

Greedy Forwarding in Dynamic Scale-Free Networks Embedded in Hyperbolic Metric Spaces

Fragkiskos Papadopoulos Dmitri Krioukov
Cooperative Association for Internet Data Analysis
University of California, San Diego
Email: {frag, dima}@caida.org

Marián Boguñá
Universitat de Barcelona
Email: marian.boguena@ub.edu

Amin Vahdat
University of California, San Diego
Email: vahdat@cs.ucsd.edu

Abstract

In this paper we show that complex (scale-free) network topologies naturally emerge from hyperbolic metric spaces. The hyperbolic geometry can be used to facilitate maximally efficient greedy forwarding on these topologies where packets can find their destinations with 100% probability following almost optimal, i.e., shortest paths, without the need for global topology knowledge. We demonstrate that this remarkable efficiency is robust under dynamic network conditions. Our findings suggest that forwarding information through complex networks, e.g., like the Internet, may be possible without the current overhead of routing protocols, and may also find practical applications in overlay networks for tasks such as application-level routing, information sharing, and data distribution.

I. INTRODUCTION

Routing information is the most basic and, perhaps, the most complicated function that networks perform. Conventional wisdom states that to find paths to destinations through the complex network maze, nodes must communicate and exchange information about the status of their connections to other nodes. This communication overhead is considered one of the most serious scaling limitations of our primary communication technologies today, including the Internet [20] and emerging wireless and sensor networks [22]. Even though other networks, such as P2P overlay networks, may not suffer from such overhead, they still rely on various techniques to find intended communication targets. These techniques include flooding-based mechanisms, random walks, etc., whose efficiency may be unpredictable and overhead costs unbounded [19].

Many networks in nature however, can somehow “route traffic” efficiently, i.e., their nodes can efficiently find intended communication targets even though they do not possess any global view of the system. Milgram’s 1969 experiment [25] showed a classic demonstration of this effect. Milgram asked some random individuals—sources—to send a letter to a specific person—the destination, described by name, occupation, age, and city of residence. The sources were asked to pass the letter to friends chosen to maximize the probability of the letter reaching its destination. The results were surprising: many of the letters reached their destination by making only a small number of hops, even though nodes had no global knowledge of the human acquaintance network topology, except their local connections and some characteristics (e.g., occupation, age, city of dwelling) of their connections.

Much later, Jon Kleinberg offered the first popular explanation of this surprising effect [14]. In his model, each node, in addition to being part of the graph representing the global network topology, resides in a coordinate space—a grid embedded in the Euclidean plane. The coordinates of a node in the plane, its address, abstracts information about the destination in Milgram’s experiments. Each node knows: 1) its coordinates; 2) the coordinates of its neighbors; and 3) the coordinates of the destination written on the packet. Given these three pieces of information, the node can route greedily by selecting its direct neighbor closest to the destination in the plane.

Clearly, the described greedy forwarding strategy can be efficient *only if* the network topology is in some way congruent with the underlying space. But the Kleinberg model does not (try to) reproduce the basic topological properties of social networks through which messages were traveling in Milgram’s experiments. For instance, the model produces only k -regular graphs while social networks, the Internet, and many other complex networks [23] are known to be *scale-free*, meaning that: i) the distribution $P(k)$ of node degrees k in a network follows power laws $P(k) \sim k^{-\gamma}$ with exponent γ often lying between 2 and 3; and ii) the network has strong clustering, i.e., a large number of triangular subgraphs [9].

Our work follows Kleinberg’s formalism. We assume that nodes in complex networks exist in some spaces that underlie the observed network topologies. We call these spaces *hidden metric spaces*. The observed network topology is coupled to the hidden space geometry in the following way: a link between two nodes in the topology exists with a certain probability that depends on the distance between the two nodes in the hidden geometry. One possible and plausible explanation for the Kleinberg model’s inability to naturally produce scale-free topologies is that the spaces hidden beneath such topologies are not Euclidean planes.

A primary contribution of this paper is the demonstration that a simple embedding of dynamically evolving networks in a *hyperbolic hidden metric space* naturally leads to the emergence of scale-free topologies. One attractive property of such

topologies is that greedy forwarding based simply on node coordinates in the metric space results in 100% reachability with near optimal path lengths even under dynamic network conditions, with link failures and node arrivals and departures. Our work raises the possibility of achieving near-optimal forwarding in complex networks, e.g. like the Internet, without the need for continuously running potentially expensive and brittle routing protocols to discover global topology knowledge. Our results may also find several practical applications in overlay networks for tasks such as application-level routing, information sharing, and data distribution.

The rest of the paper is organized as follows. In Section II we discuss related work. In Section III we reveal the connection between scale-free network topologies and hyperbolic geometries, and present a simple model that builds scale-free networks using such geometries. The whole topology of our networks in this section is constructed at once. We further demonstrate the remarkable efficiency of greedy forwarding strategies and their robustness in dynamic scenarios with link failures. As in many practical applications nodes may arrive to the system gradually, in Section IV we extend our first model and present a second model of scale-free networks that grow in hyperbolic spaces. We demonstrate that greedy forwarding strategies are still extremely efficient, even under highly dynamic network conditions where nodes randomly arrive and depart the system. We discuss practical applications of our findings in Section V, and conclude with directions for future research in Section VI.

II. RELATED WORK

The most relevant earlier work is the groundbreaking result by Robert Kleinberg, who shows in [16] how any given graph can be embedded in the hyperbolic plane such that greedy forwarding can achieve 100% reachability. However, to construct the embedding one needs to know the graph topology in advance. R. Kleinberg’s work has been recently complemented by the work in [7], where the authors propose a simple technique for online embedding of any given graph in the hyperbolic plane. The graph can also be dynamic, in the sense that once the initial graph is embedded, its topology can change, and greedy forwarding strategies can still be efficient.

While [16] and [7] show how any *given* graph can be embedded in a hyperbolic space so that greedy forwarding is efficient, here we approach the problem from the opposite angle. We fix the hyperbolic space and construct graphs in it in the simplest possible manner. We show that the resulting graphs are *not any* graphs, but are scale-free graphs. We further show that such graphs are *naturally* congruent with the underlying hyperbolic geometry, i.e., we do not enforce them by hand to be congruent, and because of this congruency, greedy forwarding strategies are maximally efficient. In terms of practical applications, while the studies in [16] and [7] are more suitable for situations where the initial network topology is already known/formed, our work here is more applicable to cases where networks are formed dynamically, where we do not know their exact topology in advance, such as in overlay network applications, [19], [21], and to cases where such global topology awareness may be prohibitive in terms of associated routing overhead costs. See Section V for a discussion of potential applications.

As mentioned earlier, the first popularization of greedy routing as a mechanism that might be responsible for efficient forwarding “in the dark”, i.e., without the knowledge of network topology, is due to Jon Kleinberg [14]. A vast amount of literature followed this seminal work, as reviewed in [15]. Other works, dealing with hyperbolic geometry in the network context, include [17], [1], [13]. No earlier work has considered the connection between complex networks and hidden hyperbolic geometries, which can be used to efficiently guide the forwarding process on them.

Finally, there has been a great deal of research trying to explain the scale-free structure of complex networks [4], among which preferential attachment [2] appears to be the most popular. However, no existing effort has considered hyperbolic hidden metric spaces as a possible explanation.

III. SCALE-FREE NETWORKS AND HYPERBOLIC SPACES

We start this section by first giving a high-level intuition for the connection between scale-free network topologies and hyperbolic geometries. We then proceed by presenting a simple model where scale-free topologies naturally emerge from such geometries, and demonstrate the remarkable efficiency of greedy forwarding strategies that use these geometries.

A. Intuition

The main metric property of hyperbolic geometry that we use in this paper is the exponential expansion of space. For example, in the hyperbolic plane, which is the two-dimensional hyperbolic space of negative curvature -1 , the length of a circle and the area of a disc of radius R are $2\pi \sinh R$ and $2\pi(\cosh R - 1)$, both growing as $\sim e^R$ with R .¹ The hyperbolic plane is thus metrically equivalent to an e -ary tree, i.e., a tree with the average branching factor equal to e . Indeed, in a b -ary tree, the analogies of the circle length or disc area are the number of nodes at distance exactly R or not more than R hops from the root. These numbers are $(b + 1)b^{R-1}$ and $((b + 1)b^R - 2)/(b - 1)$, both growing as $\sim b^R$. Informally, hyperbolic spaces can therefore be thought of as “continuous versions” of trees. Other properties of hyperbolic geometry can be found in various (text)books, e.g., [5].

¹In this paper, symbols ‘ \sim ’ and ‘ \approx ’ mean, respectively, *proportional to* and *approximately equal*.

To see why this exponential expansion of hidden space is intrinsic to scale-free networks, observe that their topology represents the structure of connections or interactions among distinguishable, heterogeneous elements abstracted as nodes. The heterogeneity implies that nodes can be somehow classified, however broadly, into a taxonomy, i.e., nodes can be split into large groups consisting of smaller subgroups, which in turn consist of even smaller subsubgroups, etc. The relationships between such groups and subgroups, called communities [10], can be approximated by tree-like structures, in which the distance between two nodes estimates how similar they are [26], [6]. The smaller their distance, the more similar the two nodes are, and the more likely they are connected. Importantly, the node classification hierarchy need not be strictly a tree. Approximate “tree-ness,” which can be formally expressed solely in terms of the metric structure of a space [12], makes the space hyperbolic.²

B. Models of scale-free networks in hyperbolic spaces

We now put our intuitive considerations to qualitative grounds. We want to see what network topologies emerge in the simplest possible settings involving hidden hyperbolic metric spaces. Specifically, we use the following strategy to formulate a network model. We specify: 1) the hyperbolic space; 2) the distribution of nodes in it, i.e., the node density; and 3) the connection probability as a function of the hyperbolic distance between nodes, i.e., we connect a pair of nodes located at hyperbolic distance x with some probability $p(x)$.

