(R-r(v))})$. For the second part of the sum we, again, first evaluate the inner integral and substitute $\theta(r(v), r) = \Theta(e^{(R-r(v)-r)/2})$ (see Equation (2)). We obtain

$$\mathbb{E}[Z_v] = \Theta \left(n^2 e^{-\alpha R} \cdot \left(e^{(\alpha-1/2)(R-r(v))} + e^{(R-r(v))/2} \int_{R-r(v)}^R e^{-(1-\alpha)r} dr \right) \right).$$

The last integral evaluates to $\mathcal{O}(e^{-(1-\alpha)(R-r(v))})$, which multiplied by the factor $e^{(R-r(v))/2}$ yields asymptotically the same expression as the first summand and we get

$$\mathbb{E}[Z_v] = \Theta \left(n^2 e^{-(\alpha-1/2)r(v)} \cdot e^{-R/2} \right).$$

Finally, we can substitute $R = 2 \log(n) + C$ to obtain the claimed bound $\mathbb{E}[Z_v] = \Theta(ne^{-(\alpha-1/2)r(v)})$. \square

For c_1 and c_2 , which both have radius ρ , the degrees of their neighbors thus sum to $\tilde{\mathcal{O}}(n^{1/(2\alpha)})$ in expectation. However, to actually prove Theorem 3.1, we need a bound that holds with high probability for all possible angular coordinates of c_1 and c_2 . As with the sum of the degrees in a sector, we prove a slightly weaker bound that matches the one in Theorem 3.1 and holds with high probability. We obtain the following lemma.

LEMMA 3.9. *Let G be a hyperbolic random graph and let v be a hypothetical vertex with radius $\rho = 1/\alpha(\log n - \log \log n)$ and arbitrary angular coordinate. The degrees of neighbors of v sum to $\tilde{\mathcal{O}}(n^{2-1/\alpha} + n^{1/(2\alpha)} + \delta_{\max})$ with high probability.*

Again, the proof is rather technical and thus deferred to Section 4. Together with the bounds on the sum of degrees in a sector of width $\varphi = \tilde{\mathcal{O}}(n^{-(1/\alpha-1)})$ (Theorem 3.5), we obtain the following corollary, which concludes the proof of Theorem 3.1.

COROLLARY 3.10. *On a hyperbolic random graph, the second phase of the bidirectional search explores with high probability only $\tilde{\mathcal{O}}(n^{2-1/\alpha} + n^{1/(2\alpha)} + \delta_{\max})$ many edges.*

4 CONCENTRATION BOUNDS FOR THE SUM OF VERTEX DEGREES

Here, we prove the concentration bounds that were announced in the previous section. For the first phase, we already know that the search space is contained within a sector S of sublinear width (Lemma 3.4). Thus, the running time in the first phase is bounded by the sum of vertex degrees in this sector. Moreover, all edges explored in the second phase also lie within the same sector S or are incident to neighbors of the two hypothetical vertices c_1 and c_2 (Lemma 3.7). Thus, the running time of the second phase is bounded by the sum of vertex degrees in S and in the neighborhood of c_1 and c_2 .

In both cases, we have to bound the sum of vertex degrees in certain areas of the disk, which can be done as follows. For each degree, we want to compute the number of vertices of this degree in the considered area and multiply it with the degree. As all vertices with a certain degree have roughly the same radius, we can separate the disk into small bands, one for each degree. Then summing over all degrees comes down to summing over all bands and multiplying the number of vertices

in this band with the corresponding degree. If we can prove that each of these values is highly concentrated (i.e., holds with probability $1 - \mathcal{O}(n^{-2})$), we obtain that the sum is concentrated as well (using the union bound). Unfortunately, this fails in two situations. For small radii, the number of vertices within the corresponding band (i.e., the number of high degree vertices) is too small to be concentrated. Moreover, for large radii the degree is too small to be concentrated around its expected value.

To overcome this issue, we partition the disk D_R into three parts. An inner part $D_R|_0^{\rho_I(\varphi)}$, containing all points of radius at most $\rho_I(\varphi)$, an outer part $D_R|_{\rho_O}^R$, containing all points of radius at least ρ_O , and a central part $D_R|_{\rho_I(\varphi)}^{\rho_O}$, containing all points in between. We choose $\rho_I(\varphi)$ such that the number of vertices with maximum degree in a sector part $S|_{\rho_I(\varphi)}^{\rho_O}$ of angular width φ is $\Omega(\log n)$, which ensures that for each vertex degree, the number of vertices with this degree is concentrated. Moreover, we choose ρ_O in such a way that the vertex degrees in $S|_{\rho_I(\varphi)}^{\rho_O}$ are sufficiently concentrated. To achieve this, we set

$$\rho_I(\varphi) = R - \frac{1}{\alpha} \left(\log(\varphi/(2\pi)) + \log n - \log \log n \right) \text{ and } \rho_O = R - (2 + \varepsilon) \log \log(n),$$

for any constant $\varepsilon \in (0, 1)$, and show concentration for the sum of the degrees in a sector and in the neighborhood of a vertex with radius ρ , separately for the three parts of the disk.

4.1 The Inner Part of the Disk

The inner part $D_R|_0^{\rho_I(\varphi)}$ contains vertices of high degree. It is not hard to see that there are only logarithmically many vertices with radius at most $\rho_I(\varphi)$.

LEMMA 4.1. *Let G be a hyperbolic random graph, let $\varphi \in [0, 2\pi]$ be an angle, and let $\xi > 0$ be a constant. A sector $S|_0^{\rho_I(\varphi)}$ of angular width $\xi\varphi \in [0, 2\pi]$ contains $\mathcal{O}(\log(n))$ vertices, with probability $1 - \mathcal{O}(n^{-c})$ for any constant c .*

PROOF. By Equation (3) the expected number of vertices in the disk $D_R|_0^{\rho_I(\varphi)}$ is given by $\mathbb{E}[|\{v \in D_R|_0^{\rho_I(\varphi)}\}|] = ne^{-\alpha(R-\rho_I(\varphi))}(1 + o(1))$. Since the angular coordinates of the vertices are distributed uniformly in $[0, 2\pi]$, the expected number of vertices in a sector portion $S|_0^{\rho_I(\varphi)}$ of angular width $\xi\varphi$ is

$$\begin{aligned} \mathbb{E}[|\{v \in S|_0^{\rho_I(\varphi)}\}|] &= \frac{\xi\varphi}{2\pi} ne^{-\alpha(R-\rho_I(\varphi))}(1 + o(1)) \\ &= \frac{\xi\varphi}{2\pi} ne^{-(\log(\varphi/2\pi) + \log n - \log \log n)}(1 + o(1)) \\ &= \xi \log(n)(1 + o(1)). \end{aligned}$$

Since $\xi > 0$ is constant, this bound is in $\Omega(\log n)$ and we can apply Corollary 2.3 to conclude that $|\{v \in S|_0^{\rho_I(\varphi)}\}| = \mathcal{O}(\log(n))$ holds with probability $1 - \mathcal{O}(n^{-c})$ for any constant c . \square

Note that, if $\varphi \in \Omega(1/n)$, we can choose at most $\mathcal{O}(n)$ sectors of width 2φ such that any sector of width φ lies completely in one of them. Thus, the probability that there exists a sector portion $S|_0^{\rho_I(\varphi)}$ where the number of vertices is super-logarithmic, is bounded by the probability that it is too large in at least one of these $\mathcal{O}(n)$ sectors (of twice the width). By choosing $\xi = 2$, we can apply Lemma 4.1 to conclude that a single sector $S|_0^{\rho_I(\varphi)}$ of twice the angular width contains at most $\mathcal{O}(\log(n))$ vertices with probability $1 - \mathcal{O}(n^{-2})$. Applying the union bound and incorporating the fact that the maximum degree in the graph is δ_{\max} , we can bound the number of edges in every such sector portion and obtain the following corollary.

COROLLARY 4.2. *Let G be a hyperbolic random graph. For every sector S of angular width $\varphi \in \Omega(1/n)$, the degrees of the vertices in $S|_0^{\rho_I(\varphi)}$ sum to $\tilde{O}(\delta_{\max})$ with high probability.*

Note that, in particular, the statement holds for the previously determined angle $\varphi = \tilde{O}(n^{-(1/\alpha-1)})$ for $\alpha \in (1/2, 1)$. Additionally, by setting $\varphi = 2\pi$, we can use Lemma 4.1 to bound the sum of the degrees of the high degree vertices in the neighborhood of a vertex with radius ρ .

