

On the Hyperbolicity of Small-World and Treelike Random Graphs

Wei Chen , Wenjie Fang , Guangda Hu & Michael W. Mahoney

To cite this article: Wei Chen , Wenjie Fang , Guangda Hu & Michael W. Mahoney (2013) On the Hyperbolicity of Small-World and Treelike Random Graphs, Internet Mathematics, 9:4, 434-491, DOI: [10.1080/15427951.2013.828336](https://doi.org/10.1080/15427951.2013.828336)

To link to this article: <http://dx.doi.org/10.1080/15427951.2013.828336>



Accepted author version posted online: 30 Sep 2013.
Published online: 30 Sep 2013.



Submit your article to this journal [↗](#)



Article views: 76



View related articles [↗](#)



Citing articles: 6 View citing articles [↗](#)

On the Hyperbolicity of Small-World and Treelike Random Graphs

Wei Chen, Wenjie Fang, Guangda Hu, and Michael W. Mahoney

Abstract. Hyperbolicity is a property of a graph that may be viewed as a “soft” version of a tree, and recent empirical and theoretical work has suggested that many graphs arising in Internet and related data applications have hyperbolic properties. Here we consider Gromov’s notion of δ -hyperbolicity and establish several positive and negative results for small-world and treelike random graph models. First, we study the hyperbolicity of the class of Kleinberg small-world random graphs $KSW(n, d, \gamma)$, where n is the number of vertices in the graph, d is the dimension of the underlying base grid B , and γ is the small-world parameter such that each node u in the graph connects to another node v in the graph with probability proportional to $1/d_B(u, v)^\gamma$, with $d_B(u, v)$ the grid distance from u to v in the base grid B . We show that when $\gamma = d$, the parameter value allowing efficient decentralized routing in Kleinberg’s small-world network, the hyperbolic δ is $\Omega((\log n)^{1/(1.5(d+1)+\varepsilon)})$ with probability $1 - o(1)$ for every $\varepsilon > 0$ independent of n . We see that hyperbolicity is not significantly improved in relation to graph diameter even when the long-range connections greatly improve decentralized navigation. We also show that for other values of γ , the hyperbolic δ is very close to the graph diameter, indicating poor hyperbolicity in these graphs as well. Next we study a class of treelike graphs called ringed trees that have constant hyperbolicity. We show that adding random links among the leaves in a manner similar to the small-world graph constructions may easily destroy the hyperbolicity of the graphs, except for a class of random edges added using an exponentially decaying probability function based on the ring distance among the leaves. Our study provides one of the first significant analytic

results on the hyperbolicity of a rich class of random graphs, which sheds light on the relationship between hyperbolicity and navigability of random graphs, as well as on the sensitivity of hyperbolic δ to noises in random graphs.

1. Introduction

Hyperbolicity, a property of metric spaces that generalizes the idea of Riemannian manifolds with negative curvature, has received considerable attention in both mathematics and computer science. When applied to graphs, one may think of hyperbolicity as characterizing a “soft” version of a tree: trees have hyperbolicity zero, and graphs that “look like” trees in terms of their metric structure have “small” hyperbolicity. Since trees are an important class of graphs and since treelike graphs arise in numerous applications, the idea of hyperbolicity has received attention in a range of applications. For example, it has found usefulness in the visualization of the Internet, the Web, and other large graphs [Lamping and Rao 94, Lamping et al. 95, Munzner and Burchard 95, Munzner 98, Walter and Ritter 02]; it has been applied to questions of compact routing, navigation, and decentralized search in Internet graphs and small-world social networks [Gavoille and Ly 05, Chepoi et al. 12, Kleinberg 07, Abraham et al. 07, Krioukov et al. 07, Boguñá et al. 09, Papadopoulos et al. 10]; and it has been applied to a range of other problems such as distance estimation, network security, sensor networks, and traffic flow and congestion minimization [Baryshnikov 02, Jonckheere and Lohsoonthorn 02, Jonckheere and Lohsoonthorn 04, Jonckheere et al. 08, Narayan and Saniee 11, de Montgolfier et al. 11].

The hyperbolicity of graphs is typically measured by Gromov’s hyperbolic δ [Gromov 87, Bridson and Haefliger 99] (see Section 2). The hyperbolic δ of a graph measures the “treelikeness” of the graph in terms of the graph distance metric. It can range from 0 to half the graph diameter, with trees having $\delta = 0$, in contrast to “circle graphs” and “grid graphs” having large δ equal to roughly half their diameters.

In this paper, we study the δ -hyperbolicity of families of random graphs that intuitively have some sort of treelike or hierarchical structure. Our motivation comes from two angles. First, although there are quite a few empirical studies on the hyperbolicity of real-world and random graphs [Baryshnikov 02, Jonckheere and Lohsoonthorn 02, Lou 08, Lohsoonthorn 03, Narayan and Saniee 11, de Montgolfier et al. 11], there are essentially no systematic analytic studies of the hyperbolicity of popular random graphs. Thus, our work is intended to fill this gap. Second, a number of algorithmic studies show that good graph

hyperbolicity leads to efficient distance labeling and routing schemes [Chepoi and Dragan 00, Gavaille and Ly 05, Chepoi and Estellon 07, Chepoi et al. 08, Krioukov et al. 10, Chepoi et al. 12], and the routing infrastructure of the Internet has been empirically shown to be hyperbolic [Baryshnikov 02]. Thus, it is of interest to investigate whether efficient routing capability implies good graph hyperbolicity.

To achieve our goal, we first provide a fine-grained characterization of δ -hyperbolicity of graph families relative to the graph diameter: a family of random graphs is

- (a) *constantly hyperbolic* if their hyperbolic δ 's are constant, regardless of the size or diameter of the graphs;
- (b) *logarithmically (or polylogarithmically) hyperbolic* if their hyperbolic δ 's are of order the logarithm (or polylogarithm) of the graph diameters;
- (c) *weakly hyperbolic* if their hyperbolic δ 's grow asymptotically slower than the graph diameters; and
- (d) *not hyperbolic* if their hyperbolic δ 's are of the same order as the graph diameters.

We study two families of random graphs. The first family is Kleinberg's grid-based small-world random graphs [Kleinberg 00], which build random long-range edges among pairs of nodes with probability inversely proportional to the γ th power of the grid distance of the pairs. Kleinberg shows that when γ equals the grid dimension d , the number of hops for decentralized routing can be improved from $\Theta(n)$ in the grid to $O(\text{polylog}(n))$, where n is the number of vertices in the graph. Contrary to the improvement in decentralized routing, we show that when $\gamma = d$, then with high probability, the small-world graph is not polylogarithmically hyperbolic. We further show that when $0 \leq \gamma < d$, the random small-world graphs are not hyperbolic, and when $\gamma > 3$ and $d = 1$, the random graphs are not polylogarithmically hyperbolic. Although there still exists a gap between hyperbolic δ and graph diameter at the sweet spot of $\gamma = d$, our results already indicate that long-range edges that enable efficient navigation do not significantly improve the hyperbolicity of the graphs.

Our second family of graphs is random *ringed trees*. A ringed tree is a binary tree with nodes in each level of the tree connected by a ring (Figure 1(d)). Ringed trees can be viewed as an idealized version of hierarchical structure with local peer connections, such as the Internet autonomous system (AS) topology. We show that a ringed tree is quasi-isometric to the Poincaré disk, the well-known hyperbolic space representation, and thus it is constantly hyperbolic. We then

study how random additions of long-range links on the leaves of a ringed tree affect the hyperbolicity of random ringed trees. Note that due to the tree base structure, random ringed trees allow efficient routing within $O(\log n)$ steps using tree branches. Our results show that if the random long-range edges between leaves are added according to a probability function that decreases exponentially fast with the ring distance between leaves, then the resulting random graph is logarithmically hyperbolic, but if the probability function decreases only as a power law with ring distance, or based on another tree distance measure similar to that in [Kleinberg 02], the resulting random graph is not hyperbolic. Furthermore, if we use binary trees instead of ringed trees as base graphs, none of the above augmentations is hyperbolic. Taken together, our results indicate that δ -hyperbolicity of graphs is quite sensitive to both base graph structures and probabilities of long-range connections.

To summarize, we provide one of the first significant analytic results on the hyperbolicity properties of important families of random graphs. Our results demonstrate that efficient routing performance does not necessarily mean good graph hyperbolicity (such as logarithmic hyperbolicity).

1.1. Related Work

There has been a considerable amount of work on search and decentralized search subsequent to Kleinberg's original work [Kleinberg 00, Kleinberg 02], much of which has been summarized in the review [Kleinberg 06].

In parallel with this, there has been empirical and theoretical work on hyperbolicity of real-world complex networks as well as simple random graph models. On the empirical side, [Baryshnikov 02] showed that measurements of the Internet are negatively curved; [Jonckheere and Lohsoonthorn 02, Jonckheere and Lohsoonthorn 04, Jonckheere et al. 08, Lou 08, Lohsoonthorn 03] provided empirical evidence that randomized scale-free and Internet graphs are more hyperbolic than other types of random graph models; [Narayan and Saniee 11] measured the average δ and related curvature to congestion; and [de Montgolfier et al. 11] measured tree-width and hyperbolicity properties of the Internet. On the theoretical side, one has [Papazian and Rémila 00, Jonckheere et al. 08, Bermudo et al. 10, Narayan et al. 12, Tucci 12, Shang 12], among which [Narayan et al. 12, Tucci 12, Shang 12] study Gromov hyperbolicity of random graphs and are most relevant to our work. In [Narayan et al. 12], the authors study δ -hyperbolicity of sparse Erdős–Rényi random graphs $G(n, p)$, where n is the number of vertices in the graph and each pair of nodes is connected by an edge with probability p independent from every other edge, with $p = c/n$ for some constant $c > 1$. The authors prove that with positive probability, these

graphs are not δ -hyperbolic for any positive constant δ (i.e., not constantly hyperbolic in our definition). In [Tucci 12], the author shows that random d -regular graphs are almost surely not constantly hyperbolic. In [Shang 12], it is shown that with nonzero probability, the Newman–Watts small-world model [Newman and Watts 99] is not constantly hyperbolic. These studies investigate only constant hyperbolicity on random graphs, while our study moves beyond constant hyperbolicity and shows whether certain random graph classes are logarithmically hyperbolic, or not hyperbolic at all, compared with the graph diameters. Moreover, the one-dimensional Newman–Watts small-world model studied in [Shang 12] is a special case of the Kleinberg small-world model that we study in this paper (with dimension $d = 1$ and small-world parameter $\gamma = 0$). As given by Theorem 3.1(2), we show that with probability $1 - o(1)$, the hyperbolic δ of these random graphs is $\Omega(\log n)$, where n is the number of vertices in the graph. Therefore, our result is stronger than the result in [Shang 12] for this particular case.

More generally, we see two approaches connecting hyperbolicity with efficient routing in graphs. One approach studies efficient computation of graph properties, such as diameters, centers, approximating trees, and packings and coverings for low-hyperbolic- δ graphs and metric spaces [Chepoi et al. 08, Chepoi and Dragan 00, Gavaille and Ly 05, Chepoi et al. 12, Chepoi and Estellon 07]. In large part, the reason for this interest is that there are often direct consequences for navigation and routing in these graphs [Gavaille and Ly 05, Chepoi et al. 12, Kleinberg 07, Abraham et al. 07]. While these results are of interest for general low-hyperbolic- δ graphs, they can be less interesting when applied to small-world and other low-diameter random models of complex networks. To take one example, [Chepoi et al. 08] provides a simple construction of a distance-approximating tree for δ -hyperbolic graphs on n vertices; but the $O(\log n)$ additive-error guarantee is clearly less interesting for models in which the diameter of the graph is $O(\log n)$. Unfortunately, this $O(\log n)$ arises for a very natural reason in the analysis, and it is nontrivial to improve it for popular treelike complex network models.

Another approach taken by several recent papers is to build random graphs from hyperbolic metric spaces and then show that such random graphs lead to several common properties of small-world complex networks, including good navigability properties [Boguñá et al. 09, Papadopoulos et al. 10, Krioukov et al. 10, Krioukov et al. 09]. While assuming a low-hyperbolicity metric space to build random graphs in these studies makes intuitive sense, it is difficult to prove nontrivial results on Gromov's δ of these random graphs even for simple random graph models that are intuitively treelike.

Understanding the relationship between these two approaches was one of the original motivations of our research. In particular, the difficulties in the above

two approaches led us to study hyperbolicity of small-world and treelike random graphs.

Finally, ideas related to hyperbolicity have been applied in numerous other network applications, such as distance estimation, network security, sensor networks, and traffic flow and congestion minimization [Shavitt and Tankel 08, Jonckheere et al. 11, Jonckheere et al. 07, Jonckheere and Lohsoonthorn 04, Narayan and Saniee 11, Baryshnikov and Tucci 10], as well as large-scale data visualization [Munzner 98]. The latter applications typically take important advantage of the idea that data are often hierarchical or treelike and that there is “more room” in hyperbolic spaces of dimension 2 than in Euclidean spaces of any finite dimension.

1.2. Organization of This Paper

In Section 2, we provide basic concepts and terminologies on hyperbolic spaces and graphs that are needed in this paper. In Sections 3 and 4, we study the hyperbolicity of small-world random graphs and ringed-tree-based random graphs. For ease of reading, in each of Sections 3 and 4 we first summarize our technical results together with their implications (Sections 3.1 and 4.1), then provide the outline of the analyses (Sections 3.2 and 4.2), followed by the detailed technical proofs (Sections 3.3 and 4.3), and finally discuss extensions of our results to other related models (Sections 3.4 and 4.4). We discuss open problems and future directions related to our study in Section 5.

2. Preliminaries on Hyperbolic Spaces and Graphs

Here, we provide basic concepts concerning hyperbolic spaces and graphs used in this paper; for more comprehensive coverage on hyperbolic spaces, see, e.g., [Bridson and Haefliger 99].

2.1. Gromov’s δ -Hyperbolicity

In [Gromov 87], there is defined a notion of hyperbolic metric space; and then hyperbolic groups are defined to be finitely generated groups with a Cayley graph that is hyperbolic. There are several equivalent definitions (up to a multiplicative constant) of Gromov’s hyperbolic metric space [Bowditch 91]. In this paper, we will mainly use the following.

Definition 2.1. (Gromov's four-point condition.) In a metric space (X, d) , given u, v, w, x with

$$d(u, v) + d(w, x) \geq d(u, x) + d(w, v) \geq d(u, w) + d(v, x)$$

in X , we note that

$$\delta(u, v, w, x) = \frac{1}{2}(d(u, v) + d(w, x) - d(u, x) - d(w, v)).$$

The space (X, d) is called δ -hyperbolic for some nonnegative real number δ if for every set of four points $u, v, w, x \in X$, we have $\delta(u, v, w, x) \leq \delta$. Let $\delta(X, d)$ be the smallest possible value of such a δ , which can also be defined as

$$\delta(X, d) = \sup_{u, v, w, x \in X} \delta(u, v, w, x).$$

Given an undirected, unweighted, and connected graph $G = (V, E)$, one can view it as a metric space (V, d_G) , where $d_G(u, v)$ denotes the (geodesic) graph distance between two vertices u and v . Then, one can apply the above four-point condition to define its δ -hyperbolicity, which we denote by $\delta = \delta(G) = \delta(V, d_G)$ (and which we sometimes refer to simply as the hyperbolicity or the δ of the graph). Trees are 0-hyperbolic; and 0-hyperbolic graphs are exactly clique trees (also called block graphs), which can be viewed as cliques connected in a treelike fashion [Howorka 79]. Thus, it is often helpful to view graphs with a low hyperbolic δ as “thickened” trees, or in other words, as treelike when viewed at large size scales.

If we let $D(G)$ denote the diameter of the graph G . Then by the triangle inequality, we have $\delta(G) \leq D(G)/2$. We will use the asymptotic difference between the hyperbolicity $\delta(G)$ and the diameter $D(G)$ to characterize the hyperbolicity of the graph G .

Definition 2.2. (Hyperbolicity of a graph.) For a family of graphs \mathcal{G} with diameter $D(G)$, with $G \in \mathcal{G}$ going to infinity as the size of G grows to infinity, we say that \mathcal{G} is *constantly* (respectively *logarithmically*, *polylogarithmically*, or *weakly*) *hyperbolic* if $\delta(G) = O(1)$ (respectively $O(\log D(G))$, $O((\log D(G))^c)$ for some constant $c > 0$, or $o(D(G))$) as $D(G)$ goes to infinity; and \mathcal{G} is *not hyperbolic* if $\delta(G) = \Theta(D(G))$, where $G \in \mathcal{G}$.

The above definition provides a more fine-grained characterization of hyperbolicity of graph families than one typically sees in the literature, which generally discusses only whether a graph family is constantly hyperbolic. This definition does not address the hyperbolicity of graph families in which the diameter

remains bounded as the size of the graph becomes unbounded. For such graph families, one may likely need tight analysis on the constant factor between the hyperbolic δ and the graph diameter, and it is beyond the scope of this paper.

2.2. Rips Condition

The Rips condition [Gromov 87, Bridson and Haefliger 99] is a technically equivalent condition to Gromov's four-point condition up to a constant factor. We use the Rips condition in analyzing the δ -hyperbolicity of ringed trees. In a metric space (X, d) , we define a *geodesic segment* $[u, v]$ between two points u, v to be the image of a function $\rho : [0, d(u, v)] \rightarrow [u, v]$ satisfying

$$\rho(0) = u, \quad \rho(d(u, v)) = v, \quad d(\rho(s), \rho(t)) = |s - t|,$$

for every $s, t \in [0, d(u, v)]$. We say that a metric space is *geodesic* if every pair of its points has a geodesic segment, not necessarily unique. In a geodesic metric space (X, d) , given u, v, w in X , we denote by

$$\Delta(u, v, w) = [u, v] \cup [v, w] \cup [w, u]$$

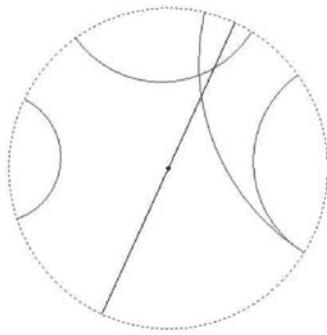
a *geodesic triangle*. We call $[u, v]$, $[v, w]$, $[w, u]$ *sides* of $\Delta(u, v, w)$. We should note that in general, geodesic segments and geodesic triangles are not unique up to their endpoints.

In a metric space, it is sometimes convenient to consider distances between point sets in the following way. We say that a set S is within distance d of another set T if S is contained in the ball $B(T, d)$ of all points within distance d of some point in T . We say that S and T are within distance d of each other if S is within distance d of T and conversely.

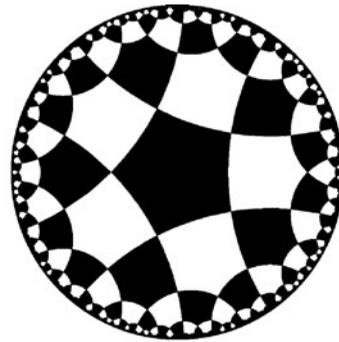
Definition 2.3. (Rips condition.) A geodesic triangle $\Delta(u, v, w)$ in a geodesic metric space (X, d) is called δ -*slim* for some nonnegative real number δ if every point on a side is within distance δ of the union of the other two sides. The space (X, d) is called *Rips δ -hyperbolic* if every geodesic triangle in (X, d) is δ -slim. We denote by $\delta_{\text{Rips}}(X, d)$ the smallest possible value of such δ (it could be infinity).

It is known (see, e.g., [Ghys and de La Harpe 90, Bridson and Haefliger 99, Chepoi et al. 08]) that $\delta(X, d)$ and $\delta_{\text{Rips}}(X, d)$ differ only within a multiplicative constant. In particular, $\delta(X, d) \leq 8\delta_{\text{Rips}}(X, d)$ and $\delta_{\text{Rips}}(X, d) \leq 4\delta(X, d)$. Since we are concerned only with the asymptotic growth of $\delta(X, d)$, the Rips condition can be used in place of Gromov's four-point condition.

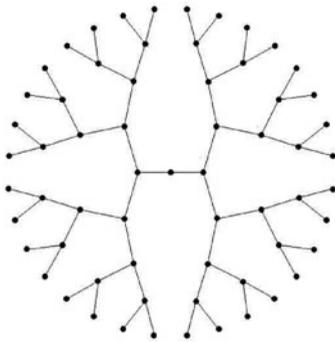
For an undirected unweighted graph $G = (V, E)$, we can also treat it as a geodesic metric space with every edge interpreted as a segment of length 1,



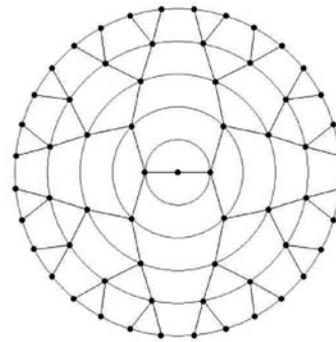
(a) Poincaré disk



(b) Tessellation of (a)



(c) Binary tree



(d) Ringed tree

Figure 1. Poincaré disk, its tessellation, a binary tree, and a ringed tree.

and thus use the Rips condition to define its hyperbolicity, which we denote by $\delta_{\text{Rips}}(G)$. Note that in the case of an unweighted graph, when considering the distance between two geodesics on the graph, we consider only the distance between vertices, since other points on the edges can add at most 2 to the distance between vertices.