The simplest hyperbolic space is the two-dimensional hyperbolic plane of constant negative curvature -1 we discussed earlier. The simplest way to place N nodes on the hyperbolic plane is to distribute them uniformly over a disc of radius R . The hyperbolically uniform node density implies that we assign the angular coordinates $\theta \in [0, 2\pi]$ to nodes with the uniform density $f(\theta) = 1/(2\pi)$, while the density for the radial coordinate $r \in [0, R]$ is exponential $f(r) = \sinh r / (\cosh R - 1) \approx e^{r-R} \sim e^r$, as the circle length at distance r from the disc center is $2\pi \sinh r$ (vs. $f(r) \sim r$ in the Euclidean plane, where the circle length is $2\pi r$). We can also generalize the model by distributing nodes non-uniformly on the disc using:

$$f(r) = \frac{\alpha \sinh \alpha r}{\cosh \alpha R - 1} \approx \alpha e^{\alpha(r-R)} \sim e^{\alpha r}, \quad (1)$$

with $\alpha = 1$ corresponding to the hyperbolically uniform node density.

The simplest connection probability we could think of is the step function $p(d) = \Theta(R - d)$, meaning that we connect a pair of nodes with polar coordinates (r, θ) and (r', θ') by a link only if the hyperbolic distance between them is $d \leq R$, where d is given by the hyperbolic law of cosines: $\cosh d = \cosh r \cosh r' - \sinh r \sinh r' \cos \Delta\theta$, with $\Delta\theta = \min(|\theta - \theta'|, 2\pi - |\theta - \theta'|)$. The following theorem states that the node degree distribution in the resulting network is a power law.

Theorem 1: The described model produces graphs with power law node degree distribution:

$$P(k) \sim k^{-\gamma}, \quad \text{with } \gamma = \begin{cases} 2\alpha + 1 & \text{if } \alpha \geq \frac{1}{2}, \\ 2 & \text{if } \alpha \leq \frac{1}{2}. \end{cases} \quad (2)$$

Proof:

First note that the volume form in polar coordinates (r, θ) in the two-dimensional hyperbolic plane of constant negative curvature -1 is $d \text{vol} = \sinh r \, dr \, d\theta$. (This gives the length of a circumference of radius r : $v(r) = 2\pi \sinh r$. And, the surface of the disc $B(r)$ of radius r : $|B(r)| = \int_{B(r)} d \text{vol} = \int_0^r \sinh x \, dx \int_0^{2\pi} d\theta = 2\pi(\cosh r - 1)$.)

We first compute the average degree $\bar{k}(r)$ of nodes located at distance r from our hyperbolic disc center. Let's consider the general case, where the connection probability is the step function $p(d) = \Theta(\chi - d)$, meaning that we connect a pair of nodes with polar coordinates (r, θ) and (r', θ') by a link only if the hyperbolic distance between them is $d \leq \chi$ where χ is some constant. First, let $\alpha = 1$. Since nodes in this case are distributed uniformly in the disc, the average degree of a node at radius r is:

$$\bar{k}(r) = \delta S(r),$$

where δ is the node density, i.e. $\delta = \frac{N}{2\pi(\cosh R - 1)}$, and $S(r)$ is the area of the intersection of discs of radii R and χ whose centers are located at distance r from each other. Let us call these discs the R - and χ -disc respectively.

To compute $S(r)$, we move the origin from the center of the R -disc to the center of the χ -disc. If $r \leq R - \chi$, then the χ -disc is fully contained in the R -disc and thus $S(r) = 2\pi(\cosh \chi - 1)$. If $R - \chi \leq r \leq R$, then using the notations in Figure 1, we get:

$$\begin{aligned} S(r) &= 2 \left\{ \int_0^{R-r} \sinh x \, dx \int_0^\pi d\theta + \int_{R-r}^\chi \sinh x \, dx \int_0^{\theta_x} d\theta \right\} \\ &= 2 \left\{ \pi(\cosh(R-r) - 1) + \int_{R-r}^\chi \sinh x \arccos \frac{\cosh r \cosh x - \cosh R}{\sinh r \sinh x} \, dx \right\}, \end{aligned} \quad (3)$$

²We call the space *hidden* to emphasize that the distance between two nodes in it is a measure of how similar they are; it is *not* their shortest path distance in the observable network graph.

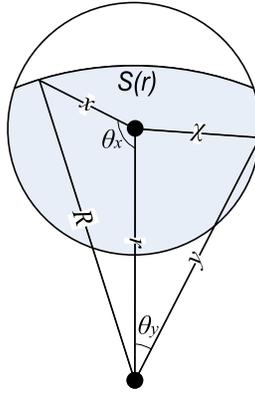


Fig. 1. Disc intersection.

where $\chi \in [0, R]$, while $\theta_x \in [0, \pi]$ is the solution of the following equation:

$$\cosh R = \cosh r \cosh x - \sinh r \sinh x \cos \theta_x.$$

The integral $\int \sinh x \arccos \frac{\cosh r \cosh x - \cosh R}{\sinh r \sinh x} dx$ can be easily computed (e.g., using substitution $t \equiv \cosh x$). The only subtlety is that the resulting expression has a singularity at the lower integration bound $R - r$. However, its right limit exists and equals to $\pi(\cosh(R - r) - \frac{1}{2} \cosh R)$. After tedious simplifications, the final answer reads:

$$\begin{aligned} \bar{k}(r) = & \delta \left\{ (\cosh R - 2)\pi + 2 \left(\cosh \chi \arccos \frac{\cosh r \cosh \chi - \cosh R}{\sinh r \sinh \chi} \right. \right. \\ & + \left. \left. \cosh R \arctan \frac{\cosh \chi - \cosh R \cosh r}{\sqrt{(\cosh r - \cosh(R - \chi))(\cosh(R + \chi) - \cosh r)}} \right) \right. \\ & \left. - \arctan \frac{8(\cosh r - \cosh R \cosh \chi) \sqrt{(\cosh r - \cosh(R - \chi))(\cosh(R + \chi) - \cosh r)}}{16(\cosh r - \cosh R \cosh \chi) \cosh r - 8 \cosh^2 r + 4(\cosh^2 R + \cosh^2 R \cosh^2 \chi + \cosh^2 \chi) - 1} \right\}. \end{aligned} \quad (4)$$

It is easy to see that only the second and third terms (the ones in the big round braces) can strongly depend on r and cause an exponential decrease of $\bar{k}(r)$. The first term is a large constant and the absolute value of the last term is bounded by $\pi/2$, so that we can safely neglect it.

One can now easily estimate the asymptotic form of $\bar{k}(r)$ in Eq. (4). Indeed, neglecting the first and last terms in the expression, we are interested only in the scaling of:

$$\begin{aligned} \xi &= \arccos \frac{\cosh r \cosh \chi - \cosh R}{\sinh r \sinh \chi}, \\ \eta &= \arctan \frac{\cosh \chi - \cosh R \cosh r}{\sqrt{(\cosh r - \cosh(R - \chi))(\cosh(R + \chi) - \cosh r)}}, \end{aligned}$$

as functions of r for $R - \chi \ll r \ll R$. Note that such region of values of r exists only if R is relatively large and χ is not too small, but comparable to R , $\chi \lesssim R$.

We first consider the more complicated second term η . In the denominator, we can neglect $\cosh(R - \chi)$ compared to $\cosh r$, and $\cosh r$ compared to $\cosh(R + \chi)$:

$$\eta \approx \arctan \frac{\cosh \chi - \cosh R \cosh r}{\sqrt{\cosh(R + \chi) \cosh r}}.$$

Since $\chi \lesssim R$, we next neglect the first term, which is infinitesimally small compared to the second term when r is large:

$$\eta \approx \tilde{\eta} = \arctan \left(-a \sqrt{\cosh r} \right).$$

where $a = \cosh R / \sqrt{\cosh(R + \chi)}$. But:

$$(\ln \tilde{\eta})'_r = \frac{a \sinh r}{2 \arctan \left(a \sqrt{\cosh r} \right) \sqrt{\cosh r} (1 + a^2 \cosh r)} \xrightarrow{r \rightarrow \infty} 0,$$

which means that for large r :

$$\eta \approx e^{0 \cdot r},$$

independently of a and, consequently, of R and χ .

The first term ξ is simpler. Since both r and χ are large, it is:

$$\xi \approx \tilde{\xi} = \arccos \left(1 - \frac{a}{\sinh r} \right),$$

where $a = \cosh R / \sinh \chi$. But:

$$(\ln \tilde{\xi})'_r = - \frac{a \cosh r}{\arccos \left(1 - \frac{a}{\sinh r} \right) \sinh^2 r \sqrt{1 - \left(1 - \frac{a}{\sinh r} \right)^2}} \xrightarrow{r \rightarrow \infty} -\frac{1}{2},$$

which means that for large r :

$$\xi \approx e^{-\frac{1}{2}r},$$

independently of a and, consequently, of R and χ .

Summing it all up, we thus observe that in the case with the uniform node density, the average node degree $\bar{k}(r) \sim e^{-\beta r}$ with $\beta = 1/2$.

If the node density is not uniform, then exact computations become rather complicated because we cannot easily move the origin to the center of the χ -disc. It appears to be easier to keep the origin at the center of the R -disc, but we have to retreat to approximate expressions.