COROLLARY 4.3. *Let G be a hyperbolic random graph. For every vertex v of radius ρ , the degrees of the neighbors of v in $D_R|_0^{\rho_I(2\pi)}$ sum to $\tilde{O}(\delta_{\max})$ with high probability.*

4.2 The Central Part of the Disk

For each possible vertex degree k , we want to compute the number of vertices with this degree in the central part $D_R|_{\rho_I(\varphi)}^{\rho_O}$. First note, that by Equation (4) a vertex with fixed radius has expected degree $\Theta(k)$ if this radius is $2 \log(n/k)$. Motivated by this, we define $r_k = 2 \log(n/k)$. To bound the sum of degrees in the central part $D_R|_{\rho_I(\varphi)}^{\rho_O}$, we use that vertices with radius significantly larger than r_k also have a smaller degree. To this end, we first prove that a vertex with degree k can actually not have a radius much larger than r_k . This has the advantage, that we can bound the number of degree- k vertices by bounding the number of vertices with these radii.

LEMMA 4.4. *Let G be a hyperbolic random graph. Then, for every constant $c > 0$, there exist constants $\kappa, \tau > 0$, such that all vertices with degree at least $k \geq \kappa \log n$ have radius at most $r_k + \tau$ with probability $1 - O(n^{-c})$.*

PROOF. To prove this lemma, it suffices to show that there exist constants $\kappa, \tau > 0$, such that the probability of a vertex with radius greater than $r_k + \tau$ having degree at least k , i.e., $\Pr[\exists v \in V: \deg(v) \geq k \wedge r(v) \geq r_k + \tau]$, is small. To obtain the following sequence of inequalities, we first use the union bound, then apply the definition of conditional probabilities, and finally use Lemma 2.1.

$$\begin{aligned} \Pr[\exists v \in V: \deg(v) \geq k \wedge r(v) \geq r_k + \tau] &\leq n \cdot \Pr[\deg(v) \geq k \wedge r(v) \geq r_k + \tau] \\ &\leq n \cdot \Pr[\deg(v) \geq k \mid r(v) \geq r_k + \tau] \\ &\leq n \cdot \Pr[\deg(v) \geq k \mid r(v) = r_k + \tau]. \end{aligned}$$

To prove the statement of the lemma, it remains to show that $\Pr[\deg(v) \geq k \mid r(v) = r_k + \tau]$ is sufficiently small, i.e., in $O(n^{-(c+1)})$.

Recall that, by Equation (4), the expected degree of a vertex with radius r is in $\Theta(ne^{-r/2})$. For a vertex v with radius $r_k + \tau$, we obtain $ne^{-(r_k + \tau)/2} = e^{-\tau/2}k$. It follows that there exists a constant $c' > 0$, such that $\mathbb{E}[\deg(v)] \leq c'e^{-\tau/2}k$. By choosing τ large enough we can ensure that $k \geq 2e\mathbb{E}[\deg(v)]$, allowing us to apply the Chernoff–Hoeffding bound in Theorem 2.2. We obtain $\Pr[\deg(v) \geq k] \leq 2^{-k}$. Finally, since $k \geq \kappa \log n$, we can choose κ such that this probability is bounded by $O(n^{-(c+1)})$. \square

We are now ready to bound the number of vertices in a sector that have degree at least k . As mentioned earlier, this bound only works for large k as the degree is not sufficiently concentrated otherwise. Moreover, the degree cannot be too large, as otherwise the number of vertices of this degree is not concentrated. The upper bound on k in the following lemma directly corresponds to our choice for $\rho_I(\varphi)$. Additionally, ρ_O is chosen such that the degrees of vertices with radii smaller than ρ_O meets the lower bound on k , i.e., the lemma holds for the central part $S|_{\rho_I(\varphi)}^{\rho_O}$.

LEMMA 4.5. *Let G be a hyperbolic random graph and let S be a sector of angular width φ . If $k = \omega(\log n)$ and $k = O((\varphi n / \log n)^{1/(2\alpha)})$, then the number of vertices in S with degree at least k is in $O(\varphi nk^{-2\alpha})$ with probability $1 - O(n^{-c})$ for any constant $c > 0$.*

PROOF. By Lemma 4.4 we know that, for any constant $c' > 0$, there are constants $\kappa, \tau > 0$ such that all vertices of degree at least $k \geq \kappa \log n$ have radius at most $r_k + \tau$, with probability $1 - \mathcal{O}(n^{-c'})$. Since $k = \omega(\log n)$ we have $k \geq \kappa \log n$ for large enough n and obtain that, with the same probability, all vertices of degree at least k that are in S are in $S|_0^{r_k + \tau}$. Since the angular width of S is φ and since the angular coordinates of the vertices are distributed uniformly, the expected number of vertices in $S|_0^{r_k + \tau}$ is given by $\varphi/(2\pi) \cdot n\mu(D_R|_0^{r_k + \tau})$. Now we can apply Equation (3), which states that a disk of radius $r_k + \tau$ centered at the origin has measure $e^{-\alpha(R-(r_k + \tau))}(1 + o(1))$ and obtain

$$\begin{aligned} \mathbb{E}[|\{v \in S|_0^{r_k + \tau}\}|] &= \frac{\varphi}{2\pi} n\mu(D_R|_0^{r_k + \tau}) \\ &= \frac{\varphi}{2\pi} n e^{-\alpha(R-(r_k + \tau))}(1 + o(1)) \\ &= \frac{\varphi}{2\pi} n e^{-2\alpha \log k - \alpha(C-\tau)}(1 + o(1)) \\ &= \Theta(\varphi n k^{-2\alpha}). \end{aligned}$$

Note that $k = \mathcal{O}((\varphi n / \log n)^{1/(2\alpha)})$ (which is a precondition of this lemma) implies that $\varphi n k^{-2\alpha} = \Omega(\log n)$. Thus, we can apply the Chernoff–Hoeffding bound in Corollary 2.3 to conclude that $|\{v \in S|_0^{r_k + \tau}\}| = \mathcal{O}(\varphi n k^{-2\alpha})$ holds with probability $1 - \mathcal{O}(n^{-c})$ for any constant $c > 0$. \square

Using these results, we can now bound the size of the search space in the central part $S|_{\rho_I(\varphi)}^{\rho_O}$ of our sector S , yielding the following lemma. (We note that the lower bound on φ that is a requirement of the following lemma, is weaker than the one we need for Theorem 3.5.)

LEMMA 4.6. *Let G be a hyperbolic random graph. For every sector S of angular width $\varphi \in \Omega(\log(n)^{2\alpha+1}/n)$, the degrees of the vertices in $S|_{\rho_I(\varphi)}^{\rho_O}$ sum to $\mathcal{O}(\varphi n)$ with high probability.*

PROOF. First note that, analogous to the argumentation about sectors in the inner part of the disk, we can choose at most $\mathcal{O}(n)$ sectors of width 2φ such that any sector of width φ lies completely in one of them. Thus, the probability that there exists a sector where the sum of the vertex degrees in the central part of the disk is too large, is bounded by the probability that it is too large in at least one of these $\mathcal{O}(n)$ sectors (of twice the width). In the following, we show for a single sector S that the probability that the sum is too large is $\mathcal{O}(n^{-2})$. The union bound then yields the claim, that the bound holds for every sector.