2.3. Poincaré Disk

The Poincaré disk (see Figure 1(a) for an illustration) is a well-studied hyperbolic metric space. Although in this paper we touch upon it only briefly when we study ringed-tree graphs, it is useful to convey some intuition about hyperbolicity and treelike behavior.

Definition 2.4. Let $D = B(0, 1)$ be an open disk in the complex plane with origin 0 and radius 1, with the following distance function:

$$d(u, v) = \operatorname{arccosh} \left(1 + \frac{2\|u - v\|^2}{(1 - \|u\|^2)(1 - \|v\|^2)} \right).$$

Then (D, d) is a metric space, which we call the Poincaré disk.

Visually, a (hyperbolic) line in the Poincaré disk is the segment of a circle in the disk that is perpendicular to the circular boundary of the disk, and thus all lines bend inward toward the origin. The hyperbolic distance between two points in the disk of fixed distance in the complex plane increases exponentially as the points move toward the boundary of the disk, meaning that there is much more “space” toward the boundary than around the origin. This can be seen from a tessellation of the Poincaré disk, as shown in Figure 1(b).

2.4. Quasi-Isometry

Quasi-isometry, defined as follows, is a concept used to capture the large-scale similarity between two metric spaces.

Definition 2.5. (Quasi-isometry.) For two metric spaces $(X, d_X), (Y, d_Y)$, we say that $f : X \rightarrow Y$ is a (λ, ϵ) -quasi-isometric embedding from X to Y if for all $u, v \in X$, we have

$$\frac{1}{\lambda} d_X(u, v) - \epsilon \leq d_Y(f(u), f(v)) \leq \lambda d_X(u, v) + \epsilon.$$

Furthermore, if the ϵ neighborhood of $f(Y)$ covers X , then we say that f is a (λ, ϵ) -quasi-isometry. Moreover, we say that X, Y are quasi-isometric if such a (λ, ϵ) -quasi-isometry exists for some constants λ and ϵ .

If two metric spaces are quasi-isometric with some constant, then they have the same “large-scale” behavior. For example, the d -dimensional grid \mathbb{Z}^d and the d -dimensional Euclidean space \mathbb{R}^d are quasi-isometric, realized by the $(\sqrt{d}, \sqrt{d}/2)$ -quasi-isometric embedding $(x, y) \mapsto (x, y)$.

As a second example, consider an infinite ringed tree: start with a binary tree (illustrated in Figure 1(c)) and then connect all vertices at a given tree level into a ring. This is defined more formally in Section 4, but an example is illustrated in Figure 1(d). As we prove in Section 4, the infinite ringed tree is quasi-isometric to the Poincaré disk; thus it may be equivalently viewed as a “softened” binary tree or as a “coarsened” Poincaré disk.

Quasi-isometric embeddings have the important property of preserving hyperbolicity, up to a constant factor, as given by the following proposition.

Proposition 2.6. [Bridson and Haefliger 99, Theorem 1.9, Chapter III.H] *Let X and X' be two metric spaces, and let $f : X' \rightarrow X$ be a (λ, ϵ) -quasi-isometric embedding. If X is δ -hyperbolic, then X' is δ' -hyperbolic, where δ' is a function of δ , λ , and ϵ .*

3. δ -Hyperbolicity of Grid-Based Small-World Graphs

In this section, we consider the δ -hyperbolicity of graphs constructed according to the small-world graph model as formulated in [Kleinberg 00], in which long-range edges are added on top of a base grid, which is a discretization of a low-dimensional Euclidean space.

The model starts with n vertices forming a d -dimensional base grid (with wraparound). More precisely, given positive integers n and d such that $n^{1/d}$ is also an integer, let $B = (V, E)$ be the base grid, with

$$V = \{(x_1, x_2, \dots, x_d) \mid x_i \in \{0, 1, \dots, n^{1/d} - 1\}, i \in [d]\},$$

and

$$E = \{((x_1, x_2, \dots, x_d), (y_1, y_2, \dots, y_d)) \mid \exists j \in [d], y_j = x_j + 1 \bmod n^{1/d} \\ \text{or } y_j = x_j - 1 \bmod n^{1/d}, \forall i \neq j, y_i = x_i\}.$$

Let d_B denote the graph distance metric on the base grid B . We then build a random graph G on top of B such that G contains all vertices and all edges (referred to as grid edges) of B , and for each node $u \in V$, it has one long-range edge (undirected) connected to some node $v \in V$, with probability proportional to $1/d_B(u, v)^\gamma$, where $\gamma \geq 0$ is a parameter. We refer to the probability space of these random graphs as $\text{KSW}(n, d, \gamma)$, and we let $\delta(\text{KSW}(n, d, \gamma))$ denote the random variable of the hyperbolic δ of a randomly chosen graph G in $\text{KSW}(n, d, \gamma)$. Recall that Kleinberg showed that the small-world graphs with $\gamma = d$ allow efficient decentralized routing (with $O(\log^2 n)$ routing hops in expectation), whereas graphs with $\gamma \neq d$ do not allow any efficient decentralized routing (with $\Omega(n^c)$ routing hops for some constant c) [Kleinberg 00], and note that the base grid B has large hyperbolic δ , i.e., $\delta(B) = \Theta(n^{1/d}) = \Theta(D(B))$. Intuitively, the structural reason for the efficient routing performance at $\gamma = d$ is that long-range edges are added “hierarchically” such that each node’s long-range edges are nearly uniformly distributed over all “distance scales.”

3.1. Results and Their Implications

The following theorem summarizes our main technical results on the hyperbolicity of small-world graphs for different combinations of d and γ .

Theorem 3.1. *With probability $1 - o(1)$ (as n goes to infinity), we have*

- (1) $\delta(\text{KSW}(n, d, \gamma)) = \Omega((\log n)^{1/(1.5(d+1)+\varepsilon)})$ when $d \geq 1$ and $\gamma = d$, for every $\varepsilon > 0$ independent of n ;
- (2) $\delta(\text{KSW}(n, d, \gamma)) = \Omega(\log n)$ when $d \geq 1$ and $0 \leq \gamma < d$; and
- (3) $\delta(\text{KSW}(n, d, \gamma)) = \Omega(n^{\frac{\gamma-2}{\gamma-1}-\epsilon})$ when $d = 1$ and $\gamma > 3$, for every $\epsilon > 0$ independent of n .

This theorem, together with the results of [Kleinberg 00] on the navigability of small-world graphs, has several implications. The first result shows that when $\gamma = d$, then with high probability, the hyperbolic δ of the small-world graphs is at least $c(\log n)^{1/1.5(d+1)}$ for some constant c . We know that the diameter is $\Theta(\log n)$ in expectation when $\gamma = d$ [Martel and Nguyen 04]. Thus the small-world graphs at the sweet spot for efficient routing are not polylogarithmically hyperbolic, i.e., δ is not $O(\log^c \log n)$ -hyperbolic for any constant $c > 0$.

However, there is still a gap between our lower bound and the upper bound provided by the diameter, and thus it is still open whether small-world graphs are weakly hyperbolic or not hyperbolic. Overall, though, our result indicates no drastic improvement on hyperbolicity (relative to the improvement of the diameter) for small-world graphs at the sweet spot (where a dramatic improvement was obtained for the efficiency of decentralized routing).

The second result shows that when $\gamma < d$, then $\delta = \Omega(\log n)$. The diameter of the graph in this case is $\Theta(\log n)$ [Martel and Nguyen 04]. Thus, we see that when $\gamma < d$, the hyperbolic δ is asymptotically the same as the diameter, i.e., although δ decreases as edges are added, small-world graphs in this range are not hyperbolic.

The third result concerns the case $\gamma > d$, whereby the random graph degenerates toward the base grid (in the sense that most of the long-range edges are very local), which itself is not hyperbolic. For the general γ , we show that for the case of $d = 1$, the hyperbolic δ is bounded below by a (low-degree) polynomial of n ; this also implies that the graphs in this range are not polylogarithmically hyperbolic. Note that our polynomial exponent $\frac{\gamma-2}{\gamma-1} - \epsilon$ matches the diameter lower bound proven in [Nguyen and Martel 05].

3.2. Outline of the Proof of Theorem 3.1

In this subsection, we provide a summary of the proof of Theorem 3.1. In our analysis, we use two different techniques, one for the first two results in Theorem 3.1, and the other for the last result; in addition, for the first two results, we further divide the analysis into the two cases $d \geq 2$ and $d = 1$.

When $d \geq 2$ and $0 \leq \gamma \leq d$, the main idea of the proof is to pick a square grid of size ℓ_0 (it does not matter from which dimension the square is chosen). We know that when only grid distance is considered, the four corners of the square grid have Gromov δ value equal to ℓ_0 . We will show that as long as ℓ_0 is not very large (to be exact, $O((\log n)^{1/(1.5(d+1)+\varepsilon)})$ when $\gamma = d$ and $O(\log n)$ when $0 \leq \gamma < d$), the probability that at least one pair of vertices on this square grid have a shortest path shorter than their grid distance after adding long-range edges is close to zero (as n tends to infinity). Therefore, with high probability, the four corners selected have Gromov δ as desired in the lower-bound results.

To prove this result, we study the probability that a randomly chosen pair of vertices u and v at grid distance ℓ are connected with a path that contains at least one long-range edge and has length at most ℓ . We bound such ℓ 's above so that this probability is close to zero. To do so, we first classify such paths into a number of categories, based on the pattern of paths connecting u and v : how a path alternates between grid edges and long-range edges, and the direction on each dimension of the grid edges and long-range edges (i.e., whether it is the same direction as from u to v in this dimension, or the opposite direction, or no move in this dimension). We then bound the probability of a path in each category existing and finally bound all such paths in the aggregate. The most difficult part of the analysis is the bounding of the probability of a path existing in each category.

For the case $d = 1$ and $0 \leq \gamma \leq d$, the general idea is similar to the above. The difference is that we do not have a base square to start with. Instead, we find a base ring of length $\Theta(\ell_0)$ using one long-range edge e_0 , where ℓ_0 is fixed to be as in the case of $d \geq 2$. We show that with high probability, (a) such an edge e_0 exists, and (b) the distance between any two vertices on the ring is simply their ring distance. This is enough to prove the lower bound on the hyperbolic δ .

For the case of $\gamma > 3$ and $d = 1$, a different technique is used to prove the lower bound on hyperbolic δ . We first show that in this case, with high probability all long-range edges connect only two vertices with ring distance at most some $\ell_0 = o(\sqrt{n})$. Next, on the one-dimensional ring, we first find two vertices A and B at two opposite ends of the ring. Then we argue that there must be a path \mathcal{P}_{AB}^+ that goes only through the clockwise side of ring from A to B , while there is another path \mathcal{P}_{AB}^- that goes through only the counterclockwise side of the ring

from A to B , and importantly, the shorter length of these two paths is at most $O(\ell_0)$ longer than the distance between A and B . We then choose the midpoints C and D of \mathcal{P}_{AB}^+ and \mathcal{P}_{AB}^- , respectively, and argue that the δ values of the four points $A, B, C,$ and D give the desired lower bound.

3.3. Detailed Proof of Theorem 3.1

3.3.1. The Case $d \geq 2$ and $0 \leq \gamma \leq d$. For this case, let $n' = n^{1/d}$ be the number of vertices on one side of the grid. For convenience, our main analysis for this case uses n' instead of n .

We first provide a couple of lemmas used in our probability calculation.

Lemma 3.2. *There exists a constant c_1 such that for all $k, m \in \mathbb{Z}^+$, we have*

$$\sum_{\substack{y_1 + \dots + y_k = m \\ y_1, \dots, y_k \in \mathbb{Z}^+}} \frac{1}{y_1 y_2 \dots y_k} \leq \frac{(c_1 \ln m)^{k-1}}{m},$$

where the left-hand side is considered to be 0 for $k > m$; and for $k = m = 1$, the right-hand side 0^0 is considered to be equal to 1.

Proof. For $k = 1$, the result is trivial. For $k = 2$, we have

$$\begin{aligned} \sum_{\substack{y_1 + y_2 = m \\ y_1, y_2 \in \mathbb{Z}^+}} \frac{1}{y_1 y_2} &\leq 2 \left(\frac{1}{\lfloor m/2 \rfloor \cdot \lceil m/2 \rceil} + \dots + \frac{1}{(m-1) \cdot 1} \right) \\ &\leq \frac{2}{\lfloor m/2 \rfloor} \left(\frac{1}{\lfloor m/2 \rfloor} + \dots + \frac{1}{1} \right) < c_1 \frac{\ln m}{m}, \end{aligned}$$

where c_1 is roughly 4.

Suppose the lemma holds for $k - 1$, with $k \geq 3$. The induction hypothesis is

$$\sum_{\substack{y_1 + \dots + y_{k-1} = m \\ y_1, \dots, y_{k-1} \in \mathbb{Z}^+}} \frac{1}{y_1 y_2 \dots y_{k-1}} \leq \frac{(c_1 \ln m)^{k-2}}{m}.$$

Since the logarithm function is increasing, we have

$$\begin{aligned} \sum_{\substack{y_1 + \dots + y_k = m \\ y_1, \dots, y_k \in \mathbb{Z}^+}} \frac{1}{y_1 y_2 \dots y_k} &\leq \frac{1}{1} \frac{(c_1 \ln(m-1))^{k-2}}{m-1} + \dots + \frac{1}{m-1} \frac{(c_1 \ln 1)^{k-2}}{1} \\ &\leq (c_1 \ln m)^{k-2} \cdot \sum_{\substack{y_1 + y_2 = m \\ y_1, y_2 \in \mathbb{Z}^+}} \frac{1}{y_1 y_2} \leq (c_1 \ln m)^{k-2} \cdot \frac{c_1 \ln m}{m}. \end{aligned}$$

Therefore, the inequality holds for all k . □

Lemma 3.3. For any constant $\theta \in \mathbb{R}$ with $0 \leq \theta < 1$, there exists a constant c_2 (which depends only on θ) such that for all constants $k, n' \in \mathbb{Z}^+$, $m \in \mathbb{R}$, and nonzero $\lambda_1, \lambda_2, \dots, \lambda_k \in \mathbb{R}$, we have

$$\sum_{\substack{\lambda_1 y_1 + \dots + \lambda_k y_k = m \\ y_1, \dots, y_k \in \{1, 2, \dots, n'\}}} \frac{1}{y_1^\theta y_2^\theta \cdots y_k^\theta} \leq (c_2 n')^{(k-1)(1-\theta)},$$

where the left-hand side is considered to be 0 if there is no $y_1, y_2, \dots, y_k \in \{1, 2, \dots, n'\}$ satisfying $\lambda_1 y_1 + \lambda_2 y_2 + \dots + \lambda_k y_k = m$.

Proof. For each tuple $(y_1, y_2, \dots, y_{k-1}) \in \{1, 2, \dots, n'\}^{k-1}$, there is at most one $y_k \in \{1, 2, \dots, n'\}$ satisfying $\lambda_1 y_1 + \lambda_2 y_2 + \dots + \lambda_k y_k = m$. Since $1/y_k^\theta \leq 1$, because $0 \leq \theta < 1$, we have

$$\begin{aligned} \sum_{\substack{\lambda_1 y_1 + \dots + \lambda_k y_k = m \\ y_1, \dots, y_k \in \{1, 2, \dots, n'\}}} \frac{1}{y_1^\theta y_2^\theta \cdots y_k^\theta} &\leq \sum_{y_1, \dots, y_{k-1} \in \{1, 2, \dots, n'\}} \frac{1}{y_1^\theta y_2^\theta \cdots y_{k-1}^\theta} \\ &= \left(\sum_{i=1}^{n'} \frac{1}{i^\theta} \right)^{k-1} \leq (c_2 n')^{(k-1)(1-\theta)}, \end{aligned}$$

where c_2 is roughly $(\frac{1}{1-\theta})^{\frac{1}{1-\theta}}$. The lemma is proved. \square

3.3.2. Classification of Paths. In a d -dimensional random graph $\text{KSW}(n, d, \gamma)$, there are two kinds of edges: *grid edges*, which are edges on the grid, and *long-range edges*, which are randomly added.

Fix two vertices u and v . A path from u to v may contain some long-range edges and some grid edges. We divide the path into several *segments* along the way from u to v : (a) each segment is either one long-range edge (called a *long-range segment*) or several consecutive grid edges (called a *grid segment*); and (b) two consecutive segments cannot both be grid segments (otherwise, combine them into one segment).

We use a d -dimensional vector to denote each edge, so that the source coordinate plus this vector equals the destination coordinate modulo n' . For a grid with wraparound, there may be multiple vectors corresponding to one edge. We choose the vector in which every element is from $\{-\lfloor \frac{n'}{2} \rfloor, -\lfloor \frac{n'}{2} \rfloor + 1, \dots, \lfloor \frac{n'-1}{2} \rfloor\}$. In this way, the vector representation of each edge is unique, and the absolute value of every dimension is the smallest. We call this the *edge vector* of that edge. For every segment in the path, we call the summation (not modulo n') of all edge vectors the *segment vector*. For a vector (x_1, x_2, \dots, x_d) , define its *sign pattern* as $(\text{sgn}(x_1), \text{sgn}(x_2), \dots, \text{sgn}(x_d))$.

We say that two paths from u to v belong to the same *category* if

- (a) they have the same number of segments;
- (b) their corresponding segments are of the same type (long-range or grid segments);
- (c) for every pair of corresponding long-range segments in the two paths, the sign patterns of their segment vectors are the same;
- (d) for every pair of corresponding grid segments in the two paths, their segment vectors are equal; and
- (e) the summations (not modulo n') of all segment vectors in the two paths are equal.

The last condition is used to enforce that the two paths must wrap around along each dimension exactly the same number of times on the grid with wraparound.

In one category, there exist paths whose long-range edges are identical but the grid edges may be different. To compute the probability of a path existing in a category, we need only consider one path among the paths with identical long-range edges, since grid edges do not change probabilistic events, and thus one such path exists if and only if other such paths exist.

We also assume that there are no repeated long-range edges in every path. For a path that has repeated long-range edges, we can obtain a shorter subpath without any repeated long-range edges, so that the original path exists if and only if the new one exists. Since we are going to calculate the probability of paths not exceeding some length, it is safe to consider only paths without repeated long-range edges.

Lemma 3.4. *There exists a constant c_3 (depending on d) such that for every fixed ℓ , the number of categories of paths from u to v of length ℓ is at most c_3^ℓ .*

Proof. For each edge on the path, if it is a grid edge, it could be in one of the d dimensions, and in each dimension, it could be in one of the two opposite directions, and thus a grid edge has $2d$ possibilities. If the edge is a long-range edge, in each dimension its sign has three possibilities, $(+1, 0, -1)$, and so 3^d possibilities altogether for the sign pattern of the long-range segment vector.

Moreover, in the grid with wraparound, we consider each wraparound of the path in some dimension to be one *circuit* in that dimension. Then the path can have at most $2\ell + 1$ different numbers of circuits in each dimension (ranging from ℓ circuits in one direction up to ℓ circuits in the other direction), so the summation of all segment vectors has at most $(2\ell + 1)^d$ different values. The choice of each edge out of $2d + 3^d$ possibilities and the total summation of segment

vectors determine a category. Therefore, the number of categories is bounded by $(2d + 3^d)^\ell (2\ell + 1)^d < c_3^\ell$ for some c_3 . \square

The above bound on the number of categories is not tight enough to be useful for later analysis, when the number of long-range segments is small. Thus, we further bound the number of categories in the following way.

Lemma 3.5. *There exists a constant c_4 (depending on d) such that for every fixed $\ell < n'$ and k with $1 \leq k \leq \ell$, the number of categories of paths from u to v of length ℓ and having k long-range segments is at most $c_4^k \ell^{(k+1)(d+1)} / k^{kd}$.*

Proof. Similar to the proof of Lemma 3.4, we consider each wraparound of the path in some dimension to be one *circuit* in that dimension. For a path from u to v of length ℓ and containing k long-range segments ($1 \leq k \leq \ell < n'$), the summation of all segment vectors has at most $(2k + 1)^d$ choices. This is because the path can have at most $2k + 1$ different number of circuits on each dimension (k circuits in one direction to k circuits in the other direction). We consider the number of categories for a fixed summation of segment vectors first.

