

On the Hyperbolicity of Small-World and Tree-Like Random Graphs

Wei Chen¹, Wenjie Fang², Guangda Hu³, and Michael W. Mahoney⁴

¹ Microsoft Research Asia

weic@microsoft.com

² Ecole Normale Supérieure de Paris

Wenjie.Fang@ens.fr

³ Princeton University

guangdah@cs.princeton.edu

⁴ Stanford University

mmahoney@cs.stanford.edu

Abstract. Hyperbolicity is a property of a graph that may be viewed as being 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 we establish several positive and negative results for small-world and tree-like random graph models. In particular, we show that small-world random graphs built from underlying grid structures do not have strong improvement in hyperbolicity, even when the rewiring greatly improves decentralized navigation. On the other hand, for a class of tree-like graphs called ringed trees that have constant hyperbolicity, 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 the first significant analytical results on the hyperbolicity of a rich class of random graphs, which shed 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.

Keywords: Graph hyperbolicity, complex networks, small-world networks, random graphs, decentralized navigation.

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 tree-like 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 [22,26,31]; it has been applied to questions of compact routing, navigation, and decentralized search in Internet graphs and small-world social networks [11,19,1,20,8]; and it has been applied to a range of other problems such as distance estimation, sensor networks, and traffic flow and congestion minimization [2,13,27,10].

The hyperbolicity of graphs is typically measured by Gromov's hyperbolic δ [12,4] (see Section 2). The hyperbolic δ of a graph measures the "tree-likeness" of the graph in terms of the graph distance metric. It can range from 0 up to the half of the graph diameter, with trees having $\delta = 0$, in contrast of "circle graphs" and "grid graphs" having large δ equal to roughly half of their diameters.

In this paper, we study the δ -hyperbolicity of families of random graphs that intuitively have some sort of tree-like or hierarchical structure. Our motivation comes from two angles. First, although there are a number of empirical studies on the hyperbolicity of real-world and random graphs [2,13,24,23,27,10], there are essentially no systematic analytical study on 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 [6,11,9,7,21,8], and the routing infrastructure of the Internet is also empirically shown to be hyperbolic [2]. Thus, it is interesting to further investigate if efficient routing capability implies good graph hyperbolicity.

To achieve our goal, we first provide 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 in the order of 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 at 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 [16], which build random long-range edges among pairs of nodes with probability inverse proportional to the γ -th power of the grid distance of the pairs. Kleinberg shows that when γ equals to the grid dimension d , decentralized routing can be improved from $\Theta(n)$ in 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 $\gamma = d$, 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 is not hyperbolic and when $\gamma > 3$ and $d = 1$, the random graphs is not polylogarithmically hyperbolic. Although there still exists a gap between hyperbolic δ and graph diameter at the sweetspot of $\gamma = d$, our results already indicate that long-range edges that enable efficient navigation do not significantly improve the hyperbolicity of the graphs.

The 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 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 [17], the resulting random graph is not hyperbolic. Furthermore, if we use binary trees instead of ringed trees as base graphs, none of the above versions 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 the first significant analytical 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).

Related Work. There has been a lot of work on decentralized search subsequent to Kleinberg’s original work [16,17], much of which has been summarized in the review [18]. In a 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, [2] showed that measurements of the Internet are negatively curved; [13,24,23] provided empirical evidence that randomized scale-free and Internet graphs are more hyperbolic than other types of random graph models; [27] measured the average δ and related curvature to congestion; and [10] measured treewidth and hyperbolicity properties of the Internet. However, on theoretical analysis of δ -hyperbolicity, the only prior work we are aware of is [28], which proves that with non-zero probability extremely sparse Erdős-Rényi random graphs are not δ -hyperbolic for any positive constant δ .

There are a number of works that connect graph hyperbolicity with efficient distance labeling and routing schemes [6,11,9,7,21,8]. Understanding the relationship between graph hyperbolicity and the ability of efficient routing is one motivation of our research. Our analytical results show, however, that the ability of efficient routing does not necessarily mean low hyperbolicity δ .