Denote the node density by $f(y)$, where y is the distance from the center of the R -disc (to distinguish it from the distance from the center of the χ -disc, see Fig. 1). Keeping the origin at the center of the R -disc somewhat complicates the disc intersection picture. We assume that $\chi > R/2$, which means that it is never the case that the χ -disc: 1) is fully contained in the R -disc and, at the same time, 2) does not contain its center. The remaining three cases are:

1) $0 \leq r \leq R - \chi$: the χ -disc: 1) is fully contained in the R -disc and 2) contains its center:

$$\bar{k}_1(r) = N \left\{ \int_0^{\chi-r} f(y) dy + \frac{1}{\pi} \int_{\chi-r}^{\chi+r} f(y) \theta_y dy \right\}; \quad (5)$$

2) $R - \chi \leq r \leq \chi$: the χ -disc: 1) is not fully contained in the R -disc but 2) still contains its center:

$$\bar{k}_2(r) = N \left\{ \int_0^{\chi-r} f(y) dy + \frac{1}{\pi} \int_{\chi-r}^R f(y) \theta_y dy \right\}; \quad (6)$$

3) $\chi \leq r \leq R$: the χ -disc: 1) is not fully contained in the R -disc and 2) does not contain its center (the case shown in Fig. 1):

$$\bar{k}_3(r) = \frac{N}{\pi} \int_{r-\chi}^R f(y) \theta_y dy. \quad (7)$$

All the three cases can be compactly written in one expression:

$$\bar{k}(r) = N \left\{ \int_0^{\max(0, \chi-r)} f(y) dy + \frac{1}{\pi} \int_{|\chi-r|}^{\min(R, \chi+r)} f(y) \theta_y dy \right\}, \quad (8)$$

where

$$\theta_y = \arccos \frac{\cosh r \cosh y - \cosh \chi}{\sinh r \sinh y}.$$

For large r , y , and χ , we have the following approximation for θ_y :

$$\theta_y \approx \arccos \left(1 - 2a^2 e^{-y} \right) \approx 2a e^{-\frac{y}{2}},$$

where

$$a = \sqrt{\frac{\cosh \chi}{\sinh r}} \approx e^{\frac{\chi-r}{2}}.$$

We are now ready to compute the approximate scaling of $\bar{k}(r)$ for the exponential node density:

$$f(y) \approx \alpha e^{\alpha(y-R)}.$$

Substituting

$$\int f(y) dy \approx e^{\alpha(y-R)},$$

$$\int f(y)\theta_y dy \approx \frac{2\alpha}{\alpha - \frac{1}{2}} e^{\frac{\chi-r}{2} - \alpha R} e^{(\alpha - \frac{1}{2})y},$$

into Eqs.(5-7), we get:

$$\bar{k}_1(r) \approx N \left\{ e^{\alpha(\chi-R-r)} + \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}} e^{\frac{\chi-r}{2} - \alpha R} \left(e^{(\alpha - \frac{1}{2})(\chi+r)} - e^{(\alpha - \frac{1}{2})(\chi-r)} \right) \right\}, \quad (9)$$

$$\bar{k}_2(r) \approx N \left\{ e^{\alpha(\chi-R-r)} + \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}} e^{\frac{\chi-r}{2} - \alpha R} \left(e^{(\alpha - \frac{1}{2})R} - e^{(\alpha - \frac{1}{2})(\chi-r)} \right) \right\}, \quad (10)$$

$$\bar{k}_3(r) \approx N \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}} e^{\frac{\chi-r}{2} - \alpha R} \left(e^{(\alpha - \frac{1}{2})R} - e^{(\alpha - \frac{1}{2})(r-\chi)} \right). \quad (11)$$

We thus see that $\bar{k}(r)$ is always an exponential function of r and the values of the exponents do not depend on R or χ . Specifically, extracting the leading terms (i.e., the terms with the largest exponent), we get:

- if $\alpha \geq 1/2$, then:

$$\begin{aligned} \bar{k}_1(r) &\sim e^{(\alpha-1)r}, \\ \bar{k}_2(r) &\sim e^{-\frac{1}{2}r}, \\ \bar{k}_3(r) &\sim e^{(\alpha-1)r}; \end{aligned}$$

- if $\alpha \leq 1/2$, then:

$$\begin{aligned} \bar{k}_1(r) &\sim e^{-\alpha r}, \\ \bar{k}_2(r) &\sim e^{-\alpha r}, \\ \bar{k}_3(r) &\sim e^{-\frac{1}{2}r}. \end{aligned}$$

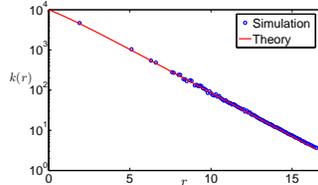
Since in our case $\chi \lesssim R$, the length $R - \chi$ of the first and last intervals (corresponding to $\bar{k}_1(r)$ and $\bar{k}_3(r)$) of values of r is small, and $\bar{k}_2(r)$ dominates the overall behavior of $\bar{k}(r)$. We finally obtain that:

$$\bar{k}(r) \sim \begin{cases} e^{-\frac{1}{2}r} & \text{if } \alpha \geq \frac{1}{2}, \\ e^{-\alpha r} & \text{if } \alpha \leq \frac{1}{2}. \end{cases}$$

Therefore, when $\chi = R$, then the regimes $\bar{k}_1(r)$ and $\bar{k}_3(r)$ in Eqs. (9,11) vanish and the overall scaling of $\bar{k}(r)$ is given by $\bar{k}_2(r)$ in Eq. (10). Therefore, $\bar{k}(r)$ simplifies to:

$$\bar{k}(r) \approx N \left\{ \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}} e^{-\frac{1}{2}r} + \left(1 - \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}} \right) e^{-\alpha r} \right\}. \quad (12)$$

We see that in this case, $\bar{k}(r)$ does not depend on R , which is not surprising, of course. We verify that the above expression perfectly matches simulations:



The limit $\alpha \rightarrow 1/2$ is well-defined:

$$\bar{k}(r) \xrightarrow{\alpha \rightarrow 1/2} N \left(1 + \frac{r}{\pi} \right) e^{-\frac{1}{2}r}.$$

Therefore, $\bar{k}(r)$ decreases exponentially, i.e. $\bar{k}(r) \sim e^{-\beta r}$, with $\beta = \frac{1}{2}$ if $\alpha \geq \frac{1}{2}$ and $\beta = \alpha$ if $\alpha \leq \frac{1}{2}$. Hence, $\bar{r}(k) \sim -\frac{1}{\beta} \ln k$. Given $f(r)$ as in Equation (1), it is easy to see that $P(k) \approx f(\bar{r}(k)) |\bar{r}'(k)| \sim k^{-\gamma}$ with $\gamma = \frac{\alpha}{\beta} + 1$. Thus, $\gamma = 2\alpha + 1$ if $\alpha \geq \frac{1}{2}$ and $\gamma = 2$ if $\alpha \leq \frac{1}{2}$. ■

We thus see that by changing α , which according to our tree analogy regulates the average branching factor of the hidden tree-like hierarchy, we can construct power-law graphs with any exponent $\gamma \geq 2$, as observed in a majority of known complex networks, including the Internet [9].

Given $f(r)$ from Equation (1) and $\bar{k}(r)$ from Equation (12), the average node degree $\bar{k} = \int_0^R \bar{k}(r) f(r) dr$, is:

$$\bar{k} \approx N \frac{\left\{ 2\alpha^2 e^{-\frac{1}{2}R} + \left[((\pi - 2)\alpha^3 - (\pi - 1)\alpha^2 + \frac{\alpha\pi}{4}) R - 2\alpha^2 \right] e^{-\alpha R} \right\}}{\pi \left(\alpha - \frac{1}{2} \right)^2},$$

where the limit $\alpha \rightarrow \frac{1}{2}$ is $\bar{k} \rightarrow N \frac{R}{2} \left(1 + \frac{R}{2\pi} \right) e^{-\frac{1}{2}R}$. Therefore, given a target \bar{k} , a target exponent γ , which is related to α via Equation (2), and the number of nodes N , the right value for the hyperbolic disc radius R is (numerically) computed by the above formula. From the formula, we can also see that the relationship between R and N is approximately logarithmic, i.e., $R \sim \ln N$.

Notice that while Theorem 1 states that the degree distribution is a power law, it does not give its expression. This is given in the following Theorem.

Theorem 2: The described model produces graphs with the power law node degree distribution

$$P(k) = 2\alpha\xi^{2\alpha} \frac{\Gamma(k - 2\alpha, \xi)}{k!}, \quad (13)$$

where $\xi = \bar{k}(2\alpha - 1)/(2\alpha)$ and Γ is the incomplete gamma function.

Proof: To find the expression of the degree distribution $P(k)$ we use the hidden variable approach described in [3]. In particular, from Equation (11) in [3], we can write:

$$P(k) = \int_0^R g(k|r) f(r) dr.$$

where $g(k|r)$ is the probability that a node with radial coordinate r is connected to other k nodes. For sparse (scale-free) networks $g(k|r)$ has the Poisson distribution, see Equation (23) in [3]:

$$g(k|r) = \frac{e^{-\bar{k}(r)} \bar{k}(r)^k}{k!},$$

where $\bar{k}(r)$ the average node degree at distance r from the disc center, as defined and computed earlier. Therefore:

$$P(k) \approx \int_0^R \frac{e^{-\bar{k}(r)} \bar{k}(r)^k}{k!} \alpha e^{\alpha(r-R)} dr. \quad (14)$$

We treat the more interesting case $\alpha > \frac{1}{2}$ where we can create networks with various exponents $\gamma > 2$. (The limiting case $\gamma \rightarrow 2$ is well defined.) In this case, from Equation (12) we can write:

$$\bar{k}(r) \approx N \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}} e^{-\frac{1}{2}r}, \quad r \approx -2 \ln \frac{\bar{k}(r)}{N \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}}}.$$

Now, let's use the substitution $t \equiv \bar{k}(r)$ in Equation (14). Notice that:

$$\frac{dt}{dr} = -\frac{1}{2} \bar{k}(r) = -\frac{1}{2} t.$$

Thus, we have:

$$P(k) \approx -2\alpha \left(N \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}} e^{-\frac{1}{2}R} \right)^{2\alpha} \frac{1}{k!} \int t^{k-2\alpha-1} e^{-t} dt,$$

where the limits of the above integration are going to be defined shortly. Now, using $\bar{k}(r)$ from above, the average degree \bar{k} is approximated by:

$$\bar{k} \approx \int_0^R N \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}} e^{-\frac{1}{2}r} \alpha e^{\alpha(r-R)} dr \approx \alpha N \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}} \frac{1}{\alpha - \frac{1}{2}} e^{-\frac{1}{2}R}.$$

