Efficient Navigation in Scale-Free Networks Embedded in Hyperbolic Metric Spaces

Dmitri Krioukov,[1] Fragkiskos Papadopoulos,[1] Marián Boguñá,[2] and Amin Vahdat[3]

[1]*Cooperative Association for Internet Data Analysis (CAIDA), University of California, San Diego (UCSD), La Jolla, CA 92093, USA*

[2]*Departament de Física Fonamental, Universitat de Barcelona, Martí i Franquès 1, 08028 Barcelona, Spain*

[3]*Department of Computer Science and Engineering, University of California, San Diego (UCSD), La Jolla, CA 92093, USA*

In this work we show that: i) the roughly hierarchical structure of complex networks is congruent with negatively curved geometries hidden beneath the observed topologies; ii) the most straightforward mapping of nodes to spaces of negative curvature naturally leads to the emergence of scale-free topologies; and iii) greedy routing on this embedding is efficient for these topologies, achieving both 100% reachability and optimal path lengths, even under dynamic network conditions. The critical important question left by this work is whether the topologies of real networks can be mapped into appropriate hidden hyperbolic metric spaces.

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, since without some knowledge of changing network connectivity, it is not possible to successfully route information through the network.

In the Internet, this required inter-node communication makes routing both expensive and fragile. The recent Internet Architecture Board report on routing and addressing [35] identifies *convergence* costs of deployed routing protocols as one of the most serious scaling issues with the existing Internet routing architecture, aggravated by explosive rates of routing table size growth. Worse yet, a recent review of compact routing [26] shows that the required number of messages for routing state to converge on small-world networks cannot scale better than linearly with network size for *any* routing algorithm.

In many other real networks however, nodes can efficiently communicate, even though they do not exchange any information about the current global state of the network topology. In 1969, Stanley Milgram performed the following experiment [53]. He asked some random individuals—sources—to send a letter to a specific person—the destination, described by his name, occupation, age, and the city he lived in. The sources were asked to pass the letter to friends chosen to maximize the probability of the letter reaching its destination. Approximately 30% of the letters reached their destination, even though nodes in this routing example 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 [22] of this surprising effect. In his model, each node, in addition to being a 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 the 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 the direct neighbor closest to the destination in the plane.

Clearly, the described greedy routing strategy can be efficient only if the network topology is in some way congruent with the underlying space. Kleinberg indeed finds that the paths produced by greedy routing are polylogarithmically short only if the probability that there is a link between a pair of nodes depends in a very specific way on the distance between the two nodes in the plane. This finding implies that greedy routing cannot be equally efficient on any arbitrary network topology built over a given underlying space. Only certain topologies, congruent with the underlying space, will perform well.

Given that the network topology is so critically important, the Kleinberg model stands closer to the beginning of an explanation for Milgram's experiment than to its end. The model does not (try to) reproduce the basic topological properties of social networks through which letters were traveling in Milgram's experiments. For instance, the Kleinberg model produces only k-regular graphs while social networks, the Internet, and many other complex networks [39] 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 lying between 2 and 3 (see Table I); and ii) the network has strong clustering, i.e., a large number of triangular subgraphs.

In this work we assume that nodes in the Internet and other 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 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 the Internet and other real networks are not Euclidean planes. In this work, we attempt to identify the most basic geometric properties of these hidden spaces.

Specifically, the main results in this paper are that if we model hidden spaces as non-Euclidean *hyperbolic spaces*, i.e., spaces of negative curvature, then this negatively curved geometry leads to:

1. naturally emerging scale-free topologies constructed over such hidden spaces; and
2. extremely efficient greedy routing on these topologies—maximally efficient, in fact, across all topologies we consider.

We have to emphasize the importance of the first result. We make no effort whatsoever to enforce either power laws or clustering in our modeled networks. These two properties emerge naturally. Therefore, hidden hyperbolic geometries appear as a new possible explanation of the scale-free structure of complex networks. To the best of our expertise, this explanation is not equivalent to the preferential attachment [3] or any other known mechanisms [39] of power-law emergence in complex networks.

In this paper however, we focus on the second result. In more detail, we find that in scale-free networks with small exponents γ, greedy routing successfully finds paths between 99.99% source-destination pairs, while simple boosting techniques bring this success ratio to 100%. More strikingly, all successful paths follow shortest paths, resulting in a maximum stretch of 1. These findings indicate that static scale-free topologies are extremely congruent with their underlying hyperbolic geometries in our models.

Real networks are dynamic however, with link and node failures the common case. We find, most remarkably, that the efficiency of greedy routing and, especially, of its modifications is extremely stable even under dynamic conditions. For example, the success ratio degrades by less than 1% when as many as 10% of all links in the network fail. Once again, our simple extensions can boost the success ratio back to 100%. These findings indicate that scale-free topologies are not only congruent with hidden hyperbolic geometries, but also that this congruency is resilient to network dynamics.

In the next section we recall the basic facts on hyperbolic geometry. They let us outline, in Section III, the main motivations that led us to our hyperbolic hidden space conjecture. We describe some details of our models of networks on hyperbolic spaces in Section IV. Leaving all the analytic derivations concerning our models for future publication, we move directly to Section V where we report our simulation results confirming our analytic predictions and focusing primarily on the efficiency characteristics of greedy routing and its modifications in modeled networks. After a review of related work in Section VI, we conclude, in Section VII, with an outline of the main results, their applicability, and directions for future work.

II. HYPERBOLIC SPACES

In this section we review the basic facts about hyperbolic geometry. More detailed accounts can be found in various (text)books [2, 8, 9, 19].

TABLE I: Values of power-law exponent γ observed in some complex networks. The values for the AS Internet and PGP trust network come from [14, 45] and [5]. All other values are from [39]. In many networks, γ is close to 2.

Network	γ
AS Internet	2.1
WWW	2.1
P2P	2.1
PGP trust relationships	2.5
Film actor collaboration	2.3
Metabolic reactions	2.2
Protein interactions	2.4

In two dimensions, there are only three types of isotropic[1] spaces: Euclidean (flat), spherical (positively curved), and hyperbolic (negatively curved). Most readers are familiar with the first two. Hyperbolic spaces of constant negative curvature are more difficult to envisage because they cannot be isometrically[2] embedded in any Euclidean spaces. The reason is, informally, that the former are "larger," have more "space" than the latter.

One can measure the curvature of a surface at a given point by the deviation of the lengths of circles centered at the point from their Euclidean values. Formally, the curvature of a surface at a given point can be defined as

$$K = \frac{3}{\pi} \lim_{R \to 0} \frac{2\pi R - l(R)}{R^3}, \qquad (1)$$

where $l(R)$ is the length of the circle of radius R lying on the surface and centered at this point. If there is no deviation from $2\pi R$, then the surface is flat; if circles are shorter or longer than in the Euclidean case, then the surface is positively or negatively curved. A classic example of hyperbolic surfaces is the one-sheeted hyperboloid, obtained by rotating a hyperbola around one of its symmetry axis. The hyperboloid curvature is not constant: it is -1 at its narrowest part, but approaches 0 at infinity. Note for comparison that at any point on a sphere of radius 1, the curvature is 1.

Because of the fundamental difficulties in representing spaces of *constant* negative curvature as subsets of Euclidean spaces, there are not one but many models for the hyperbolic plane, that is the 2-dimensional hyperbolic space of constant curvature -1. Each model emphasizes different aspects of hyperbolic geometry, but no model simultaneously represents all its properties. For illustration purposes, we consider just one model—the Poincaré disc model. It is conformal, meaning that Euclidean angles between hyperbolic lines in the model are equal to their hyperbolic values, but it does not preserve distances or areas.

---

[1] The space is isotropic if it "looks the same" at every point and in every direction.

[2] An isometric embedding preserves distances.