Suppose there are t grid segments, each having a_1, a_2, \dots, a_t edges ($t \leq k + 1$, $a_i \geq 1$, $a_1 + a_2 + \dots + a_t < \ell$). If $t = 0$, then $k = \ell$, and it is easy to see that there are at most $(3^d)^k$ categories. Suppose now $t \geq 1$. For the i th grid segment with a_i grid edges, its segment vector is such that in each dimension, the only possible values are $-a_i, -a_i + 1, \dots, 0, \dots, a_i - 1, a_i$. Thus, the number of possible segment vectors is $(2a_i + 1)^d \leq 3^d a_i^d$. Since each long-range edge has 3^d possible sign patterns, the number of categories for fixed t and a_1, a_2, \dots, a_t is at most $(3^d)^k \prod_{i=1}^t 3^d a_i^d < 9^{(k+1)d} (\ell/t)^{td}$, where the arithmetic-geometric mean inequality is used.

The tuple (a_1, a_2, \dots, a_t) has fewer than ℓ^t possibilities. Considering that there are $(2k + 1)^d$ different possibilities of segment vector summations, the total number of categories from u to v with length ℓ and containing k long-range edges is at most

$$\begin{aligned} & (2k + 1)^d \left\{ (3^d)^k + \sum_{t=1}^{k+1} \ell^t \cdot 9^{(k+1)d} (\ell/t)^{td} \right\} \\ & < (2k + 1)^d \left\{ 3^{dk} + (k + 1) 9^{(k+1)d} \max_{1 \leq t \leq k+1} \{ \ell^t (\ell/t)^{td} \} \right\} \\ & = (2k + 1)^d \left\{ 3^{dk} + (k + 1) 9^{(k+1)d} \ell^{k+1} \left(\frac{\ell}{k+1} \right)^{(k+1)d} \right\} \\ & < c_4^k \cdot \frac{\ell^{(k+1)(d+1)}}{k^{kd}}, \end{aligned}$$

where c_4 is a constant depending only on d . \square

3.3.3. Probability Calculation. We first give a lemma to calculate the probability of the existence of a specific edge. For an integer x , we define \bar{x} to be $|x|$ if $x \neq 0$ and 1 if $x = 0$. We also define the function $f(n')$ as follows:

$$f(n') = \begin{cases} \ln n', & \gamma = d, \\ (n')^{d-\gamma}, & 0 \leq \gamma < d. \end{cases}$$

Lemma 3.6. *For two vertices u and v , the probability of the existence of a long-range undirected edge between u and v is at most $c_5(\bar{x}_1 \cdot \bar{x}_2 \cdots \bar{x}_d)^{-\gamma/d} / f(n')$, where c_5 is a constant depending only on d and γ , and (x_1, x_2, \dots, x_d) is the edge vector if there exists a long-range edge from u to v that depends only on u and v .*

Proof. Suppose the nonzero elements of (x_1, x_2, \dots, x_d) are $(x_{i_1}, x_{i_2}, \dots, x_{i_{d'}})$ ($d' \leq d$). Let p be the probability of adding an edge from u to v . Then

$$\begin{aligned} p &= \frac{(|x_1| + |x_2| + \cdots + |x_d|)^{-\gamma}}{\Theta(\sum_{i=1}^{n'} \frac{i^{d-1}}{i^\gamma})} = O\left(\frac{|x_{i_1} \cdot x_{i_2} \cdots x_{i_{d'}}|^{-\gamma/d'}}{\sum_{i=1}^{n'} i^{d-1-\gamma}}\right) \\ &\leq O\left(\frac{|x_{i_1} \cdot x_{i_2} \cdots x_{i_{d'}}|^{-\gamma/d}}{\sum_{i=1}^{n'} i^{d-1-\gamma}}\right) = O\left(\frac{(\bar{x}_1 \cdot \bar{x}_2 \cdots \bar{x}_d)^{-\gamma/d}}{\sum_{i=1}^{n'} i^{d-1-\gamma}}\right) \\ &= O\left(\frac{(\bar{x}_1 \cdot \bar{x}_2 \cdots \bar{x}_d)^{-\gamma/d}}{f(n')}\right). \end{aligned}$$

Moreover, the edge may also be from v to u , which also has probability p . By union bound, the probability of the undirected edge (u, v) is $O(p)$. \square

The following lemma gives the probability that one edge jumps within a local area.

Lemma 3.7. *For a vertex u and a long-range edge (u, v) from u , the probability that the grid distance between u and v is less than s is at most $c_6 f(s) / f(n')$, where c_6 is a constant depending only on d and γ .*

Proof. By union bound, the probability is at most

$$O\left(\sum_{i=1}^s i^{d-1} \frac{i^{-\gamma}}{f(n')}\right) = O\left(\sum_{i=1}^s \frac{i^{d-1-\gamma}}{f(n')}\right) \leq c_6 \frac{f(s)}{f(n')},$$

where c_6 is a constant. \square

Given a path category \mathcal{C} , we now calculate the probability of a path in \mathcal{C} existing.

Lemma 3.8. *Given a path category \mathcal{C} with length ℓ and k long-range edges, the probability that there exists a path in \mathcal{C} is at most*

$$\begin{cases} c_5^k (c_7^k k^k)^d / (\ln n')^k, & \gamma = d, \\ c_5^k c_2^{(k-1)(d-\gamma)} / (n')^{d-\gamma}, & 0 \leq \gamma < d, \end{cases}$$

where c_2 and c_5 are the constants given in Lemma 3.6, and c_7 is another constant.

Proof. Let the segment vectors of the long-range edges in a path $P \in \mathcal{C}$ be

$$(x_{11}, x_{12}, \dots, x_{1d}), \quad (x_{21}, x_{22}, \dots, x_{2d}), \quad \dots, \quad (x_{k1}, x_{k2}, \dots, x_{kd}).$$

By our definition of a category, all paths in \mathcal{C} have the same sign patterns on the corresponding long-range segment vectors. Thus, we can define the following sets for the category \mathcal{C} :

$$\begin{aligned} A_i^+ &= \{j \mid x_{ji} > 0, 1 \leq j \leq k\}, \\ A_i^- &= \{j \mid x_{ji} < 0, 1 \leq j \leq k\}, \\ A_i^0 &= \{j \mid x_{ji} = 0, 1 \leq j \leq k\}, \end{aligned}$$

for all $1 \leq i \leq d$. For fixed u and v , there is a fixed vector (t_1, t_2, \dots, t_d) such that the summation of the segment vectors of all long-range segments in any $P \in \mathcal{C}$ is the vector (t_1, t_2, \dots, t_d) . This is because the summation of all segment vectors from u to v is fixed, and all grid segments have fixed segment vectors. Therefore, a path in \mathcal{C} can be characterized by kd integers x_{11}, \dots, x_{kd} satisfying the following for all $1 \leq i \leq d$:

$$\begin{cases} x_{ji} \in \{1, \dots, n'\} & \text{for } j \in A_i^+, \\ x_{ji} \in \{-n', \dots, -1\} & \text{for } j \in A_i^-, \\ x_{ji} = 0 & \text{for } j \in A_i^0, \\ \sum_{j \in A_i^+} |x_{ji}| - \sum_{j \in A_i^-} |x_{ji}| = t_i. \end{cases} \tag{3.1}$$

The probability that some path exists is the product of the probability of the first edge, the probability of the second edge conditioned on the existence of the first edge, the probability of the third edge conditioned on the existence of the first two edges, etc. In our model, the probability of an (undirected) edge conditioned on the existence of other (undirected) edges is less than or equal to the probability without condition, because each vertex can connect to exactly one other vertex (when the edge direction is taken into consideration). Hence we can use the product of the probabilities of all edges as an upper bound for the probability of a path. By union bound and Lemma 3.6, the probability that a

path exists in \mathcal{C} is at most

$$\begin{aligned}
 & \sum_{\text{all paths in } \mathcal{C}} \prod_{j=1}^k \frac{c_5(\overline{x}_{j1} \cdot \overline{x}_{j2} \cdots \overline{x}_{jd})^{-\gamma/d}}{f(n')} \\
 & \leq \frac{c_5^k}{f^k(n')} \sum_{\substack{x_{11}, \dots, x_{kd}: \\ \text{satisfying (3.1)}}} \prod_{j=1}^k (\overline{x}_{j1} \cdot \overline{x}_{j2} \cdots \overline{x}_{jd})^{-\gamma/d} \\
 & = \frac{c_5^k}{f^k(n')} \cdot \prod_{i=1}^d \left\{ \sum_{\substack{\overline{x}_{1i}, \dots, \overline{x}_{ki}: \\ \sum_{j \in A_i^+} \overline{x}_{ji} - \sum_{j \in A_i^-} \overline{x}_{ji} = t_i; \\ \text{for } j \in A_i^0, \overline{x}_{ji} = 1; \text{ for } j \in A_i^+ \cup A_i^-, \overline{x}_{ji} \in \{1, \dots, n'\}}} \left(\prod_{j=1}^k \overline{x}_{ji} \right)^{-\gamma/d} \right\} \\
 & = \frac{c_5^k}{f^k(n')} \cdot \prod_{i=1}^d \left\{ \sum_{\substack{\overline{x}_{ji} \text{ (} j \text{ ranges in } A_i^+ \cup A_i^-): \\ \sum_{j \in A_i^+} \overline{x}_{ji} - \sum_{j \in A_i^-} \overline{x}_{ji} = t_i; \text{ and } \overline{x}_{ji} \in \{1, \dots, n'\}}} \left(\prod_{j \in A_i^+ \cup A_i^-} \overline{x}_{ji} \right)^{-\gamma/d} \right\}. \tag{3.2}
 \end{aligned}$$

The inner brace summation is considered to be 1 for $A_i^+ = A_i^- = \emptyset$. Now we consider the following sum for disjoint sets $A^+, A^- \subseteq \{1, 2, \dots, k\}$ (at least one is not empty) and numbers $t \in \mathbb{Z}, \gamma \geq 0, d \geq 0$:

$$\sum_{\substack{y_j \text{ (} j \text{ ranges in } A^+ \cup A^-): \\ \sum_{j \in A^+} y_j - \sum_{j \in A^-} y_j = t; \text{ and } y_j \in \{1, 2, \dots, n'\}}} \left(\prod_{j \in A^+ \cup A^-} y_j \right)^{-\gamma/d}. \tag{3.3}$$

Case $\gamma = d$. If one of A^+ and A^- is \emptyset , we need only consider the case that $t \neq 0$, because the sum (3.3) is 0 for $t = 0$. By Lemma 3.2, the sum (3.3) is bounded by

$$c_1^k \frac{(\ln |t|)^k}{|t|} \leq c_1^k \max_{x \geq 1} \left\{ \frac{(\ln x)^k}{x} \right\} = \left(\frac{c_1 k}{e} \right)^k.$$

If neither A^+ nor A^- is \emptyset , we assume that $t \geq 0$. (The calculation for $t < 0$ is similar if we exchange A^+ and A^- .) By Lemma 3.2, the sum (3.3) is bounded by

$$\begin{aligned} & \sum_{s=1}^{kn'} \left(\sum_{j \in A^-} y_j = s \frac{1}{\prod_{j \in A^-} y_j} \right) \left(\sum_{j \in A^+} y_j = s+t \frac{1}{\prod_{j \in A^+} y_j} \right) \\ & \leq \sum_{s=1}^{kn'} \frac{(c_1 \ln s)^{|A^-|-1}}{s} \cdot \frac{(c_1 \ln(s+t))^{|A^+|-1}}{s+t} \leq c_1^k \sum_{s=1}^{kn'} \frac{(\ln(s+t))^k}{s(s+t)} \\ & = c_1^k \sum_{s=1}^{kn'} \left(\frac{(\ln(s+t))^k}{\sqrt{s+t}} \cdot \frac{1}{s\sqrt{s+t}} \right) \leq c_1^k \cdot \max_{x \geq 1} \left\{ \frac{(\ln x)^k}{x^{0.5}} \right\} \cdot \sum_{s=1}^{\infty} \frac{1}{s^{1.5}} \\ & = c_1^k \cdot \frac{(2k)^k}{e^k} \cdot O(1). \end{aligned}$$

In either case, the sum (3.3) is bounded by $c_7^k k^k$ for some constant c_7 . By (3.2), the probability of a path in \mathcal{C} is at most $c_5^k (c_7^k k^k)^d / (\ln n')^k$.

Case $0 \leq \gamma < d$. By Lemma 3.3, the sum (3.3) is bounded by

$$(c_2 n')^{(|A^+|+|A^-|-1)(1-\gamma/d)} \leq (c_2 n')^{(k-1)(1-\gamma/d)}.$$

Therefore, by (3.2), the probability of a path in \mathcal{C} is at most

$$\frac{c_5^k}{(n')^{k(d-\gamma)}} \cdot \left((c_2 n')^{(k-1)(1-\gamma/d)} \right)^d = \frac{c_5^k c_2^{(k-1)(d-\gamma)}}{(n')^{d-\gamma}}.$$

□

Finally, we apply Lemmas 3.4, 3.5, 3.7, and 3.8 to show that the probability that at least one pair of vertices are connected by a short path with at least one long-range edge is vanishingly small.

Lemma 3.9. *For any two vertices u and v , the probability that a path exists connecting u and v with at least one long-range edge and total length at most ℓ is $o(1)$ when $\ell \leq (\log n')^{1/(1.5(d+1)+\varepsilon)}$ ($\varepsilon > 0$) for $\gamma = d$; and $\ell < c \log n'$ (c is some constant depending only on d and γ) for $0 \leq \gamma < d$.*

Proof. We first study the probability of a path with exact length ℓ , and we divide it into the following two cases.

Case $\gamma = d$. If $k = 1$, the probability that a path from u to v with length ℓ exists is at most $O(\ell^d c_6 \ln(3\ell) / \ln n') = O(\ell^d \ln \ell / \ln n')$. To see this, we divide the path into three segments: a grid segment followed by one long-range edge, then followed by another grid segment. The first grid segment can reach at most $O(\ell^d)$ destinations. For each such destination w , the long-range edge has to reach

some vertex within grid distance ℓ of vertex v . By the triangle inequality, it must be within grid distance 3ℓ of vertex w , since the distances between v and u and between u and w are both at most ℓ .

By Lemma 3.7, we know that this probability is $c_6 \ln(3\ell) / \ln n'$. Therefore, the above statement holds.

For $k \geq 2$, we simply combine Lemmas 3.5 and 3.8. Thus, the probability that a path with length ℓ exists is at most

$$\begin{aligned} & O\left(\frac{\ell^d \ln \ell}{\ln n'}\right) + \sum_{k=2}^{\ell} \frac{c_4^k \ell^{(k+1)(d+1)}}{k^{kd}} \cdot \frac{c_5^k (c_7^k k^k)^d}{(\ln n')^k} \\ &= O\left(\frac{\ell^d \ln \ell}{\ln n'}\right) + \sum_{k=2}^{\ell} \left(\frac{c_4 c_5 c_7^d \cdot \ell^{(d+1)(1+\frac{1}{k})}}{\ln n'}\right)^k \\ &\leq O\left(\frac{\ell^d \ln \ell}{\ln n'}\right) + \sum_{k=2}^{\ell} \left(\frac{c_4 c_5 c_7^d \cdot \ell^{1.5(d+1)}}{\ln n'}\right)^k. \end{aligned}$$

This probability is $o(1)$ when $\ell < (\log n')^{1/(1.5(d+1)+\varepsilon)}$ for any $\varepsilon > 0$.

Case $0 \leq \gamma < d$. In this case, we combine Lemmas 3.4 and 3.8. The probability that a path with length ℓ exists is at most

$$c_3^\ell \cdot \max_{1 \leq k \leq \ell} \left\{ \frac{c_5^k c_2^{(k-1)(d-\gamma)}}{(n')^{d-\gamma}} \right\} \leq \frac{(c_3 c_5 c_2^{d-\gamma})^\ell}{(n')^{d-\gamma}},$$

where the inequality is based on $c_2, c_5 \geq 1$, which is obviously the case. This probability is $o(1)$ when $\ell < c \log n'$ for some properly chosen constant c , which depends only on d and γ .

We now consider the case of path length less than ℓ . Let w be a grid neighbor of v . For any path connecting u and v with length less than ℓ , we can add grid edges (v, w) followed by (w, v) , to increase the path length to either $\ell - 1$ or ℓ . Thus, the probability that a path of length at most ℓ exists is the same as the probability that a path of length $\ell - 1$ or ℓ exists. The case of $\ell - 1$ is exactly like the case of ℓ argued above. Therefore, the lemma holds. \square

With Lemma 3.9, we are ready to prove Theorem 3.1 for the case $d \geq 2$ and $0 \leq \gamma \leq d$.

Proof of Theorem 3.1 for the case $d \geq 2$ and $0 \leq \gamma \leq d$. We define

$$\ell = \begin{cases} \lfloor (\log n')^{1/(1.5(d+1)+\varepsilon)} / 2 \rfloor & \text{when } \gamma = d, \\ \lfloor c \log n' / 2 \rfloor & \text{when } 0 \leq \gamma < d, \end{cases}$$

where c is the constant determined by Lemma 3.9. In the base grid, find any square in any dimension with one side of length equal to ℓ . Consider the four vertices of the square. By Lemma 3.9, the probability that a pair of vertices of

this square are connected by a path with at least one long-range edge and total length at most 2ℓ is $o(1)$. Thus by union bound, the distance between any pair of these four vertices is exactly their grid distance, with probability $1 - o(1)$. This means that the δ value of these four vertices is ℓ , with probability $1 - o(1)$. Therefore, we know that with probability $1 - o(1)$,

$$\delta(\text{KSW}(n, d, \gamma)) = \Omega\left((\log n')^{1/(1.5(d+1)+\varepsilon)}\right) = \Omega\left((\log n)^{1/(1.5(d+1)+\varepsilon)}\right)$$

when $d \geq 2$ and $\gamma = d$, and

$$\delta(\text{KSW}(n, d, \gamma)) = \Omega(\log n') = \Omega(\log n)$$

when $d \geq 2$ and $0 \leq \gamma < d$. □

Remark 3.10. (The limitation of this approach.) We already have a tight lower bound for $0 \leq \gamma < d$. However, the lower bound $\Omega((\log n)^{1/(1.5(d+1)+\varepsilon)})$ for $\gamma = d$ does not match the upper bound $O(\log n)$. We show below that the lower bound can never be improved to $\Omega(\log n)$ by the above technique of proving that the grid distance is the shortest on every small square (with high probability).

Consider a KSW graph with $\gamma = d$. For an $\ell \times \ell$ square S , let u be its upper left vertex and let v be its lower right vertex. Similar to the proof in Lemma 3.7, the probability of existence of an edge linking any w to the $\ell/2 \times \ell/2$ square on the lower right side of w within the square S (but excluding w 's two grid neighbors in the square) is

$$q = \Theta\left(\sum_{i=2}^{\ell/2} i \cdot \frac{i^{-d}}{\ln n'}\right) = \begin{cases} \Theta\left(\frac{\log \ell}{\log n'}\right) & d = 2, \\ \Theta\left(\frac{1}{\log n'}\right) & d \geq 3. \end{cases}$$

Consider the $\ell/2 \times \ell/2$ square on the lower right side of u , which is the upper left quadrant of the original square. The probability that at least one vertex w in this quadrant links to its lower right $\ell/2 \times \ell/2$ square is $1 - (1 - q)^{\ell/2 \times \ell/2}$. This probability is almost 1 when

$$\ell = \omega\left(\sqrt{\frac{\log n'}{\log \log n'}}\right), \quad \text{for } d = 2,$$

or $\ell = \omega(\sqrt{\log n'})$, for $d \geq 3$. If there exists such a vertex w in the upper right quadrant, suppose it links to a vertex x in its lower right $\ell/2 \times \ell/2$ square. Then x must also be in the original square S , and we can have a path from u to w following the grid path, then the long-range edge from w to x , and then the grid path from x to v . This path connects u and v and must be shorter than the grid

paths from u to v . Therefore, our technique cannot improve the lower bound to

$$\omega\left(\sqrt{\frac{\log n'}{\log \log n'}}\right) = \omega\left(\sqrt{\frac{\log n}{\log \log n}}\right), \quad \text{for } d = 2,$$

or $\omega(\sqrt{\log n'}) = \omega(\sqrt{\log n})$, for $d \geq 3$.

3.3.4. The Case $d = 1$ and $0 \leq \gamma \leq 1$. In this section, we give lower bounds on δ for the one-dimensional KSW model (based on an n -vertex ring). Let the n vertices be v_0, \dots, v_{n-1} . Let

$$\ell_0 = \begin{cases} \lfloor (\log n)^{1/(1.5(d+1)+\varepsilon)} \rfloor & \text{when } \gamma = 1, \\ \lfloor c \log n \rfloor & \text{when } 0 \leq \gamma < 1, \end{cases}$$

where c is the number determined in Lemma 3.9.

The idea is to find a long-range edge e_0 between two vertices with grid distance ℓ_0 . Since e_0 forms a local ring with the original grid, we can give a lower bound of δ such that the ring distances (with respect to the grid edges and e_0) are minimal even after adding the long-range edges. We first calculate the probability of ring distances being minimal under the condition of existing e_0 .