To sum the degrees of all vertices in S , think of a vertex v of degree $\deg(v)$ as a rectangle of height 1 and width $\deg(v)$. For a small graph, Figure 6 shows all such rectangles stacked on top of each other, sorted by their degree. Note that the sum of degrees is equal to the area under the function $g(x) = |V_x^S|$ where $V_x^S = \{v \in S \mid \deg(v) \geq x\}$ is the set of vertices in S that have degree at least x . Also note that the above considerations do not take into account that we sum only the degrees of vertices in the central part $S|_{\rho_I(\varphi)}^{\rho_O}$ of S . To resolve this, let k_{\min} and k_{\max} be the minimum and maximum degree of vertices in $S|_{\rho_I(\varphi)}^{\rho_O}$, respectively. One can see in Figure 6 that summing only those degrees that are larger than k_{\min} is equivalent to integrating over $|V_{\max(k_{\min}, x)}^S|$ instead of $|V_x^S|$. Thus, we can compute the sum of all degrees as

$$\sum_{v \in S|_{\rho_I(\varphi)}^{\rho_O}} \deg(v) \leq \sum_{\substack{v \in S, \\ k_{\min} \leq \deg(v) \leq k_{\max}}} \deg(v) = \int_0^{k_{\max}} |V_{\max(k_{\min}, x)}^S| dx = k_{\min} |V_{k_{\min}}^S| + \int_{k_{\min}}^{k_{\max}} |V_x^S| dx.$$

To compute this integral, we first calculate the minimum and maximum degrees k_{\min} and k_{\max} . Afterwards, we apply Lemma 4.5 to bound $|V_x^S|$. For the minimum degree k_{\min} , assume that vertex v has radius $\rho_O = R - (2 + \varepsilon) \log \log(n)$ for any constant $\varepsilon \in (0, 1)$. Using Equation (4) the expected




Fig. 6. Visualization of how the sum over the degrees can be turned into an integral (left). The same visualization but only the degrees of vertices with degree at least 3 are summed up (right).

degree of v is $\mathbb{E}[\deg(v)] = \Theta(ne^{-R/2+(1+\varepsilon/2)\log\log(n)}) = \Theta(\log(n)^{1+\varepsilon/2})$. Since $\varepsilon > 0$, this bound is $\omega(\log n)$, allowing us to apply the Chernoff–Hoeffding bounds in Corollaries 2.3 and 2.4 to conclude that $\deg(v) = \Theta(\log(n)^{1+\varepsilon/2})$ with high probability. Note that this only holds under the assumption that v has radius exactly ρ_O . However, by Lemma 2.1 all vertices with smaller radius have larger expected degree. Therefore, $\Theta(\log(n)^{1+\varepsilon/2})$ is a lower bound on the expected degrees of all such vertices, allowing us to apply Corollary 2.4 together with a union bound, to conclude that, with high probability, no vertex with smaller radius has smaller degree. Thus, with high probability, the minimum degree in $S_{\rho_I(\varphi)}^{\rho_O}$ is $k_{\min} = \Theta(\log(n)^{1+\varepsilon/2})$. Analogously, the bound on the maximum degree k_{\max} of a vertex in $S_{\rho_I(\varphi)}^{\rho_O}$ can be obtained as follows. Let v be a vertex with radius $\rho_I(\varphi) = R - 1/\alpha(\log(\varphi/2\pi) + \log n - \log\log n)$. The expected degree of v is $\mathbb{E}[\deg(v)] = \Theta((\varphi n / \log n)^{1/(2\alpha)})$ (Equation (4)). Since $\varphi \in \Omega(\log(n)^{2\alpha+1}/n)$, which is a precondition of this lemma, we can conclude that this bound on the expected degree of v is $\Omega(\log n)$, allowing us to apply Corollary 2.3 to conclude that $\mathbb{E}[\deg(v)] = O((\varphi n / \log n)^{1/(2\alpha)})$ holds with high probability. Again, this only holds under the assumption that v has radius exactly $\rho_I(\varphi)$. However, by Lemma 2.1 all vertices with larger radius have smaller expected degree. Therefore, $O((\varphi n / \log n)^{1/(2\alpha)})$ is a valid upper bound on all their expected degrees, allowing us to apply Corollary 2.3 together with a union bound, to conclude that no vertex with larger radius has larger degree. Thus, the maximum degree in $S_{\rho_I(\varphi)}^{\rho_O}$ is $k_{\max} = O((\varphi n / \log n)^{1/(2\alpha)})$ with high probability.

Using Lemma 4.5 we obtain $|V_x^S| = O(\varphi n x^{-2\alpha})$ with probability $1 - O(n^{-c})$ for any constant $c > 0$. Note that the requirements $x = \omega(\log n)$ and $x = O((\varphi n / \log n)^{1/(2\alpha)})$ in Lemma 4.5 are satisfied as $k_{\min} \leq x \leq k_{\max}$. By choosing $c = 2$ and applying the union bound over all degrees, we can conclude that, with high probability

$$\begin{aligned} \sum_{v \in S_{\rho_I(\varphi)}^{\rho_O}} \deg(v) &= O\left(\varphi n k_{\min}^{-(2\alpha-1)}\right) + O\left(\varphi n \cdot \int_{k_{\min}}^{k_{\max}} x^{-2\alpha} dx\right) \\ &= O\left(\varphi n k_{\min}^{-(2\alpha-1)}\right) + O\left(\varphi n \cdot k_{\min}^{-(2\alpha-1)} \left(1 - (k_{\min}/k_{\max})^{2\alpha-1}\right)\right). \end{aligned}$$

As $k_{\min} \leq k_{\max}$, this can be simplified to $O(\varphi n k_{\min}^{-(2\alpha-1)})$, which is $O(\varphi n)$ since $k_{\min} = \omega(\log n)$. \square

It remains to bound the sum of the degrees of vertices in the central part of the disk $D_R|_{\rho_I(2\pi)}^{\rho_O}$ that lie in the neighborhood of a vertex v with radius ρ , i.e., vertices lying in $D_R(v)$. Similar to

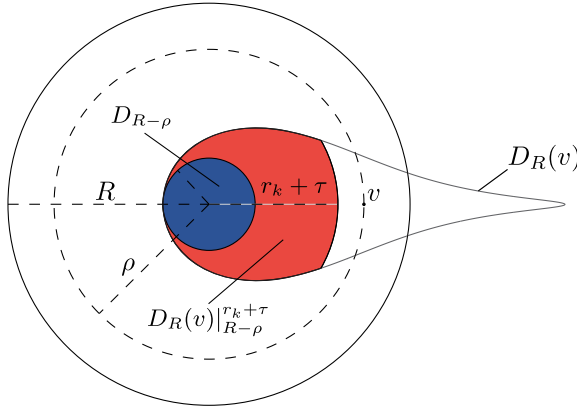


Fig. 7. Determining the sum of degrees of the neighbors of vertex v that are all contained in $D_R(v)$. To compute the measure of $D_R(v)|_0^{r_k+\tau}$ we divide it into two regions $D_R(v)|_0^{R-\rho} = D_{R-\rho}$ (blue) and $D_R(v)|_{R-\rho}^{r_k+\tau}$ (red).

the bounds for a sector S , we bound the sum of degrees in $D_R(v)$ by bounding the number of vertices with a fixed degree k for every possible value of k . If all these bounds hold with probability $1 - \mathcal{O}(n^{-3})$, then the union bound shows that the sum is concentrated with probability $1 - \mathcal{O}(n^{-2})$. To obtain a bound that holds for every possible angular coordinate of v (as claimed in Section 3.3.2), we apply Lemma 2.5. There, we choose the random variables X_w to represent the degrees of the vertices. Our bound on the sum that holds with probability $1 - \mathcal{O}(n^{-2})$ at a fixed angular coordinate, can then be translated to the same asymptotic bound that holds with probability $1 - \mathcal{O}(n^{-1})$ at every possible angular coordinate.

For a fixed degree $k = \omega(\log n)$, all vertices with degree at least k have radius at most $r_k + \tau$ with high probability due to Lemma 4.4, where $r_k = 2 \log(n/k)$ and τ is constant. Thus, all vertices of degree at least k in $D_R(v)$ lie in $D_R(v)|_0^{r_k+\tau}$, with high probability. In analogy to Lemma 4.5, we obtain the following bound on the number of vertices in $D_R(v)|_0^{r_k+\tau}$.

LEMMA 4.7. *Let G be a hyperbolic random graph and let v be a vertex with radius $\rho = 1/\alpha(\log n - \log \log n)$. If $k = \omega(\log n)$, the number of neighbors of v with degree at least k is*

$$|\{w \in N(v) \mid \deg(w) \geq k\}| = \mathcal{O}(n^{1-1/(2\alpha)} \log(n)^{1/(2\alpha)} k^{-(2\alpha-1)} + \log n),$$

with probability $1 - \mathcal{O}(n^{-c})$ for any constant $c > 0$.