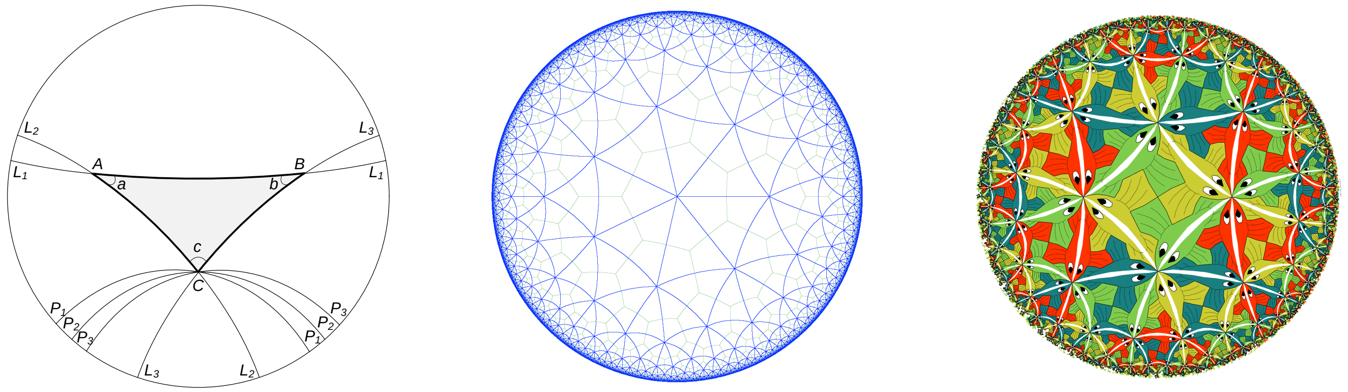

(a)Lines and triangles (b)Tessellation (c)Hyperbolic art

FIG. 1: Poincaré disc model. Fig. (a) shows examples of hyperbolic lines ($L_{1,2,3}, P_{1,2,3}$). Lines $L_{1,2,3}$ intersect to form triangle ABC. The sum of its angles $a + b + c < \pi$. There are infinitely many lines (examples are $P_{1,2,3}$) that are parallel to line L_1 and go through a point C that does not belong to L_1. Fig. (b) shows a tessellation of the hyperbolic plane by equilateral triangles, and the dual tessellation by regular heptagons. All triangles and heptagons are of the same hyperbolic size but the size of their Euclidean representations exponentially decreases as a function of the distance from the center, while their number exponentially increases. Fig. (c) is a tessellation-inspired artistic visualization of the hyperbolic plane by Silvio Levy, based on M. C. Escher's *Circle Limit III*. Printed with the permission from the Geometry Center, University of Minnesota.

In the model, the whole infinite hyperbolic plane is represented by the interior of the Euclidean disc of radius 1, see Fig. 1. The boundary of the disc is not a part of the hyperbolic plane, but represents its infinitely remote points. Hyperbolically straight, infinite lines, i.e., geodesics, are disc diameters and Euclidean arcs orthogonal to its boundary. Euclidean and hyperbolic distances, r_e and r_h, from the disc center, or the origin of the hyperbolic plane, are related by $r_e = \tanh(r_h/2)$. Euclidean distances in the radial direction thus correspond to exponentially longer hyperbolic distances as we move closer to the disc boundary. For example, the hyperbolic lengths of radial intervals of Euclidean length 0.01 located at Euclidean distances 0.0, 0.5, and 0.9 from the center are 0.02, 0.03, and 0.70:

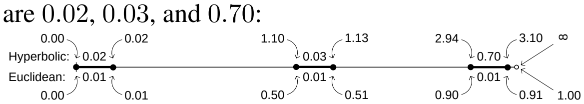

More generally, by definition of the negative curvature, hyperbolic spaces expand faster than Euclidean spaces. Specifically, while Euclidean spaces expand polynomially, hyperbolic spaces expand exponentially, as illustrated in Fig. 1. In the hyperbolic plane, for example, the length of the circle and the area of the disc of hyperbolic radius R are

$$l(R) = 2\pi \sinh R, \tag{2}$$
$$s(R) = 2\pi(\cosh R - 1). \tag{3}$$

In the Poincaré disc model, these relations imply that Euclidean distances in the tangential direction also correspond to exponentially longer hyperbolic distances as we move closer to the disc boundary. For example, the hyperbolic lengths of tangential intervals of Euclidean length 0.01 located at Euclidean distances 0.0, 0.5, and 0.9 from the center are 0.02, 0.03, and 0.10:

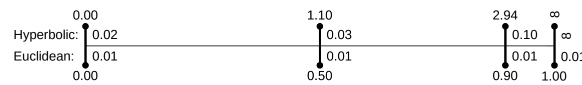

Note that substituting Eq. (2) in Eq. (1) yields curvature $K = -1$ at every point of the hyperbolic plane. Eqs. (2,3) mean that both circle lengths and disc areas grow as e^R with radius R.

Some find these properties of hyperbolic space bizarre. However, certain more familiar objects, including trees, display these properties. In a b-ary tree (a tree with branching factor b), 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 b^R with R. We thus see that the metric structures of the hyperbolic plane and trees are extremely similar: in the former case circle lengths and disc areas grow as e^R, while in the latter—as b^R. In other words, from the purely metric perspective, the hyperbolic plane is equivalent to a tree with average branching factor e. Informally, trees can therefore be thought of as "discreet hyperbolic spaces." Formally, trees, even infinite ones, allow isometric embeddings into the hyperbolic plane. For example, any tessellation of the hyperbolic plane (see Fig. 1) naturally defines isometric embeddings for a class of trees formed by certain subsets of polygon sides. Numerous software packages for visualization of massive tree-like graphs utilize the fact that trees embed isometrically into hyperbolic spaces. Note that in general, trees do not isometrically embed into Euclidean spaces. Informally, trees need an exponential amount of space for branching, and only hyperbolic geometry has it.

Table II collects these and other characteristic properties of hyperbolic geometry and juxtaposes them against the corresponding properties of Euclidean and spherical geometries.

TABLE II: Characteristic properties of Euclidean, spherical, and hyperbolic geometries. *Parallel lines* is the number of lines that are parallel to a line and that go through a point not belonging to this line.

Property	Euclid.	Spherical	Hyperbolic
Curvature	0	1	-1
Parallel lines	1	0	∞
Triangles are	normal	thick	thin
Shape of triangles	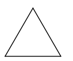	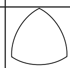	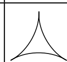
Sum of angles	π	$> \pi$	$< \pi$
Circle length	$2\pi R$	$2\pi \sin R$	$2\pi \sinh R$
Disc area	$2\pi R^2/2$	$2\pi(1-\cos R)$	$2\pi(\cosh R - 1)$

III. MOTIVATION

If not all properties of hyperbolic spaces are necessarily easy to comprehend, then why suggest them as models for hidden metric spaces underlying the real Internet and other networks? There are at least two reasons, discussed below.

A. Node taxonomies imply negative curvature of hidden spaces

Complex networks connect distinguishable, heterogenous elements abstracted as nodes. Understood broadly, this heterogeneity implies that there is at least some taxonomy of elements, meaning that all nodes can be somehow classified. In most general settings, this classification implies that nodes can be split in large groups consisting of smaller subgroups, which in turn consist of even smaller subsubgroups, etc. The relationships between such groups and subgroups can thus be approximated by tree-like structures, sometimes called *dendrograms* [12], that represent hidden hierarchies in networks. But as discussed in Section II, the geometries of trees and hyperbolic spaces are intimately related, and they both are negatively curved. We emphasize that we do not assume that the node classification hierarchy among a particular dimension is strictly a tree, but that it is approximately a tree. As soon as it is at least approximately a tree, it is negatively curved [19].

The above discussion obviously applies only to a snapshot of a network taken at some moment of time. A logical question is how these taxonomies emerge. Clearly, when the first node of a future network appears, the node classification is degenerate, but as more and more nodes join the network and evolve in it, they necessarily diversify and specialize, thus deepening their classification hierarchy.

With these observations, the distance between nodes in those underlying tree-like geometries is a rough approximation of how similar two nodes are [27]. The more similar a pair of nodes, the more likely they are connected.