We divide the construction of a KSW graph into two stages: (1) every vertex links to exactly one other vertex according to some distribution; and (2) we ignore the edge direction and consider the graph undirected. Let \mathcal{E}_i ($0 \leq i \leq n - 1$) be the event that v_i links to $v_{(i+\ell_0 \bmod n)}$ in the first stage. Under the condition that \mathcal{E}_i happens, $v_{(i+\ell_0 \bmod n)}$ is still free to link to any vertex. The events $\mathcal{E}_0, \mathcal{E}_1, \dots, \mathcal{E}_{n-1}$ are independent.

Lemma 3.11. *Under the condition that \mathcal{E}_i happens (v_i links to $v_{(i+\ell_0 \bmod n)}$), for any two vertices u, w on the curve from v_i to $v_{(i+\ell_0 \bmod n)}$ (exclusive), the conditional probability of a path existing that connects u and w with at least one long-range edge other than $(v_i, v_{(i+\ell_0 \bmod n)})$ and path length at most ℓ_0 is $o(1)$.*

Proof. Let e_0 be the edge $(v_i, v_{(i+\ell_0 \bmod n)})$. We still follow the arguments for $d \geq 2$ and $0 \leq \gamma \leq d$, but we change the classification of paths slightly: now e_0 is a type of edge by itself, and thus together with grid edges and long-range edges, we have three types of edges and three types of corresponding segments. For each original category of paths with length ℓ and k long-range segments defined in the previous section, we further divide it into at most $k + 1$ categories, based on whether e_0 was the first, second, \dots , or k th long-range segment, or e_0 does not appear in the path. The number of long-range edges in a category with e_0 is decreased by 1. For the categories with only grid edges and e_0 , those paths now

have no long-range edge, and they exist with conditional probability 1; hence they will not be calculated.

For Lemma 3.4, the number of categories is increased by a factor of at most $\ell + 1$. For Lemma 3.5, the number of categories is increased by a factor of at most $k + 1$. Thus, we have only to properly adjust the constants c_3 and c_4 in these two lemmas to make them still hold. Lemma 3.6 is not changed except that e_0 is no longer a long-range edge and the lemma is not applicable on e_0 . Lemma 3.7 is not affected. In the proof of Lemma 3.8, a category without e_0 can be calculated normally. For a category with e_0 , the number of long-range edges is decreased by 1, as stated above. One can see that all the arguments still hold by changing the summation of long-range edges (t_i in (3.1)) to contain e_0 .

For Lemma 3.9, only the case $\gamma = d$ and $k = 1$ (one long-range edge) needs some modification. If the path does not contain e_0 , the calculation still holds. Otherwise, we can divide the path into five segments: grid, long-range, grid, e_0 , grid (or grid, e_0 , grid, long-range, grid). We just consider the three consecutive segments grid, e_0 , grid as a whole, which contains $O(\ell)$ edges and can reach $O(\ell^d) = O(\ell)$ destinations. One can see that the previous argument still holds. Therefore, the consequence of Lemma 3.9 still holds. \square

With Lemma 3.11, we give the proof of Theorem 3.1 for the case $d = 1$ and $0 \leq \gamma \leq 1$.

Proof of Theorem 3.1 for the case $d = 1$ and $0 \leq \gamma \leq 1$. Let \mathcal{F} be the event $\delta \geq \frac{1}{4}\ell_0 - 3$. We first show that $\Pr\{\mathcal{F} \mid \mathcal{E}_i\} = 1 - o(1)$. Pick four vertices

$$\begin{aligned} A_i &= v_{(i + \lfloor \frac{1}{8}\ell_0 \rfloor \bmod n)}, & B_i &= v_{(i + \lfloor \frac{3}{8}\ell_0 \rfloor \bmod n)}, \\ C_i &= v_{(i + \lfloor \frac{5}{8}\ell_0 \rfloor \bmod n)}, & D_i &= v_{(i + \lfloor \frac{7}{8}\ell_0 \rfloor \bmod n)}. \end{aligned}$$

By Lemma 3.11 and union bound, we know that with probability $1 - o(1)$, the distances between every pair of vertices are the ring (grid edges plus $(v_i, v_{(i+\ell_0 \bmod n)})$) distances. That is, the distances between $A_i B_i$, $B_i C_i$, $C_i D_i$, $D_i A_i$ are roughly $\frac{1}{4}\ell_0$ (off by at most 2), and the distances between $A_i C_i$ and $B_i D_i$ are roughly $\frac{1}{2}\ell_0$ (off by at most 1). Therefore, by considering the four vertices A_i, B_i, C_i, D_i , we have δ at least

$$\frac{1}{2} \left(\frac{1}{2}\ell_0 - 1 + \frac{1}{2}\ell_0 - 1 - \frac{1}{4}\ell_0 - 2 - \frac{1}{4}\ell_0 - 2 \right) = \frac{1}{4}\ell_0 - 3$$

with conditional probability $1 - o(1)$.

For every \mathcal{E}_i , the probability that v_i links to $v_{(i+\ell_0 \bmod n)}$ is

$$\Pr\{\mathcal{E}_i\} = \Theta \left(\frac{\ell_0^{-\gamma}}{f(n)} \right) = \begin{cases} \Theta \left((\log n)^{-\frac{2+\epsilon}{3+\epsilon}} \right) & \gamma = 1, \\ \Theta (\log n / n^{1-\gamma}) & 0 \leq \gamma < 1. \end{cases}$$

Define this probability as q . We have

$$\Pr\{\mathcal{F} \text{ and } \mathcal{E}_i\} = \Pr\{\mathcal{F} \mid \mathcal{E}_i\} \Pr\{\mathcal{E}_i\} = (1 - o(1))q.$$

Let K be the random variable denoting the number of \mathcal{E}_i 's that occur. We define $m = E[K]$, and we have $m = nq$. One can check that $m = \text{poly}(n)$ for both cases $\gamma = 1$ and $0 \leq \gamma < 1$. By the Chernoff bound,

$$\Pr\{|K - m| \leq m^{0.6}\} > 1 - 2e^{-(m^{-0.4})^2/4m} = 1 - 2e^{-m^{0.2}/4}.$$

Hence with very high probability, K is close to m . Let \mathcal{G} denote the event that $m - m^{0.6} \leq K \leq m + m^{0.6}$. We show that $\Pr\{\mathcal{F} \text{ and } \mathcal{E}_i \text{ and } \mathcal{G}\}$ is very close to $\Pr\{\mathcal{F} \text{ and } \mathcal{E}_i\}$. In fact,

$$\Pr\{\mathcal{F} \text{ and } \mathcal{E}_i \text{ and } \mathcal{G}\} \geq \Pr\{\mathcal{F} \text{ and } \mathcal{E}_i\} - \Pr\{\text{not } \mathcal{G}\} = (1 - o(1)) \Pr\{\mathcal{F} \text{ and } \mathcal{E}_i\},$$

because

$$\Pr\{\text{not } \mathcal{G}\} < 2e^{-m^{0.2}/4} = 2e^{-\text{poly}(n)}$$

is much smaller than

$$\Pr\{\mathcal{F} \text{ and } \mathcal{E}_i\} = (1 - o(1))q = \omega(1/n).$$

On the other hand, it is straightforward that

$$\Pr\{\mathcal{F} \text{ and } \mathcal{E}_i \text{ and } \mathcal{G}\} \leq \Pr\{\mathcal{F} \text{ and } \mathcal{E}_i\}.$$

Hence $\Pr\{\mathcal{F} \text{ and } \mathcal{E}_i \text{ and } \mathcal{G}\}$ differs from $\Pr\{\mathcal{F} \text{ and } \mathcal{E}_i\}$ by a factor at most $(1 - o(1))$, and

$$\Pr\{\mathcal{F} \text{ and } \mathcal{E}_i \text{ and } \mathcal{G}\} = (1 - o(1))q.$$

For every $r = r_0 r_1 \cdots r_{n-1} \in \{0, 1\}^n$, define \mathcal{H}_r to be the event that \mathcal{E}_j happens iff $r_j = 1$ for all $j = 0, 1, \dots, n - 1$. We can see that the \mathcal{H}_r are mutually exclusive for $r \in \{0, 1\}^n$. Hence

$$\Pr\{\mathcal{F} \text{ and } \mathcal{E}_i \text{ and } \mathcal{G}\} = \sum_{k=m-m^{0.6}}^{m+m^{0.6}} \sum_{\substack{r \in \{0,1\}^n \\ r_i=1 \text{ and } r \text{ has } k \text{ 1's}}} \Pr\{\mathcal{F} \text{ and } \mathcal{H}_r\}.$$

Therefore,

$$\begin{aligned}
 \Pr\{\mathcal{F} \text{ and } \mathcal{G}\} &= \sum_{k=m-m^{0.6}}^{m+m^{0.6}} \sum_{\substack{r \in \{0,1\}^n \\ r \text{ has } k \text{ 1's}}} \Pr\{\mathcal{F} \text{ and } \mathcal{H}_r\} \\
 &= \sum_{k=m-m^{0.6}}^{m+m^{0.6}} \frac{1}{k} \cdot \sum_{i=0}^{n-1} \sum_{\substack{r \in \{0,1\}^n \\ r_i=1 \text{ and } r \text{ has } k \text{ 1's}}} \Pr\{\mathcal{F} \text{ and } \mathcal{H}_r\} \\
 &> \frac{1}{m+m^{0.6}} \cdot \sum_{i=0}^{n-1} \sum_{k=m-m^{0.6}}^{m+m^{0.6}} \sum_{\substack{r \in \{0,1\}^n \\ r_i=1 \text{ and } r \text{ has } k \text{ 1's}}} \Pr\{\mathcal{F} \text{ and } \mathcal{H}_r\} \\
 &= \frac{1}{m+m^{0.6}} \sum_{i=0}^{n-1} \Pr\{\mathcal{F} \text{ and } \mathcal{E}_i \text{ and } \mathcal{G}\} \\
 &= \frac{n(1-o(1))q}{m+m^{0.6}} = \frac{m(1-o(1))}{m+m^{0.6}} = 1-o(1).
 \end{aligned}$$

The probability that $\delta \geq \frac{1}{4}\ell_0 - 3 = \Omega(\ell_0)$ is

$$\Pr\{\mathcal{F}\} \geq \Pr\{\mathcal{F} \text{ and } \mathcal{G}\} = 1 - o(1).$$

The proof is complete. □

3.3.5. The Case $d = 1$ and $\gamma > 3$. We first show that with high probability, all long-range edges connect two vertices with grid distance $o(n)$, for general d and $\gamma > 2d$.

Lemma 3.12. *In a random graph from $\text{KSW}(n, d, \gamma)$ with $\gamma > 2d$, with probability $1 - o(1)$ there is no long-range edge that connects two vertices with grid distance larger than $n^{\frac{1}{\gamma-d} + \varepsilon}$, where ε is any positive number.*

Proof. For a vertex u , the probability that the long-range edge from u links to somewhere with distance longer than $\ell_0 = n^{\frac{1}{\gamma-d} + \varepsilon}$ from u is

$$O\left(\sum_{i=\ell_0}^{n^{1/d}} i^{d-1} \frac{i^{-\gamma}}{\sum_{j=1}^{n^{1/d}} j^{d-1} j^{-\gamma}}\right) = O\left(\sum_{i=\ell_0}^{\infty} i^{d-1-\gamma}\right) = O\left(\ell_0^{d-\gamma}\right).$$

By union bound, the probability that such u exists is

$$O\left(n\ell_0^{d-\gamma}\right) = O\left(n \cdot n^{(\frac{1}{\gamma-d} + \varepsilon)(d-\gamma)}\right) = O\left(n^{\varepsilon(d-\gamma)}\right) = o(1).$$

Therefore, with probability $1 - o(1)$, such u does not exist. □

Given a graph G in $\text{KSW}(n, d, \gamma)$, let $\ell_0(G)$ be the largest grid distance of two vertices connected by a long-range edge in G . From the above result, we know that when $\gamma > 3d$, $\ell_0(G) < n^{\frac{1}{\gamma-d} + \varepsilon} = o(\sqrt{n^{1/d}})$ with high probability. For the rest of this section, with $d = 1$, we fix G to be any graph in $\text{KSW}(n, 1, \gamma)$ with $\ell_0(G) < n^{\frac{1}{\gamma-1} + \varepsilon}$, and show that $\delta(G) = \Omega(n^c)$ for some constant c . Since G is fixed, we will use ℓ_0 to denote $\ell_0(G)$.

We now return to the one-dimensional grid with wraparound, which is a ring with vertices v_0, v_1, \dots, v_{n-1} . Note that in the one-dimensional case, the edge vector defined in Section 3.3.1 degenerates to a scalar value from $\{-\lfloor \frac{n}{2} \rfloor, -\lfloor \frac{n}{2} \rfloor + 1, \dots, \lfloor \frac{n-1}{2} \rfloor\}$. We arrange v_0, v_1, \dots, v_{n-1} clockwise on the ring. Then a positive edge scalar corresponds to a clockwise hop, while a negative edge scalar corresponds to a counterclockwise hop.

Let $A = v_0$ and $B = v_{\lfloor n/2 \rfloor}$ be two specific vertices. We define two kinds of paths between A and B : a *positive path* is one in which the summation of edge scalars (not taken modulo n) is positive, while a *negative path* is one in which the summation of edge scalars (not taken modulo n) is negative.

Lemma 3.13. *There exists a positive path from A to B that does not go through $v_{\lfloor n/2 \rfloor + 1}, v_{\lfloor n/2 \rfloor + 2}, \dots, v_{n-1}$, and the length is at most $2\ell_0$ longer than the shortest positive path from A to B . Similarly, there exists a negative path from A to B that does not go through $v_1, v_2, \dots, v_{\lfloor n/2 \rfloor - 1}$, and the length is at most $2\ell_0$ longer than the shortest negative path from A to B .*

Proof. We give only the proof for the positive-path case. Consider any shortest positive path \mathcal{P} from A to B . We first establish the following claim. Let $S_A = \{v_1, \dots, v_{\ell_0}\}$ be the set of ℓ_0 consecutive vertices clockwise to A , and $S_B = \{v_{\lfloor n/2 \rfloor - \ell_0}, v_{\lfloor n/2 \rfloor - \ell_0 + 1}, \dots, v_{\lfloor n/2 \rfloor - 1}\}$ the set of ℓ_0 consecutive vertices counterclockwise to B .

Claim 3.14. *There must exist a subpath in \mathcal{P} from a vertex $u \in S_A$ to a vertex $w \in S_B$ that does not go through $v_{\lfloor n/2 \rfloor + 1}, v_{\lfloor n/2 \rfloor + 2}, \dots, v_{n-1}$.*

Suppose there are m edges in \mathcal{P} . Let s_i ($0 \leq i \leq m$) denote the sum of the first i edge scalars in \mathcal{P} (not taken modulo n). Initially, we have $s_0 = 0$, and the final value s_m is $kn + \lfloor n/2 \rfloor$ for some integer $k \geq 0$. One can also see that the position after going through the first i edges in \mathcal{P} is at $v_{(s_i \bmod n)}$. Let i_1 be the smallest integer such that $s_{i_1} \geq \lfloor n/2 \rfloor$ (such exists, because $s_m \geq \lfloor n/2 \rfloor$), and let i_2 be the largest integer such that $i_2 < i_1$ and $s_{i_2} \leq 0$ (such exists, because $s_0 = 0$).

We consider the $(i_2 + 1)$ th edge, which begins at $v_{(s_{i_2} \bmod n)}$ for some $s_{i_2} \leq 0$ and ends at $v_{(s_{i_2+1} \bmod n)}$ for some $s_{i_2+1} > 0$. The number s_{i_2+1} is at most

$s_{i_2} + \ell_0$, since no edge is longer than ℓ_0 . Therefore, s_{i_2+1} must be a number in $(0, \ell_0]$, and $v_{(s_{i_2+1} \bmod n)}$ must be in S_A . We choose $u = v_{(s_{i_2+1} \bmod n)} \in S_A$. Similarly, pick $w = s_{i_1-1}$, which is the beginning point of the i_1 th edge. We have $w \in S_B$. The intermediate values $s_{i_2+1}, s_{i_2+2}, \dots, s_{i_1-1}$ are all in the interval $(0, \lfloor n/2 \rfloor]$ by the definitions of i_1 and i_2 . That is, for all j such that $i_2 + 1 \leq j \leq i_1 - 1$, we have $s_j \bmod n = s_j$, and the corresponding vertex $v_{(s_j \bmod n)}$ is $v_{s_j} \in \{v_1, v_2, \dots, v_{\lfloor n/2 \rfloor - 1}\}$. Therefore, the subpath from u to w does not go through $v_{\lfloor n/2 \rfloor + 1}, v_{\lfloor n/2 \rfloor + 2}, \dots, v_{n-1}$, and the claim holds.

With the claim, we can construct a positive path \mathcal{P}' that uses ring edges from A to u , then uses the subpath in the claim from u to w , and then from w to B using ring edges. The length of \mathcal{P}' is at most $2\ell_0$ longer than \mathcal{P} , the shortest positive path from A to B . \square

We use \mathcal{P}_{AB}^+ and \mathcal{P}_{AB}^- to denote the two paths stated in the above lemma. According to this lemma, one of \mathcal{P}_{AB}^+ and \mathcal{P}_{AB}^- is at most $2\ell_0$ longer than the shortest path between A and B .

Let C be the midpoint of \mathcal{P}_{AB}^+ (take a vertex nearest the middle if the path has an odd number of edges), and let $\mathcal{P}_{AC}^+, \mathcal{P}_{CB}^+$ be the two subpaths from A to C and C to B . Similarly, let D be the midpoint of \mathcal{P}_{AB}^- , and let $\mathcal{P}_{AD}^-, \mathcal{P}_{DB}^-$ be the two subpaths.

Lemma 3.15. *The paths $\mathcal{P}_{AC}^+, \mathcal{P}_{CB}^+, \mathcal{P}_{AD}^-$, and \mathcal{P}_{DB}^- are at most $3\ell_0 + 1$ longer than the shortest paths between corresponding pairs of vertices.*

Proof. We give the proof only for \mathcal{P}_{AC}^+ . Suppose that it is not true, and the shortest path \mathcal{P}_{AC}^* from A to C is at least $3\ell_0 + 2$ shorter than \mathcal{P}_{AC}^+ . The path \mathcal{P}_{AC}^* must be at least $3\ell_0 + 1$ shorter than \mathcal{P}_{CB}^+ , because C is the point nearest the middle of \mathcal{P}_{AB}^+ .

Consider the last time that the path \mathcal{P}_{AC}^* gets into the range $\{v_0, v_1, \dots, v_{\lfloor n/2 \rfloor}\}$. The subpath of \mathcal{P}_{AC}^* from that point to C must be one of the following cases:

- It is a subpath from a vertex $A' \in \{v_0, v_1, \dots, v_{\ell_0}\}$ to C not going through $v_{\lfloor n/2 \rfloor + 1}, v_{\lfloor n/2 \rfloor + 2}, \dots, v_{n-1}$. Replace \mathcal{P}_{AC}^+ by the path from A to A' through ring edges concatenated with the subpath of \mathcal{P}_{AC}^* from A' to C . This will cause the length of \mathcal{P}_{AB}^+ to decrease by at least $3\ell_0 + 2 - \ell_0 > 2\ell_0$, which is impossible by Lemma 3.13.
- It is a subpath from a vertex $B' \in \{v_{\lfloor n/2 \rfloor - \ell_0}, v_{\lfloor n/2 \rfloor - \ell_0 + 1}, \dots, v_{\lfloor n/2 \rfloor}\}$ to C that does not go through $v_{\lfloor n/2 \rfloor + 1}, v_{\lfloor n/2 \rfloor + 2}, \dots, v_{n-1}$. Replace \mathcal{P}_{CB}^+ by the reverse of this subpath of \mathcal{P}_{AC}^* from C to B' concatenated with ring edges

from B' to B . This will cause the length of \mathcal{P}_{AB}^+ to decrease by at least $3\ell_0 + 1 - \ell_0 > 2\ell_0$, which is impossible by Lemma 3.13.

Therefore the lemma holds. □

Next we consider the shortest path between C and D .

Lemma 3.16. *Either the concatenation of \mathcal{P}_{CB}^+ and reversed \mathcal{P}_{DB}^- , or the concatenation of reversed \mathcal{P}_{AC}^+ and \mathcal{P}_{DA}^- , is at most $8\ell_0 + 2$ longer than the shortest path between C and D .*

Proof. The shortest path from C to D (say \mathcal{P}_{CD}^*) must go through either B 's neighborhood $v_{\lfloor n/2 \rfloor}, v_{\lfloor n/2 \rfloor + 1}, \dots, v_{\lfloor n/2 \rfloor + \ell_0}$ or A 's neighborhood $v_0, v_1, \dots, v_{\ell_0}$. Without loss of generality, we assume that it goes through the point B' in B 's neighborhood. Use $\mathcal{P}_{CB'}^*$ and $\mathcal{P}_{B'D}^*$ to denote the two subpaths from C to B' and B' to D respectively. They must also be shortest paths of CB' and $B'D$.