Ideas related to hyperbolicity have been applied in numerous other networks applications, e.g., to problems such as distance estimation, sensor networks, and traffic flow and congestion minimization [30,14,15,27,3], as well as large-scale data visualization [22,26,31]. The latter applications typically take important advantage of the idea that data are often hierarchical or tree-like and that there is “more room” in hyperbolic spaces of a given dimension than corresponding Euclidean spaces.

The full version of this conference paper, including detailed proofs and additional results, is available as the technical report [5].

2 Preliminaries on Hyperbolic Spaces and Graphs

We provide basic concepts concerning hyperbolic spaces and graphs used in this paper. For more comprehensive coverage on hyperbolic spaces, see, e.g., [4].




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

2.1 Gromov’s δ -Hyperbolicity

In this paper, we use the following four-point condition originally introduced by Gromov [12] as the hyperbolicity measure of a metric space.

Definition 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 denote $\delta(u, v, w, x) = (d(u, v) + d(w, x) - d(u, x) - d(w, v))/2$. (X, d) is called δ -hyperbolic for some non-negative real number δ if for any four points $u, v, w, x \in X$, $\delta(u, v, w, x) \leq \delta$. Let $\delta(X, d)$ be the smallest possible value of such δ , which can also be defined as $\delta(X, d) = \sup_{u, v, w, x \in X} \delta(u, v, w, x)$.*

An undirected, unweighted and connected graph $G = (V, E)$ can be viewed as a metric space (V, d_G) with the standard graph distance metric d_G . We then apply the four-point condition defined above to define the δ -hyperbolicity of graph G , denoted as $\delta(G) = \delta(V, d_G)$. Trees are 0-hyperbolic, and it is often helpful to view graphs with a low hyperbolic δ as tree-like when viewed at large-size scales.

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

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

2.2 Poincaré Disk

The Poincaré disk is a well-studied hyperbolic metric space. In this paper, we use the Poincaré disk to mainly convey some intuition about hyperbolicity and tree-like behaviors, and thus we defer its technical definition to [5]. Visually, the Poincaré disk is an

open disk with unit radius, and 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 towards the origin (Figure 1(a)). For two points maintaining the same Euclidean distance on the disk, their hyperbolic distance increases exponentially fast when they move from the center to the boundary of the disk, meaning that there are “more room” towards the boundary. This can be seen from a tessellation of the Poincaré disk shown in Figure 1(b).

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

In this section, we consider the δ -hyperbolicity of random graphs constructed according to the small-world graph model of Kleinberg [16], 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 wrap-around). 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 $V = \{(x_1, x_2, \dots, x_d) \mid x_i \in \{0, 1, \dots, n^{1/d} - 1\}, i \in [d]\}$, $E = \{((x_1, x_2, \dots, x_d), (y_1, y_2, \dots, y_d)) \mid \exists j \in [d], y_j = x_j + 1 \pmod{n^{1/d}} \text{ or } y_j = x_j - 1 \pmod{n^{1/d}}, \forall i \neq j, y_i = x_i\}$. 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 $KSW(n, d, \gamma)$; and we let $\delta(KSW(n, d, \gamma))$ denote the random variable of the hyperbolic δ of a randomly picked graph G in $KSW(n, d, \gamma)$. 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) [16]; 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”.

Results and Their Implications. The following theorem summarizes our main technical results on the hyperbolicity of small-world graphs.

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

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

This theorem, together with the results of [16] on the navigability of small-world graphs, have several implications. The first result shows that when $\gamma = d$, with high probability the hyperbolic δ of the small-world graphs is at least $c(\log n)^{\frac{1}{1.5(d+1)}}$ for some

constant c . We know that the diameter is $\Theta(\log n)$ in expectation when $\gamma = d$ [25]. Thus the small-world graphs at the sweetspot for efficient routing is 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 the hyperbolicity (relative to the improvement of the diameter) for small-world graphs at the sweetspot (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 $\Theta(\log n)$ [25]; 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 $\gamma > d$, in which case the random graph degenerates towards 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 lower bounded by a (low-degree) polynomial of n ; this also implies that the graphs in this range are not polylogarithmically hyperbolic. Our polynomial exponent $\frac{\gamma-2}{\gamma-1} - \epsilon$ matches the diameter lower bound proven in [29].