We consider several examples suggesting that these general considerations are applicable to different real networks. Social networks form the most straightforward class of examples, where network community structures [5, 17] represent hidden hierarchies [54]. More concretely, in paper citation networks—nodes are scientific papers and links are references—the underlying geometries can approximately be the relationships between scientific subject categories, and the closer the subjects of two papers, the more likely they cite each other [6, 44]. Conceptually similar subject-based classification of web pages (or more specifically, of the Wikipedia pages [37]) also shows the same effect: the closer the subjects of a pair of web pages, the more likely there is a hyperlink between them [34]. In biology, the distance between two species on the phylogenetic tree is a widely used measure of similarity between the species [38]. Note that this example emphasizes both the existing taxonomy of elements and their evolution.

Finally, the evolution of the Internet can be also mapped to this general framework. In the beginning, there were only two IMPs, then ARPANET that grew and inspired MFENet, HEPNet, CSNET, USENET, BITNET, etc., and finally NSFNET whose decommission eventually led to the appearance of a collection of interconnected ASs, and their growth in number and diversity [28]. Currently, ASs can be classified based on their geographic position and coverage, size, number and type of customers, business role, and many other factors [15].

B. Power laws as a consequence of exponential expansion of hidden space

We have seen in Section II that hyperbolic spaces expand exponentially. In particular, if nodes are distributed approximately uniformly in a hidden hyperbolic space, then the number of nodes $n(r)$ at distance r from any reference point, e.g., an abstraction of the root or origin of hidden hierarchies described in Section III A, grows exponentially with r,[3]

$$n(r) \sim e^{\alpha r}, \quad \alpha > 0, \tag{4}$$

with $\alpha = 1$ for the hyperbolic plane.

At the same time, the average degree $k(r)$ of nodes located at distance r from the hierarchy origin should decrease with r. Indeed, nodes farther from the origin are, on average, newer (as discussed in Section III A), and therefore have had less time to establish connections to other nodes, whatever the specific mechanisms of the network evolution.

Using this kind of high-level consideration, we can speculate that $k(r)$ decreases exponentially with increasing r. Indeed, we can always define any characteristic hyperbolic shape span by a node's connections, e.g., the minimum-size hyperbolic disc containing $x\%$ of its neighbors in the network. Independent of the connection shape for nodes at distance r from the origin, the average area $S(r)$ of the intersection of their connection shapes with the minimum-sized disc centered at the origin and containing all nodes in the network decreases exponentially with r, which is an obvious consequence of the

[3] In this paper, symbols '$\sim$' and '$\approx$' mean, respectively, *proportional to* and *approximately equal*.

exponential dependence of hyperbolic area on its characteristic size, cf. Section II. Since by definition $k(r)$ is proportional to $S(r)$, the former also decreases exponentially:

$$k(r) \sim e^{-\beta r}, \quad r(k) \sim -\frac{1}{\beta}\ln k, \quad \beta > 0. \tag{5}$$

The combination of exponentials [40] in Eqs. (4,5) yields the power-law node degree distribution: the number of nodes $n(k)$ of degree k is given by:

$$n(k) \approx n[r(k)]|r'(k)| \sim k^{-\gamma} \tag{6}$$

with the power-law exponent:

$$\gamma = 1 + \frac{\alpha}{\beta}. \tag{7}$$

This way, the power laws ubiquitously observed in complex networks, including the Internet, emerge as a simple and natural consequence of the exponential expansion of space in hidden hyperbolic geometries.

IV. MODELS OF NETWORKS EMBEDDED IN HYPERBOLIC SPACES

In this section we develop network models with nodes residing in hyperbolic spaces.

Given our main premise that Internet nodes (ASs or routers) and nodes of other real networks exist in hidden spaces of negative curvature, we would ultimately like to specify exactly what these spaces are for each given real network. Further, each node should be able to compute its hidden coordinates based solely on the information accessible to it locally. For example, each AS has access to the information on the identity of its customers, providers, and peers. It also knows the details of its peering agreements, its business role in the Internet economy, the size of IP address space allocated to it, its geographic spread, i.e., the number, size, and geographic location of PoPs it is present at, etc. We thus see that there are numerous, non-unique ways to combine (some function of) these and other pieces of locally-known information into an ultimate formula that would compute the AS's coordinates in the underlying space, which is yet unknown to us.

We leave the difficult problem of reconstructing the exact structure of hidden spaces underlying the Internet and other networks for future work. In this paper, we consider the simplest hyperbolic space possible, a hyperbolic plane, and see if we can construct simple network models where nodes mapped to this space naturally reproduce the main topological characteristics of the Internet and other real networks. Specifically, according to [32], reproducing the Internet's power-law node degree distributions, correlations, and clustering, one can capture many other metrics as well. As a bare minimum, we want to construct network models with nodes lying in a hyperbolic space and with power-law node degree distributions arising naturally, e.g., without our designing them into the modeled networks.

We achieve this task using the *hidden variable* approach described in [4]. This approach works as follows: given a network size N, each node i is first assigned a hidden variable h_i drawn from some probability distribution, and then a link between each pair of nodes (i, j) is created with a *connection probability* $p(h_i, h_j)$, which is a function of the hidden variable values at the two nodes. In our models, hidden variables are node coordinates in a hidden hyperbolic space and the connection probability depends on the hyperbolic distance between them. We can thus fully describe each network model by specifying: 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.

The simplest hyperbolic space is the hyperbolic plane discussed in Section II. The simplest way to place N nodes on the hyperbolic plane is to distribute them uniformly over a disc of radius R. Since the disc area is given by Eq. (3), we have $N \sim e^R$. In view of the analogy between hyperbolic spaces and trees discussed in Section II, the radius R of the disc is an abstraction of the depth of the hidden network hierarchy due to node taxonomies. When the network is small this hierarchy is shallow, and as the network grows, its hidden node hierarchy deepens, as discussed in Section III A. In our model, the relationship between the disc radius and network size, $R \sim \ln N$, is qualitatively the same as the relationship between the depth of a balanced tree and its size.

We have thus far selected the simplest hyperbolic space and the simplest node distribution. We have next to select the simplest node connection probability function. To do so, we observe that the disc is a compact set, and R is its characteristic scale. A natural candidate for the simplest connection probability function is then a function proportional to the maximum-entropy distribution with the compact support $[0, R]$, which is again the uniform distribution on this interval. The simplest such function is the step function on $[0, R]$. In other words, we connect each pair of nodes by a link if the hyperbolic distance between them is $x \leqslant R$. It turns out that the described model produces graphs with the power-law distribution $P(k) \approx n(k)/N$ of node degrees k:

$$P(k) \sim k^{-3}. \tag{8}$$

Note that we have done nothing to enforce this power law. It appears as a simple consequence of the negative curvature of the underlying space. To understand why, recall (Section III B) that a uniform spatial node distribution in hyperbolic spaces implies that the node density grows exponentially as a function of the distance from a point. Specifically, if this point is the center of our disc, then the normalized node density $\rho(r)$, where $r \in [0, R]$ is the distance from the disc center, is:

$$\rho(r) = \frac{\sinh(r)}{\cosh R - 1} \approx e^{r-R} \sim e^r, \tag{9}$$

i.e., the number of nodes $n(r) = N\rho(r)$ at distance r from the disc center is proportional to e^r. In the Euclidean case, the number of nodes is proportional to r.

It is then a matter of simple (but lengthy) analytic calculations to find the average degree $k(r)$ of nodes located at distance r from the disc center. The result is that the step-function connection probability yields:

$$k(r) \sim e^{-\frac{1}{2}r}, \tag{10}$$

i.e., the average node degree decreases exponentially with the distance from the disc center. In fact, we can calculate the exact analytic expression for $k(r)$. This expression is rather long. We omit it for brevity, but show in Fig. 2 that it perfectly matches simulations. Taken together, Eqs. (9,10,7) yield that

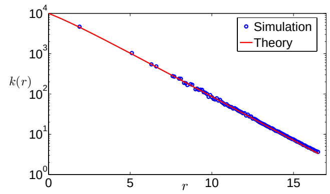

FIG. 2: Average degree at distance r from the origin.

graphs constructed by this model have a power-law node degree distribution with exponent $\gamma = 1+\alpha/\beta = 1+1/\left(\frac{1}{2}\right) = 3$.