PROOF. Since $k = \omega(\log n)$, we can apply Lemma 4.4 stating that all vertices of degree at least k in $D_R(v)$ lie within $D_R(v)|_0^{r_k+\tau}$ with high probability. To bound the number of neighbors of v with degree at least k we first compute the measure $\mu(D_R(v)|_0^{r_k+\tau})$. To do this, we separate $D_R(v)|_0^{r_k+\tau}$ into the disk $D_R(v)|_0^{R-\rho} = D_{R-\rho}$ and $D_R(v)|_{R-\rho}^{r_k+\tau}$; see Figure 7. Due to Equation (3), we have $\mu(D_{R-\rho}) = \mathcal{O}(e^{-\alpha(R-(R-\rho))}) = \mathcal{O}(\log n/n)$, which is already an upper bound on $\mu(D_R(v)|_0^{r_k+\tau})$ for the case where $r_k + \tau \leq R - \rho$. When $r_k + \tau > R - \rho$, we need to add the measure of $D_R(v)|_{R-\rho}^{r_k+\tau}$, which is given by

$$\mu(D_R(v)|_{R-\rho}^{r_k+\tau}) = \int_{R-\rho}^{r_k+\tau} 2 \int_0^{\theta(\rho, r)} f(r, \phi) d\phi dr = \mathcal{O}\left(\int_{R-\rho}^{r_k+\tau} \theta(\rho, r) f(r) dr\right).$$

Since we consider $r \in [R - \rho, r_k + \tau]$ in the integral, we have $r \geq R - \rho$, allowing us to apply Equation (2) to conclude that $\theta(\rho, r) = \mathcal{O}(e^{(R-\rho-r)/2})$. Furthermore, we can substitute the probability

density $f(r) = O(e^{-\alpha(R-r)})$ (Equation (1)) to obtain

$$\begin{aligned} \mu(D_R(v)|_{R-\rho}^{r_k+\tau}) &= O\left(\int_{R-\rho}^{r_k+\tau} e^{(R-\rho-r)/2} \cdot e^{-\alpha(R-r)} dr\right) \\ &= O\left(e^{(R-\rho)/2} \cdot e^{-\alpha R} \cdot \int_{R-\rho}^{r_k+\tau} e^{(\alpha-1/2)r} dr\right) \\ &= O\left(e^{-(\alpha-1/2)R} \cdot e^{-\rho/2} \cdot \left[e^{(\alpha-1/2)(r_k+\tau)} - e^{(\alpha-1/2)(R-\rho)}\right]\right). \end{aligned}$$

Dropping the negative term in the brackets and substituting $R = 2 \log n + C$, $\rho = 1/\alpha(\log n - \log \log n)$, and $r_k = 2 \log(n/k)$, we obtain

$$\begin{aligned} \mu(D_R(v)|_{R-\rho}^{r_k+\tau}) &= O\left(e^{-(\alpha-1/2)R} \cdot e^{-\rho/2} \cdot e^{(\alpha-1/2)(r_k+\tau)}\right) \\ &= O\left(n^{-(2\alpha-1)} \cdot n^{-1/(2\alpha)} \log(n)^{1/(2\alpha)} \cdot n^{2\alpha-1} \cdot k^{-(2\alpha-1)}\right) \\ &= O\left((\log(n)/n)^{1/(2\alpha)} \cdot k^{-(2\alpha-1)}\right). \end{aligned}$$

The expected number of vertices in $D_R(v)|_0^{r_k+\tau}$ is now obtained by reversing the previous split and adding the measures of $D_{R-\rho}$ and $D_R(v)|_{R-\rho}^{r_k+\tau}$, which yields

$$\mathbb{E}[|\{v \in D_R(v)|_0^{r_k+\tau}\}|] = n \cdot \left(\mu(D_{R-\rho}) + \mu(D_R(v)|_{R-\rho}^{r_k+\tau})\right) = O(\log n + n^{-1/(2\alpha)} \log(n)^{1/(2\alpha)} k^{-(2\alpha-1)}),$$

and it remains to show that this bound holds with large enough probability. Clearly, this bound is at least logarithmic. Thus, we can apply Corollary 2.3 to conclude that it holds with probability $1 - O(n^{-c})$ for any constant c . \square

With this, we are now ready to bound the sum of the degrees of the vertices in the central part of the disk that are in the neighborhood of a vertex with radius ρ . The proof of the following lemma is analogous to the one of Lemma 4.6.

LEMMA 4.8. *Let G be a hyperbolic random graph and let v be a hypothetical vertex with radius $\rho = 1/\alpha(\log n - \log \log n)$ and arbitrary angular coordinate. The degrees of neighbors of v in $D_R|_{\rho_I(2\pi)}^{\rho_O}$ sum to $\tilde{O}(n^{1/(2\alpha)})$ with high probability.*

PROOF. Recall that $D_R(v)$ is the disk containing all neighbors of v . To bound the sum of the degrees of the vertices in $D_R(v)|_{\rho_I(2\pi)}^{\rho_O}$, we use basically the same proof as in Lemma 4.6 except we use Lemma 4.7 instead of Lemma 4.5. Thus,

$$\sum_{w \in D_R(v)|_{\rho_I(2\pi)}^{\rho_O}} \deg(w) \leq k_{\min} |V_{k_{\min}}^{D_R(v)}| + \int_{k_{\min}}^{k_{\max}} |V_x^{D_R(v)}| dx,$$

where $V_x^{D_R(v)}$ is the set of vertices of degree at least x in $D_R(v)$ and k_{\min} and k_{\max} are the maximum and minimum degree in $D_R(v)|_{\rho_I(2\pi)}^{\rho_O}$, respectively.

We start with computing k_{\min} and k_{\max} . Using Equation (4) and Corollaries 2.3 and 2.4, we obtain that a vertex of radius $\rho_O = R - (2 + \varepsilon) \log \log n$, for any $\varepsilon \in (0, 1)$, has degree $k_{\min} = \Theta((\log n)^{1+\varepsilon/2})$ with high probability. Moreover, by the same argumentation as in the proof of Lemma 4.6 no vertex with smaller radius has smaller degree, with high probability. Additionally, a vertex with radius $\rho_I(2\pi) = R - 1/\alpha(\log n - \log \log n)$ has degree $k_{\max} = O((n/\log n)^{1/(2\alpha)})$ and no vertex with larger

radius has larger degree, with high probability. It follows that we can use the bound shown in Lemma 4.7 for $|V_x^{D_R(v)}|$. Thus, we obtain

$$\sum_{w \in D_R(v)|_{\rho_I(2\pi)}^{\rho_O}} \deg(w) = \tilde{O}(k_{\min} \cdot n^{1-1/(2\alpha)} k_{\min}^{-(2\alpha-1)}) + \tilde{O}\left(n^{1-1/(2\alpha)} \int_{k_{\min}}^{k_{\max}} x^{-(2\alpha-1)} dx\right).$$

Replacing k_{\min} and simplifying the first term in the sum yields $\tilde{O}(n^{1-1/(2\alpha)})$, which is smaller than the claimed bound. For the second term, we obtain

$$\tilde{O}\left(n^{1-1/(2\alpha)} \int_{k_{\min}}^{k_{\max}} x^{-(2\alpha-1)} dx\right) = \tilde{O}\left(n^{1-1/(2\alpha)} [k_{\max}^{2-2\alpha} - k_{\min}^{2-2\alpha}]\right).$$

Dropping the negative term and replacing $k_{\max} = \tilde{O}(n^{1/(2\alpha)})$, we obtain $\tilde{O}(n^{1-1/(2\alpha)+1/\alpha-1}) = \tilde{O}(n^{1/(2\alpha)})$. \square

4.3 The Outer Part of the Disk

At this point we have bounded the sum of the degrees of the vertices with radius at most $\rho_O = R - (2+\varepsilon) \log \log n$ (for any constant $\varepsilon \in (0, 1)$) that lie in a sector of angular width $\varphi \in \Omega(\log(n)^{2\alpha+1}/n)$ or in the neighborhood of a vertex with radius ρ . It remains to bound the sums when considering vertices with radii larger than ρ_O .

To bound the sum of the vertex degrees in the outer part of a sector $S|_{\rho_O}^R$, we start by computing the expected value.