Outline of the Proof of Theorem 1. In our analysis, we use two different techniques, one for the first two results in Theorem 1, and the other for the last result. For the first two results, we further divide the analysis into two cases $d \geq 2$ and $d = 1$.

When $d \geq 2$ and $0 \leq \gamma \leq d$, we first pick an arbitrary square grid with ℓ_0 nodes on each side. We know that when only grid distance is considered, the four corners of the square grid have the Gromov δ value equal to ℓ_0 . We will show that, as long as ℓ_0 is not very large (to be exact, $O((\log n)^{\frac{1}{1.5(d+1)+\epsilon}})$ when $\gamma = d$ and $O(\log n)$ when $0 \leq \gamma < d$), the probability that any 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. 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 any 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 upper bound such ℓ 's 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 it 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 existing a path in each category and finally bound all such paths in aggregate. The most difficult part of the analysis is the bounding of the probability of existing a path in each category.

For the case of $d = 1$ 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 edges e_0 , where ℓ_0 is fixed to be the same as the case of $d \geq 2$. We show that with high probability, (a) such an edge e_0 exists, and (b)

the distance of any two vertices on the ring is simply their ring distance. This is enough to show 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 only connect 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 the two opposite ends on the ring. Then we argue that there must be a path \mathcal{P}_{AB}^+ that only goes through the clockwise side of ring from A to B , while another path \mathcal{P}_{AB}^- that only goes through the counter-clockwise side of the ring from A to B , and importantly, the shorter length of these two paths are at most $O(\ell_0)$ longer than the distance between A and B . We then pick the middle point C and D of \mathcal{P}_{AB}^+ and \mathcal{P}_{AB}^- , respectively, and argue that the δ value of the four points $A, B, C,$ and D give the desired lower bound.

Extensions to Other Models. We further study several extensions of the *KSW* model, including base grid without wrap-around, constant number of long-range links per node, and independent linking probabilities of each edge. We show that Theorem 1 still holds in all these models (except the case of $d = 1$ and $\gamma > 3$ for the grid with no wrap-around extension) and their combinations.

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 binary tree or tree-like low- δ graph. In particular, we consider as based graphs both binary trees (Figure 1(c)) and ringed trees (Figure 1(d)), which contain concentric rings connecting all nodes in the same level of the binary tree, and adding long range links on these base graphs. The ringed tree is formally defined as follows.

Definition 3 (Ringed tree). A ringed tree of level k , denoted $RT(k)$, 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 child and the right child of a vertex with string σ are represented as $\sigma 0$ and $\sigma 1$, 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.

Note that the diameter of the ringed tree $RT(k)$ is $\Theta(\log n)$, where $n = 2^k - 1$ is the number of vertices in $RT(k)$, and we will use $RT(\infty)$ to denote the infinite ringed tree when k in $RT(k)$ goes to infinity. Thus, a ringed tree may be thought of as a soft version of a binary tree. To some extent, a ringed tree can also be viewed as an idealized picture reflecting the hierarchical structure in real networks coupled with local neighborhood connections, such as Internet autonomous system (AS) networks, which has both a hierarchical structure of different level of AS'es, and peer connections based on geographical proximity.

Results and Their Implications. A visual comparison of the ringed tree of Figure 1(d) with the tessellation of Poincaré disk (Figure 1(b)) suggests that the ringed tree can be

seen as an approximate tessellation or coarsening of the Poincaré disk, just as a two-dimensional grid can be seen as a coarsening of a two dimensional Euclidean space. Quasi-isometry is a technical concept making it precise what coarsening means. We show that the infinite ringed tree $RT(\infty)$ is indeed quasi-isometric to the Poincaré disk. This also implies that ringed tree $RT(k)$ for any k is constantly hyperbolic (technical definition of quasi-isometry and the above results are included in [5]).