We can easily alter the described model. For example, we can distribute nodes non-uniformly on the disc. The most natural generalization of node density in Eq. (9) appears to be:

$$\rho(r) \approx \alpha e^{\alpha(r-R)} \sim e^{\alpha r}, \tag{11}$$

with $\alpha = 1$ corresponding to the hyperbolically uniform node distribution. It turns out—we omit these calculations—that the average node degree $k(r)$ decreases as:

$$k(r) \sim \begin{cases} e^{-\frac{1}{2}r} & \text{if } \alpha \geqslant \frac{1}{2}, \\ e^{-\alpha r} & \text{if } \alpha \leqslant \frac{1}{2}, \end{cases} \tag{12}$$

and, consequently, the node degree distribution $P(k)$ is a power law. Specifically, one can show—we omit these calculations as well—that the degree distribution is:

$$P(k) = 2\alpha\xi^{2\alpha}\Gamma\left(k-2\alpha,\xi\right), \tag{13}$$

where $\xi = \bar{k}(2\alpha-1)/(2\alpha)$, $\bar{k} = \sum_k kP(k)$ is the average degree, and Γ is the incomplete gamma function. For large k, Eq. (13) scales as:

$$P(k) \sim k^{-\gamma} \tag{14}$$

with:

$$\gamma = \begin{cases} 2\alpha+1 & \text{if } \alpha \geqslant \frac{1}{2}, \\ 2 & \text{if } \alpha \leqslant \frac{1}{2}. \end{cases} \tag{15}$$

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 $\gamma \geqslant 2$, as observed in a majority of known complex networks.

V. ROUTING ON THE MODELED NETWORKS

In this section we first briefly describe the properties of networks generated by our model, and then focus on the efficiency characteristics of greedy routing and its modifications in these networks.

A. Modeled networks

In all our simulations, unless mentioned otherwise, we fix the target number of nodes in the network to $N = 10^4$ and its average degree to $\bar{k} = 6.5$, which is roughly the same as in Internet's AS topology. Given a target number of nodes N and average degree $\bar{k}$, we generate our networks as follows:

- Fix the radius R of the hyperbolic disc according to $N = \kappa e^{R/2}$, where parameter κ is used to tune the average degree to target $\bar{k}$. This relationship between N and R ensures that the network remains sparse in the large-graph limit $N \to \infty$.
- Assign to each node an angular coordinate $\theta \in [0, 2\pi)$ distributed uniformly.
- Assign to each node a radial coordinate $r \in [0, R]$ with probability $\rho(r) = \alpha e^{\alpha r}(e^{\alpha R}-1)^{-1}$, $\alpha \in [1/2, 1]$.
- Connect every pair of nodes whenever the hyperbolic distance between them is smaller than R. The hyperbolic distance x between two nodes with coordinates (r, θ) and (r', θ') is given by the hyperbolic law of cosines

$$\cosh x = \cosh r \cosh r' - \sinh r \sinh r' \cos \Delta\theta \tag{16}$$

where $\Delta\theta = \min\left(|\theta-\theta'|, 2\pi - |\theta-\theta'|\right)$.

In Fig. 3(a), we visualize one network instance of small size. We notice that hyperbolic geometry prevents peripheral nodes from connecting to each other, even if the Euclidean distance between them is small. This effect is due to the visualization settings in this figure. We set Euclidean radial distances to hyperbolic ones, but hyperbolic distances in the tangential direction grow exponentially with the distance from the disc center. Therefore a Euclidean distance in the tangential direction corresponds to a longer hyperbolic distance than the same Euclidean distance in the radial direction. The farther from the disc center, the exponentially stronger this disproportion. This explains why most links appear radially oriented.

In Fig. 4, we show the degree distribution, correlations, and clustering, i.e., the dK-properties [32] in our modeled networks. We observe agreement between simulation results and the analytical prediction for the degree distribution in Eq. (13).

We see that all our networks possess strong clustering. Strong clustering, or large numbers of triangles in generated networks, is a simple consequence of the triangle inequality in the hyperbolic plane. 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 our model. We also observe that the smaller the γ, the stronger the clustering of high-degree nodes, which means that in networks with low γ's, high-degree hubs participate in

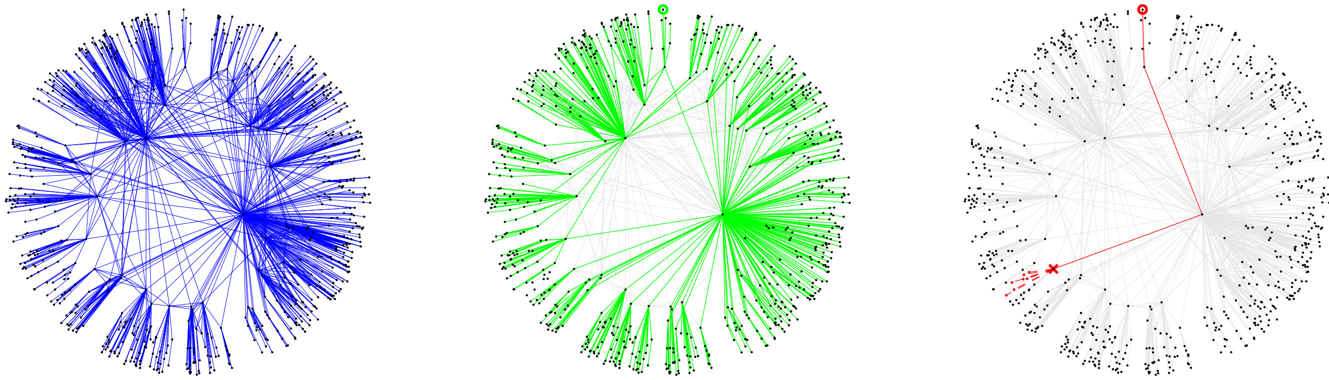

(a)Graph instance (b)Successful paths (c)Unsuccessful paths

FIG. 3: Visualization of a modeled network and greedy routing on it. Fig. (a) shows a modeled network with $N = 740$ nodes and $\gamma = 2.2$ embedded in the hyperbolic plane. For visualization purposes, the hyperbolic plane is not represented as in the Poincaré disc model. Instead, the Euclidean distances between all nodes and the disc center are their hyperbolic distances from the origin. Fig. (b) and (c) show all the links span by the successful and unsuccessful paths from the top node to all other nodes. For all the seven unsuccessful destinations, the last hop on the unsuccessful paths to them, i.e., the *local minimum* marked by the cross, is the same. The dashed lines in Fig. (c) show the unsuccessful destinations by connecting them to their local minimum.

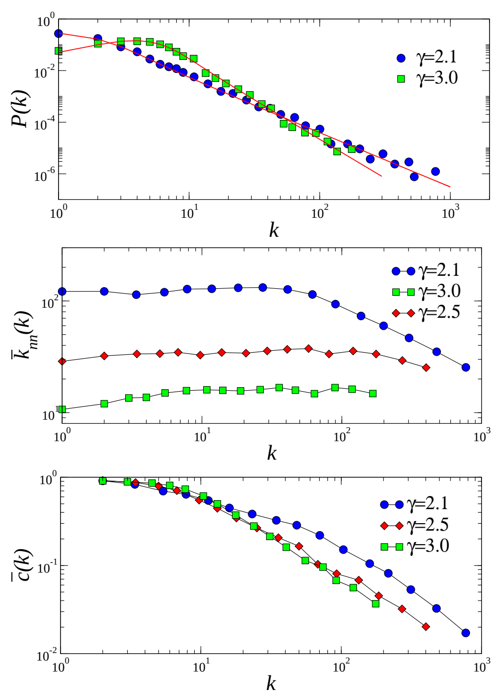

FIG. 4: The dK-properties [32] of simulated networks. From top to bottom: degree distribution $P(k)$, degree correlation as the average degree $\bar{k}_{nn}(k)$ of neighbors of k-degree nodes, and average clustering $\bar{c}(k)$ of k-degree nodes. The degree distribution for $\gamma = 2.5$ is not shown for clarity. Solid lines are the theoretical prediction given by Eq. (13).