LEMMA 4.9. *Let G be a hyperbolic random graph. For a sector S of angular width φ , the degrees of vertices in $S|_{\rho_O}^R$ sum to $\Theta(\varphi n)$ in expectation.*

PROOF. Let $\deg(v)$ be the random variable describing the degree of a vertex v . Moreover, let X_v be the indicator variable that is 1 if $v \in S|_{\rho_O}^R$ and 0 otherwise. Then the expected sum of the degrees of vertices in $S|_{\rho_O}^R$ is given by

$$\mathbb{E}\left[\sum_{v \in V} X_v \cdot \deg(v)\right] = \sum_{v \in V} \mathbb{E}[X_v \cdot \deg(v)] = n \cdot \Pr[v \in S|_{\rho_O}^R] \cdot \mathbb{E}[\deg(v) \mid v \in S|_{\rho_O}^R].$$

Note that $\Pr[v \in S|_{\rho_O}^R]$ is simply the measure $\mu(S|_{\rho_O}^R)$. As the angular coordinate is uniformly distributed, the whole sector S has measure $\Theta(\varphi)$. Moreover, the region of the disk containing the points with constant distance to the boundary has constant measure. Thus, the measure of $S|_{\rho_O}^R$ is also in $\Theta(\varphi)$. For the sake of completeness, the measure of $S|_{\rho_O}^R$ can be formally computed as

$$\begin{aligned} \mu(S|_{\rho_O}^R) &= \mu(S \setminus S|_0^{\rho_O}) \\ &= \frac{\varphi}{2\pi} (1 - \mu(D_{\rho_O})) \\ &= \frac{\varphi}{2\pi} (1 - e^{-\alpha(R-\rho_O)} (1 + o(1))) \\ &= \frac{\varphi}{2\pi} (1 - O((\log n)^{-\alpha(2+\varepsilon)})) \\ &= \Theta(\varphi). \end{aligned}$$

It remains to determine $\mathbb{E}[\deg(v) \mid v \in S_{\rho_O}^R]$, which can be done as follows:

$$\begin{aligned}
 \mathbb{E}[\deg(v) \mid v \in S_{\rho_O}^R] &= \iint_{S_{\rho_O}^R} \mathbb{E}[\deg(v) \mid r(v) = r] \frac{f(r, \phi)}{\mu(S_{\rho_O}^R)} d\phi dr \\
 &= \frac{1}{\mu(S_{\rho_O}^R)} \cdot \int_{\rho_O}^R \int_0^\varphi \mathbb{E}[\deg(v) \mid r(v) = r] f(r, \phi) d\phi dr \\
 &= \Theta(1) \cdot \int_{\rho_O}^R \mathbb{E}[\deg(v) \mid r(v) = r] f(r) dr \\
 &= \Theta(1) \cdot n \cdot e^{-\alpha R} \int_{\rho_O}^R e^{(\alpha-1/2)r} dr \\
 &= \Theta(1) \cdot n \cdot e^{-\alpha R} \left[e^{(\alpha-1/2)R} - e^{(\alpha-1/2)\rho_O} \right] \\
 &= \Theta(1) \cdot n \cdot e^{-R/2} \left[1 - e^{-(\alpha-1/2)(R-\rho_O)} \right].
 \end{aligned}$$

Note that the part in brackets is bounded by a constant. Moreover, as $R = 2 \log n + C$, $n \cdot e^{-R/2}$ is constant as well. Thus, $\mathbb{E}[\deg(v) \mid v \in S_{\rho_O}^R]$ is in $\Theta(1)$. It follows that the expected sum of the degrees is $\Theta(\varphi n)$. \square

Unfortunately, the sum of the vertex degrees in $S_{\rho_O}^R$ is not concentrated sufficiently well around its expectation to conclude that this bound also holds with high probability. The problem lies with the high-degree vertices in the graph, which can be adjacent to none or all vertices in $S_{\rho_O}^R$ depending on their positions. That is, small perturbations of the position of a single high-degree vertex can change the sum by too much. To overcome this issue, we consider the impact of high-degree vertices separately. To this end, we partition the edge set that contributes to the degrees of the vertices in $S_{\rho_O}^R$ into two sets E_I and E_O , denoting the *inner edges* where the other endpoint is in $D_R|_0^{\rho_O}$ and the *outer edges* where the other endpoint is in $D_R|_{\rho_O}^R$. The sum of the degrees of the vertices in $S_{\rho_O}^R$ can then be bounded by taking the number of inner edges and adding them to twice the number of outer edges. That is,

$$\sum_{v \in S_{\rho_O}^R} \deg(v) \leq |E_I| + 2|E_O|.$$

Since E_I denotes all edges with one endpoint in $S_{\rho_O}^R$ and the other in the inner or central part of the disk, we can obtain an upper bound on the first summand by summing the degrees of the vertices in $D_R|_0^{\rho_O}$ that are adjacent to any vertex in $S_{\rho_O}^R$. Since $\rho \leq \rho_O$, we have $S_{\rho_O}^R \subseteq S_{\rho}^R$, allowing us to apply Lemma 3.7 to conclude that all such vertices are contained in S or are neighbors of the two hypothetical corner vertices c_1 and c_2 , which both have radius ρ . Thus, $|E_I|$ can be bounded by the sum of the degrees of vertices in a sector and in the neighborhood of a vertex with radius ρ , but constrained to vertices in the inner and central parts of the disk. Corresponding bounds that hold with high probability have been determined above. For the sector we obtain an upper bound of $\tilde{O}(\delta_{\max})$ for the inner part (Corollary 4.2) and $\tilde{O}(\varphi n)$ for the central part (Lemma 4.6). For the neighborhood of a vertex with radius ρ we have $\tilde{O}(\delta_{\max})$ for the inner part (Corollary 4.3) and $\tilde{O}(n^{1/(2\alpha)})$ for the central part (Lemma 4.8). Taking them together, we obtain the following corollary.

COROLLARY 4.10. *Let G be a hyperbolic random graph. For every sector S of angular width $\varphi \in \Omega(\log(n)^{2\alpha+1}/n)$, the number of edges with one endpoint in $S_{\rho_O}^R$ and the other in $D_R|_0^{\rho_O}$ is in $\tilde{O}(\varphi n + n^{1/(2\alpha)} + \delta_{\max})$, with high probability.*

To obtain an upper bound on the second part of the above sum, we aim at applying a *method of typical bounded differences* based on the fact that changing the position of a single vertex has *typically* only little impact on the number of outer edges. The idea is as follows. We consider $|E_O|$ as a function that only depends on the positions P_1, \dots, P_n of the vertices in the graph and we ask ourselves: How much can $|E_O|$ change, if we alter the position of a single vertex i ? Clearly, this change can be large in the worst case. Assume that we move i from outside $D_R|_{\rho_O}^R$ into $S|_{\rho_O}^R$. Then, i does not contribute anything to $|E_O|$ before the move and the increase in $|E_O|$ depends on the number of outer edges that are incident to i after the move, which can be $n - 1$ in the worst case. However, it is very unlikely that a vertex in $S|_{\rho_O}^R$ has this many neighbors that lie in the outer part of the disk. In fact, its degree is typically much smaller. To formalize this, we represent the typical case using an event A , denoting that the degree of such a vertex is at most a constant factor larger than the expected degree of a vertex with radius $\rho_O = R - (2 + \varepsilon) \log \log n$ for any constant $\varepsilon \in (0, 1)$. More precisely, A denotes the event in which all disks of radius R with center in $D_R|_{\rho_O}^R$ contain at most $O(\log(n)^{1+\varepsilon/2})$ vertices. In this case, moving a vertex i in the same way as before leads to a much smaller increase in the number of outer edges. Assuming that A holds before the move, there are at most $O(\log(n)^{1+\varepsilon/2})$ outer edges incident to i after the move, which corresponds to the increase of $|E_O|$. The following lemma defines the event A formally and shows that it holds with high probability.

LEMMA 4.11. *Let G be a hyperbolic random graph and let $\rho_O = R - (2 + \varepsilon) \log \log(n)$ for any constant $\varepsilon \in (0, 1)$. Then, all disks D with radius R and center in $D_R|_{\rho_O}^R$ contain at most $|\{v \in D\}| = O(\log(n)^{1+\varepsilon/2})$ vertices, with probability $1 - O(n^{-c})$ for any constant c .*

PROOF. Let D be a disk of radius R and center $P \in D_R|_{\rho_O}^R$. By Lemma 2.1, a valid upper bound on the expected number of vertices in D can be obtained by considering the disk D' at center P' instead, which has the same angular coordinate as P and radius ρ_O . Thus, $\mathbb{E}[|\{v \in D\}|] \leq \mathbb{E}[|\{v \in D'\}|] = O(ne^{-\rho_O/2}) = O(\log(n)^{1+\varepsilon/2})$ (see Lemma (4)). Moreover, since $\varepsilon > 0$, this bound is $\omega(\log n)$ and we can apply Corollary 2.3 to conclude that $|\{v \in D\}| \in O(\log(n)^{1+\varepsilon/2})$ holds with probability $1 - O(n^{-c'})$ for any constant c' . To obtain a bound that holds for every possible angular coordinate for P , we apply Lemma 2, which allows us to translate our bound that holds for any given disk D with probability $1 - O(n^{-c'})$ to the same asymptotic bound that holds with probability $1 - O(n^{-c'+1})$ for all possible angular coordinates. Choosing $c' = c + 1$ then yields the claim. \square

So while moving a single vertex leads to a large change in the number of outer edges $|E_O|$ in the worst case, we observe only small changes in the typical case A . Formally, we say that a function $f: \Omega^n \rightarrow \mathbb{R}$ satisfies the *typical bounded differences condition* with respect to an event $A \subseteq \Omega^n$ if for all $i \in \{1, \dots, n\}$ there exist $\Delta_i^A \leq \Delta_i$ such that

$$|f(\mathbf{x}) - f(\mathbf{x}')| \leq \begin{cases} \Delta_i^A, & \text{if } \mathbf{x} \in A, \\ \Delta_i, & \text{otherwise,} \end{cases}$$

for all $\mathbf{x}, \mathbf{x}' \in \Omega^n$ that differ only in the i th component.