We now consider random ringed trees constructed by adding random edges between two vertices at the outermost level, i.e., level $k - 1$, such that the probability 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 $k - 1$. Given a real-valued positive function $g(u, v)$, let $RRT(k, g)$ denote a random ringed tree constructed as follows. We start with the ringed tree $RT(k)$, 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: the first family uses an exponential decay function $g_1(u, v) = e^{-\alpha d_R(u, v)}$, and 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 as $h(u, v)$, 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 in a tree-based small-world model of Kleinberg [17]. The following theorem summarizes the hyperbolicity behavior of these three families of random ringed trees.

Theorem 2. *Considering the follow families of functions (with u and v as the variables of the function) for random ringed trees $RRT(k, g)$, for any positive integer k and positive real number α , with probability $1 - o(1)$ (when n tends to infinity), we have*

1. $\delta(RRT(k, e^{-\alpha d_R(u, v)})) = O(\log \log n)$;
2. $\delta(RRT(k, d_R(u, v)^{-\alpha})) = \Theta(\log n)$;
3. $\delta(RRT(k, 2^{-\alpha h(u, v)})) = \Theta(\log n)$;

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

Theorem 2 states that, when the random long-range edges are selected using 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. Intuitively, when it is more likely for a long-range edge to connect two far-away 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 through tree nodes. In Internet routing paths going through such shortcuts are referred to as *valley routes*.

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

Theorem 3. *Considering the follow families of functions (with u and v as the variables of the function) for random binary trees $RBT(k, g)$, for any positive integer k and positive real number α , with probability $1 - o(1)$ (when n tends to infinity), we have*

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

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

Thus, in this case, the original hyperbolicity of the base graph ($\delta = 0$ for the binary tree) is destroyed. Comparing with Theorem 2, 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.

Extensions of the Random Ringed Tree Model. We further show that all our results in this section apply to extended models that allow a constant number of long-range edges per node, or independent selection of long-range edges for each node, or both.

5 Discussions 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 “sweetspot”, as well as extending 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, exact computation of δ by its 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 tree-like-ness that can be quite sensitive to noise in graphs, and in particular to randomness as it is implemented in common network generative models. Moreover, our results for the δ hyperbolicity of rewired trees versus rewired low- δ tree-like metrics 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 it is used in popular network generation models; is it possible to incorporate in a meaningful way nontrivial randomness in other low δ -hyperbolicity graph families; and can we construct non-trivial random graph families that contain as much randomness as possible while having low δ -hyperbolicity comparing to graph diameter?

References

1. Abraham, I., Balakrishnan, M., Kuhn, F., Malkhi, D., Ramasubramanian, V., Talwar, K.: Reconstructing approximate tree metrics. In: PODC (2007)
2. Baryshnikov, Y.: On the curvature of the Internet. In: Workshop on Stochastic Geometry and Teletraffic, Eindhoven, The Netherlands (April 2002)
3. Baryshnikov, Y., Tucci, G.H.: Asymptotic traffic flow in an hyperbolic network I: Definition and properties of the core. Technical Report Preprint: arXiv:1010.3304 (2010)