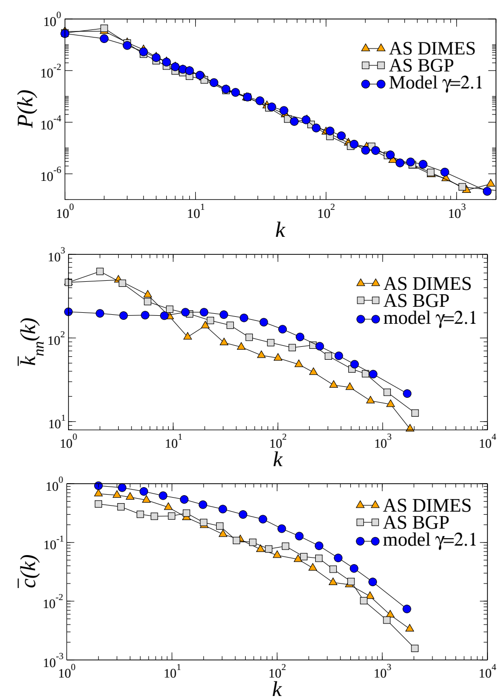

FIG. 5: Simulated networks with $\gamma = 2.1$ vs. AS topologies from RouteViews BGP tables [45] and DIMES traceroute data [14].

more triangles than in networks with large γ's. This fact turns out to have important consequences for the performance of greedy routing in Section V B.

In Fig. 5, we compare the Internet AS topology and our networks with $\gamma = 2.1$ (see Table I). We use two different sources of Internet topology data: BGP tables from RouteViews [45] and traceroute data from the DIMES project [14]. The degree distribution in our networks is remarkably close to the empirical AS degree distribution, as expected. Although our networks do not show the same degree correlations as in the Internet for small degrees k, the slope of the average neighbor's degree function $\bar{k}_{nn}(k)$ for large k is the same. Surprisingly, the shape of the clustering curve $\bar{c}(k)$ for our networks is similar to the Internet's. We currently do not have any satisfactory explanation for these coincidences between our toy model and the real Internet. We see that although the clustering shape is the same, the absolute values of clustering in our model are larger than in the Internet. Note that the DIMES clustering at small degrees is larger than BGP's because DIMES finds more "missing" links between small-degree ASs. It is therefore quite plausible that the clustering in the real Internet—once all the "missing" links are added—is even stronger and thus closer to what our model produces.

B. Greedy routing

We now evaluate the performance of greedy routing 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.

Below we report simulation results for two forms of greedy routing, *original* and *modified*. In both algorithms, a current hop selects as the next hop the neighbor that is closest to the destination in the hyperbolic space. The subtle difference between the two is in the failure detection mechanism. The original 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 modified algorithm does not always do so. The current hop excludes itself from any distance comparison operations, and finds the neighbor closest to the destination. The packet is dropped only if this neighbor is the same as the previous hop from which the current hop receives the packet.

We have also experimented with a variety of other greedy routing modifications, most notably with *geodesic routing*, which selects, among all the downstream neighbors of the current hop, the one which is closest to the hyperbolic geodesic connecting the source and destination. All these modifications deliver quite similar results, with only subtle differences across all the metrics we consider. Some algorithms perform slightly better w.r.t. some metrics but slightly worse w.r.t. others.

We compute the following metrics: (i) the percentage of successful paths, p_s, which is the proportion of paths that reach their destinations; (ii) the average hop-length $\bar{h}$ of successful paths; and (iii) the average and maximum stretch of successful paths.

Since we have not only graphs, but also hyperbolic spaces underneath, we compute not one, but three types of stretch. The first stretch is the standard hop stretch defined as the ratio between the hop-lengths of greedy routing paths and the corresponding shortest paths in the graph. We denote its average and maximum by s_1 and $\max(s_1)$. The other two stretches are *hyperbolic*. They measure the deviation of the hyperbolic length, traveled by a packet along either the greedy or shortest path, from the hyperbolic distance between the source and destination. Formally, let (s,t) be a source-destination pair and let $s = h_0, h_1, ..., h_\tau = t$ be the greedy or shortest path between s and t, and τ its hop length. Further, let $d_i, i = 1...\tau$, be the hyperbolic distance between h_i and h_{i-1}. The hyperbolic stretch is the ratio $\sum_i d_i/d_{st}$, where d_{st} is the hyperbolic distance between s and t. For greedy routing paths, we denote the average and maximum of this stretch by s_2 and $\max(s_2)$; for shortest paths—by s_3 and $\max(s_3)$. The lower these two stretches, the closer the greedy and shortest paths stay to the hyperbolic geodesics, and the more congruent the network topology is with the underlying geometry.

We first focus on static networks, where the network topology does not change, and then move to dynamic networks, where we emulate link failures by randomly selecting and removing one or more links from the topology. For each generated network instance, we extract the giant connected component (GCC), and perform greedy routing between 10^4 random source-destination pairs belonging to the GCC. In Fig. 3(b,c) we visualize the greedy routing performance in a small network instance.

1. *Static networks*

Fig. 6(i) shows the success ratio (p_s) of our greedy routing in networks with different γ's, while Fig. 6(ii) shows the corresponding average number of hops ($\bar{h}$) for the successful paths. For each value of γ, we average the results over 5 different network instances. First, we observe that p_s decreases as γ increases, and that modified greedy routing performs slightly better than the original one for all values of γ. Second, we can see that smaller values of γ, e.g., $\gamma \leqslant 2.4$, observed in many complex networks including the Internet, maximize the efficiency of greedy routing, yielding remarkably high success ratios close to 1. For example, when $\gamma = 2.1$, i.e., equal to γ observed in the AS Internet, original greedy routing yields $p_s = 0.99920$ and the modified one gives $p_s = 0.99986$. Further, we observe an increasing trend in $\bar{h}$ as we increase γ, with the two greedy routing algorithms performing approximately the same.

Fig. 7 shows the stretch results. For each value of γ, the maximum stretch corresponds to the maximum observed value across the 5 different network instances, while the average stretch is taken as the average across these instances. We observe that the average hop stretch s_1 for both greedy routing algorithms is approximately 1 for all values of γ, implying that all paths are approximately optimal. While the difference between the two algorithms is not notable in terms of s_1, we can observe some differences in terms of $\max(s_1)$. We see that original greedy routing never performs worse than the modified one, which is expected because the original algorithm never sends the packet to a next hop that is hyperbolically farther from the destination than the current hop, while the modified algorithm sometimes does so to increase the success ratio.

Remarkably, the original algorithm yields both $s_1 = 1$ and $\max(s_1) = 1$ for $\gamma \leqslant 2.2$, which means that *all greedy routing paths are shortest paths*. The modified algorithm for the same range of γ, gives $s_1 \approx 1$ and $\max(s_1) \leqslant 1.25$. There-

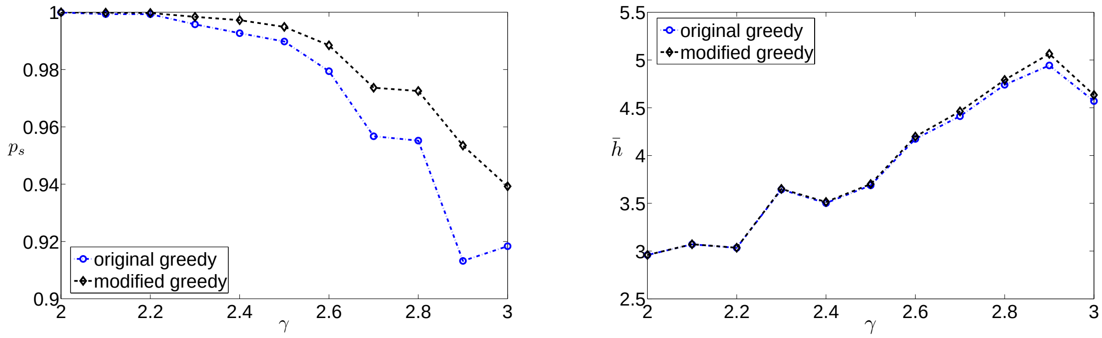

(i) (ii)

FIG. 6: (i) Percentage p_s of successful greedy routing paths, and (ii) their average hop-length $\bar{h}$. (Static networks.)