The path $\mathcal{P}_{CB'}^*$ is at most ℓ_0 shorter than the shortest path between C and B . Otherwise, the path $\mathcal{P}_{CB'}^*$ concatenated with ring edges from B' to B would be shorter than the shortest path. Similarly, $\mathcal{P}_{B'D}^*$ is at most ℓ_0 shorter than the shortest path between B and D . Therefore, the shortest path between C and D is at most $2\ell_0$ shorter than the concatenation of shortest paths of CB and BD . Then by Lemma 3.15, the summation of \mathcal{P}_{CB}^+ and \mathcal{P}_{DB}^- is at most $2(3\ell_0 + 1) + 2\ell_0 = 8\ell_0 + 2$ longer than the shortest path between C and D . □

We have the following corollary, since the lengths of the two paths in this lemma differ by at most 2.

Corollary 3.17. *The concatenation of \mathcal{P}_{CB}^+ and reversed \mathcal{P}_{DB}^- , and the concatenation of reversed \mathcal{P}_{AC}^+ and \mathcal{P}_{DA}^- , are both at most $8\ell_0 + 4$ longer than the shortest path between C and D .*

Now we can prove the lower bound of δ for this case.

Proof of Theorem 3.1 for the case $d = 1$ and $\gamma > 3$. Consider the four points $A, B, C,$ and D defined above. Let $d(x, y)$ denote the distance between vertices x and y . We can see the following consequences about pairwise distances:

- (a) $d(A, B) \geq \lfloor n/2 \rfloor / \ell_0$;
- (b) $d(C, D) \geq |\mathcal{P}_{CB}^+| + |\mathcal{P}_{DB}^-| - (8\ell_0 + 4) \geq |\mathcal{P}_{AC}^+| - 1 + |\mathcal{P}_{DB}^-| - (8\ell_0 + 4) \geq d(A, C) + d(D, B) - (8\ell_0 + 5)$, where the first inequality is due to Corollary 3.17;
- (c) similarly, $d(C, D) \geq d(A, D) + d(C, B) - (8\ell_0 + 5)$.

Therefore, we have both

$$d(A, B) + d(C, D) \geq d(A, C) + d(D, B) + \lfloor n/2 \rfloor / \ell_0 - (8\ell_0 + 5)$$

and

$$d(A, B) + d(C, D) \geq d(A, D) + d(C, B) + \lfloor n/2 \rfloor / \ell_0 - (8\ell_0 + 5).$$

For $\ell_0 < n^{\frac{1}{\gamma-1} + \varepsilon}$ with any sufficiently small $\varepsilon > 0$ and sufficiently large n , we have

$$\frac{\lfloor n/2 \rfloor}{\ell_0} \gg 8\ell_0 + 5,$$

and thus $d(A, B) + d(C, D)$ is the largest distance pair. In this case,

$$\delta \geq \frac{\lfloor n/2 \rfloor}{\ell_0} - (8\ell_0 + 5).$$

By Lemma 3.12, with probability $1 - o(1)$ we have $\ell_0 < n^{\frac{1}{\gamma-1} + \varepsilon}$. Therefore, with probability $1 - o(1)$,

$$\delta(\text{KSW}(n, 1, \gamma)) = \Omega\left(n/n^{\frac{1}{\gamma-1} + \varepsilon}\right) = \Omega\left(n^{\frac{\gamma-2}{\gamma-1} - \varepsilon}\right)$$

for $d = 1$, $\gamma > 3$, and any sufficiently small $\varepsilon > 0$. Since for every $\varepsilon' > \varepsilon > 0$, we have

$$n^{\frac{\gamma-2}{\gamma-1} - \varepsilon} = \Omega\left(n^{\frac{\gamma-2}{\gamma-1} - \varepsilon'}\right),$$

it follows that

$$\delta(\text{KSW}(n, 1, \gamma)) = \Omega\left(n^{\frac{\gamma-2}{\gamma-1} - \varepsilon}\right)$$

for every $\varepsilon > 0$. □

3.4. Extensions of the KSW Model

Our analysis also holds for some variants of the KSW model. In this section, we study one variant of the underlying structure, grid without wraparound, and two variants of edge linking: multiple edges for each vertex and linking edges independently.

3.4.1. Grid without Wraparound. We modify our analysis so that the first two results of Theorem 3.1 still hold. For the case of $d \geq 2$ and $0 \leq \gamma \leq d$ (Section 3.3.1), the changes are as follows. For a path from u to v , we divide it into segments as before. Elements in an edge vector are now in $\{-n', -n' + 1, \dots, n' - 1, n'\}$. Recall the last condition by which we define that two paths from u to v belong to the same category: the summations (not modulo n') of all segment vectors in the two paths are equal. This is always satisfied for grid without wraparound,

because the summation of all segment vectors depends only on the positions of u and v .

In the proofs of Lemmas 3.4 and 3.5, the summation of all segment vectors is fixed, rather than there being $(2\ell + 1)^d$ or $(2k + 1)^d$ choices respectively. Hence the upper bounds given in Lemmas 3.4 and 3.5 still hold. In Lemma 3.6, an edge between u and v can be from u to v or from v to u . The probabilities of the two cases may differ by a constant factor in the case of grid without wraparound. Hence the probability of there existing an edge between u and v is changed by at most a constant factor, and Lemma 3.6 still holds. One can verify that the rest of the analysis in Section 3.3.1 still holds. For the case $d = 1$ and $0 \leq \gamma \leq 1$ (Section 3.3.4), the only change is that the event \mathcal{E}_i , which is the event that v_i links to $v_{i+\ell_0}$, applies only when $i = 0, 1, \dots, n - \ell_0 - 1$. Since $\ell_0 = O(\log n)$ is much smaller than n , there are still almost n events \mathcal{E}_i , and the argument requires no significant change. Hence Theorem 3.1 still holds for the cases $d \geq 1$ and $0 \leq \gamma \leq d$.

3.4.2. Multiple Edges for Each Vertex. In this model, each vertex links a constant, say d_0 , number of edges according to the same distribution with which u links to v with probability $d_B(u, v)^{-\gamma} / \sum_{v'} d_B(u, v')^{-\gamma}$. We show that all our previous analysis still holds with some slight changes. For the case $d \geq 2$ and $1 \leq \gamma \leq d$ (Section 3.3.1), Lemma 3.6 still holds, since the probability of the edge (u, v) is increased by at most d_0 times using union bound. And the proof of Lemma 3.7 still works, because $O(i^{d-1}i^{-\gamma}/f(n'))$ is an upper bound on the probability that u links to some vertex at distance i , again by union bound. For the case $d = 1$ and $0 \leq \gamma \leq 1$ (Section 3.3.4), we define \mathcal{E}_i to be the event that at least one of v_i 's edges links to $v_{(i+\ell_0 \bmod n)}$. One can see that $\Pr\{\mathcal{E}_i\}$ is still $\Theta(\ell_0^{-\gamma}/f(n))$, and the remainder of the argument also holds. For the case $d = 1$ and $\gamma > 3$ (Section 3.3.5), Lemma 3.12 still holds for the same reason as Lemma 3.7, namely that

$$O\left(i^{d-1} \frac{i^{-\gamma}}{\sum_{j=1}^{n^{1/d}} j^{d-1} j^{-\gamma}}\right)$$

is an upper bound on the probability that u links to some vertex at distance i . One can verify that all the results of Theorem 3.1 still hold under this change.

3.4.3. Linking Edges Independently. In this model, all edges exist independently. The edge between (u, v) exists with probability

$$\frac{d_0 \cdot d_B(u, v)^{-\gamma}}{\sum_{i=1}^{n^{1/d}} i^{d-1-\gamma}},$$

where d_0 is some constant. We also give the changes in our analysis. Lemma 3.6 is straightforward in this model. Lemmas 3.7 and 3.12 still hold for the same reason as above. In the proof of Lemma 3.8, the first line of (3.2), which uses the product of edges' probabilities for an upper bound of the path's probability, still holds, because it is now just the product of independent events. For the case $d = 1$ and $0 \leq \gamma \leq 1$ (Section 3.3.4), we define \mathcal{E}_i to be the event that the edge $(v_i, v_{(i+\ell_0 \bmod n)})$ exists. One can see that the analysis remains valid. All the results of Theorem 3.1 still hold under this change.

In summary, Theorem 3.1 for the case $d \geq 1$ and $0 \leq \gamma \leq d$ still holds for grid without wraparound, and Theorem 3.1 for all cases still holds for both variants of edge linking. The variants of edge linking can be combined with grid without wraparound, for which Theorem 3.1 still holds for the case $d \geq 1$ and $0 \leq \gamma \leq d$.

4. δ -Hyperbolicity of Ringed Trees

In this section, we consider the δ -hyperbolicity of graphs constructed according to a variant of the small-world graph model, in which long-range edges are added on top of a base graph that is a binary tree or treelike low- δ graph. In particular, we will analyze the effect on the δ -hyperbolicity of adding long-range links to a ringed-tree base graph; and then we will consider several related extensions, including an extension to the binary tree.

Definition 4.1. (Ringed tree.) A *ringed tree* of level k , denoted by $\text{RT}(k)$, is a fully binary tree with k levels (counting the root as a level), in which all vertices at the same level are connected by a ring. More precisely, we can use a binary string to represent each vertex in the tree such that the root (at level 0) is represented by an empty string, and the left and right children of a vertex with string σ are represented by $\sigma 0$ (concatenating 0 after string σ) and $\sigma 1$ (concatenating 1 after string σ), respectively. Then, at each level $i = 1, 2, \dots, k - 1$, we connect two vertices u and v represented by binary strings σ_u and σ_v if $(\sigma_u + 1) \bmod 2^i = \sigma_v$, where the addition treats the binary strings as the integers they represent. As a convention, we say that a level is higher if it has a smaller level number and thus is closer to the root.

Figure 1(d) illustrates the ringed tree $\text{RT}(6)$. Note that the diameter of the ringed tree $\text{RT}(k)$ is $\Theta(\log n)$, where $n = 2^k - 1$ is the number of vertices in $\text{RT}(k)$, and we will use $\text{RT}(\infty)$ to denote the infinite ringed tree as k in $\text{RT}(k)$ goes to infinity. Thus, a ringed tree may be thought of as a soft version of a binary tree, and to some extent, one can view a ringed tree as an idealized picture reflecting the hierarchical structure in real networks coupled with local

neighborhood connections, such as Internet autonomous system (AS) networks, which have both a hierarchical structure of different levels of autonomous systems and peer connections based on geographical proximity.

4.1. Results and Their Implications

A visual comparison of the ringed tree of Figure 1(d) with the tessellation of the Poincaré disk (Figure 1(b)) suggests that the ringed tree can be seen as an approximate tessellation or coarsening of the Poincaré disk. Our first result in this section makes this precise; in particular, we show that the infinite ringed tree and the Poincaré disk are quasi-isometric.

Theorem 4.2. *The infinite ringed tree $\text{RT}(\infty)$ and the Poincaré disk are quasi-isometric.*

Thus, by Proposition 2.6, we immediately have the following result.

Corollary 4.3. *There exists a constant c such that for all k , the ringed tree $\text{RT}(k)$ is c -hyperbolic.*

Alternatively, we also provide a direct proof of this corollary (Section 4.3.3) to show that the ringed tree $\text{RT}(k)$ is Rips 5-hyperbolic, and Gromov 40-hyperbolic in terms of the four-point condition. Our direct analysis also provides important properties of ringed trees that are used in later analyses.

Next, we address the question whether long-range edges added at each level of the ring maintain or destroy the hyperbolicity of the base graph. Given two vertices u and v at some level t of the ringed tree, we define the ring distance between u and v , denoted by $d_R(u, v)$, to be the length of the shorter path connecting u and v purely through the ring edges at the level t .

Given any function f from positive integers to positive integers, let $\text{RT}(k, f)$ denote the class of graphs constructed by adding long-range edges on the ringed tree $\text{RT}(k)$ such that for each long-range edge (u, v) connecting vertices u and v at the same level, $d_R(u, v) \leq f(n)$, where $n = 2^k - 1$ is the number of vertices in the ringed tree $\text{RT}(k)$. Since long-range edges do not reduce distances from the root to any other vertices, the diameter of every graph in $\text{RT}(k, f)$ is still $\Theta(\log n)$. Define $\delta(\text{RT}(k, f)) = \max_{G \in \text{RT}(k, f)} \delta(G)$.

Our second result (used in the proof of the first part of our next result, but explicitly stated here, since it is also of independent interest) is the following.

Theorem 4.4. *$\delta(\text{RT}(k, f)) = O(\log f(n))$ for every positive function f and positive integer k , where $n = 2^k - 1$ is the number of vertices in the ringed tree $\text{RT}(k)$.*

This result indicates that if the long-range edges added do not span distant vertices, then the graph should have good hyperbolicity. In particular, if we take $f(n) = \log n$, then the theorem implies that the class $\text{RT}(k, f)$ is logarithmically hyperbolic. The theorem covers all (deterministic) graphs in the class $\text{RT}(k, f)$. We can extend it to random graphs such that if we can show that with high probability, the random graph is in the class $\text{RT}(k, f)$, then we know that the hyperbolic δ of the random graph is $O(\log f(n))$ with high probability. The first result in the next theorem is proved using this approach.

Next, we consider adding random edges between two vertices at the outermost level, i.e., level $k - 1$, such that the probability of connecting two vertices u and v is determined by a function $g(u, v)$. Let V_{k-1} denote the set of vertices at level $k - 1$, i.e., the leaves of the original binary tree. Given a real-valued positive function $g(u, v)$, let $\text{RRT}(k, g)$ denote a random graph constructed as follows. We start with the ringed tree $\text{RT}(k)$, and then for each vertex $v \in V_{k-1}$, we add one long-range edge to a vertex u with probability proportional to $g(u, v)$, that is, with probability $g(u, v)\rho_v^{-1}$, where $\rho_v = \sum_{u \in V_{k-1}} g(u, v)$.

We study three families of functions g , each of which has the characteristic that vertices closer to one another (by some measure) are more likely to be connected by a long-range edge. The first two families use the ring distance $d_R(u, v)$ as the closeness measure. In particular, the first family uses an exponential decay function $g_1(u, v) = e^{-\alpha d_R(u, v)}$. The second family uses a power-law decay function $g_2(u, v) = d_R(u, v)^{-\alpha}$, where $\alpha > 0$. The third family uses the height of the lowest common ancestor of u and v , denoted by $h(u, v)$, as the closeness measure, and the function is $g_3 = 2^{-\alpha h(u, v)}$.

Note that this last probability function matches the function used by Kleinberg in a small-world model based on the tree structure [Kleinberg 02]. Moreover, although g_3 and g_2 are similar, in the ringed tree they are not the same, since for two leaf nodes u and v , $d_R(u, v)$ may not be the same as $2^{h(u, v)}$. For example, let u be the rightmost leaf of the left subtree of the root (i.e., u is represented as the string $01 \dots 1$), and let v be the leftmost leaf of the right subtree of the root (i.e., v is represented as the string $10 \dots 0$). Then $d_R(u, v) = 1$, while $h(u, v) = \Theta(\log n)$.

The following theorem summarizes the hyperbolicity behavior of these three families of random ringed trees.

Theorem 4.5. *Consider the following families of functions (of variables u and v) for random ringed trees $\text{RRT}(k, g)$, for a positive integer k and positive real number α . With probability $1 - o(1)$ (as n tends to infinity), we have*

$$(1) \quad \delta(\text{RRT}(k, e^{-\alpha d_R(u, v)})) = O(\log \log n),$$

$$(2) \delta(\text{RRT}(k, d_R(u, v)^{-\alpha})) = \Theta(\log n),$$

$$(3) \delta(\text{RRT}(k, 2^{-\alpha h(u, v)})) = \Theta(\log n),$$

where $n = 2^k - 1$ is the number of vertices in the ringed tree $\text{RT}(k)$.

This theorem states that when the random long-range edges are selected using an exponential decay function based on the ring distance measure, the resulting graph is logarithmically hyperbolic, i.e., the constant hyperbolicity of the original base graph is degraded only slightly; but when a power-law decay function based on the ring distance measure or an exponential decay function based on common ancestor measure is used, then hyperbolicity is destroyed, and the resulting graph is not hyperbolic. One may notice that the function forms in (1) and (3) above are similar, but the results are different. This is because with height $h(u, v)$, the subtree actually covers $\Theta(2^{h(u, v)})$ leaves, and thus (3) is naturally closer to the power-law function of (2). Intuitively, when it is more likely for a long-range edge to connect two distant vertices, such an edge creates a shortcut for many internal tree nodes, so that many shortest paths will go through this shortcut instead of traversing tree nodes. (In Internet routing, such paths are referred to as *valley routes*).

Finally, as a comparison, we also study the hyperbolicity of random binary trees $\text{RBT}(k, g)$, which are the same as random ringed trees $\text{RRT}(k, g)$ except that we remove all ring edges.

Theorem 4.6. *Consider the following families of functions (of the variables u and v) for random binary trees $\text{RBT}(k, g)$, for a positive integer k and positive real number α . Then with probability $1 - o(1)$ (as n tends to infinity), we have*

$$\begin{aligned} \delta\left(\text{RBT}\left(k, e^{-\alpha d_R(u, v)}\right)\right) &= \delta\left(\text{RBT}\left(k, d_R(u, v)^{-\alpha}\right)\right) \\ &= \delta\left(\text{RBT}\left(k, 2^{-\alpha h(u, v)}\right)\right) \\ &= \Theta(\log n), \end{aligned}$$

where $n = 2^k - 1$ is the number of vertices in the binary tree $\text{RBT}(k, g)$.

Thus in this case, the original hyperbolicity of the base graph ($\delta = 0$ for the binary tree) is destroyed. Comparing with Theorem 4.5, we see that our results above suggest that the “softening” of the hyperbolicity provided by the rings is essential in maintaining good hyperbolicity: with rings, random ringed trees with exponential decay function (depending on the ringed distance) are logarithmically hyperbolic, but without the rings, the resulting graphs are not hyperbolic.

4.2. Outline of the Analysis

In this subsection, we provide a summary of the proofs of the four theorems in Section 4.1. For Theorem 4.2, we provide an embedding of the ringed tree in the Poincaré disk, intuitively similar to the picture we show in Figure 1(d), and prove that it is a quasi-isometry.

For the analysis of δ -hyperbolicity, we apply the Rips condition, which is equivalent to Gromov's four-point condition up to a constant factor.

For any two vertices u and v on the ringed tree $\text{RT}(k)$, we define the canonical geodesic $\langle u, v \rangle$ to be the geodesic from u to v such that the geodesic always first goes upward, then follows ring edges, and then goes downward (any of these segments may be omitted). We show that the canonical geodesic $\langle u, v \rangle$ and any other geodesic $[u, v]$ are within distance 1 of each other, and any triangle $\Delta(u, v, w)$ formed by three canonical geodesics $\langle u, v \rangle$, $\langle u, w \rangle$, and $\langle v, w \rangle$ (called *canonical triangles*) are 3-slim. This immediately implies that every geodesic triangle in $\text{RT}(k)$ is 5-slim, which is a direct proof that ringed trees are constantly hyperbolic.

For Theorem 4.4, we inductively prove that every geodesic $[u, v]$ in $\text{RT}(k, f)$ is within $O(\log f(n))$ distance of the canonical geodesic $\langle u, v \rangle$, and conversely. Together with the result that every canonical triangle is 3-slim, it follows that every geodesic triangle is $O(\log f(n))$ -slim.

For Theorem 4.5, part (1), we show that with high probability, the long-range edges connect only vertices within ring distance $O(\log n)$, and then we can apply Theorem 4.4 to achieve the $O(\log \log n)$ bound.

For Theorem 4.5, part (2), the key is to show that (a) with high probability, some long-range link connects two vertices at ring distance $\Theta(n^c)$ for some constant c , and (b) if such a long-range edge (u, v) exists, then we consider the geodesic triangle $\Delta(u, v, r)$, where r is a point with lowest layer number on the canonical geodesic between u and v , and show that the midpoint of $[r, u]$ is $\Theta(\log n)$ away from the union of $[r, v]$ and $[v, u]$.

For Theorem 4.5, part (3), we first show that with high probability, there is some pair of vertices u, v with $h(u, v) \geq c \log_2 n$ for some constant $c > 0$, and then we observe that in such a configuration, the ring distance $d_R(u, v)$ has high probability of being $\Omega(n^{c/2})$, and the argument follows exactly the same analysis in the previous part.

For Theorem 4.6, parts (2) and (3) follow a similar strategy as that of Theorem 4.5. For part (1), we know that two prospective ring neighbors u and v have constant probability of having a long-range connection. However, since we do not have ring edges, the alternative path between u and v through the tree may be $\Theta(\log n)$ in length. We show that there are at least $\Omega(\sqrt{n})$ such pairs, so

that with high probability, at least one pair is connected, generating a bad δ of $\Omega(\log n)$.