4. Bridson, M.R., Haefliger, A.: *Metric Spaces of Non-Positive Curvature*. Springer (1999)
5. Chen, W., Fang, W., Hu, G., Mahoney, M.W.: On the hyperbolicity of small-world and tree-like random graphs. Technical Report Preprint: arXiv:1201.1717v2 (2012)
6. Chepoi, V., Dragan, F.: A note on distance approximating trees in graphs. *European Journal of Combinatorics* 21(6), 761–766 (2000)
7. Chepoi, V., Dragan, F.F., Estellon, B., Habib, M., Vaxès, Y.: Diameters, centers, and approximating trees of δ -hyperbolic geodesic spaces and graphs. In: SoCG (2008)
8. Chepoi, V., Dragan, F.F., Estellon, B., Habib, M., Vaxès, Y., Xiang, Y.: Additive spanners and distance and routing labeling schemes for hyperbolic graphs. *Algorithmica* 62(3-4), 713–732 (2012)
9. Chepoi, V., Estellon, B.: Packing and Covering δ -Hyperbolic Spaces by Balls. In: Charikar, M., Jansen, K., Reingold, O., Rolim, J.D.P. (eds.) APPROX and RANDOM 2007. LNCS, vol. 4627, pp. 59–73. Springer, Heidelberg (2007)
10. de Montgolfier, F., Soto, M., Viennot, L.: Treewidth and hyperbolicity of the internet. In: *IEEE Networks Computing and Applications 2011*. IEEE (2011)
11. Gavoille, C., Ly, O.: Distance Labeling in Hyperbolic Graphs. In: Deng, X., Du, D.-Z. (eds.) ISAAC 2005. LNCS, vol. 3827, pp. 1071–1079. Springer, Heidelberg (2005)
12. Gromov, M.: Hyperbolic groups. *Essays in Group Theory* 8, 75–263 (1987)
13. Jonckheere, E., Lohsoonthorn, P.: Hyperbolic geometry approach to multipath routing. In: MED (2002)
14. Jonckheere, E., Lou, M., Bonahon, F., Baryshnikov, Y.: Euclidean versus hyperbolic congestion in idealized versus experimental networks. Technical Report Preprint: arXiv:0911.2538 (2009)
15. Jonckheere, E.A., Lou, M., Hespanha, J., Barooah, P.: Effective resistance of Gromov-hyperbolic graphs: Application to asymptotic sensor network problems. In: CDC (2007)
16. Kleinberg, J.: The small-world phenomenon: an algorithm perspective. In: STOC (2000)
17. Kleinberg, J.: Small-world phenomena and the dynamics of information. In: NIPS (2002)
18. Kleinberg, J.: Complex networks and decentralized search algorithms. In: ICM (2006)
19. Kleinberg, R.: Geographic routing using hyperbolic space. In: Infocom (2007)
20. Krioukov, D., Claffy, K.C., Fall, K., Brady, A.: On compact routing for the Internet. *Computer Communication Review* 37(3), 41–52 (2007)
21. Krioukov, D., Papadopoulos, F., Kitsak, M., Vahdat, A., Boguñá, M.: Hyperbolic geometry of complex networks. *Physical Review E* 82, 036106 (2010)
22. Lamping, J., Rao, R.: Laying out and visualizing large trees using a hyperbolic space. In: UIST (1994)
23. Lohsoonthorn, P.: *Hyperbolic Geometry of Networks*. PhD thesis, University of Southern California (2003)
24. Lou, M.: *Traffic pattern in negatively curved network*. PhD thesis, University of Southern California (2008)
25. Martel, C.U., Nguyen, V.: Analyzing Kleinberg’s (and other) small-world models. In: PODC (2004)
26. Munzner, T., Burchard, P.: Visualizing the structure of the World Wide Web in 3D hyperbolic space. In: *Web3D-VRML* (1995)
27. Narayan, O., Sanjeev, I.: The large scale curvature of networks. Technical Report Preprint: arXiv:0907.1478 (2009)

28. Narayan, O., Saniee, I., Tucci, G.H.: Lack of spectral gap and hyperbolicity in asymptotic Erdős-Rényi random graphs. Technical Report Preprint: arXiv:1009.5700 (2010)
29. Nguyen, V., Martel, C.U.: Analyzing and characterizing small-world graphs. In: SODA (2005)
30. Shavitt, Y., Tankel, T.: Hyperbolic embedding of Internet graph for distance estimation and overlay construction. *IEEE/ACM Transactions on Networking* 16(1), 25–36 (2008)
31. Walter, J.A., Ritter, H.: On interactive visualization of high-dimensional data using the hyperbolic plane. In: KDD (2002)