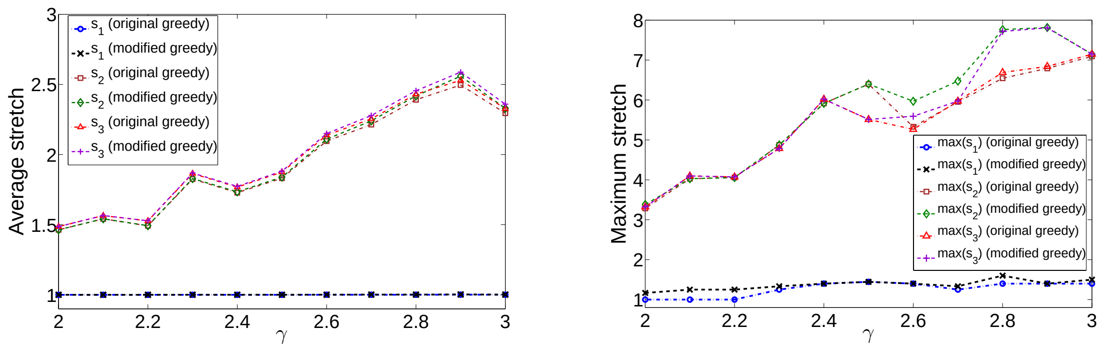

(i) (ii)

FIG. 7: (i) Average and (ii) maximum stretch. (Static networks.)

fore, while the modified algorithm performs slightly better in terms of the success ratio than the original one, it performs slightly worse in terms of stretch.

Interestingly, we see that for each algorithm, the hyperbolic stretch of shortest paths (s_3 and $\max(s_3)$) is slightly worse (larger) than of greedy paths (s_2 and $\max(s_2)$). In that sense, we can say, informally, that for small γ's, *greedy routing paths are "shorter than shortest"* as for small γ they are shortest hop-wise and shorter hyperbolically. This effect is expected because shortest path computation algorithms (e.g., Dijkstra) do not care about hyperbolic distances, while greedy routing does. At the same time, the differences between s_2 and s_3, and between $\max(s_2)$ and $\max(s_3)$ are minimal, which is also expected because the hop stretch s_1 and $\max(s_1)$ is so low. Finally, we see that the increase of stretches s_2 and s_3 with γ is similar to the one of the path hop length $\bar{h}$ in Fig. 6(ii), as also expected since longer paths are more likely to travel away from the the hyperbolic geodesic.

Summarizing, both greedy routing algorithms are extremely efficient in static networks, especially for the smaller values of γ observed in a vast majority of complex networks, including the Internet. The algorithms yield high success ratios close to 1 and optimal (or almost optimal) path lengths.

2. *Dynamic networks*

We now evaluate the greedy routing performance in dynamic scenarios with link failures. For each value of γ, we randomly select a network instance from the ones considered earlier. For this particular instance, we consider the same source-destination pairs as before, and emulate failures by randomly selecting and removing one or more links from the topology. We consider the following two link-failure scenarios:

Scenario 1. In this scenario we study the degradation of the success ratio and stretch under removal of a percentage p_r of randomly selected links from the topology. After link removal, we compute the new GCC, and then, for all of the source-destination pairs that belong to the new GCC we compute the new success ratio p_s^{new}, and the new average and maximum stretch of the successful paths, denoted by s_1^{new} and $\max(s_1^{new})$ respectively. We vary p_r from 0% to 30%.

Scenario 2. In this scenario we provide a finer-grain view focusing on the paths that used a removed link. Specifically, we select one link at random, remove it from the topology, and compute the new GCC. Then, for all previously successful paths traversing this link, we compute the percentage of paths that remain successful, denoted by p_s^l. For these paths, we also compute the average and maximum stretch as before, which we now denote by s_1^l and $\max(s_1^l)$ respectively. We repeat this procedure over 1000 different links, and report the average p_s^l and the average s_1^l, and the maximum observed value of $\max(s_1^l)$.

Fig. 8 presents results for Scenario 1. For smaller values of γ, the success ratio p_s^{new} remains remarkably high, for all meaningful values of p_r. For example, modified greedy routing on networks with $\gamma = 2.1$ and $p_r \leqslant 0.1$, yields $p_s^{new} > 0.99$. In fact, this remains true for all networks with $\gamma \leqslant 2.5$, not shown in the figure for clarity. Note that $p_r = 0.1$ corresponds to removal of 10% of all links in the network. For comparison, the simultaneous failure of 10% links in the Internet is a rare catastrophe. Fig. 8(i) also shows that modified greedy routing outperforms the original algorithm. We also observe that smaller values of γ yield a higher success ratio. Fig. 8(ii) shows that the average routing stretch slightly increases as we increase p_r, for all values of γ. However, it still remains quite low. Also for clarity, we do not show results for maximum stretch in the figure. We report that for any of the two algorithms and for any value of p_r, $\max(s_1^{new}) \leqslant 2$ for $\gamma = 2.1$, and $\max(s_1^{new}) \leqslant 2.5$ for $\gamma = 2.6, 2.8$.

Fig. 9 presents results for Scenario 2. For the modified algorithm, the percentage p_s^l of paths that used a removed link and that found a by-pass after its removal is approximately equal to 1 for small γ's. This percentage decays slowly as γ increases. We also see in Fig. 9(i) that for all values of γ, modified greedy routing outperforms original greedy routing in terms of the success ratio p_s^l. In Fig. 9(ii) we see that the average routing stretch for both algorithms remains low, below 1.1, and the maximum routing stretch never exceeds 1.5. However, comparing Fig. 9(ii) with Fig. 7(i) and 7(ii), we detect a slight increase in stretch, as expected.

Summarizing, greedy routing strategies (e.g., our modified greedy routing algorithm) can be quite efficient and robust in dynamic network conditions. In particular, for the range of γ's we are mostly interested in, they maintain remarkably high success ratios and low stretch.

3. Random and exponential networks

We provide further evidence that the efficiency of greedy routing on scale-free network topologies is due to their congruency with the underlying hyperbolic geometries. We consider other network topologies, also embedded in a hyperbolic space, and find that their success ratio is significantly worse.

Specifically, we place nodes on the hyperbolic plane as before, and construct two types of networks: (i) classical random graphs [16], where links exist between any two nodes with a constant probability p, independent of the hyperbolic distance between them; and (ii) exponential graphs, in which links exist between any two nodes with probability e^{-x}, where x is the hyperbolic distance between the two nodes. It is known that the node degree distribution is binomial in the first case, and we can show (but skip for brevity) that this distribution is exponential in the second case. For random networks, the average success ratio is $p_s = 0.00254$, which is remarkably low. For exponential networks, $p_s = 0.68$, still significantly lower than the success ratios on scale-free topologies. These results underscore that topologies of other, non-scale-free networks are not naturally congruent with hyperbolic geometry.

C. Closer look at success ratio

Although the success ratios in scale-free networks with small γ's are extremely close to 1 (and we have encountered graph instances where it is 1), it is not exactly 1 on average—and it could not be, thanks to randomness of graph construction in our models. We believe that if the hyperbolic space underlying the real Internet is reconstructed exactly, then the success ratio will be 1. This belief is supported, in part, by the fact that any graph can be embedded in the hyperbolic plane such that the success ratio is 1 [24]. However, since the real hyperbolic space hidden beneath the Internet may not be reconstructed exactly but only approximately, and since there may be different applications of our work in this paper (see Section VII), we discuss some techniques that can boost the success ratio to 1, and report basic statistics relevant for the performance of these techniques on our modeled networks. These statistics also shed some light on the nature of the unsuccessful paths in our networks.