4.3. Detailed Analysis on Ringed Trees

4.3.1. Properties of Ringed Trees. We begin with some properties of ringed trees that will be used repeatedly in the following analysis on ringed-tree-related graphs and which may be of independent interest.

We define the *ring distance* $d_R(u, v)$ of u and v on the same level to be their distance on the ring. Ringed trees have the following fundamental property.

Lemma 4.7. *Let u and v be two vertices on the same level, and let u' and v' be their respective parents. Then $d_R(u', v') \leq (d_R(u, v) + 1)/2$.*

Proof. On the ring, there are $d_R(u, v) + 1$ vertices on the segment between u and v , which belong to at most $(d_R(u, v) + 1)/2 + 1$ parents, which correspond to at most $(d_R(u, v) + 1)/2 + 1$ vertices on the ring segment between u' and v' . This concludes the proof. \square

For a geodesic $[u, v]$ on the ringed tree $\text{RT}(k)$, we call the sequence of levels through which it passes from u to v its *level sequence*. Lemma 4.7 implies that the level sequence of any geodesic must be reversed unimodal: it first decreases, and then increases (but the increasing or decreasing segment may be omitted). The following lemma further characterizes geodesics in ringed trees.

Lemma 4.8. *Let u, v be two vertices, u on level ℓ , and u' the parent of u at level $\ell - 1$. Suppose $[u, v]$ intersects level $\ell - 1$, and let t be the intersection closest to u . Then $d(u, t) \leq 2$, and the segment $[u, t]$ of $[u, v]$ and $\{u, u'\}$ are within distance 1 of each other.*

Proof. Let t' be the node just before t to u on $[t, u]$. Then

$$d(t, u) \leq 1 + d_R(t, u') \leq 1 + \frac{d_R(t', u) + 1}{2}$$

by Lemma 4.7. Since t is the closest node on level $\ell - 1$ on the geodesic to u , it follows that $d(t, u) \geq 1 + d_R(t', u)$. We get $d(t', u) \leq 1$ and $d(t, u) \leq 2$ by combining these two inequalities. The segment $[u, t]$ of $[u, v]$ and $\{u, u'\}$ are within distance 1 of each other because $d(t, u') \leq 1$. \square

For two vertices u and v , we define the *canonical geodesic* $\langle u, v \rangle$ in a recursive fashion:

1. For u, v on the same level and $d_R(u, v) \leq 3$, $\langle u, v \rangle$ is the path on the ring from u to v .
2. For u, v on the same level but $d_R(u, v) > 3$, let u', v' be the respective parents of u, v . Then $\langle u, v \rangle = [u, u'] \cup \langle u', v' \rangle \cup [v', v]$.
3. For u, v on different levels, supposing u on upper level, let v' be the parent of v . Then $\langle u, v \rangle = \langle u, v' \rangle \cup [v', v]$.

This is a well-formed definition. At each level of recursion, either the difference of levels of nodes decreases, or in the case of nodes on the same level, the ring distance decreases by Lemma 4.7, until we reach the base case, where we have $d_R(u, v) \leq 3$.

We prove now that canonical geodesics are really geodesics.

Lemma 4.9. *For all u, v , $\langle u, v \rangle$ is a geodesic between u and v .*

Proof. For u, v on the same level and $d_R(u, v) \leq 3$, we can check that $\langle u, v \rangle$ is a geodesic between u and v .

For u, v on the same level but $d_R(u, v) > 3$, let u', v' be the respective parents of u, v . Let t, s be the closest node to u, v in an upper level on $[u, v]$. If $t \neq u'$, then $d(u, t) \geq 2$, because the only way to go up one level in one step is to go to the parent. By Lemma 4.8, $[u, t]$ and $\{u, u'\}$ are within distance 1 of each other. Therefore, $d(u', t) = 1$, since $d(u, t) \geq 2$, and $\{u, u'\} \cup [t, v]$ is also a geodesic from u to v . This is also correct for $t = u'$. The same can be proved for s . By combining, we have that $\{u, u'\} \cup [t, s] \cup \{v', v\}$ is also a geodesic between u and v for any geodesics $[t, s]$. If we pick the canonical geodesic $\langle t, s \rangle$, this will always be the canonical geodesic $\langle u, v \rangle$. Therefore, $\langle u, v \rangle$ is a geodesic.

For u, v on different levels, the induction is essentially the same as in the previous case, but instead of reasoning on both sides of u and v , we only need to reason on one side of u or v , depending on which side is at the lower level.

This concludes our induction. \square

4.3.2. Proof of Theorem 4.2: Quasi-Isometry from Infinite Ringed Tree to the Poincaré Disk. In this subsection, we will exhibit and prove a quasi-isometry from the infinite ringed tree $RT(\infty)$ to the Poincaré disk. We denote its distance by d_{RT} . We denote by (D, d_P) the Poincaré disk, where D is the open disk of radius 1 in the complex plane.

Here is a brief summary of our approach. First, we propose a candidate quasi-isometry, and then all possible cases of images of two points of the ringed tree are divided into four categories, each of which is separately analyzed. We then

proceed with an analysis of the metric of the ringed tree and show that the candidate quasi-isometry is effective, and thus the ringed tree and the Poincaré disk are quasi-isometric.

The following inequalities will be used:

$$\begin{aligned} \ln(x) &\leq \cosh^{-1}(x) \leq \ln(2x) \quad \text{for } x \geq 1, \\ \frac{2}{\pi}x &\leq \sin(x) \leq x \quad \text{for } 0 \leq x \leq \frac{\pi}{2}, \\ 1 - \frac{x}{2} &\leq \sqrt{1-x} \leq 1 - \frac{3x}{5} \quad \text{for } 0 \leq x \leq \frac{1}{2}. \end{aligned}$$

To state the promised quasi-isometry, we give coordinates to the nodes in $\text{RT}(\infty)$. We know that we can number the nodes with binary strings, each of which can be regarded as a number. For a node on the k th level and numbered by m , its coordinates are (k, m) , with $0 \leq m \leq 2^k - 1$. The root is level 0.

Definition 4.10. Let the following mapping be the candidate for the quasi-isometry:

$$f : \text{RT}(\infty) \rightarrow D, (k, m) \mapsto \sqrt{1 - 2^{-k}} e^{2i\pi m/2^k}.$$

For $0 \leq k \leq \ell$, $m < 2^{\ell-1}$, we define $D(k, \ell, m) = d_P(f(k, 0), f(\ell, m))$. This is a distance in the Poincaré disk. Following is its full expression in k and m :

$$\begin{aligned} D(k, \ell, m) &= \cosh^{-1} \left(1 + 2 \frac{\|\sqrt{1 - 2^{-k}} - \sqrt{1 - 2^{-\ell}} e^{2i\pi m/2^\ell}\|^2}{2^{-k-\ell}} \right) \\ &= \cosh^{-1} \left(1 + 2 \left(\sqrt{2^\ell (2^k - 1)} - \sqrt{2^k (2^\ell - 1)} \right)^2 \right. \\ &\quad \left. + 8\sqrt{2^\ell (2^\ell - 1) 2^k (2^k - 1)} \sin^2 \left(\pi \frac{m}{2^\ell} \right) \right). \end{aligned}$$

We also define

$$D'(k, \ell, m) = d_{\text{RT}}(f(k, 0), f(\ell, m)),$$

with $0 \leq k \leq \ell$, $m < 2^{\ell-1}$. This is a distance in the ringed tree.

We will now try to bound $D(k, \ell, m)$ with the following lemma.

Lemma 4.11. *We have the following bounds on $D(k, \ell, m)$.*

(1) For $k = 0$,

$$\frac{\ln(2)}{2} \ell + \frac{\ln(2)}{2} \leq D(0, \ell, m) \leq \frac{\ln(2)}{2} \ell + \ln(6).$$

(2) For $0 < k = \ell$, $m > 0$,

$$\ln(2)(4 + 2\lfloor \log_2 m \rfloor) \leq D(k, k, m) \leq \ln(2)(4 + 2\lfloor \log_2 m \rfloor) + \ln\left(\frac{5}{4\pi^2}\right)$$

for $k \geq 1$.

(3) For $m = 0$, $1 \leq k < \ell$,

$$\ln(2)(\ell - k) - \ln(50) \leq D(k, \ell, 0) \leq \ln(2)(\ell - k).$$

(4) For $0 < k < \ell$ and $0 < m < 2^{\ell-k}$,

$$\ln(2)(\ell - k) - \ln(100) \leq D(k, \ell, m) \leq \ln(2)(\ell - k) + \ln(66\pi^2).$$

(5) For $0 < k < \ell$ and $2^{\ell-k} \leq m < 2^{\ell-1}$,

$$\begin{aligned} \ln(2)(k - \ell + 2\lfloor \log_2 m \rfloor + 4) \\ \leq D(k, \ell, m) \leq \ln(2)(k - \ell + 2\lfloor \log_2 m \rfloor + 6) + \ln(\pi^2 + 1). \end{aligned}$$

Proof. 1. Case $k = 0$:

$$D(0, \ell, m) = \cosh^{-1}\left(1 + 2\sqrt{2^\ell - 1}\right).$$

If $\ell = 0$, then $D(0, 0, m) = 0$. For $\ell \geq 1$, we have

$$D(0, \ell, m) \geq \cosh^{-1}\left(2\sqrt{2^{\ell-1}}\right) \geq \frac{\ln(2)}{2}\ell + \frac{\ln(2)}{2}$$

and

$$D(0, \ell, m) \leq \cosh^{-1}\left(3\sqrt{2^\ell}\right) \leq \frac{\ln(2)}{2}\ell + \ln(6).$$

2. Case $0 < k = \ell$, $m > 0$:

$$D(k, k, m) = \cosh^{-1}\left(1 + 2^{k+3}(2^k - 1)\sin^2\left(\frac{m}{2^k}\pi\right)\right).$$

Let $a = \lfloor \log_2 m \rfloor$. We have

$$\begin{aligned} D(k, k, m) &\geq \cosh^{-1}\left(2^{2k+2}\sin^2(2^{a-k}\pi)\right) \geq \cosh^{-1}\left(2^{2k+2}(2^{a+1-k})^2\right) \\ &\geq \ln(2)(4 + 2a) \end{aligned}$$

and

$$\begin{aligned} D(k, k, m) &\leq \cosh^{-1}\left(\frac{5}{4}2^{2k+2}\sin^2(2^{a+1-k}\pi)\right) \\ &\leq \cosh^{-1}\left(\frac{5}{4}2^{2k+2}(2^{a+1-k}\pi)^2\right) \leq \ln(2)(4 + 2a) + \ln\left(\frac{5}{4\pi^2}\right). \end{aligned}$$

3. Case $0 < k < \ell, m = 0$:

$$\begin{aligned} D(k, \ell, 0) &= \cosh^{-1} \left(1 + 2 \left(\sqrt{2^\ell (2^k - 1)} - \sqrt{2^k (2^\ell - 1)} \right)^2 \right) \\ &= \cosh^{-1} \left(1 + 2^{k+\ell+1} \left(\sqrt{1 - 2^{-k}} - \sqrt{1 - 2^{-\ell}} \right)^2 \right). \end{aligned}$$

Since $1 \leq k < \ell$, it follows that $\sqrt{1 - 2^{-\ell}} - \sqrt{1 - 2^{-k}} > 0$, and we have

$$\sqrt{1 - 2^{-\ell}} - \sqrt{1 - 2^{-k}} \geq 1 - 2^{-\ell-1} - 1 + \frac{3}{5} 2^{-k} \geq \frac{1}{5} 2^{-k-1}$$

and

$$\sqrt{1 - 2^{-\ell}} - \sqrt{1 - 2^{-k}} \leq 1 - \frac{3}{5} 2^{-\ell} - 1 + 2^{-k-1} \leq 2^{-k-1}.$$

Therefore, we have

$$\ln(2)(\ell - k) - \ln(50) \leq D(k, \ell, 0) \leq \ln(2)(\ell - k).$$

4. Case $0 < k < \ell, 0 < m < 2^{\ell-1}$. We are now dealing with the general case of $m > 0$ and $0 < k < \ell$:

$$D(k, \ell, m) = \cosh^{-1} \left(\cosh(D(k, \ell, 0)) + 8\sqrt{2^\ell (2^\ell - 1) 2^k (2^k - 1)} \sin^2 \left(\pi \frac{m}{2^\ell} \right) \right).$$

We note that $a = \lfloor \log_2 m \rfloor$, and we have

$$\sqrt{2^\ell (2^\ell - 1) 2^k (2^k - 1)} \sin^2 \left(\pi \frac{m}{2^\ell} \right) \geq 2^{\ell+k-1} 2^{2a-2\ell+2} = 2^{k-\ell+2a+1}$$

and

$$\sqrt{2^\ell (2^\ell - 1) 2^k (2^k - 1)} \sin^2 \left(\pi \frac{m}{2^\ell} \right) \leq 2^{\ell+k} 2^{2a+2-2\ell} \pi^2 \leq 2^{k-\ell+2a+2} \pi^2.$$

The previous bound on $D(k, \ell, 0)$ transforms into the following by applying cosh:

$$\frac{2^{\ell-k}}{100} \leq \cosh(D(k, \ell, 0)) \leq 2^{\ell-k}.$$

A suitable substitution of $\cosh(D(k, \ell, 0))$ gives

$$\cosh^{-1} \left(\frac{2^{\ell-k}}{100} + 2^{k-\ell+2a+4} \right) \leq D(k, \ell, m) \leq \cosh^{-1} \left(2^{\ell-k} + 2^{k-\ell+2a+5} \pi^2 \right).$$

For $0 \leq a \leq \ell - k$, therefore $0 < m < 2^{\ell-k}$,

$$D(k, \ell, m) \geq \cosh^{-1} \left(\frac{2^{\ell-k}}{100} + 2^{k-\ell+2a+4} \right) \geq \ln \left(\frac{2^{\ell-k}}{100} \right) = \ln(2)(\ell - k) - \ln(100),$$

and

$$\begin{aligned} D(k, \ell, m) &\leq \cosh^{-1} (2^{\ell-k} + 2^{k-\ell+2a+5} \pi^2) \\ &\leq \ln (2^{\ell-k} 33\pi^2) + \ln (2) = \ln (2) (\ell - k + 1) + \ln (33\pi^2). \end{aligned}$$

For $\ell - k + 1 \leq a < m - 1$, therefore $2^{\ell-k} \leq m < 2^{\ell-1}$,

$$\begin{aligned} D(k, \ell, m) &\geq \cosh^{-1} \left(\frac{2^{\ell-k}}{100} + 2^{k-\ell+2a+4} \right) \geq \ln (2^{k-\ell+2a+4}) \\ &= \ln (2) (k - \ell + 2a + 4), \end{aligned}$$

and

$$\begin{aligned} D(k, \ell, m) &\leq \cosh^{-1} (2^{\ell-k} + 2^{k-\ell+2a+5} \pi^2) \\ &\leq \ln ((\pi^2 + 1) 2^{k-\ell+2a+5}) + \ln (2) \\ &= \ln (2) (k - \ell + 2a + 6) + \ln (\pi^2 + 1). \end{aligned}$$

The proof is complete. \square

We will now try to relate $D(k, \ell, m)$ and $D'(k, \ell, m)$ in the following lemma.

Lemma 4.12. For $0 \leq k \leq \ell$, $0 \leq m < 2^{\ell-1}$, we have

$$\frac{\ln(2)}{2} D'(k, \ell, m) - \ln(200) \leq D(k, \ell, m) \leq \ln(2) D'(k, \ell, m) + \ln(66\pi^2).$$

Proof. Consider the canonical geodesic. If $k = \ell$ and $m = 1$, then it is an edge from $(k, 0)$ to $(k, 1)$. In other cases, it goes up first from (ℓ, m) to a certain ancestor, then makes two or three moves on the ring, and then finishes by going straight down to $(k, 0)$.

In the first case, $D'(k, k, 1) = 1$. In the second case, since the ancestor of any node (k, m) is $(k - 1, \lfloor \frac{m}{2} \rfloor)$, by the form of canonical geodesics, we should first go up $\ell - k$ steps to reach level k . If $m \leq 2^{\ell-k+1}$, it reaches $(k, 0)$ by at most an extra step. In this case,

$$\ell - k \leq D'(k, \ell, m) \leq \ell - k + 1.$$

If $m > 2^{\ell-k+1}$, we go up $\lfloor \log_2 m \rfloor - 1$ steps from (ℓ, m) to reach an ancestor numbered 2 or 3, then go two or three steps on the ring to the node numbered 0 on the same level, and finish by going down to $(k, 0)$. Thus we have

$$2 \lfloor \log_2 m \rfloor + k - \ell \leq D'(k, \ell, m) \leq 2 \lfloor \log_2 m \rfloor + k - \ell + 1$$

in this case.

We conclude by comparing to bounds in Lemma 4.11. \square

We want to bound the distance between any two points in $\text{RT}(\infty)$ with the following lemma.

Lemma 4.13. *For $u_1, v_1, u_2, v_2 \in \text{RT}(\infty)$ with u_1, v_1, u_2, v_2 on the same level and $d_R(u_1, v_1) = d_R(u_2, v_2)$, we have $|d(u_1, v_1) - d(u_2, v_2)| \leq 3$.*

Proof. Let $A = d_R(u_1, v_1)$. The statement of the lemma clearly holds when $A \leq 3$ by the structure of canonical geodesics.

For $A > 3$, when we consider the canonical geodesics, we take successive ancestors of u_1 and v_1 until their distance is 2 or 3. This takes at least $\lfloor \log_2 A \rfloor - 1$ generations, but at most $\lfloor \log_2(A - 1) \rfloor$. These bounds differ by at most 1. Platforms differ by at most 1, and therefore $|d(u_1, v_1) - d(u_2, v_2)| \leq 3$. \square

Proof of Theorem 4.2. By symmetry, Lemma 4.13, and Lemma 4.12, for all $u, v \in \text{RT}(\infty)$, we have

$$\frac{\ln(2)}{2} d_{\text{RT}}(u, v) - \ln(200) \leq d_P(f(u), f(v)) \leq \ln(2) d_{\text{RT}}(u, v) + \ln(66\pi^2).$$

There is only one thing left in order to prove quasi-isometry. We will now prove that for some constant ϵ , $B(f(\text{RT}(\infty)), \epsilon)$ covers the Poincaré disk. Images of each level are all on concentric circles, and the difference in radii between successive levels can be bounded by $\ln(6)$. The distance between images of neighboring nodes on the same level can be bounded by $\ln(16)$. For every point in the Poincaré disk, its distance to the nearest image of nodes is bounded by $\ln(96)$, by first moving straight away from 0 until reaching a concentric circle of images and then taking the shortest path to reach the image of a certain node.

This proves f to be a $\left(\frac{2}{\ln(2)}, \ln(66\pi^2)\right)$ -quasi-isometry from $\text{RT}(\infty)$ to the Poincaré disk. Thus $\text{RT}(\infty)$ and the Poincaré disk are quasi-isometric. The constants we found here are not tight. \square

4.3.3. A Direct Proof of Corollary 4.3. We begin with a lemma about the distance between a general geodesic and the corresponding canonical geodesic.

Lemma 4.14. *For a geodesic $[u, v]$, we have that $[u, v]$ and $\langle u, v \rangle$ are within distance 1 of each other.*

Proof. We perform an induction on the structure of $\langle u, v \rangle$.

For u, v on the same level and $d_R(u, v) \leq 3$, we can check that $[u, v]$ and $\langle u, v \rangle$ are within distance 1 of each other.

For u, v on the same level but $d_R(u, v) > 3$, let u', v' be the respective parents of u, v . We have $\langle u, v \rangle = \langle u, u' \rangle \cup \langle u', v' \rangle \cup \langle v', v \rangle$ by construction. By Lemma

4.8, the parts of $[u, v]$ that are on the same level with u, v already satisfy the condition. We have only to deal with the part on levels with lower numbering.

Let t, s be the closest nodes to u, v that are at an upper level on $[u, v]$. If $t \neq u'$, then $d(u, t) \geq 2$, because the only way to go up one level in one step is to go to the parent. By Lemma 4.8, $[u, t]$ and $\langle u, u' \rangle$ are within distance 1 of each other. Therefore, $d(u', t) = 1$, and since $d(u, t) \geq 2$, $\langle u, u' \rangle \cup \langle u', t \rangle \cup [t, v]$ is also a geodesic from u to v . This is also correct for $t = u'$. The same can be proved for s . By combining the above, we have that $\langle u, u' \rangle \cup \langle u', t \rangle \cup [t, s] \cup \langle s, v' \rangle \cup \langle v', v \rangle$ is also a geodesic between u and v . Thus the section $\langle u', t \rangle \cup [t, s] \cup \langle s, v' \rangle$ is also a geodesic from u' to v' . By the induction hypothesis, it is within distance 1 of $\langle u', v' \rangle$. However, this geodesic contains the part of $[u, v]$ on levels with lower numbering than that of u, v , which is precisely $[t, s]$. We conclude that $[u, v]$ and $\langle u, v \rangle$ are within distance 1 of each other.