THEOREM 4.12 (METHOD OF TYPICAL BOUNDED DIFFERENCES, [27, THEOREM 2³]). *Let $X_1, \dots, X_n \in \Omega$ be independent random variables and let $A \subseteq \Omega^n$ be an event. Furthermore, let $f: \Omega^n \rightarrow \mathbb{R}$ be a function that satisfies the typical bounded differences condition with respect to A and with parameters $\Delta_i^A \leq \Delta_i$ for $i \in \{1, \dots, n\}$. Then, for all $\varepsilon_1, \dots, \varepsilon_n \in (0, 1]$ there exists an*

³We state a slightly simplified version in order to facilitate understandability. The original theorem allows for the random variables X_1, \dots, X_n to be defined in different sample spaces.

event B satisfying $\bar{B} \subseteq A$ and $\Pr[B] \leq \Pr[\bar{A}] \cdot \sum_i 1/\varepsilon_i$, such that for $\Delta = \sum_i (\Delta_i^A + \varepsilon_i(\Delta_i - \Delta_i^A))^2$ and $t \geq 0$ it holds that

$$\Pr[f > \mathbb{E}[f] + t \wedge \bar{B}] \leq e^{-t^2/(2\Delta)}.$$

Intuitively, the choice of the values for ε_i has two effects. On one hand, choosing ε_i small allows us to compensate for a potentially large worst-case change Δ_i . On the other hand, this also increases the bound on the probability of the event B that represents the atypical case. However, in that case one can still obtain meaningful bounds if the typical event A occurs with high enough probability. In the following, we show that an upper bound on the expected value $\mathbb{E}[f]$ is sufficient to apply the method of typical bounded differences, before applying it to bound the number of outer edges in a sector.

COROLLARY 4.13. *Let $X_1, \dots, X_n \in \Omega$ be independent random variables and let $A \subseteq \Omega^n$ be an event. Furthermore, let $f: \Omega^n \rightarrow \mathbb{R}$ be a function that satisfies the typical bounded differences condition with respect to A and with parameters $\Delta_i^A \leq \Delta_i$ for $i \in \{1, \dots, n\}$ and let $g(n)$ be an upper bound on $\mathbb{E}[f]$. Then for all $\varepsilon_1, \dots, \varepsilon_n \in (0, 1]$, $\Delta = \sum_i (\Delta_i^A + \varepsilon_i(\Delta_i - \Delta_i^A))^2$, and $c \geq 1$ it holds that*

$$\Pr[f > cg(n)] \leq e^{-((c-1)g(n))^2/(2\Delta)} + \Pr[\bar{A}] \sum_i 1/\varepsilon_i.$$

PROOF. Let $h(n) \geq 0$ be a function with $f' = f + h(n)$ such that $\mathbb{E}[f'] = g(n)$. Note that $h(n)$ exists since $g(n) \geq \mathbb{E}[f]$. As a consequence, we have $f \leq f'$ for all outcomes of X_1, \dots, X_n and it holds that

$$|f'(\mathbf{x}) - f'(\mathbf{x}')| = |f(\mathbf{x}) + h(n) - f(\mathbf{x}') - h(n)| = |f(\mathbf{x}) - f(\mathbf{x}')|,$$

for all $\mathbf{x}, \mathbf{x}' \in \Omega^n$. Consequently, f' satisfies the typical bounded differences condition with respect to A with the same parameters $\Delta_i^A \leq \Delta_i$ as f . Since $f \leq f'$ it holds that

$$\Pr[f > cg(n)] \leq \Pr[f' > cg(n)] = \Pr[f' > c\mathbb{E}[f']].$$

By choosing $t = (c-1)\mathbb{E}[f']$ this can be written as

$$\Pr[f' > c\mathbb{E}[f']] = \Pr[f' > \mathbb{E}[f'] + t].$$

Theorem 4.12 now guarantees the existence of an event B with $\Pr[B] \leq \Pr[\bar{A}] \cdot \sum_i 1/\varepsilon_i$ and $\bar{B} \subseteq A$, such that $\Pr[f' > \mathbb{E}[f'] + t \wedge \bar{B}] \leq e^{-t^2/(2\Delta)}$. To bound $\Pr[f' > \mathbb{E}[f'] + t]$ we apply the law of total probability and consider the events B and \bar{B} separately

$$\Pr[f' > \mathbb{E}[f'] + t] = \Pr[f' > \mathbb{E}[f'] + t \mid \bar{B}] \cdot \Pr[\bar{B}] + \Pr[f' > \mathbb{E}[f'] + t \mid B] \cdot \Pr[B].$$

The first part of the sum can be simplified using the definition of conditional probabilities. Moreover, it holds that $\Pr[f' > \mathbb{E}[f'] + t \mid B] \leq 1$. Thus, we can bound the above term by

$$\Pr[f' > \mathbb{E}[f'] + t] \leq \Pr[f' > \mathbb{E}[f'] + t \wedge \bar{B}] + \Pr[B].$$

Both remaining summands can now be bounded using the upper bounds that we previously obtained by applying Theorem 4.12, i.e., $\Pr[f' > \mathbb{E}[f'] + t \wedge \bar{B}] \leq e^{-t^2/(2\Delta)}$ and $\Pr[B] \leq \Pr[\bar{A}] \cdot \sum_i 1/\varepsilon_i$. Thus,

$$\Pr[f' > \mathbb{E}[f'] + t] \leq e^{-t^2/(2\Delta)} + \Pr[\bar{A}] \cdot \sum_i 1/\varepsilon_i.$$

Finally, since t was chosen as $t = (c-1)\mathbb{E}[f']$ and since $\mathbb{E}[f'] = g(n)$, we obtain the claimed bound. \square

We are now ready to bound the number $|E_O|$ of outer edges, i.e., edges that are incident to vertices in a sector $S_{\rho_O}^R$ and have their other endpoint in $D_R|_{\rho_O}^R$.

LEMMA 4.14. *Let G be a hyperbolic random graph. For every sector S of angular width $\varphi \in \Omega(\log(n)^2/n^{1/2})$, the number of edges with one endpoint in $S|_{\rho_O}^R$ and the other in $D_R|_{\rho_O}^R$ is in $\mathcal{O}(\varphi n)$, with high probability.*

PROOF. First note that, analogous to the proof of Lemma 4.6, we can cover the disk with $\mathcal{O}(n)$ sectors of angular width 2φ such that any sector of angular width φ lies completely in one of them. In the following, we show that the claimed bound holds with probability $\mathcal{O}(n^{-2})$ for a single sector S . Applying union bound then yields the claim.

We consider $|E_O|$, the number of edges with one endpoint in $S|_{\rho_O}^R$ and the other in $D_R|_{\rho_O}^R$, as a function that only depends on the positions P_1, \dots, P_n of the vertices in the graph. To show that $|E_O|$ does not exceed an upper bound with high probability, we aim at applying the method of typical bounded differences (Corollary 4.13). We represent the typical case with an event A , denoting that all disks D of radius R and center in $D_R|_{\rho_O}^R$ contain at most $\mathcal{O}(\log(n)^{1+\varepsilon/2})$ vertices for any constant $\varepsilon \in (0, 1)$. In order to determine the parameters $\Delta_i^A \leq \Delta_i$ for $i \in \{1, \dots, n\}$ with which $|E_O|$ fulfills the typical bounded differences condition with respect to A , we have to bound the maximum change in $|E_O|$ obtained by moving a single vertex. As argued before, this change is at most $\Delta_i = n - 1$ for all $i \in \{1, \dots, n\}$ in the worst case. To bound the Δ_i^A , we start with a configuration of vertex coordinates in which the event A holds. In this case, it is easy to see that moving a single vertex i changes $|E_O|$ by at most $\Delta_i^A = \mathcal{O}(\log(n)^{1+\varepsilon/2})$ for all $i \in \{1, \dots, n\}$, since the degree of i is at most this large after the move and so is the number of outer edges it contributes to $|E_O|$.