Hence:

$$P(k) \approx -2\alpha \left(\bar{k} \frac{2\alpha - 1}{2\alpha} \right)^{2\alpha} \frac{1}{k!} \int t^{k-2\alpha-1} e^{-t} dt.$$

Now, when $r = 0$, $t = \bar{k}(0) = N \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}}$. When $r = R$, $t = \bar{k}(R) = N \frac{2}{\pi} \frac{\alpha}{\alpha - \frac{1}{2}} e^{-\frac{1}{2}R} = \bar{k} \frac{2\alpha - 1}{2\alpha}$. Let $\xi \equiv \bar{k} \frac{2\alpha - 1}{2\alpha}$. Performing the integration with the N large limit (technically letting $N \rightarrow \infty$), we get:

$$P(k) \approx 2\alpha \xi^{2\alpha} \frac{1}{k!} \int_{\xi}^{\infty} t^{k-2\alpha-1} e^{-t} dt = 2\alpha \xi^{2\alpha} \frac{\Gamma(k - 2\alpha, \xi)}{k!}. \quad (15)$$

■

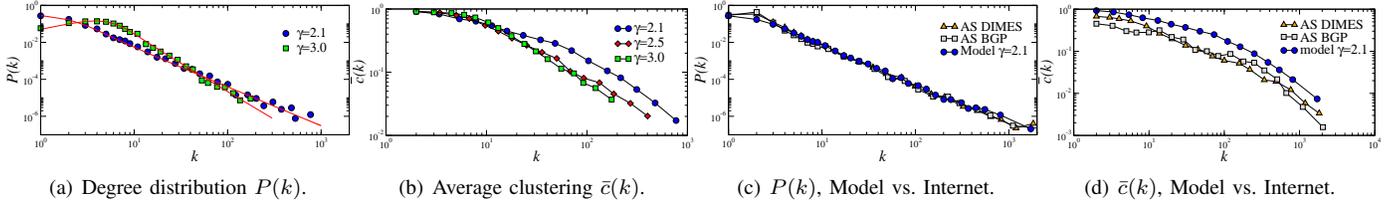


Fig. 2. The first two plots show the degree distribution $P(k)$, and average clustering $\bar{c}(k)$ of k -degree nodes. The degree distribution for $\gamma = 2.5$ is not shown for clarity. Solid lines in the first plot are the theoretical prediction given by Equation (13). The last two plots show the same statistics for simulated networks with $\gamma = 2.1$ vs. AS topologies from RouteViews BGP tables [24] and DIMES traceroute data [8].

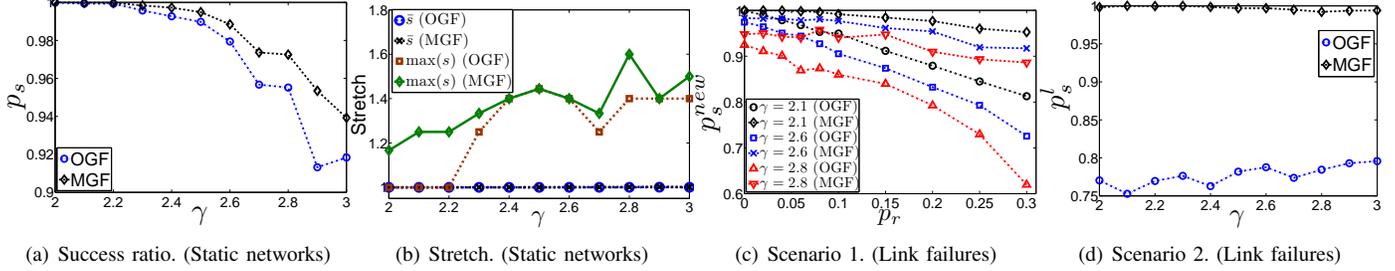


Fig. 3. Performance of greedy forwarding (GF).

Our networks also possess strong clustering. Strong clustering, or large numbers of triangles in generated networks, is a simple consequence of the triangle inequality in the metric space. Indeed, if node a is close to node b in the plane, and b is close to a third node c , then a is also close to c because of the triangle inequality. Since all three nodes are close to each other, links between all of them forming triangle abc exist.

In Figures 2(a) and 2(b) we show, in log-log scale, the degree distribution $P(k)$ and average clustering as a function of the node degree $\bar{c}(k)$ [27], for modeled networks with $N = 10000$ and $\bar{k} = 6.5$. We observe agreement between simulation results and the analytical prediction for the degree distribution in Equation (13). In Figures 2(c) and 2(d) we compare the same statistics between the above modeled networks with $\gamma = 2.1$ and the AS Internet topologies from RouteViews BGP tables [24] and DIMES traceroute data [8]. The degree distribution in our networks is remarkably close to the empirical AS degree distribution. Further, interestingly enough, the shape of the clustering curve $\bar{c}(k)$ for our networks is similar to the Internet's. In [18] we also show how clustering can be matched exactly.

C. Greedy forwarding

We now evaluate the performance of greedy forwarding (GF) strategies on our modeled networks. A node's address is its hyperbolic coordinates, and each node knows only its own address, the addresses of its neighbors, and the destination address written in the packet. GF forwards a packet at each hop to the neighbor closest to the destination in the hyperbolic space. We present simulation results for two simple forms of GF, *original* (OGF) and *modified* (MGF). The OGF algorithm drops the packet if the current hop is a *local minimum*, meaning that it does not have any neighbor closer to the destination than itself. The MGF algorithm excludes the current hop from any distance comparisons, and finds the neighbor closest to the destination. The packet is dropped only if this neighbor is the same as the packet's previous hop. We report the following metrics: (i) the percentage of successful paths, p_s , which is the proportion of paths that reach their destinations; and (ii) the average and maximum stretch of successful paths, denoted by \bar{s} and $\max(s)$ respectively. The stretch is defined as the ratio between the hop-lengths of greedy paths and the corresponding shortest paths in the graph.

We initially focus on static networks, where the network topology does not change, and then emulate network topology dynamics by randomly removing one or more links from the topology. As before, we fix the target number of nodes in the network to $N = 10000$ and its average degree to $\bar{k} = 6.5$, which roughly matches the Internet's AS topology. For each generated network, we extract the Giant Connected Component (GCC), and perform GF between 10000 random source-destination pairs.

Static networks. Figures 3(a) and 3(b) show the results for static networks of different degree exponent γ . We see that the success ratio p_s increases and the stretch decreases as we decrease γ to 2. For example, for $\gamma = 2.1$, i.e., equal to γ observed in the AS Internet, OGF and MGF yield $p_s = 0.99920$ and $p_s = 0.99986$, with the OGF's maximum stretch of 1, meaning that *all greedy paths are shortest paths*. In summary, GF is exceptionally efficient in static networks, especially for the small γ 's observed in the vast majority of complex networks [9]. The two GF algorithms yield high success ratios close to 1 and optimal (or almost optimal) path lengths, i.e., stretch close to 1. The reason for this remarkable GF performance is the congruency between the network topology and the underlying hyperbolic geometry, as visually demonstrated in Figure 4.

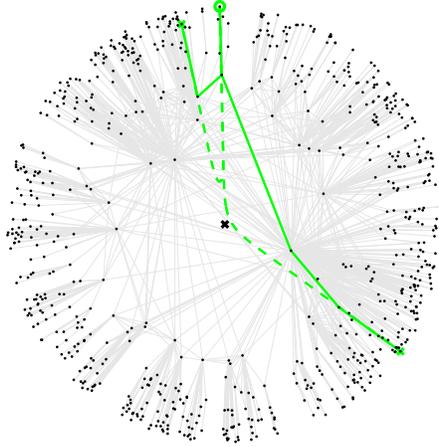


Fig. 4. Visualization of a modeled network embedded in the hyperbolic plane, and greedy forwarding in it. The figure shows two hyperbolically straight lines, i.e., geodesics, the dashed curves, vs. the greedy paths, the solid curves, between the same source-destination nodes (the source is the top circled node and the destinations are marked by a cross). The geodesics and greedy paths follow the same pattern, i.e., they are approximately congruent.

Link failures. We next study the GF performance in dynamic scenarios with link failures. We consider the following two link-failure scenarios. In Scenario 1 we remove a percentage p_r , ranging from 0% to 30%, of all links in the network, recompute the GCC, and compute the new success ratio p_s^{new} . In Scenario 2, we provide a finer-grain view focusing on paths that used a removed link. We remove one link from the network, recompute and GCC, and find the percentage p_s^l of successful paths, only among those previously successful paths that traversed the removed link and belong to the GCC. We repeat the procedure for 1000 random links, and report the average value for p_s^l . Figures 3(c) and 3(d) present the results. We see that for small γ 's, the success ratio p_s^{new} remains remarkably high, for all meaningful values of p_r . For example, MGF on networks with $\gamma = 2.1$ and $p_r \leq 0.1$, yields $p_s^{new} > 0.99$. Note that the simultaneous failure of 10% of the links in networks like the Internet is a rare catastrophe, but even in this case GF is still efficient. The percentage p_s^l of MGF paths that used a removed link and that found a by-pass after its removal is also remarkably high, close to 100% for small γ 's. In both scenarios and for all γ , the average stretch remains remarkably low, below 1.1.

In summary, GF is not only efficient in static networks, but its efficiency can be also very robust in the presence of network topology dynamics. Thanks to high path diversity, there are many shortest paths, disjoint over some links or nodes, between the same source and destination, which all closely follow their geodesics. Link removals affect some shortest paths, but others remain, and greedy forwarding can use the underlying hyperbolic “guidance system” to find them.

More sophisticated greedy forwarding strategies. Notice that although the success ratios in scale-free networks with small γ 's are extremely close to 1, they are not exactly 1. Further, the performance of our simple GF strategies deteriorates very slowly in the presence of topology dynamics. These observations suggest that more sophisticated, yet still simple GF strategies, may be able to achieve 100% success ratio preserving the extremely small stretch even under more extreme dynamic network conditions than the ones we studied above. We will explore such a strategy in the next section, in scenarios where nodes randomly arrive and depart the network.