For u, v on different levels, the induction is essentially the same as in the previous case, but we only need to reason on one side of u or v , depending on which one is at the lower level.

This concludes our induction. \square

Lemma 4.15. *Let u and v be two vertices at the same level ℓ such that $\langle u, v \rangle$ stays at level ℓ . Then for any vertex w , $\langle u, w \rangle$ and $\langle v, w \rangle$ are within distance 3 of each other.*

Proof. Let ℓ_0 and ℓ'_0 be the highest levels reached by $\langle u, w \rangle$ and $\langle v, w \rangle$, respectively. Without loss of generality, we assume that $\ell_0 \geq \ell'_0$. Thus $\ell \geq \ell_0$. We prove the lemma by an induction on ℓ .

Consider the base case $\ell = \ell_0$. Let w' be the ancestor of w at level ℓ_0 . Thus we know that both $\langle w', u \rangle$ and $\langle u, v \rangle$ stay at level ℓ_0 . By definition, $d_R(w', u) \leq 3$ and $d_R(u, v) \leq 3$. We can enumerate all the possible cases of arrangements of u, v, w' to check that $\langle w', v \rangle$ is at most one level higher than ℓ_0 (i.e., $\ell'_0 \geq \ell_0 - 1$), and $\langle w', v \rangle$ and $\langle w', u \rangle$ are within distance 3 of each other. Since $\langle u, w \rangle$ and $\langle v, w \rangle$ share the portion $\langle w, w' \rangle$, the lemma holds for the case $\ell = \ell_0$.

For the induction step, consider $\ell > \ell_0$. Let u' and v' be the respective parents of u and v . By Lemma 4.7 and $d_R(u, v) \leq 3$, we know that $d_R(u', v') \leq 2$. Then $\langle u', v' \rangle$ must stay at level $\ell - 1$. By the induction hypothesis, $\langle u', w \rangle$ and $\langle v', w \rangle$ are within distance 3 of each other. Since $\langle u, w \rangle = \langle u, u' \rangle \cup \langle u', w \rangle$, $\langle v, w \rangle = \langle v, v' \rangle \cup \langle v', w \rangle$, and $d_R(u, v) \leq 3$, we know that the lemma holds in this case. \square

We define a *canonical triangle* $\hat{\Delta}(u, v, w)$ to be a geodesic triangle in which all sides are canonical geodesics.

Lemma 4.16. *Every canonical triangle $\hat{\Delta}(u, v, w)$ in $\text{RT}(k)$ is 3-slim.*

Proof. Without loss of generality, let u be the vertex at the lowest level, which is ℓ . We prove the lemma by induction on ℓ . The base case $\ell = 0$ is trivial.

Consider the induction step with $\ell > 0$. If neither v nor w is at level ℓ , then both $\langle u, v \rangle$ and $\langle u, w \rangle$ go through u 's parent u' . By the induction hypothesis, we know that $\hat{\Delta}(u', v, w)$ is 3-slim. Adding u in this case does not change the distances among the sides, and thus $\hat{\Delta}(u, v, w)$ is also 3-slim. Suppose now that v or w or both are at level ℓ . In the first case, suppose that at least one pair, say u and v , is such that $\langle u, v \rangle$ stays at level ℓ . By Lemma 4.15, we know that $\langle u, w \rangle$ and $\langle v, w \rangle$ are within distance 3 of each other. Since $\langle u, v \rangle$ has length at most 3, we know that $\hat{\Delta}(u, v, w)$ is 3-slim.

In the second case, suppose that all pairs at level ℓ have their canonical geodesics in level $\ell - 1$. For each vertex at level ℓ , we take its parent, together with perhaps another vertex already within level $\ell - 1$, and we can apply the induction hypothesis and show that their canonical triangle is 3-slim.

Since for every vertex at level ℓ , its canonical geodesics to the other two vertices all go through its parent, the vertices at level ℓ do not change the distance between any pair of sides of the canonical triangle. Therefore, $\hat{\Delta}(u, v, w)$ is 3-slim. \square

With Lemmas 4.14 and 4.16, we can provide a direct proof of Corollary 4.3.

Proof of Corollary 4.3. We use the Rips condition here. For all u, v, w , By Lemma 4.14, we have $[u, v] \subseteq B(\langle u, v \rangle, 1)$. By Lemma 4.16, we have $\langle u, v \rangle \subseteq B(\langle u, w \rangle \cup \langle v, w \rangle, 3)$. By Lemma 4.14 again, we have $\langle u, w \rangle \cup \langle v, w \rangle \subseteq B([u, w] \cup [v, w], 1)$. Therefore, we conclude that $[u, v] \subseteq B([u, w] \cup [v, w], 5)$, and thus $\delta_{\text{Rips}}(\text{RT}(k)) \leq 5$. Since $\delta(\text{RT}(k)) \leq 8\delta_{\text{Rips}}(\text{RT}(k))$, we have $\delta(\text{RT}(k)) \leq 40$. \square

4.3.4. Proof of Theorem 4.4. We now apply the Rips condition to analyze the hyperbolicity of ringed trees with limited long-range edges $\text{RT}(k, f)$. First we show that a geodesic in $\text{RT}(k, f)$ cannot make too many hops at the same level of the ringed tree. For two vertices in a graph $G \in \text{RT}(k, f)$, let $[u, v]$ denote any one of the geodesics between u and v in G , while $\langle u, v \rangle$ still denotes the canonical geodesic between u and v in the base ringed tree $\text{RT}(k)$.

Lemma 4.17. *For u, v on the same level in $\text{RT}(k)$ with $d(u, v) > 1$, we have $2 \log_2 d_R(u, v) \leq d(u, v) \leq 2 \log_2(d_R(u, v) - 1) + 2$.*

Proof. We carry out an induction on $d_R(u, v)$. If $1 < d_R(u, v) \leq 3$, we can check that the assertion is correct. If $d_R(u, v) > 3$, let u', v' be the respective parents

of u, v . By the induction hypothesis,

$$2 \log_2 d_R(u', v') \leq d(u', v') \leq 2 \log_2 (d_R(u', v') - 1) + 2.$$

By the triangle inequality,

$$d(u, v) \leq d(u', v') + 2 \leq 2 \log_2 (d_R(u', v') - 1) + 4.$$

But $d_R(u', v') - 1 \leq (d_R(u, v) - 1)/2$ by Lemma 4.7. And we have

$$d(u, v) \leq 2 \log_2 (d_R(u', v') - 1) + 4 \leq 2 \log_2 (d_R(u, v) - 1) + 2.$$

On the other hand,

$$d(u, v) \geq 2 + d(u', v') \geq 2 \log_2 (2d_R(u', v')).$$

Since $d_R(u', v') \geq d_R(u, v)/2$, we have $d(u, v) \geq 2 \log_2 d_R(u, v)$. Combining these two inequalities concludes the induction. \square

Corollary 4.18. *For u, v on the same level in $\text{RT}(k)$, we have*

$$d(u, v) \leq 2 \log_2 d_R(u, v) + 2.$$

Proof. We check the case $d(u, v) = 1$ and see that the assertion is satisfied. For $d(u, v) > 1$, the result follows directly from

$$d(u, v) \leq 2 \log_2 (d_R(u, v) - 1) + 2 \leq 2 \log_2 d_R(u, v) + 2.$$

\square

Lemma 4.19. *For a graph in the class $\text{RT}(k, f)$ and two vertices u and v at the same level j , if $[u, v]$ never goes into vertices at level $i < j$, then $d(u, v) \leq \max(32, 4 \log_2 f(n)) - 1$ and $d_R(u, v) \leq f(n) \max(32, 4 \log_2 f(n)) - 1$.*

Proof. First, we will prove $d_R(u, v)/f(n) \leq d(u, v) \leq 2 \log_2 d_R(u, v) + 2$. The first inequality holds because the shortest distance from u to v without going into vertices in level $i < j$ in the base ringed tree $\text{RT}(k)$ is $d_R(u, v)$, and each long-range edge can jump at most $f(n)$ hops on the ring at any level. The second inequality follows directly from Corollary 4.18.

Let $x = d(u, v)$. From the above two inequalities, we have $x \leq 2 \log_2 x f(n) + 2$. Then $f(n) \geq 2^{(x-2)/2}/x \geq 2^{(x+1)/4}$ for $x \geq 32$, and thus $x = d(u, v) \leq \max(32, 4 \log_2 f(n)) - 1$. It follows that $d_R(u, v) \leq f(n) \max(32, 4 \log_2 f(n)) - 1$.

\square

Lemma 4.20. *For a graph in the class $\text{RT}(k, f(n))$ and two vertices u and v , $[u, v]$ and $\langle u, v \rangle$ are within distance $2 \max(32, 4 \log_2 f(n))$ of each other.*

Proof. The case in which u and v are ancestor and descendant to each other on the tree is trivial. Thus we consider the case that u and v are not ancestor and descendant to each other. Let ℓ_0^\diamond and ℓ_0^\square be respectively the innermost levels that geodesics $\langle u, v \rangle$ and $[u, v]$ reach. If $\ell_0^\square < \ell_0^\diamond$, then $[u, v]$ uses at least two more tree edges than $\langle u, v \rangle$. This means that $\langle u, v \rangle$ uses at least two ring edges.

If $\langle u, v \rangle$ uses exactly two ring edges, all edges in $[u, v]$ must be tree edges, which is impossible for a geodesic, given that u and v are not ancestor and descendant to each other. Thus, $\langle u, v \rangle$ uses exactly three ring edges. Then $[u, v]$ uses exactly one ring or long-range edge.

Let u' and v' be respective ancestors of u and v at level ℓ_0^\diamond . In this case, the only possible situation is (a) $\ell_0^\square = \ell_0^\diamond - 1$, and (b) the parents of u' and v' are connected either by a ring edge or a long-range edge, which is used by $[u, v]$. Hence, $[u, v]$ and $\langle u, v \rangle$ share the tree edges from u to u' and v' to v , and $\langle u, v \rangle$ goes through ring edges from u' to v' , while $[u, v]$ goes through the edge connecting the parents of u' and v' . We thus have that $[u, v]$ and $\langle u, v \rangle$ are within distance 1 of each other. Therefore, from now on, we consider $\ell_0^\square \geq \ell_0^\diamond$.

Let ℓ_u and ℓ_v be the levels of u and v respectively. Without loss of generality, we can suppose that $\ell_u \geq \ell_v \geq \ell_0^\square$.

For any level ℓ reachable by both $[u, v]$ and $\langle u, v \rangle$, i.e., $\ell_0^\square \leq \ell \leq \ell_u$, let x_ℓ be the first level- ℓ vertex on the geodesic $[u, v]$ starting from u , and let y_ℓ be the first level- ℓ vertex on the geodesic $\langle u, v \rangle$ starting from u . We claim that $d_R(x_\ell, y_\ell) \leq f(n) \max(32, 4 \log_2 f(n))$.

For level ℓ_u , this is trivial. Suppose that our claim is correct for some level $\ell > \ell_0^\square$. We shall inductively prove the claim for level $\ell - 1$.

By the induction hypothesis, we know that

$$d_R(x_\ell, y_\ell) \leq f(n) \max(32, 4 \log_2 f(n)).$$

Let x'_ℓ be the level- ℓ vertex just before $x_{\ell-1}$ on $[u, v]$ starting from u . By Lemma 4.19 and the fact that the portion $[x_\ell, x'_\ell]$ never goes to level $i < \ell$, we have $d_R(x_\ell, x'_\ell) \leq f(n) \max(32, 4 \log_2 f(n)) - 1$. Therefore,

$$d_R(x'_\ell, y_\ell) \leq d_R(x_\ell, y_\ell) + d_R(x_\ell, x'_\ell) \leq 2f(n) \max(32, 4 \log_2 f(n)) - 1.$$

Since $x_{\ell-1}$ and $y_{\ell-1}$ are the parents of x'_ℓ and y_ℓ respectively, by Lemma 4.7 we have $d_R(x_{\ell-1}, y_{\ell-1}) \leq f(n) \max(32, 4 \log_2 f(n))$. Our claim holds for level $\ell - 1$. By induction, our claim stands.

We thus have

$$d(x_\ell, y_\ell) \leq 2 \log_2(d_R(x_\ell, y_\ell)) + 2 \leq \max(32, 4 \log_2 f(n))$$

for all $\ell_0^{[1]} \leq \ell \leq \ell_u$. For any vertex x between x_ℓ and x'_ℓ on $[u, v]$ (note that x may be at a level $\ell' \geq \ell$), by Lemma 4.19 we have

$$d(x, y_\ell) \leq d(x, x_\ell) + d(x_\ell, y_\ell) \leq d(x'_\ell, x_\ell) + d(x_\ell, y_\ell) \leq 2 \max(32, 4 \log_2 f(n)).$$

Hence all such vertices x are within distance $2 \max(32, 4 \log_2 f(n))$ of vertex y_ℓ in $\langle u, v \rangle$, and conversely.

Similarly, we can define z_ℓ to be the first level- ℓ vertex on the geodesic $[v, u]$ starting from v , and w_ℓ to be the first level- ℓ vertex on the geodesic $\langle v, u \rangle$ starting from v , for all $\ell_0^{[1]} \leq \ell \leq \ell_v$. By a symmetric argument, we can show that the w_ℓ are within distance $2 \max(32, 4 \log_2 f(n))$ of all vertices in the segment of $[u, v]$ from $z_{\ell-1}$ to z_ℓ for $\ell > \ell_0^{[1]}$, and $d(z_{\ell_0^{[1]}}, w_{\ell_0^{[1]}}) \leq \max(32, 4 \log_2 f(n))$.

The only portion left to argue is from $x_{\ell_0^{[1]}}$ to $z_{\ell_0^{[1]}}$ in the geodesic $[u, v]$ and from $y_{\ell_0^{[1]}}$ to $w_{\ell_0^{[1]}}$ in the geodesic $\langle u, v \rangle$. By Lemma 4.19 and the definition of $\ell_0^{[1]}$, we know that $d(x_{\ell_0^{[1]}}, z_{\ell_0^{[1]}}) \leq \max(32, 4 \log_2 f(n))$. Therefore, all vertices in the segment from $x_{\ell_0^{[1]}}$ to $z_{\ell_0^{[1]}}$ in the geodesic $[u, v]$ are within $2 \max(32, 4 \log_2 f(n))$ of both $y_{\ell_0^{[1]}}$ and $w_{\ell_0^{[1]}}$. Now, for any vertex x in the segment from $y_{\ell_0^{[1]}}$ to $w_{\ell_0^{[1]}}$ in the geodesic $\langle u, v \rangle$, we need to bound the distance from x to $[u, v]$. Since

$$\begin{aligned} d_R(y_{\ell_0^{[1]}}, w_{\ell_0^{[1]}}) &\leq d_R(y_{\ell_0^{[1]}}, x_{\ell_0^{[1]}}) + d_R(x_{\ell_0^{[1]}}, z_{\ell_0^{[1]}}) + d_R(z_{\ell_0^{[1]}}, w_{\ell_0^{[1]}}) \\ &\leq 3f(n) \max(32, 4 \log_2 f(n)), \end{aligned}$$

we have

$$d(y_{\ell_0^{[1]}}, w_{\ell_0^{[1]}}) \leq 2 \log_2 d_R(y_{\ell_0^{[1]}}, w_{\ell_0^{[1]}}) + 2 \leq \max(32, 4 \log_2 f(n)).$$

Thus, for any vertex x in the segment from $y_{\ell_0^{[1]}}$ to $w_{\ell_0^{[1]}}$ in the geodesic $\langle u, v \rangle$, x can reach either $x_{\ell_0^{[1]}}$ or $z_{\ell_0^{[1]}}$ in at most $2 \max(32, 4 \log_2 f(n))$ hops. Therefore, x is within distance $2 \max(32, 4 \log_2 f(n))$ of $[u, v]$. \square

Proof of Theorem 4.4. We use the Rips condition here. For any choice of u, v, w , by Lemma 4.20, we have $[u, v] \subseteq B(\langle u, v \rangle, 2 \max(32, 4 \log_2 f(n)))$. By Lemma 4.16, $\langle u, v \rangle \subseteq B(\langle u, w \rangle \cup \langle v, w \rangle, 3)$. Again by Lemma 4.20,

$$\langle u, w \rangle \cup \langle v, w \rangle \subseteq B([u, w] \cup [v, w], 2 \max(32, 4 \log_2 f(n))).$$

Combining these, we have

$$[u, v] \subseteq B([u, w] \cup [v, w], 4 \max(32, 4 \log_2 f(n)) + 3).$$

Therefore, since δ and δ_{Rips} differ by a constant factor, we have $\delta(\text{RT}(k, f)) \leq c \log f(n)$ for some constant c . \square

4.3.5. Proof of Theorem 4.5. We first analyze the δ -hyperbolicity of $\text{RRT}(k, e^{-\alpha d_R(u,v)})$.

Proof of Theorem 4.5, part 1. We have $e^\alpha \leq \rho = \sum_{v \in V \setminus \{u\}} e^{-d_R(u,v)\alpha} \leq 2 \sum_{i=1}^{+\infty} e^{i\alpha}$. Therefore, $\rho = \Theta(1)$. A vertex u on the leaves of $\text{RT}(k)$ has a long-range edge with ring distance greater than k with probability $\Theta(e^{-k\alpha})/\rho$. Let $k = \frac{2}{\alpha} \log n$. We know that a vertex has a long-range edge of ring distance greater than $\frac{2}{\alpha} \log n$ with probability $\Theta(1/n^2) = o(1/n)$.

Therefore, with probability $1 - o(1)$, long-range edges never exceed the ring distance $\frac{2}{\alpha} \log n$. From Theorem 4.4, it follows that

$$\delta(\text{RRT}(k, e^{-\alpha d_R(u,v)})) = O(\log \log n)$$

for every $\alpha > 0$. \square

For the case $\delta(\text{RRT}(k, d_R(u,v)^{-\alpha}))$, we look first at a lemma about the effect of long-range edges with ring distance $\Omega(n^c)$ on the δ -hyperbolicity of ringed trees.

Lemma 4.21. *If we add an edge between u and v on the outermost ring of a k -level ringed tree with $d_R(u,v) \geq c'n^c$, for some constants c and c' , then the resulting graph G (possibly with other edges on the outermost ring) has $\delta(G) = \Omega(\log n)$.*

Proof. Let w be a node of lowest layer number on $\langle u, v \rangle$. By Lemma 4.17 and the structure of a canonical geodesic,

$$d(w, u) \geq \frac{2 \log_2 c'n^c - 3}{2} \geq c \log_2 n + \log_2 c' - \frac{3}{2}.$$

We consider the midpoint x of $[u, w]$. For any point y in $[v, w]$, by considering the canonical geodesic $\langle x, y \rangle$, we know that

$$d(x, y) \geq \frac{c \log_2 n + \log_2 c' - 3/2}{2} - 3.$$

Therefore, $\Delta(u, v, w)$ is at best

$$\left(\frac{c \log_2 n + \log_2 c' - 3/2}{2} - 3 \right)\text{-slim,}$$

and thus $\delta(G) = \delta_{\text{Rips}}(G) = \Omega(\log n)$. \square

A probabilistic version comes naturally as the following corollary.

Corollary 4.22. *For a random graph G formed by linking edges on leaves of a ringed tree, if for some constant c with $0 < c < 1$, with high probability there exists an edge linking some u and v with $d_R(u, v) = \Theta(n^c)$, then with high probability $\delta(G) = \Theta(\log n)$.*

Proof. The diameter of a ringed tree gives an $O(\log n)$ upper bound. Lemma 4.21 gives an $\Omega(\log n)$ lower bound. \square

We can now estimate the δ -hyperbolicity of $\text{RRT}(k, d_R(u, v)^{-\alpha})$.

Proof of Theorem 4.5, part 2. Note that in a ringed tree with n vertices, at least $n/2$ of them are leaves. For a constant c with $0 < c < 1$ and a fixed vertex u , the probability that the long-range edge (u, v) has $d_R(u, v) \leq (n/4)^c$ (we shall say that it is “good”) is

$$p = 2\rho^{-1} \sum_{d=1}^{(n/4)^c} d^{-\alpha}, \quad \text{where } \rho = \Theta\left(\sum_{d=1}^n d^{-\alpha}\right).$$

For $0 < \alpha < 1$, we have $p = O(n^{(1-\alpha)(c-1)}) = o(1)$. For $\alpha = 1$, we have $p = c + o(1)$. In these two cases, all edges are good with probability at most

$$(c + o(1))^{(n/2)} = o(1).$$

For $\alpha > 1$, we have $\rho = O(1)$. We take $q = 1 - p$. We have $q = O(n^{c(1-\alpha)})$. By choosing

$$c = \min\left(1, \frac{1}{2(\alpha - 1)}\right),$$

we have $q = O(n^{-1/2})$. All edges are good with probability

$$(1 - q)^{n/2} = O\left(e^{-\sqrt{n}}\right) = o(1).$$

In all three cases, by Corollary 4.22, with probability $1 - o(1)$, we have $\delta(\text{RRT}(k, d_R(u, v)^{-\alpha})) = \Theta(\log n)$. \square

Finally, we estimate the δ -hyperbolicity of $\text{RRT}(k, 2^{-\alpha h(u, v)})$.