We are now ready to apply the method of typical bounded differences (Corollary 4.13). For an upper bound $g(n)$ on $|E_O|$, any constant $c > 1$, and all $\varepsilon_1, \dots, \varepsilon_n \in (0, 1]$ it states that

$$\Pr[|E_O| > cg(n)] \leq e^{-((c-1)g(n))^2/(2\Delta)} + \Pr[\bar{A}] \sum_i 1/\varepsilon_i,$$

where $\Delta = \sum_i (\Delta_i^A + \varepsilon_i(\Delta_i - \Delta_i^A))^2$. First note that a valid upper bound on the expected number of outer edges incident to vertices in $S|_{\rho_O}^R$ is given by the expected sum of the degrees of these vertices. Thus, by Lemma 4.9 we can choose $g(n) = \Theta(\varphi n)$. Moreover, by choosing $\varepsilon_i = 1/n$ for all $i \in \{1, \dots, n\}$ and since $\Delta_i = n - 1$ and $\Delta_i^A = \mathcal{O}(\log(n)^{1+\varepsilon/2})$ for all $i \in \{1, \dots, n\}$, we can compute Δ as

$$\begin{aligned} \Delta &= \sum_i (\Delta_i^A + \varepsilon_i(\Delta_i - \Delta_i^A))^2 \\ &= \mathcal{O}\left(n \cdot (\log(n)^{1+\varepsilon/2} + 1/n(n - \log(n)^{1+\varepsilon/2}))^2\right) \\ &= \mathcal{O}\left(n \cdot (\log(n)^{1+\varepsilon/2} + (1 - o(1)))^2\right) \\ &= \mathcal{O}\left(n \cdot \log(n)^{2+\varepsilon}\right). \end{aligned}$$

Consequently, the above probability can be bounded by

$$\begin{aligned} \Pr[|E_O| > cg(n)] &\leq \exp\left(-\Theta\left(\frac{(\varphi n)^2}{n \log(n)^{2+\varepsilon}}\right)\right) + \Pr[\bar{A}] \cdot n^2 \\ &\leq \exp\left(-\Theta\left(\frac{\varphi^2 n}{\log(n)^{2+\varepsilon}}\right)\right) + \Pr[\bar{A}] \cdot n^2. \end{aligned}$$

Since $\varphi \in \Omega(\log(n)^2/n^{1/2})$ is a precondition of this lemma and since $\varepsilon < 1$, we can conclude that the fraction is $\omega(\log n)$, which means that the first summand is $\mathcal{O}(n^{-c'})$ for any constant c' . Moreover,

by Lemma 4.11 event A holds with probability $1 - O(n^{-c'})$ for any constant c' . Choosing $c' = 3$ then yields the claim. \square

4.4 The Complete Disk

Having obtained the required bounds for the inner, central, and outer parts of the disk, we can now combine them to bound the sum of the degrees in a sector and in the neighborhoods of the hypothetical corner vertices. We start with Theorem 3.5, which bounds the sum of degrees in a sector. To improve readability, we restate the theorem here.

THEOREM 3.15. *Let G be a hyperbolic random graph. The degrees of vertices in every sector of angular width φ sum to $\tilde{O}(\varphi n + n^{1/(2\alpha)} + \delta_{\max})$ with high probability if $\varphi = \Omega(\log(n)^2/n^{1/2})$.*

PROOF. For the inner and central parts of every sector the sum of the vertex degrees is bounded by $\tilde{O}(\varphi n + \delta_{\max})$ with high probability due to Corollary 4.2 and Lemma 4.6. As argued above, the sum of the degrees of the remaining vertices, i.e., vertices with radius at least ρ_O , can be bounded by counting the number of inner edges and adding twice the number of outer edges. Since $\varphi \in \Omega(\log(n)^2/n^{1/2})$, we can apply Corollary 4.10 and Corollary 4.14 to conclude that the corresponding sum is bounded by $\tilde{O}(\varphi n + n^{1/(2\alpha)} + \delta_{\max})$, with high probability. \square

Lastly, it remains to bound the sum of the degrees of the neighbors of the hypothetical corner vertices that were used to bound the size of the search space in the second phase. Again, for the sake of readability, we restate the corresponding lemma here.

LEMMA 3.16. *Let G be a hyperbolic random graph and let v be a hypothetical vertex with radius $\rho = 1/\alpha(\log n - \log \log n)$ and arbitrary angular coordinate. The degrees of neighbors of v sum to $\tilde{O}(n^{2-1/\alpha} + n^{1/(2\alpha)} + \delta_{\max})$ with high probability.*

PROOF. For the inner and central parts of the neighborhood of a vertex with radius ρ and arbitrary angular coordinate the sum of the degrees is bounded by $\tilde{O}(\delta_{\max} + n^{1/(2\alpha)})$ with high probability, due to Corollary 4.3 and Lemma 4.8. For the sum of the degrees in the outer part of the disk, note that all neighbors of radius at least ρ have angular distance at most $\varphi = O(n^{-(1/\alpha-1)})$; see Section 3.3.1. Thus, we can use Theorem 3.5 to conclude that claimed bound holds for the sum of their degrees. Note that if φ is too small to meet the requirements of Theorem 3.5, we can choose $\varphi = \tilde{O}(n^{-1/2})$ as a valid upper bound to conclude that the sum of degrees in the outer part of the neighborhood is in $\tilde{O}(n^{1/2})$, which is $\tilde{O}(n^{1/(2\alpha)})$ for $\alpha \in (1/2, 1)$. \square

5 CONCLUSION

In the following, we briefly discuss why we think that the bound $\tilde{O}(n^{2-1/\alpha} + n^{1/(2\alpha)} + \delta_{\max})$ is rather tight; see Figure 8 (left) for a plot of the exponents. Clearly, the maximum degree of the graph is a lower bound, i.e., we cannot improve the δ_{\max} . As $\delta_{\max} = \Theta(n^{1/(2\alpha)})$ holds almost surely [17], we also cannot improve below $\tilde{O}(n^{1/(2\alpha)})$. For the term $n^{2-1/\alpha}$ we do not have a lower bound. Thus, the blue region in Figure 8 (left) is the only part where our bound can potentially be improved. However, by only making a single step from a vertex with radius $\rho = 1/\alpha(\log n - \log \log n)$, we can already reach vertices with angular distance $\Theta(n^{-(1/\alpha-1)})$. Thus, it seems likely, that there exists a start–destination pair such that all vertices within a sector of this angular width are actually explored. As such a sector contains $\Theta(n^{2-1/\alpha})$ vertices, our bound seems rather tight (at least asymptotically and up to poly-logarithmic factors). For a comparison of our theoretical bound with actual search-space sizes in hyperbolic random graphs; see Figure 8.




Fig. 8. Left: The exponent of our theoretical bound depending on α . Right: The corresponding empirically measured search spaces. The data were obtained by generating 20 hyperbolic random graphs with average degree roughly 8 for each shown α and each $n \in \{100k, 200k, 300k\}$. For each graph we sampled 300k start-destination pairs and report the maximum number of edges explored in one search. The numbers are normalized with the total number of edges m of the graph such that x is plotted for a search space of size m^x .

Finally, in order to put our results into perspective, we discuss the following question: How does a heterogeneous degree distribution impact the exponent in the running time of the bidirectional BFS? First, considering networks with no underlying geometry, the exponent is $1/2$ for homogeneous networks and $(4 - \beta)/2 = 3/2 - \alpha$ for heterogeneous networks with power-law exponent β [10]. That is, when increasing the heterogeneity by letting α go from 1 to $1/2$, the exponent increases from $1/2$ to 1. This can be explained by the fact that a heterogeneous degree distribution leads to high-degree vertices, which leads to a higher running time when they are explored.