IV. NETWORKS GROWING IN HYPERBOLIC SPACES

The model we presented in the previous section generates a whole network at once. However, in many applications, the network topology is formed by nodes gradually arriving over time. In this section, we extend our previous model and present a model of scale-free networks that grow in hyperbolic spaces. We then demonstrate the remarkable efficiency of greedy forwarding in highly dynamic conditions, where nodes randomly arrive and depart the system.

A. Growing model

We assume that the network initially consists of 0 nodes. We number each arriving node by its order of arrival. We do not consider node departures for now. A new node $i \geq 1$ that arrives to the system needs to know: (i) the current number of nodes in the network, including itself, $N(i) = i$; (ii) a system pre-specified parameter α for the node radial density, which as before, will determine the exponent of the degree distribution γ ; and (iii) a system pre-specified constant c , which will determine the average node degree as will be explained below. Then, to connect to the network, the node performs the following operations inspired by the model in Section III:

- i. select an angular coordinate θ uniformly distributed in $[0, 2\pi]$;

- ii. compute the current hyperbolic disk radius $R(i)$ according to $R(i) = \frac{1}{\alpha} \ln \frac{i}{c}$, i.e., $i = ce^{\alpha R(i)}$; ³
- iii. select a radial coordinate $r \in [0, R(i)]$, according to the probability density function $f(r|R(i)) = \frac{\alpha \sinh(\alpha r)}{\cosh(\alpha R(i)) - 1} \approx \alpha e^{\alpha(r-R(i))}$;
- iv. connect to every node $1 \leq j < i$, already in the network, for which the hyperbolic distance to it, denoted by d_{ij} , satisfies $d_{ij} \leq R(i)$. ⁴

Notice that to build the network in a fully decentralized manner, each arriving node needs to be able to *discover* the current number of nodes in it, and the nodes to connect to, i.e. its neighbors. We will present a technique for this later. Before doing so, we first need to analyze the statistical characteristics of the resulting network topologies and show that they are indeed scale-free.

As before, strong clustering is easily justified by the triangle inequality in the hidden space. However, we need to show that if we look at the network at any time instance, where the number of nodes in it is some (sufficiently large) value, say t , the node degree distribution follows a power law. The analysis becomes fairly more complicated than the one in Section III, as the hyperbolic disc radius is no longer constant, but grows as the number of nodes in the network grows. To proceed, we first need to compute the probability density function of the radial coordinate r of a randomly selected node when the number of nodes in the network is some value t , denoted by $f(r, t)$, as well as the average node degree at distance r from the disc center, denoted by $\bar{k}(r, t)$. Having the expressions for $f(r, t)$, $\bar{k}(r, t)$, we can then easily identify how these quantities scale as a function of r , and use arguments similar to those in the proof of Theorem 1 to show that the degree distribution is a power law. The intermediate calculations for the derivation of $f(r, t)$ and $\bar{k}(r, t)$ are given in the Appendix. Below, we present the final expressions and show that the degree distribution is indeed power law.

Let $R(t)$ be the hyperbolic disc radius when the number of nodes in the network is t . According to the model, $R(t) = \frac{1}{\alpha} \ln \frac{t}{c}$. The approximate, simplified expression for the node radial density $f(r, t)$ is:

$$f(r, t) \approx \alpha^2 (R(t) - r) e^{\alpha(r-R(t))}. \quad (16)$$

It is easy to see from the above formula that $f(r, t)$ scales according to the following lemma.

Lemma 1: For r not close to $R(t)$ and t sufficiently large, $f(r, t) \sim e^{\alpha r}$. ⁵

Proof: For r not close to $R(t)$, we can write:

$$\begin{aligned} f(r, t) &\approx \alpha^2 R(t) \left(1 - \frac{r}{R(t)}\right) \left(1 - \frac{r}{R(t)}\right)^{\alpha \frac{R(t)}{r} (R(t)-r)} \\ &= \alpha^2 R(t) e^{-\alpha R(t)} e^{(\alpha - \frac{1}{R(t)})r} \sim e^{(\alpha - \frac{1}{R(t)})r} \approx e^{\alpha r}, \end{aligned}$$

where the last approximation holds for sufficiently large t . ■

For $\alpha > \frac{1}{2}$, the average node degree as a function of r and t , $\bar{k}(r, t)$, is:

$$\bar{k}(r, t) \approx t \left\{ P(r) + \frac{1 - e^{-\alpha(R(t)-r)}}{\alpha(R(t)-r)} (G(r) - P(r)) \right\}, \quad (17)$$

where:

$$G(r) = \left(1 - \frac{2\alpha^2}{\pi(\alpha - \frac{1}{2})^2} + \left(1 - \frac{2\alpha}{\pi(\alpha - \frac{1}{2})}\right)\alpha r\right) e^{-\alpha r} + \frac{2\alpha^2}{\pi(\alpha - \frac{1}{2})^2} e^{-\frac{1}{2}r}, \quad (18)$$

$$P(r) = \frac{2\alpha}{\pi(\alpha - \frac{1}{2})} e^{-\frac{1}{2}r} - \frac{1}{2(\alpha - \frac{1}{2})} e^{-2\alpha \ln \frac{\pi}{2}} e^{-\alpha r}, \quad (19)$$

and the limit $\alpha \rightarrow \frac{1}{2}$ is well defined and easily computed from above. It is easy to see that $\bar{k}(r, t)$ scales as follows.

Lemma 2: For $\alpha \geq \frac{1}{2}$ and t sufficiently large, $\bar{k}(r, t) \sim e^{-\frac{1}{2}r}$.

³We can generalize the model by setting $R(i) = \delta \ln \frac{i}{c}$ for some constant $\delta \neq \frac{1}{\alpha}$. However, this generalization is out of the scope of the paper.

⁴If we wish to maintain a connected graph at any time instance, a node that picks coordinates for which $d_{ij} > R(i) \forall j$, can randomly re-select new ones, until $d_{ij} \leq R(i)$ for at least one j .

⁵It is easy to draw $f(r, t)$ and see that it increases exponentially with r , but as r approaches $R(t)$ there is a turning point, where it starts decreasing, reaching zero at $r = R(t)$. This however, turns out not to affect our construction of power-law graphs.

Proof: For $\alpha > \frac{1}{2}$ the scaling behavior of Equation (17) can be determined by looking at the leading terms, i.e., the terms multiplied by $e^{-\frac{1}{2}r}$. Therefore, we look at the sum:

$$\begin{aligned} & \left(\frac{2\alpha}{\pi(\alpha - \frac{1}{2})} e^{-\frac{1}{2}r} \right) + \frac{1 - e^{-\alpha(R(t)-r)}}{\alpha(R(t) - r)} \left(\frac{\alpha}{\pi(\alpha - \frac{1}{2})^2} e^{-\frac{1}{2}r} \right) \\ & \approx \left(\frac{2\alpha}{\pi(\alpha - \frac{1}{2})} \right) \frac{1}{e^{\frac{1}{2}r}} + \left(\frac{1}{\pi(\alpha - \frac{1}{2})^2} \right) \frac{1}{(R(t) - r)e^{\frac{1}{2}r}} \\ & \approx \left(\frac{2\alpha}{\pi(\alpha - \frac{1}{2})} \right) e^{-\frac{1}{2}r} + \left(\frac{1}{\pi(\alpha - \frac{1}{2})^2} \right) \frac{e^{-\left(\frac{1}{2} - \frac{1}{R(t)}\right)r}}{R(t)} \sim e^{-\left(\frac{1}{2} - \frac{1}{R(t)}\right)r} \approx e^{-\frac{1}{2}r}, \end{aligned}$$

where the last approximation holds for sufficiently large t . The same scaling behavior holds as $\alpha \rightarrow \frac{1}{2}$. \blacksquare

We are now ready to turn our attention to the node degree distribution, which is given by the following theorem.

Theorem 3: The described growing model with $\alpha \geq \frac{1}{2}$ produces graphs with approximately a power law node degree distribution $P(k, t) \sim k^{-\gamma}$, with $\gamma = 2\alpha + 1$.

Proof: The proof is similar to that of Theorem 1. For t sufficiently large, by Lemma 2, $\bar{k}(r, t) \sim e^{-\frac{1}{2}r}$. Hence, $\bar{r}(k, t) \sim -2 \ln k$. By Lemma 1, $f(r, t) \sim e^{\alpha r}$, approximately. Therefore, $P(k, t) \approx f(\bar{r}(k, t), t) |\bar{r}'(k, t)| \sim k^{-\gamma}$, with $\gamma = 2\alpha + 1$. \blacksquare We thus see again that we can create power law graphs with any exponent $\gamma \geq 2$.

The average node degree $\bar{k}(t) = \int_0^{R(t)} \bar{k}(r, t) f(r, t) dr$, is:

$$\bar{k}(t) \approx 2\alpha c \left(e^{(\alpha - \frac{1}{2})R(t)} C_1 - \alpha R(t)^2 C_2 - R(t) C_3 - C_4 \right), \quad (20)$$

where $C_1 = \frac{2\alpha^3}{\pi(\alpha - \frac{1}{2})^3(2\alpha - \frac{1}{2})}$, $C_2 = \frac{2\alpha - \pi(\alpha - \frac{1}{2})}{2\pi(\alpha - \frac{1}{2})}$, $C_3 = \frac{\alpha}{\pi(\alpha - \frac{1}{2})^2}$, $C_4 = \frac{2}{\pi(\alpha - \frac{1}{2})} \left(\frac{\alpha}{2(\alpha - \frac{1}{2})^2} + 1 \right)$, and the limit $\alpha \rightarrow \frac{1}{2}$ is again well defined and easily computed. From Equation (20) we can thus see that the constant c of our model can be used to set the average node degree to a target value for when the number of nodes in the network becomes t . Notice that for fixed c and $\alpha \rightarrow \frac{1}{2}$ the exponential term in Equation (20) vanishes and $\bar{k}(t)$ grows very slowly with t , as a function of $\ln t$, since $R(t) = \frac{1}{\alpha} \ln \frac{t}{c}$. This is desirable, as in many practical cases, we would like the average node degree not to depend (at least significantly) on the system size. Interestingly, by Theorem 3, $\alpha = \frac{1}{2}$ yields degree exponent $\gamma = 2$, which as we have seen in Section III also maximizes the efficiency of greedy forwarding.