The first obvious boosting technique is to forward a packet to a landmark if the source cannot reach the destination using greedy routing. Here a *landmark* is a node that the source can reach and that can reach the destination. The relevant statistics for this technique include: (i) the percentage of sources that can reach all destinations $p_{src}^{all-dsts}$, (ii) the percentage of destinations that all sources can reach $p_{dst}^{all-srcs}$, and (iii) the percentage of nodes that can reach all other nodes and be reached by all other nodes p^{all}, i.e., the intersection of the previous two sets. We compute all these statistics by performing greedy routing between all possible source-destination pairs in our modeled networks with $N = 10^3$ nodes and average degree $\bar{k} = 6.5$. We average results across 5 different network instances, and report them in Table III and IV. We see that for small γ's, there are a significant number of nodes that can serve as landmarks. We also observe a common trend that all the statistics degrade as γ increases.

γ	p_s	$p_{src}^{all-dsts}$	$p_{dst}^{all-srcs}$	p^{all}
2.1	0.99943	0.59892	0.99576	0.59651
2.4	0.98613	0.19979	0.39829	0.19937
2.7	0.96141	0	0.00480	0
3.0	0.82119	0	0	0
Exponential	0.52799	0	0	0
Random	0.01575	0	0	0

TABLE III: Reachability of original greedy routing.

The second boosting technique is applicable to cases where we can slightly alter the network topology by adding a small number of links (or virtual links, e.g., tunnels) to boost the success ratio to 1. We have designed a straightforward greedy

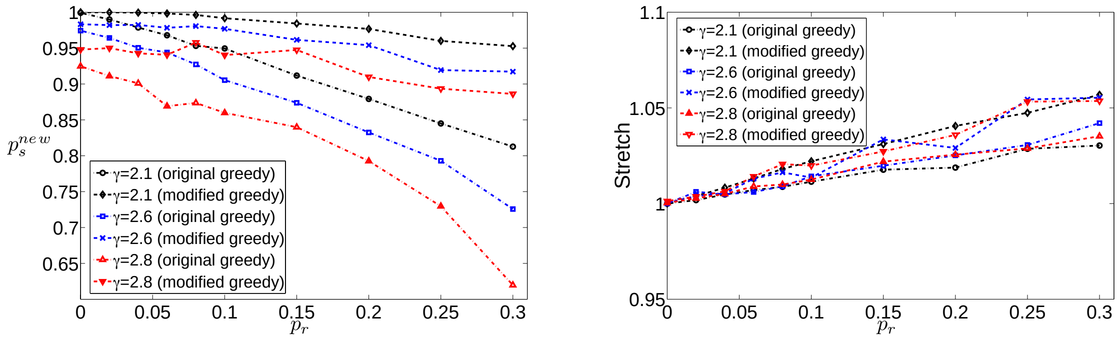

(i) (ii)

FIG. 8: (i) Success ratio p_s^{new} and (ii) stretch. (Dynamic networks, Scenario 1.)

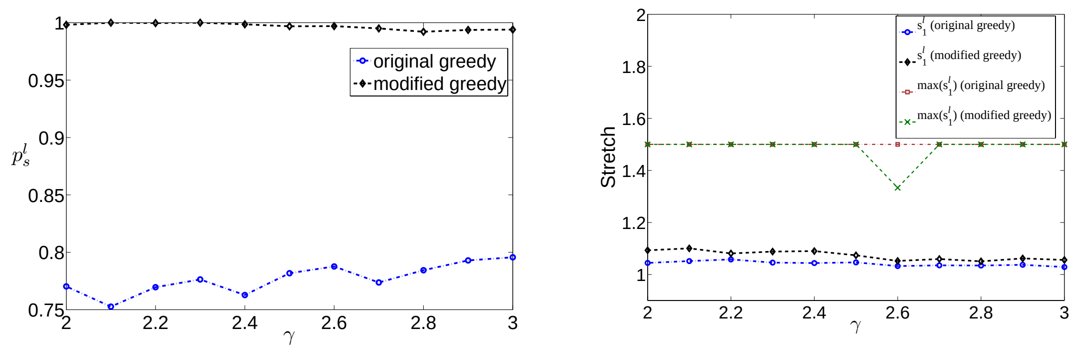

(i) (ii)

FIG. 9: (i) Success ratio p_s^l and (ii) stretch. (Dynamic networks, Scenario 2.)

γ	p_s	$p_{src}^{all-dsts}$	$p_{dst}^{all-srcs}$	p^{all}
2.1	0.99967	0.79948	0.99890	0.79923
2.4	0.99193	0.39850	0.78759	0.39849
2.7	0.97828	0.04084	0.39689	0.04084
3.0	0.85932	0	0.00250	0
Exponential	0.72622	0	0	0
Random	0.02328	0	0	0

TABLE IV: Reachability of modified greedy routing.

algorithm to compute the approximately minimum number of links to be added to a graph to boost its success ratio to 1. We omit this algorithm specification for brevity, and report its results in Fig. 10. We see that our networks are just a small number of links away from being 100%-successful. For instance, fewer than 10 links are required to make all paths successful for $\gamma \leqslant 2.4$. We also detect a strong correlation between the number of added links and the number of local minima that existed in the network before any link addition. Interestingly, we require fewer links than the number of local minima, as adding one link can sometimes eliminate several local minima. We also computed the number of added links in exponential and random graphs of the same size and average degree. In exponential graphs with original and modified greedy routing we need to add 645 and 694 links. For random graphs, these numbers are respectively 1604 and 3466. These results once again confirm incongruity between hyperbolic space and these other topologies.

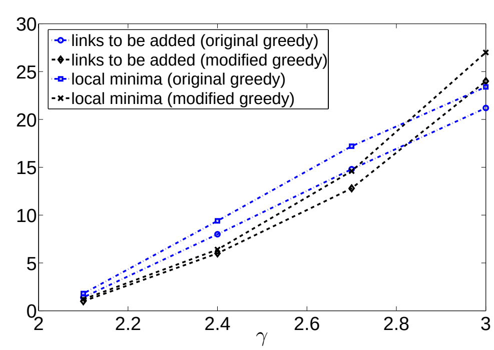

FIG. 10: Number of local minima and links to add to boost the success ratio to 1.

There are also other boosting techniques that try to avoid or escape from local minima. These techniques are based on various forms of path backtracking, lookahead, simulated annealing, etc. Their analysis is more complicated, and we omit them for brevity.

VI. RELATED WORK

The literature dealing with the design of new scalable routing for the Internet is abundant. Some recent papers in this area include [10, 11, 20, 36, 50]. There has also been recent interest in a variety of compact routing (CR) [1, 7, 52] to address scalability concerns for Internet routing. CR currently promises the best scaling among known routing alternatives. No CR algorithm can guarantee, however, convergence cost scaling better than linear with network size [26]. Assuming an appropriate mapping function between scale-free topologies with small γ's and a hyperbolic metric space, the approach outlined in this paper achieves:

- zero convergence costs (obviously) vs. linear in CR;
- constant address size (node coordinates in hyperbolic spaces) vs. polylogarithmic in CR;
- maximum stretch equal to 1 (all paths are shortest) with minor stretch degradation under network dynamics vs. different constants greater than 1 in CR;
- routing table sizes proportional to node degrees (as each node stores nothing but its coordinates and the coordinates of its neighbors) vs. polylogarithmic and polynomial in CR with topology-aware and topology-unaware addressing;
- success ratio nearly 1 (or 1 with boosting techniques) and minor success ratio degradation under network dynamics vs. 1 in CR.

We also juxtapose our approach with recent work [24] showing that any graph can be embedded in a hyperbolic space to achieve 100%-successful routing. The approach in [24] is equivalent to CR [26] in the sense that they both require global knowledge, i.e., a full view of the graph. The best way to understand the principal difference between [24, 26] and our approach is to recognize that we essentially have the opposite problem formulation. For example, in [24], the problem is: given the topology of an entire graph, find its embedding in a hyperbolic space such that greedy routing is 100%-successful. In our case, the problem is: given a hyperbolic space, check if there exists a simple graph construction procedure that would naturally yield scale-free topologies, and if greedy routing on these topologies would be efficient.