Proof of Theorem 4.5, part 3. We note that the number of leaves is $n_L = (n + 1)/2$. Fix a leaf u . There are 2^{h-1} leaves v such that $h(u, v) = h$. Therefore,

$$\rho = \sum_{h=1}^{\log_2 n_L} 2^{h-1-\alpha h}.$$

For a constant c with $0 < c < 1$, let

$$p(c) = \rho^{-1} \sum_{h=1}^{c \log_2 n_L} 2^{h-1-\alpha h}$$

be the probability that u never links to any v with $h(u, v) \geq c \log_2 n$. We have

$$\rho = \begin{cases} 2^{-\alpha} n_L^{1-\alpha} / (1 - 2^{1-\alpha}) & \text{for } \alpha \neq 1, \\ \frac{1}{2} \log_2 n_L & \text{for } \alpha = 1. \end{cases}$$

For $\alpha \neq 1$, we have $p(c) = n_L^{-(1-\alpha)(1-c)}$. For $\alpha = 1$, we have $p(c) = c$.

For the case $\alpha \leq 1$, we have $p(1/2) = O(1)$. Therefore, with probability $(1 - p(1/2))^{n_L} = o(1)$, there exist some u, v linked together with $h(u, v) \geq \frac{1}{2} \log_2 n$. For $\alpha > 1$, we take the constant $c_0 = \min(1, \frac{1}{2(\alpha-1)})$, and we have $p(c_0) = n_L^{-1/2}$. Therefore, with probability $(1 - p(c_0))^{n_L} = O(e^{-\sqrt{n_L}}) = o(1)$, there exist some u, v linked together with $h(u, v) \geq c_0 \log_2 n_L$. In any case, with $1 - o(1)$ probability, there exist u, v linked together by a long-range edge with $h(u, v) \geq c \log_2 n_L$ for some constant c . We notice that this occurs uniformly through all edges.

Given u, v with $h(u, v) = h$, we have $d_R(u, v) < 2^{h/2}$ with probability at most $2^h 2^{1-2h} = 2^{1-h}$ by simply counting pairs within ring distance $2^{h/2}$. With $h(u, v) \geq c \log_2 n_L, c > 0$, we know that with probability $1 - 2n_L^{-c} = 1 - o(1)$, we have $d_R(u, v) \geq 2^{h/2} = n_L^{c/2}$. Combining with the previous analysis, we prove that with $1 - o(1)$ probability, there exist u, v linked together by a long-range edge with $d_R(u, v) \geq n_L^{c/2}$ for some constant $c > 0$. By Corollary 4.22 and by $n_L > n/2$, with probability $1 - o(1)$ we have $\delta(\text{RRT}(k, 2^{-\alpha h(u,v)})) = \Theta(\log n)$. \square

4.3.6. Proof of Theorem 4.6. We can order a binary tree to give it a ring distance. We will suppose hereinafter that such a distance is defined. We begin with a counterpart of Lemma 4.21 in a binary tree.

Lemma 4.23. *If there is an edge between two leaves u, v of a binary tree of size n with distance to lowest common ancestor $h(u, v) = c_1 \log_2 n + c_2$ for some constants $c_1 > 0, c_2 > 0$, then the resulting graph G (possibly with other edges on the outermost ring) has $\delta(G) = \Omega(\log n)$.*

Proof. Consider w , the lowest common ancestor of u, v , and x , the midpoint of $[u, w]$. We have $d(x, w) = d(w, u)/2 = h(u, v)/2$. For every y in $[w, v]$, $d(x, y) \geq h(u, v)/2$, since the only path in the tree from y to x always passes by w , and we need to climb $h(u, v)/2$ levels if we use links on leaves. Therefore, $\Delta(u, v, w)$ is at best $h(u, v)/2 = \frac{1}{2}(c_1 \log_2 n + c_2)$, and we conclude that $\delta(G) = \Omega(\log n)$. \square

Corollary 4.24. *For a random graph G formed by linking edges on leaves of a binary tree, if for some constant c with $0 < c < 1$, with high probability there exists an edge linking some u and v with $h(u, v) = \Theta(\log n)$, then with high probability, $\delta(G) = \Theta(\log n)$.*

Proof. The diameter of a binary tree gives an $O(\log n)$ upper bound. Lemma 4.23 gives an $\Omega(\log n)$ lower bound. \square

Proof of Theorem 4.6. For $\text{RBT}(k, e^{-\alpha d_R(u,v)})$, the height of the whole tree is $h = \lceil \log_2 n \rceil$. There are $\Theta(\sqrt{n})$ subtrees of height $h/2$, with the root at level $h/2$. For every neighboring such subtree, the rightmost leaf u on the left subtree and the leftmost leaf v on the right subtree satisfy $d_R(u, v) = 1$, $h(u, v) \geq h/2$. For each leaf, we have $\rho = O(1)$. Therefore, for u, v with $d_R(u, v) = 1$, there is an extra edge between u and v with constant probability $e^{-\alpha} \rho^{-1} > 0$. Since there are \sqrt{n} such pairs, with probability $1 - (1 - e^{-\alpha} \rho^{-1})^{\sqrt{n}} = 1 - o(1)$, there is a pair of leaves u, v linked by an extra edge with $d_R(u, v) = 1$, $h(u, v) \geq h/2 = \Theta(\log n)$. By Corollary 4.24, with probability $1 - o(1)$, $\delta(\text{RBT}(k, e^{-\alpha d_R(u,v)})) = \Theta(\log n)$.

For $\text{RBT}(k, d_R(u, v)^{-\alpha})$ and $\text{RBT}(k, 2^{-\alpha h(u,v)})$, using the same analysis as that in the proof of Theorem 4.5, we know that for some constant $c > 0$, with $1 - o(1)$ probability, there is an extra edge between u and v with $d_R(u, v) = \Omega(n^c)$. We have $h(u, v) = \Omega(\log n)$, because a subtree of height h spans a ring distance at most 2^h . By Corollary 4.24, we have $\delta(\text{RBT}(k, d_R(u, v)^{-\alpha})) = \delta(\text{RBT}(k, 2^{-\alpha h(u,v)})) = \Theta(\log n)$. \square

4.4. Extensions of the Random Ringed Tree Model

We will now discuss some extensions of the random ringed tree (RRT) model and show that our results still hold for these extensions, thus extending the model's expressivity.

We begin with some observations on the proof of Theorem 4.5. In this proof, the upper bound of δ -hyperbolicity is given by Theorem 4.4, and the lower bound is given by Corollary 4.22. In the statement of Theorem 4.4, by the definition of $\text{RT}(k, f)$, only a uniform bound of ring distance $d_R(u, v)$ for each long-range edge (u, v) is considered. In the statement of Corollary 4.22, the only quantity concerning a long-range edge (u, v) is also the ringed distance between u and v , and to apply this corollary, we have only to show that a long-range edge (u, v) with $d_R(u, v) = \Theta(n^c)$ for some constant c exists with high probability. Therefore, the proof of Theorem 4.5 relies only on the ring distances of long-range edges.

To extend the RRT model while keeping similar properties on δ -hyperbolicity, we have only to show that Theorem 4.4 and Corollary 4.22 are still applicable

in these extensions. We will here discuss two extensions on choosing long-range edges.

4.4.1. A Constant Number of Long-Range Edges for Each Node. In the original RRT model, each node has only one long-range edge connecting to other nodes. We can extend the model to allow each node to have a constant number of long-range edges connecting to a constant number of other nodes. In this extension, Theorem 4.5 still holds, since Theorem 4.4 is not affected by the number of long-range edges, and Corollary 4.22 is still applicable, since the required probability only increases with extra long-range edges.

4.4.2. Independent Long-Range Edges. In the original RRT model, we choose exactly one long-range edge for each node. A variant of the model is that each node u can choose edge (u, v) independently from other edges (u, v') , with the same probability as in the original model such that in expectation, u connects with one long-range edge. In this variant, Theorem 4.5 still holds. The reason is that in expectation, at least a constant fraction of nodes issue only one edge, and the computation for applying Corollary 4.22 is similar. It is clear that the application of Theorem 4.4 remains valid.

In the two variants discussed above, we can see that Theorem 4.5 still applies, and we have exactly the same property of δ -hyperbolicity of these variants. We can also combine these two variants, and it is clear that our results are still valid.

5. Discussion and Open Problems

Perhaps the most obvious extension of our results is to close the gap in the bounds on the hyperbolicity in the low-dimensional small-world model when γ is at the “sweet spot” as well as to extend the results for large γ to dimensions $d \geq 2$. Also of interest is characterizing in more detail the hyperbolicity properties of other random graph models, in particular those that have substantial heavy-tailed properties. Finally, the exact computation of δ by definition takes $O(n^4)$ time, which is not scalable to large graphs, and thus the design of more efficient exact or approximation algorithms would be of interest.

From a broader perspective, however, our results suggest that δ is a measure of treelikeness that can be quite sensitive to noise in graphs, and in particular to randomness as it is implemented in common network generative models. For example, ringed trees have constant hyperbolicity, but once some random links are added among leaves, our results show that very likely, their hyperbolic δ reaches the level of the graph diameter and they become not hyperbolic at all.

Moreover, our results for the δ hyperbolicity of rewired trees (Theorem 4.6) versus rewired low- δ treelike metrics (Theorem 4.5 (1)) suggest that while quite appropriate for continuous negatively curved manifolds, the usual definition of δ may be somewhat less useful for discrete graphs. Thus, it would be of interest to address questions such as, does there exist a measure other than Gromov's δ that is more appropriate for graph-based data or more robust to noise/randomness as used in popular network generation models? Is it possible to incorporate in a meaningful way nontrivial randomness in other low- δ -hyperbolicity graph families? Can we construct nontrivial random graph families that contain as much randomness as possible while having low δ -hyperbolicity compared to graph diameter?

Acknowledgments. We are grateful to Yajun Wang and Xiaohui Bei for their helpful discussions on this topic.

References

- [Abraham et al. 07] I. Abraham, M. Balakrishnan, F. Kuhn, D. Malkhi, V. Ramasubramanian, and K. Talwar. "Reconstructing Approximate Tree Metrics." In *Proceedings of the 26th Annual ACM Symposium on Principles of Distributed Computing*, pp. 43–52, 2007.
- [Baryshnikov 02] Y. Baryshnikov. "On the Curvature of the Internet." In *Workshop on Stochastic Geometry and Teletraffic, Eindhoven, The Netherlands*, 2002.
- [Baryshnikov and Tucci 10] Y. Baryshnikov and G. H. Tucci. "Asymptotic Traffic Flow in an Hyperbolic Network I: Definition and Properties of the Core." arXiv:1010.3304, 2010.
- [Bermudo et al. 10] S. Bermudo, J. M. Rodríguez, J. M. Sigarreta, and J.-M. Vilaire. "Mathematical Properties of Gromov Hyperbolic Graphs." In *Proceedings of the 2010 International Conference of Numerical Analysis and Applied Mathematics*, pp. 575–578, 2010.
- [Boguñá et al. 09] M. Boguñá, D. Krioukov, and K. C. Claffy. "Navigability of Complex Networks." *Nature Physics* 5 (2009), 74–80.
- [Bowditch 91] Brian H. Bowditch. "Notes on Gromov's Hyperbolicity Criterion for Path-Metric Spaces." In *Group Theory from a Geometrical Viewpoint (Trieste, 1990)*, pp. 64–167, 1991.
- [Bridson and Haefliger 99] Martin R. Bridson and André Haefliger. *Metric Spaces of Non-positive Curvature*. Springer, 1999.
- [Chepoi and Dragan 00] V. Chepoi and F. Dragan. "A Note on Distance Approximating Trees in Graphs." *European Journal of Combinatorics* 21 (2000), 761–766.
- [Chepoi and Estellon 07] V. Chepoi and B. Estellon. "Packing and Covering δ -Hyperbolic Spaces by Balls." In *Proceedings of the 10th International Workshop on Approximation*, pp. 59–73, 2007.

- [Chepoi et al. 08] Victor Chepoi, Feodor F. Dragan, Bertrand Estellon, Michel Habib, and Yann Vaxès. “Diameters, Centers, and Approximating Trees of δ -Hyperbolic Geodesic Spaces and Graphs.” In *Proceedings of the 24th Annual Symposium on Computational Geometry*, pp. 59–68. ACM, 2008.
- [Chepoi et al. 12] V. Chepoi, F. F. Dragan, B. Estellon, M. Habib, Y. Vaxès, and Y. Xiang. “Additive Spanners and Distance and Routing Labeling Schemes for Hyperbolic Graphs.” *Algorithmica* 62 (2012), 713–732.
- [Gavoille and Ly 05] C. Gavoille and O. Ly. “Distance Labeling in Hyperbolic Graphs.” In *Proceedings of the 16th Annual International Symposium on Algorithms and Computation*, pp. 1071–1079, 2005.
- [Ghys and de La Harpe 90] Étienne Ghys and Pierre de La Harpe. “Sur les groupes hyperboliques d’après Mikhael Gromov.” Birkhäuser, 1990.
- [Gromov 87] Mikhael Gromov. “Hyperbolic Groups.” *Essays in Group Theory* 8 (1987), 75–263.
- [Howorka 79] Edward Howorka. “On Metric Properties of Certain Clique Graphs.” *Journal of Combinatorial Theory, Series B* 27 (1979), 67–74.
- [Jonckheere and Lohsoonthorn 02] E. Jonckheere and P. Lohsoonthorn. “Hyperbolic Geometry Approach to Multipath Routing.” In *Proceedings of the 10th Mediterranean Conference on Control and Automation*, 2002.
- [Jonckheere and Lohsoonthorn 04] E. Jonckheere and P. Lohsoonthorn. “Geometry of Network Security,” In *Proceedings of the 2004 American Control Conference*, pp. 976–981, 2004.
- [Jonckheere et al. 07] E. A. Jonckheere, M. Lou, J. Hespanha, and P. Barooah. “Effective Resistance of Gromov-Hyperbolic Graphs: Application to Asymptotic Sensor Network Problems.” In *Proceedings of the 46th IEEE Conference on Decision and Control*, pp. 1453–1458, 2007.
- [Jonckheere et al. 08] E. Jonckheere, P. Lohsoonthorn, and F. Bonahon. “Scaled Gromov Hyperbolic Graphs.” *Journal of Graph Theory* 57 (2008), 157–180.
- [Jonckheere et al. 11] Edmond Jonckheere, Mingji Lou, Francis Bonahon, and Yuliy Baryshnikov. “Euclidean versus Hyperbolic Congestion in Idealized versus Experimental Networks.” *Internet Mathematics* 7 (2011), 1–27.
- [Kleinberg 00] J. Kleinberg. “The Small-World Phenomenon: An Algorithm Perspective.” In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing*, pp. 163–170, 2000.
- [Kleinberg 02] J. Kleinberg. “Small-World Phenomena and the Dynamics of Information.” *Proceedings of Annual Advances in Neural Information Processing Systems*, pp. 431–438, 2002.
- [Kleinberg 06] J. Kleinberg. “Complex Networks and Decentralized Search Algorithms.” In *Proceedings of the International Congress of Mathematicians*, 2006.
- [Kleinberg 07] R. Kleinberg. “Geographic Routing Using Hyperbolic Space.” In *Proceedings of the 26th IEEE International Conference on Computer Communications*, pp. 1902–1909, 2007.

- [Krioukov et al. 07] D. Krioukov, K. C. Claffy, K. Fall, and A. Brady. “On Compact Routing for the Internet.” *Computer Communication Review* 37 (2007), 41–52.
- [Krioukov et al. 09] D. Krioukov, F. Papadopoulos, A. Vahdat, and M. Boguñá. “Curvature and Temperature of Complex Networks.” *Physical Review E* 80 (2009), 035101(R).
- [Krioukov et al. 10] D. Krioukov, F. Papadopoulos, M. Kitsak, A. Vahdat, and M. Boguñá, “Hyperbolic Geometry of Complex Networks.” *Physical Review E* 82 (2010), 036106.
- [Lamping and Rao 94] J. Lamping and R. Rao. “Laying Out and Visualizing Large Trees Using a Hyperbolic Space.” In *Proceedings of the 7th Annual ACM Symposium on User Interface Software and Technology*, pp. 13–14, 1994.
- [Lamping et al. 95] J. Lamping, R. Rao, and P. Pirolli. “A Focus + Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies.” In *Proceedings of CHI’95: Conference on Human Factors in Computing*, pp. 401–408, 1995.
- [Lohsoonthorn 03] P. Lohsoonthorn. “Hyperbolic Geometry of Networks.” PhD thesis, University of Southern California, 2003.
- [Lou 08] M. Lou. “Traffic Pattern in Negatively Curved Network.” PhD thesis, University of Southern California, 2008.
- [Martel and Nguyen 04] Charles U. Martel and Van Nguyen. “Analyzing Kleinberg’s (and Other) Small-World Models.” In *Proceedings of the 23rd Annual ACM Symposium on Principles of Distributed Computing*, pp. 179–188, 2004.
- [de Montgolfier et al. 11] Fabien de Montgolfier, Mauricio Soto, and Laurent Viennot. “Treewidth and Hyperbolicity of the Internet.” In *IEEE Networks Computing and Applications 2011*. IEEE, 2011.
- [Munzner 98] T. Munzner. “Exploring Large Graphs in 3D Hyperbolic Space.” *IEEE Computer Graphics and Applications* 18 (1998), 18–23.
- [Munzner and Burchard 95] T. Munzner and P. Burchard. “Visualizing the Structure of the World Wide Web in 3D Hyperbolic Space.” In *Proceedings of the First Symposium on Virtual Reality Modeling Language*, pp. 33–38, 1995.
- [Narayan and Saniee 11] Onuttom Narayan and Iraj Saniee. “Large-Scale Curvature of Networks.” *Phys. Rev. E* 84 (2011), 066108.
- [Narayan et al. 12] O. Narayan, I. Saniee, and G. H. Tucci. “Lack of Spectral Gap and Hyperbolicity in Asymptotic Erdős–Rényi Random Graphs.” In *Proceedings of the 5th International Symposium on Communications Control and Signal Processing (ISCCSP)*, pp. 1–4, 2012.
- [Newman and Watts 99] M. E. J. Newman and D. J. Watts. “Scaling and Percolation in the Small-World Network Model.” *Physical Review E* 60 (1999), 7332–7342.
- [Nguyen and Martel 05] Van Nguyen and Charles U. Martel. “Analyzing and Characterizing Small-World Graphs.” In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 311–320, 2005.
- [Papadopoulos et al. 10] Fragkiskos Papadopoulos, Dmitri V. Krioukov, Marián Boguñá, and Amin Vahdat. “Greedy Forwarding in Dynamic Scale-Free Networks Embedded in Hyperbolic Metric Spaces.” In *Proceedings of the 29th IEEE*

- International Conference on Computer Communications (INFOCOM)*, pp. 2973–2981, 2010.
- [Papazian and Rémila 00] C. Papazian and E. Rémila. “Some Properties of Hyperbolic Networks.” In *Proceedings of the 9th International Conference on Discrete Geometry for Computer Imagery*, pp. 149–158, 2000.
- [Shang 12] Yilun Shang. “Lack of Gromov-Hyperbolicity in Small-World Networks.” *Central European Journal of Mathematics* 10 (2012), 1152–1158.
- [Shavitt and Tankel 08] Y. Shavitt and T. Tankel. “Hyperbolic Embedding of Internet Graph for Distance Estimation and Overlay Construction.” *IEEE/ACM Transactions on Networking* 16 (2008), 25–36.
- [Tucci 12] Gabriel H. Tucci. “Random Regular Graphs Are Not Asymptotically Gromov Hyperbolic.” arXiv:1203.5069, 2012.
- [Walter and Ritter 02] J. A. Walter and H. Ritter. “On Interactive Visualization of High-Dimensional Data Using the Hyperbolic Plane.” *Proceedings of the 8th Annual ACM SIGKDD Conference*, pp. 123–132, 2002.

Wei Chen, Microsoft Research Asia, 5 Dan Ling Street, Haidian District, Beijing 100080, China (weic@microsoft.com)

Wenjie Fang, LIAFA, University Paris Diderot - Paris 7, 75205 Paris Cedex 13, France (Wenjie.Fang@ens.fr)

Guangda Hu, Department of Computer Science, Princeton University, Princeton, NJ 08540, USA (guangdah@cs.princeton.edu)

Michael W. Mahoney, Department of Mathematics, Stanford University, Stanford, CA 94305, USA (mmahoney@cs.stanford.edu)