In Figure 5, we demonstrate the accuracy of our predictions. The figure corresponds to a growing network with parameters $\alpha = 0.75$, i.e. target $\gamma = 2.5$, and $c = 0.0014$, i.e. target $\bar{k}(t) = 6.5$ when $t = 10000$. Figure 5(a) shows, in semi-log scale,

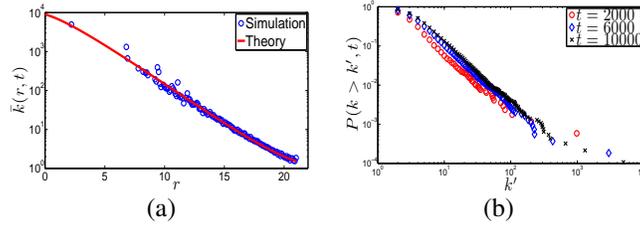


Fig. 5. (a) $\bar{k}(r, t)$ at $t = 10000$, and (b) degree ccdf at various t .

that Equation (17) closely matches simulations. Figure 5(b) shows, in log-log scale, the behavior of the degree ccdf as the network grows, i.e., at various t . The ccdf approximately follows a power law with slope, i.e., exponent, virtually independent of t , as our theoretical arguments suggest.⁶ The exponent γ is approximately 2.5. We further observe that the ccdf is slightly shifting to the right, indicating that the average node degree in the network grows slowly. At $t = 10000$ the average degree is 6.5768, which is very close to our target one. Similar results hold for other parameter values.

B. Decentralized implementation

As mentioned earlier, to build a network in a decentralized manner, each arriving node needs to be able to discover the current number of nodes in it, and its neighbors. Below, we give an example of an efficient greedy algorithm for doing so, at least approximately. Note that this is by no means the only algorithm. Other techniques and optimizations are possible and it is in our future plans to identify the best such technique(s).

Discovering the current number of nodes in the network and their coordinates: Each arriving node contacts a random node currently in the network, which acts as a *bootstrap node*. If it cannot contact a node, then the system is empty. The bootstrap node then sends an *exploration packet* to the network. In a nutshell, the lifetime of the exploration packet is as

⁶Note that if $P(k)$ is power law with exponent γ , then $P(k > k')$ is power law with exponent $\gamma - 1$.

follows. Each hop that receives the packet first registers in it the coordinates and ids of its neighbors and of itself, and then forwards it to its *highest degree neighbor* that has not seen the packet before. The process terminates when all neighbors of a node have seen the packet. Below, we describe the process in detail.

The exploration packet starts from the bootstrap node and keeps a list of the node ids it has visited, denoted by L_V . It also keeps a list of node ids along with their corresponding coordinates, denoted by L_C . Each node that receives the packet records its own id into L_V , and its id and coordinates into L_C . Further, it also records the id and the corresponding coordinates of each of its neighbors into L_C . The node then selects from its neighbors that are not included in L_V , the one with the maximum degree, and forwards the packet to it. The process terminates when all neighbors of a node are listed in L_V , in which case the exploration packet is sent back to the bootstrap node. The list L_C is then given to the arriving node.

This process, called “*search utilizing high degree nodes*”, is very efficient in power law graphs. In particular, for degree exponents $2 < \gamma < 3$, the exploration packet can discover a very large percentage of the nodes in the graph and their coordinates (recorded in L_C), by traversing only a small number of hops $h \sim N^{2-\frac{4}{\gamma}}$ (recorded in L_V), see Chapter 13 in [4]. Having an estimate of the number of nodes, the arriving node can compute the current hyperbolic disc radius, and in turn, its own coordinates. Further, knowing the hyperbolic disc radius, the coordinates of the nodes, and their ids, it can compute to which nodes to connect to.

Note that extensions of the above basic technique are also possible. For example, one can impose upper bounds on the size of the L_C list. Once these bounds are reached, the current L_C list is returned to the bootstrap node and then cleared in the exploration packet. This extension adds control on the maximum size the exploration packet can have.

C. Greedy forwarding

We now evaluate via simulation the performance of greedy forwarding in highly dynamic conditions, where nodes randomly arrive and depart the system. Our setup is as follows.

Without loss of generality, time is slotted. During each time slot, a new node arrives w.p. $p = 0.1$, and each node currently in the network departs w.p. $q = 0.0001$. Initially the network consists of 0 nodes. An arriving node joins the network according to our growing model. It also discovers the current number of nodes in it and their coordinates according to the procedure described earlier. In our experiments below, the exploration packet discovers, approximately, on average 95% of nodes in the network, by traversing only 1.5% of the nodes. If all neighbors of a node depart, the node re-initiates the join process to reconnect to the network. To ensure that the network remains connected, we assume that the first $t_{start} = 200$ nodes that arrive never depart. According to our growing model, these will be, on average, high degree nodes. High degree nodes are required to maintain connectivity in scale-free graphs [4]. Also, see discussion in Section V.

The average number of nodes in the network grows and stabilizes at the steady state value $\bar{t}_{steady} = \frac{p}{q} + t_{start} = 10200$.⁷ Notice that $\frac{t_{start}}{\bar{t}_{steady}} \approx 2\%$. Since the network grows on average until it reaches steady state we can use our earlier growing model analysis to make approximate/rough predictions.⁸ We set the system parameters to $\alpha = 0.5$ and $c = 0.01$. According to our earlier analysis, using as target number of nodes the average number of nodes in steady state \bar{t}_{steady} , these will correspond approximately to an average node degree $\bar{k}(\bar{t}_{steady}) = 6.5$ and degree exponent $\gamma = 2$. Recall from Section III that $\gamma = 2$ maximizes the efficiency of greedy forwarding. Also, as explained, for $\gamma = 2$ the average degree depends logarithmically on the network size. Therefore, even rough estimates of the \bar{t}_{steady} value would be sufficient.

Figure 6(a) shows the average node degree $\bar{k}(t)$ as a function of the current number of nodes t in the network, until and when we reach steady state. The average degree grows initially and then stays above, but close to, our target value of 6.5. Figure 6(b) shows how the degree cdf behaves. The degree exponent is $\gamma \approx 2$ and does not change as the network grows, as expected.

As mentioned in Section (III), more sophisticated yet still simple GF strategies may be able to achieve better performance than our simple OGF and MGF strategies, even under more dynamic network conditions. One such strategy that we now explore is the *Gravity-Pressure Greedy Forwarding* algorithm (GPGF), suggested in [7] and described below.

Gravity-Pressure Greedy Forwarding (GPGF). Each packet carries a bit to indicate whether the packet is in *Gravity* or *Pressure* forwarding mode. The packet starts in Gravity mode, where the forwarding procedure is exactly the same as in our OGF algorithm. However, if the packet reaches a local minimum, it is not dropped as in OGF. Instead, it first records the distance of the local minimum to the destination, which we call *current local-minimum distance*, and then enters Pressure mode. In Pressure mode, the packet maintains a list of the nodes in the network it has visited since it entered this mode, and the number of visits to each node. A node that receives the packet first determines all neighbors that the packet has visited the least number of times, and selects among those the one with the minimum distance to the destination to forward the packet to. This process

⁷Of course, in steady state we still have node arrivals and departures but the network does not grow on average. Formally, if $\bar{t}(s)$ is the average number of nodes in some time slot s excluding the first t_{start} nodes, and $\bar{t}(s+1)$ the corresponding number in slot $s+1$, then $\bar{t}(s+1) = \bar{t}(s)(1-q) + p$. The steady state value of $\bar{t}(s)$ is given as $\lim_{s \rightarrow \infty} \bar{t}(s) = \lim_{s \rightarrow \infty} \bar{t}(s+1) = \frac{p}{q}$. Adding t_{start} to this last quantity yields \bar{t}_{steady} .

⁸An exact analysis would need to account for the specific node arrival and departure processes, and it is out of the main focus of this paper.

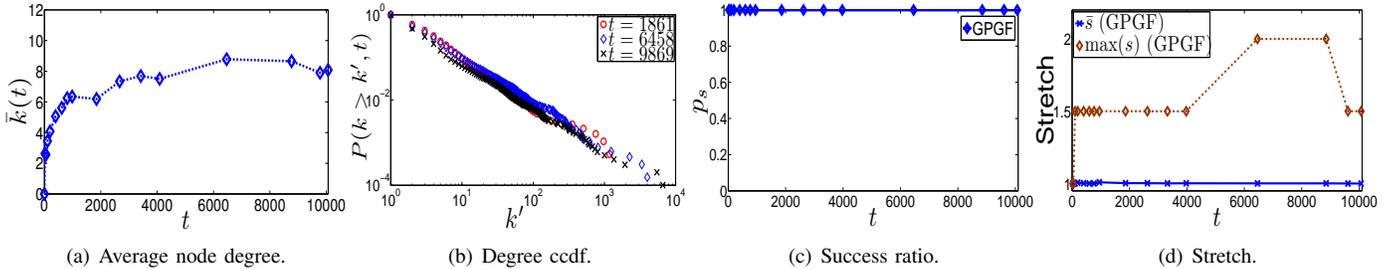


Fig. 6. Dynamic networks.

continues until the packet either reaches the destination or a node whose distance to the destination is smaller than the current local-minimum distance. In the latter case, the packet switches back to and continues in the Gravity mode.