On hyperbolic random graphs, we get the same effect. The $1/(2\alpha)$ -part of the exponent (the red function in Figure 8) is very similar to the above $3/2 - \alpha$. However, due to the underlying geometry, the heterogeneity has another effect, expressed by the $2 - 1/\alpha$ -part of the exponent (the green function in Figure 8). This can be explained as follows. The underlying geometry constrains the parts of the graph that a vertex can connect to. As a result, the search space cannot expand sufficiently fast on homogeneous networks and we only get a constant speedup, i.e., the exponent is 1. However, increasing the heterogeneity leads to high degree vertices, which accelerate the expansion of the search spaces, leading to a lower exponent.

In conclusion, we can say that heterogeneity has two effects on the bidirectional BFS:

- (1) More heterogeneity leads to higher running times as exploring high degree-vertices is costly.
- (2) More heterogeneity leads to lower running times as high degree-vertices let the search spaces expand quickly.

For networks without underlying geometry, the second effect is irrelevant, as the search space always expands quickly due to the independence of edges. Thus, the running time is better the more homogeneous the network. For networks with underlying geometry, both effects play an important role leading to the v-shape in Figure 8. For high heterogeneity ($\alpha < 0.75$), the cost of exploring high degree vertices dominates, leading to the exponent $1/(2\alpha)$. For lower heterogeneity ($\alpha > 0.75$), the slower expanding search space due to the underlying geometry dominates, leading to the exponent $2 - 1/\alpha$.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their insightful comments and suggestions, which have improved the quality of the paper.

REFERENCES

- [1] Takuya Akiba, Christian Sommer, and Ken-ichi Kawarabayashi. 2012. Shortest-path queries for complex networks: Exploiting low tree-width outside the core. In *Proceedings of the 15th International Conference on Extending Database Technology*. 144–155. DOI : <https://doi.org/10.1145/2247596.2247614>
- [2] Gregorio Alanis-Lobato, Pablo Mier, and Miguel A. Andrade-Navarro. 2016. Manifold learning and maximum likelihood estimation for hyperbolic network embedding. *Applied Network Science* 1, 10 (2016), 1–14. DOI : <https://doi.org/10.1007/s41109-016-0013-0>
- [3] Thomas Bläsius, Philipp Fischbeck, Tobias Friedrich, and Maximilian Katzmann. 2020. Solving vertex cover in polynomial time on hyperbolic random graphs. In *Proceedings of the 37th International Symposium on Theoretical Aspects of Computer Science*. 25:1–25:14. DOI : <https://doi.org/10.4230/LIPIcs.STACS.2020.25>
- [4] Thomas Bläsius, Cedric Freiberger, Tobias Friedrich, Maximilian Katzmann, Felix Montenegro-Retana, and Marianne Thieffry. 2018. Efficient shortest paths in scale-free networks with underlying hyperbolic geometry. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming*. 20:1–20:14. DOI : <https://doi.org/10.4230/LIPIcs.ICALP.2018.20>
- [5] Thomas Bläsius, Tobias Friedrich, and Anton Krohmer. 2016. Hyperbolic random graphs: Separators and treewidth. In *Proceedings of the 24th Annual European Symposium on Algorithms*. 15:1–15:16. DOI : <https://doi.org/10.4230/LIPIcs.ESA.2016.15>
- [6] Thomas Bläsius, Tobias Friedrich, and Anton Krohmer. 2018. Cliques in hyperbolic random graphs. *Algorithmica* 80, 8 (2018), 2324–2344. DOI : <https://doi.org/10.1007/s00453-017-0323-3>
- [7] Michel Bode, Nikolaos Fountoulakis, and Tobias Müller. 2013. On the giant component of random hyperbolic graphs. In *Proceedings of the 7th European Conference on Combinatorics, Graph Theory and Applications*. 425–429. DOI : https://doi.org/10.1007/978-88-7642-475-5_68
- [8] Michel Bode, N. Fountoulakis, and Tobias Müller. 2015. On the largest component of a hyperbolic model of complex networks. *Electronic Journal of Combinatorics* 22, 3 (2015), 1–46. DOI : <https://doi.org/10.37236/4958>
- [9] Marián Boguñá, Fragkiskos Papadopoulos, and Dmitri Krioukov. 2010. Sustaining the internet with hyperbolic mapping. *Nature Communications* 1 (2010), 1–8. DOI : <https://doi.org/10.1038/ncomms1063>
- [10] Michele Borassi and Emanuele Natale. 2016. KADABRA is an ADaptive algorithm for betweenness via random approximation. In *Proceedings of the 24th Annual European Symposium on Algorithms*. 20:1–20:18. DOI : <https://doi.org/10.4230/LIPIcs.ESA.2016.20>
- [11] Ralph Keusch. 2018. Geometric inhomogeneous random graphs and graph coloring games. DOI : [10.3929/ethz-b-000269658](https://doi.org/10.3929/ethz-b-000269658)
- [12] Karl Bringmann, Ralph Keusch, and Johannes Lengler. 2017. Sampling geometric inhomogeneous random graphs in linear time. In *Proceedings of the 25th Annual European Symposium on Algorithms*. 20:1–20:15. DOI : <https://doi.org/10.4230/LIPIcs.ESA.2017.20>
- [13] Karl Bringmann, Ralph Keusch, Johannes Lengler, Yannic Maus, and Anisur Rahaman Molla. 2017. Greedy routing and the algorithmic small-world phenomenon. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*. 371–380. DOI : <https://doi.org/10.1145/3087801.3087829>
- [14] Devdatt P. Dubhashi and Alessandro Panconesi. 2012. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press.
- [15] Tobias Friedrich and Anton Krohmer. 2018. On the diameter of hyperbolic random graphs. *SIAM Journal on Discrete Mathematics* 32, 2 (2018), 1314–1334. DOI : <https://doi.org/10.1137/17M1123961>
- [16] Tobias Friedrich, Thomas Sauerwald, and Alexandre Stauffer. 2013. Diameter and broadcast time of random geometric graphs in arbitrary dimensions. *Algorithmica* (2013), 65–88. DOI : <https://doi.org/10.1007/s00453-012-9710-y>
- [17] Luca Gugelmann, Konstantinos Panagiotou, and Ueli Peter. 2012. Random hyperbolic graphs: Degree sequence and clustering. In *Proceedings of the 39th International Colloquium on Automata, Languages, and Programming*. 573–585. DOI : https://doi.org/10.1007/978-3-642-31585-5_51
- [18] Marcos Kiwi and Dieter Mitsche. 2019. On the second largest component of random hyperbolic graphs. *SIAM Journal on Discrete Mathematics* 33, 4 (2019), 2200–2217. DOI : <https://doi.org/10.1137/18M121201X>
- [19] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá. 2010. Hyperbolic geometry of complex networks. *Physical Review E* 82, 3 (2010), 036106. DOI : <https://doi.org/10.1103/PhysRevE.82.036106>
- [20] Zheng Lu, Yunhe Feng, and Qing Cao. 2017. Decentralized search for shortest path approximation in large-scale complex networks. In *Proceedings of the IEEE International Conference on Cloud Computing Technology and Science*. 130–137. DOI : <https://doi.org/10.1109/CloudCom.2017.36>

- [21] Michael Luby and Prabhakar Ragde. 1989. A bidirectional shortest-path algorithm with good average-case behavior. *Algorithmica* 4, 1 (1989), 551–567. DOI : <https://doi.org/10.1007/BF01553908>
- [22] Tobias Müller and Merlijn Staps. 2019. The diameter of KPKVB random graphs. *Advances in Applied Probability* 51, 2 (2019), 358–377. DOI : <https://doi.org/10.1017/apr.2019.23>
- [23] Wei Peng, Xiaofeng Hu, Feng Zhao, and Jinshu Su. 2012. A fast algorithm to find all-pairs shortest paths in complex networks. *Procedia Computer Science* 9 (2012), 557–566. DOI : <https://doi.org/10.1016/j.procs.2012.04.060>
- [24] Mathew Penrose. 2003. *Random Geometric Graphs*. Oxford University Press.
- [25] Ueli Peter. 2014. *Random Graph Models for Complex Systems*. Ph.D. Dissertation. ETH Zürich.
- [26] Ira Sheldon Pohl. 1969. *Bi-directional and Heuristic Search in Path Problems*. Ph.D. Dissertation. Stanford University.
- [27] Lutz Warnke. 2016. On the method of typical bounded differences. *Combinatorics, Probability and Computing* 25, 2 (2016), 269–299. DOI : <https://doi.org/10.1017/S0963548315000103>

Received May 2020; revised September 2021; accepted January 2022