The success ratio of GPGF is guaranteed to be always $p_s=1$, *irrespective of the geometry of the underlying space*. A formal proof is given in [7]. We also verify this as a function of the current number of nodes t in the network, in Figure 6(c). What is *not* guaranteed by the algorithm is the *stretch*, which can be enormous, as in the worst case a packet may need to visit *all* nodes in the network to find its destination! However, in Figure 6(d) we can see that its performance in terms of stretch is also exceptional in our networks. In particular, we see that the average stretch \bar{s} remains extremely close to 1, while the maximum stretch $\max(s)$ never exceeds 2. As before, this remarkable efficiency is due to the congruency between scale-free network topology and underlying hyperbolic geometry, which persists even under highly dynamic conditions.

V. DISCUSSION

In addition to their theoretical importance, our findings can find significant applications in overlay network constructions, e.g., [19], [21]. For example, the idea of using an underlying geometry to guide the forwarding process is similar in spirit to Distributed Hash Table (DHT) overlay architectures [19]. In these overlay networks, messages searching for data content are usually greedily forwarded, based on some distance metric, to the node in the network that is responsible for the data, usually the node that is closer to the data in terms of the distance metric. The main objective is to have low average node degree, and a low search, expressed as the average number of hops until the search packet finds the responsible node [19]. Both of these metrics should grow slowly with the number of nodes in the network N . However, the best search efficiency of $O(\ln \ln N)$ with $O(1)$ average node degree, can be achieved only when the network is scale-free; it is only in scale-free networks that shortest paths grow as $\sim \ln \ln N$, independently of the average node degree [4]. Therefore, given that greedy paths are approximately shortest in our networks, as well as the extreme robustness of greedy forwarding in dynamic conditions, which does not require any updates to be exchanged into the network, our findings can be used in such contexts or in similar ones, such as in overlay routing [21], to improve routing/search efficiency.

However, note that by no means are we proposing a solution in this paper that it is better than all existing overlay architectures in *all* aspects. For example, there are cases where it is desirable that all nodes in a network have similar degrees, which is not the case in scale-free networks where the degree distribution follows a power law. As another example, while it is well known that scale-free networks are robust to random node failures/departures, they are vulnerable to failures of the highest degree nodes, which can result in network disconnection [4]. To address this, one may want, for example, to allow for a hybrid architecture with resilient super-nodes (the high degree nodes) [11], or use techniques similar to those in [28] to ensure recovery when such nodes fail. However, as we have demonstrated in the previous section, it is enough to ensure that only the first few nodes that arrive into the system, which consisted in our experiments of 2% of the total steady state number of nodes, do not fail.

Finally, notice that in addition to applications in overlay networks, our findings also suggest that forwarding information through existing scale-free networks, like the Internet, may be possible without requiring global topology knowledge and routing updates. As one of the most serious scaling limitations with the existing Internet architecture is the communication overhead of deployed routing protocols [20], our findings are also of special interest in this context, as they may lay the groundwork for potentially moving to an Internet where efficient forwarding can take place without this overhead.

VI. CONCLUSION

This paper shows that scale-free network topologies naturally emerge from, and are congruent with, hyperbolic geometries. This congruency can be used to efficiently guide the forwarding process to find destinations with 100% success probability, following almost optimal (i.e. shortest) paths, even under dynamic network conditions. Our findings complement the important work in [14], and can have several practical applications in overlay network construction for improving routing/search performance.

There are several interesting directions for future work. One is to explore other techniques for decentralized implementation. Another is the design of greedy forwarding strategies that could also be used for improving other network performance metrics, e.g., strategies that could avoid congestion areas, perform load-balancing, and so on.

Finally, this paper also suggests that forwarding information through existing scale-free networks, like the Internet, may be possible without the current overhead of routing protocols. To answer this question, one of the most interesting, yet challenging, future work direction is the following *inverse problem*: Can we embed any *given* scale-free graph, e.g., the real Internet topology, into a hyperbolic space, and greedily forward through this embedding with similar efficiency? How can each node in this graph compute its coordinates in the space *having no global knowledge of the graph topology*, so that the resulting embedding is *congruent* with the space?

APPENDIX – GROWING MODEL ANALYSIS

Here we look at a growing network at some time instance where the number of nodes in it is some value t , and derive the approximate expressions for the node radial density $f(r, t)$ and the average node degree from the hyperbolic disc center $\bar{k}(r, t)$. As before, let's number each node by its order of arrival. Recall that there are no node departures and that each node $i \geq 1$ that arrives to the system computes a hyperbolic disc radius $R(i) = \frac{1}{\alpha} \ln \frac{i}{c}$. The radial coordinate $r \in [0, R(i)]$ of node i is distributed according to the density $f(r|R(i)) \approx \alpha e^{\alpha(r-R(i))}$.

We start with $f(r, t)$. Let $\{R(1) \dots R(t)\}$ be the sequence of the hyperbolic disc radii that nodes $\{1 \dots t\}$ compute on their arrival, and let R be the random variable representing the computed disc radius of a randomly selected node from $\{1 \dots t\}$. It is easy to see that for $i \in \{1 \dots t\}$, $P(R \leq R(i)) = \frac{i}{t} = \frac{ce^{\alpha R(i)}}{ce^{\alpha R(t)}} = e^{\alpha(R(i)-R(t))}$.

We treat R as a continuous random variable. As we see in Section IV this does not affect the accuracy of the predictions. Therefore, if we denote by $F(R, t) = e^{\alpha(R-R(t))}$ the distribution function of the computed disc radius R , then the probability density function of R , denoted by $f(R, t)$, is:

$$f(R, t) = \alpha e^{\alpha(R-R(t))},$$

and it is obtained by differentiating $F(R, t)$ w.r.t. R . Since, given the value of R a node computes its radial coordinate r according to:

$$f(r|R) \approx \alpha e^{\alpha(r-R)},$$

and $r \leq R \leq R(t)$, we can write:

$$f(r, t) = \int_r^{R(t)} f(r|R)f(R, t)dR \approx \alpha^2(R(t) - r)e^{\alpha(r-R(t))}, \quad (21)$$

which is Equation (16).

We now proceed with $\bar{k}(r, t)$. We can break $\bar{k}(r, t)$ into two parts. The first is $\bar{k}_{init}(r, t)$, which is the initial average degree of a node with radial coordinate r , i.e., the average number of nodes, already in the network, that the node connects to when it arrives. And, the second is $\bar{k}_{new}(r, t)$, which is the average number of new connections to the node at r , coming from new nodes that have arrived to the system after it. Clearly:

$$\bar{k}(r, t) = \bar{k}_{init}(r, t) + \bar{k}_{new}(r, t). \quad (22)$$

We first compute $\bar{k}_{init}(r, t)$. Suppose that the node at r computed a disc radius equal to R when it arrived. According to our model, the node then connected to all other nodes in the network, for which the hyperbolic distance to it was $d \leq R$. The average number of these nodes, denoted by $\bar{k}_{init}(r|R)$, can be computed in exactly the same way as $\bar{k}(r)$ in the model of Section III. In particular, it is the average number of nodes that lie in the intersection area of the two discs of the same radius R , one in which all current nodes reside, and the other centered at distance r from the center of the first disc. The node radial density at the time that the above node arrived is given by setting $R(t) = R$ in the $f(r, t)$ formula as the current disc radius where all nodes reside was $R \leq R(t)$. Let's denote this density by $f(r, R)$ so that we do not confuse it with $f(r, t)$. In other words:

$$f(r, R) \approx \alpha^2(R - r)e^{\alpha(r-R)}.$$

Using Equation (8) with $\chi = R$ and $N = ce^{\alpha R}$, we have:

$$\bar{k}_{init}(r|R) \approx ce^{\alpha R} \left\{ \int_0^{R-r} f(r', R)dr' + \frac{1}{\pi} \int_{R-r}^R f(r', R)\theta_{r'}dr' \right\}, \quad (23)$$

where $\theta_{r'} \approx 2e^{\frac{1}{2}(R-r)}e^{-\frac{1}{2}r'}$. After performing the integration, we can deduce that:

$$\bar{k}_{init}(r|R) \approx ce^{\alpha R}G(r), \quad (24)$$

where:

$$G(r) = \left(1 - \frac{2\alpha^2}{\pi(\alpha - \frac{1}{2})^2} + \left(1 - \frac{2\alpha}{\pi(\alpha - \frac{1}{2})}\right)\alpha r\right) e^{-\alpha r} + \frac{2\alpha^2}{\pi(\alpha - \frac{1}{2})^2} e^{-\frac{1}{2}r},$$

i.e., $G(r)$ in Equation (18), whose limit $\alpha \rightarrow \frac{1}{2}$ is $G(r) = (1 + \frac{r}{2} + \frac{r^2}{4\pi})e^{-\frac{1}{2}r}$. Note the Equation (24) is approximate and is accurate for $\alpha \geq \frac{1}{2}$.

Removing the condition on R from $\bar{k}_{init}(r|R)$, by accounting for all possible R with $r \leq R \leq R(t)$, gives $\bar{k}_{init}(r, t)$. To do so, we need the conditional probability density of R given that a node's radial coordinate is r , denoted by $f(R, t|r)$. Given $f(r|R)$, $f(R, t)$ and $f(r, t)$ as above, we can write

$$f(R, t|r) = \frac{f(r|R)f(R, t)}{f(r, t)}.$$

Therefore:

$$\bar{k}_{init}(r, t) = \int_r^{R(t)} \bar{k}_{init}(r|R) f(R, t|r) dR \approx t \frac{1 - e^{-\alpha(R(t)-r)}}{\alpha(R(t)-r)} G(r). \quad (25)$$

We now proceed with $\bar{k}_{new}(r, t)$. Let N_{new} be the number of new nodes that arrived to the system after a node with radial coordinate r . Clearly, the computed hyperbolic disc radius R' of a new such node satisfies $r \leq R' \leq R(t)$. According to our model, the new node connects to the node at r only if the hyperbolic distance to it is $d \leq R'$. We next compute the probability of this event.

To do so, we first need to compute the probability $P_l(r|r', R')$, which is the probability that there is a link to a node with radial coordinate r , coming from a node with radial coordinate r' whose computed disc radius is R' . Recall that the hyperbolic distance d between two nodes with coordinates (r, θ) and (r', θ') is given as:

$$d = \operatorname{arccosh}(\cosh r \cosh r' - \sinh r \sinh r' \cos \Delta\theta),$$

where $\Delta\theta$ the angle between the two nodes. To simplify the analysis we look for approximate expressions for d . Unfortunately, there is no single one. Assuming that $R' \gg 1$, there are at least three important approximations:

$$d \approx \begin{cases} r + r' + 2 \ln \sin \frac{\Delta\theta}{2} & \text{if } r, r' \gg 1; \Delta\theta \gg 0, \\ |r - r'| + \frac{1}{4} e^{2r'} (\Delta\theta)^2 & \text{if } r > r' \gg 1; \Delta\theta \ll 1, \\ |r - r'| + \frac{\sinh r \sinh r'}{2 \cosh(r-r')} (\Delta\theta)^2 & \text{if } r' > r \gg 1; \Delta\theta \ll 1 \text{ or } r \ll R'; r' \gg 1 \text{ or } r \gg 1; r' \ll R'. \end{cases}$$

One can check (e.g., using command `Manipulate` in `Mathematica` v.6) that all these approximations nicely glue to each other at different values of $\Delta\theta$, with the gluing point position depending on specific values of r, r' . Since the node density is exponential, there are very few nodes at $r \ll R'$ and $r' \ll R'$, therefore we can neglect the approximations for $r \ll R'$ or $r' \ll R'$. For $r, r' \gg 1$, the gluing point is always at $\Delta\theta \ll 1$. Therefore we can completely neglect the second and third approximations, and use only the first one, which for sufficiently large r, r' and $\Delta\theta$, further simplifies to:

$$d \approx r + r' + 2 \ln \frac{\Delta\theta}{2}.$$

Since there is a link between the two nodes only if $d \leq R'$, we have:

$$P_l(r|r', R') = P(d \leq R') \approx P(\Delta\theta \leq 2e^{\frac{1}{2}(R'-r'-r)}).$$

Further, since the angular coordinates of nodes are uniformly distributed in $[0, 2\pi]$, $\Delta\theta$ is uniformly distributed in $[0, \pi]$. Therefore, we can write:

$$P_l(r|r', R') \approx \begin{cases} 1 & \text{if } r' \leq R' - r - 2 \ln \frac{\pi}{2}, \\ \frac{2}{\pi} e^{\frac{1}{2}(R'-r'-r)} & \text{otherwise.} \end{cases} \quad (26)$$

To find the probability that a new node with a computed disc radius R' connects to the node at r , we only need to remove the condition on r' from Equation 26. Let's denote this probability by $P_l(r|R')$. Since r' is distributed according to $f(r') \approx \alpha e^{\alpha(r'-R')}$, we have:

$$\begin{aligned} P_l(r|R') &\approx \int_0^{R'-r-2 \ln \frac{\pi}{2}} \alpha e^{\alpha(r'-R')} dr' + \int_{R'-r-2 \ln \frac{\pi}{2}}^{R'} \frac{2}{\pi} e^{\frac{1}{2}(R'-r'-r)} \alpha e^{\alpha(r'-R')} dr' \\ &= \frac{2\alpha}{\pi(\alpha - \frac{1}{2})} e^{-\frac{1}{2}r} - \frac{1}{2(\alpha - \frac{1}{2})} e^{-2\alpha \ln \frac{\pi}{2}} e^{-\alpha r} - e^{-\alpha R'} \\ &\approx \frac{2\alpha}{\pi(\alpha - \frac{1}{2})} e^{-\frac{1}{2}r} - \frac{1}{2(\alpha - \frac{1}{2})} e^{-2\alpha \ln \frac{\pi}{2}} e^{-\alpha r} \triangleq P(r), \end{aligned} \quad (27)$$

which is Equation (19). The last approximation is accurate for $\alpha > \frac{1}{2}$, as we are ignoring the term $e^{-\alpha R'}$. We thus see that the probability that a new node connects to a node with radial coordinate r is approximately independent of the exact value of the new node's computed disc radius R' and depends only the radial coordinate r of the (old) node. Further, the limit $\alpha \rightarrow \frac{1}{2}$ is $P(r) \rightarrow \frac{2}{\pi}(1 + \ln \frac{\pi}{2} + \frac{1}{2}r)e^{-\frac{1}{2}r}$.

Now, given that the node at r was the i^{th} node arrival, $N_{new} = t - i$. If $\bar{k}_{new}(r, t|i)$ denotes the average number of new connections to node i , we can write:

$$\bar{k}_{new}(r, t|i) \approx (t - i)P(r) = (t - ce^{\alpha R(i)})P(r).$$

To find $\bar{k}_{new}(r, t)$, we need to remove the condition on the index i . In other words, we need to account again for all R with $r \leq R \leq R(t)$. The conditional density of R , $f(R, t|r)$, was computed earlier. Therefore:

$$\bar{k}_{new}(r, t) \approx \int_r^{R(t)} (t - ce^{\alpha R})P(r)f(R, t|r)dR = t \left(1 - \frac{1 - e^{-\alpha(R(t)-r)}}{\alpha(R(t) - r)} \right) P(r). \quad (28)$$

Adding $\bar{k}_{init}(r, t)$ from Equation (25) and $\bar{k}_{new}(r, t)$ from Equation (28), we get:

$$\bar{k}(r, t) \approx t \left\{ P(r) + \frac{1 - e^{-\alpha(R(t)-r)}}{\alpha(R(t) - r)} (G(r) - P(r)) \right\}, \quad (29)$$

which is Equation (17), and the $\lim_{\alpha \rightarrow \frac{1}{2}} \bar{k}(r, t)$ can be easily computed.

Note that Theorem 3 states that we can create power law graphs with $\gamma = 2\alpha + 1 \geq 2$, as our analysis above holds for $\alpha \geq \frac{1}{2}$. For $\alpha < \frac{1}{2}$ our analysis above does not accurately hold. However, we have experimentally verified that we can again generate power law graphs where the exponent is $\gamma = 2$.

REFERENCES

- [1] I. Abraham, M. Balakrishnan, F. Kuhn, D. Malkhi, V. Ramasubramanian, and K. Talwar. Reconstructing approximate tree metrics. In *PODC*, 2007.
- [2] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.
- [3] M. Boguñá and R. Pastor-Satorras. Class of correlated random networks with hidden variables. *Phys Rev E*, 68:036112, 2003.
- [4] S. Bornholdt and H. G. Schuster (Eds.). *Handbook of Graph and Networks: From the Genome to the Internet*. Wiley-VCH, Berlin, 2002.
- [5] M. R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer-Verlag, Berlin, 1999.
- [6] A. Clauset, C. Moore, and M. E. J. Newman. Hierarchical structure and the prediction of missing links in networks. *Nature*, 453:98–101, 2008.
- [7] A. Cvetkovski and M. Crovella. Hyperbolic embedding and routing for dynamic graphs. In *INFOCOM*, 2009.
- [8] The DIMES project. <http://www.netdimes.org/>.
- [9] S. N. Dorogovtsev and J. F. F. Mendes. *Evolution of Networks: From Biological Nets to the Internet and WWW*. Oxford University Press, Oxford, 2003.
- [10] M. Girvan and M. E. J. Newman. Community structure in social and biological networks. *Proc Natl Acad Sci USA*, 99:7821–7826, 2002.
- [11] C. Gkantsidis and M. Mihail. Hybrid search schemes for unstructured peer-to-peer networks. In *INFOCOM*, 2005.
- [12] M. Gromov. *Metric Structures for Riemannian and Non-Riemannian Spaces*. Birkhäuser, Boston, 2007.
- [13] E. Jonckheere, P. Lohsoonthorn, and F. Bonahon. Scaled Gromov hyperbolic graphs. *J. Graph Theory*, 57(2), 2008.
- [14] J. Kleinberg. Navigation in a small world. *Nature*, 406:845, 2000.
- [15] J. Kleinberg. Complex networks and decentralized search algorithms. In *ICM*, 2006.
- [16] R. Kleinberg. Geographic routing using hyperbolic space. In *INFOCOM*, 2007.
- [17] R. Krauthgamer and J. R. Lee. Algorithms on negatively curved spaces. In *FOCS*, 2006.
- [18] D. Krioukov, F. Papadopoulos, A. Vahdat, and M. Boguna. On curvature and temperature of complex networks, 2009. <http://arxiv.org/abs/0903.2584>.
- [19] E. K. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim. A survey and comparison of peer-to-peer overlay network schemes. *IEEE Communications survey and tutorial*, 7(2), 2005.
- [20] D. Meyer, L. Zhang, and K. Fall, editors. *RFC4984*. The Internet Architecture Board, 2007.
- [21] A. Nakao, L. Peterson, and A. Bavier. Scalable routing overlay networks. *SIGOPS Oper. Syst. Rev.*, 40(1), 2006.
- [22] V. Naumov and T. Gross. Scalability of routing methods in ad hoc networks. *Performance Evaluation*, 62:193–209, 2005.
- [23] M. E. J. Newman. The structure and function of complex networks. *SIAM Rev*, 45(2):167–256, 2003.
- [24] University of Oregon RouteViews Project. <http://www.routeviews.org/>.
- [25] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32:425–443, 1969.
- [26] D. J. Watts, P. S. Dodds, and M. E. J. Newman. Identity and search in social networks. *Science*, 296:1302–1305, 2002.
- [27] D. J. Watts and S. H. Strogatz. Collective dynamics of “small-world” networks. *Nature*, 393:440–442, 1998.
- [28] R. H. Wouhaybi and A. Campbell. Phenix: supporting resilient low-diameter peer-to-peer topologies. In *INFOCOM*, 2004.