This difference in problem formulation results in a significant difference in what the two approaches can accomplish. In [24, 26], one performs traditional routing "in the light." Knowing the topology of any graph—which requires a lot of communication overhead in distributed settings—one finds a way to efficiently route on this topology. The fact that it is possible to do is by no means surprising, of course. In our case, we route "in the dark." Since we do not send any routing control messages, we cannot know the network topology. Therefore we cannot efficiently route on an arbitrary network. We can route efficiently only on the networks whose topologies are congruent with underlying spaces. We have shown that scale-free networks appear to be naturally congruent with hyperbolic spaces that are thus potentially good candidates for embeddings of the real Internet topology.

To the best of our knowledge, the first mention of hyperbolic spaces in the context of the Internet appeared in the work by Shavitt and Tankel [47, 48]. In this work, the authors consider heuristics embedding Internet distances into hyperbolic spaces. They showed that the negative curvature of target spaces improves the efficiency and accuracy of overlay network constructions and Internet delay distance estimations.

Network delay estimation services, e.g. [30, 41, 51] or Vivaldi [13], try to accurately estimate network delays between Internet hosts by embedding them into virtual coordinate spaces such that the distances between nodes in this space are approximately equal to the corresponding inter-node delays. This problem is therefore directly related to the problem of low-distortion embedding of finite metric spaces into normed spaces. This related problem has seen significant research progress [33].

More recently, Krauthgamer and Lee [25] refer to network embedding applications—and, in particular, to [47]—as one of the motivations for constructing efficient algorithms for various problems, including routing. The efficiency of the algorithms constructed in this work is rooted in the negative curvature of underlying metric spaces. In the same paper, the authors popularize the observation, originally due to Gromov [18, 19], of the intimate connections between hyperbolic spaces and trees.

The first popularization of greedy routing as a mechanism that might be responsible for efficient routing "in the dark," i.e., without the knowledge of network topology, is due to J. Kleinberg [22]. A vast amount of literature followed this seminal work, as reviewed in [23]. This work includes searching the Web [34], social [29, 54], and overlay [46] networks.

Overlay networks, such as CAN [43] or Chord [49] (see also survey [31]), use greedy routing in the overlay virtual topology, which does not have to be related to the topology of the real underlying network.

Another area in networking where greedy routing is a core element is geographic routing [21, 42]. Since, as in overlay networks, the network topology may not accurately reflect the underlying geographic space, one of the main problems in this line of research is how to deal with local minima.

VII. CONCLUSION

This paper shows that negatively curved spaces lead to natural emergence of scale-free network topologies, quite similar to those observed in the Internet and other complex networks. Yet more remarkably, greedy routing in these settings perform exceptionally well, better than all existing proposals for compact routing for instance. The reason for this dramatic improvement is a routing paradigm shift: routing no longer requires any non-local knowledge about observed network topology, but relies solely on hidden geometries to find paths to destinations. This geometric underpinning drastically simplifies the routing function that no longer performs any complicated information processing or computation, but just forwards packets in the right hidden directions towards destinations. Congruency between observed scale-free topologies and hidden hyperbolic geometries ensures the efficiency of such "dumb routing in the dark." We found that this congruency is strongest and greedy routing is maximally efficient on scale-free networks with small values of the exponent of the power-law degree distribution. These small values characterize many real complex networks, including the Internet.

We roughly split potential applications of these findings in two categories. The first category concerns synthetic networks, such as overlays. In this case, we can freely design a hidden hyperbolic space and build a congruent network topology over it. Future work in this direction includes constructing models of networks that grow over hyperbolic spaces. All the network models considered in this paper generate a whole network at once, and therefore they are not immediately applicable for overlay-like applications.

More interesting but also more challenging are the applications for real networks—the Internet in the first place. The practical challenges in this case are due to obvious difficulties in finding the exact structure of hidden hyperbolic spaces underlying real networks, and node coordinates in these spaces. Indeed, we can hardly expect that the over-simplistic abstraction of negatively curved geometries provided by the hyperbolic plane is a serious candidate for hidden spaces underlying real networks. However, quite promising is the rigidity property of hyperbolic metric spaces [19], suggesting that we do not have to always know the hyperbolic structure exactly—an approximate, coarse-grained view of a hyperbolic space often suffices for practical purposes.

The most promising future research path appears to be a step-by-step narrowing of the space of possible hidden metric spaces. One way to proceed is to compare their geometric and topological properties with the topology data available for the Internet and other real-world networks. In that context, the conceptual contribution of this work is that we have narrowed down the class of hidden metric spaces to spaces of negative curvature.

[1] I. Abraham, C. Gavoille, D. Malkhi, N. Nisan, and M. Thorup. Compact name-independent routing with minimum stretch. In *SPAA*, 2004.

[2] J. W. Anderson. *Hyperbolic Geometry*. Springer-Verlag, London, 2005.

[3] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[4] M. Boguñá and R. Pastor-Satorras. Class of correlated random networks with hidden variables. *Phys Rev E*, 68:036112, 2003.

[5] M. Boguñá, R. Pastor-Satorras, A. Díaz-Guilera, and A. Arenas. Models of social networks based on social distance attachment. *Phys Rev E*, 70:056122, 2004.

[6] K. Börner, J. T. Maru, and R. L. Goldstone. The simultaneous evolution of author and paper networks. *PNAS*, 101(Suppl. 1):5266–5273, 2004.

[7] A. Brady and L. Cowen. Compact routing on power-law graphs with additive stretch. In *ALENEX*, 2006.

[8] M. R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer-Verlag, Berlin, 1999.

[9] D. Burago, Y. Burago, and S. Ivanov. *A Course in Metric Geometry*. AMS, Providence, 2001.

[10] M. Caesar, M. Castro, E. B. Nightingale, G. O'Shea, and A. Rowstron. Virtual ring routing: Network routing inspired by DHTs. In *SIGCOMM*, 2006.

[11] M. Caesar, T. Condie, J. Kannan, K. Lakshminarayanan, I. Stoica, and S. Shenker. ROFL: Routing on flat labels. In *SIGCOMM*, 2006.

[12] A. Clauset, C. Moore, and M. E. J. Newman. Hierarchical structure and the prediction of missing links in networks. *Nature*, 453:98–101, 2008.

[13] F. Dabek, R. Cox, F. Kaashoek, and R. Morris. Vivaldi: A decentralized network coordinate system. In *SIGCOMM*, 2004.

[14] The DIMES project. `http://www.netdimes.org/`.

[15] X. Dimitropoulos, D. Krioukov, G. Riley, and kc claffy. Revealing the Autonomous System taxonomy: The machine learning approach. In *PAM*, 2006.

[16] P. Erdős and A. Rényi. On random graphs. *Publ Math*, 6:290–297, 1959.

[17] M. Girvan and M. E. J. Newman. Community structure in social and biological networks. *Proc Natl Acad Sci USA*, 99:7821–7826, 2002.

[18] M. Gromov. Hyperbolic groups. In *Essays in Group Theory*, volume 8 of *Math Sci Res Inst Publ*, pages 75–263. Springer, New York, 1987.

[19] M. Gromov. *Metric Structures for Riemannian and Non-Riemannian Spaces*. Birkhäuser, Boston, 2007.

[20] R. Gummadi, R. Govindan, N. Kothari, B. Karp, Y.-J. Kim, and S. Shenker. Reduced state routing in the Internet. In *HotNets*, 2004.

[21] Y.-J. Kim, R. Govindan, B. Karp, and S. Shenker. Geographic routing made practical. In *NSDI*, 2005.

[22] J. Kleinberg. Navigation in a small world. *Nature*, 406:845, 2000.

[23] J. Kleinberg. Complex networks and decentralized search algorithms. In *ICM*, 2006.

[24] R. Kleinberg. Geographic routing using hyperbolic space. In *INFOCOM*, 2007.

[25] R. Krauthgamer and J. R. Lee. Algorithms on negatively curved spaces. In *FOCS*, 2006.

[26] D. Krioukov, kc claffy, K. Fall, and A. Brady. On compact routing for the Internet. *Comput Commun Rev*, 37(3):41–52, 2007.

[27] E. A. Leicht, P. Holme, and M. E. J. Newman. Vertex similarity in networks. *Phys Rev E*, 73:026120, 2006.

[28] B. Leiner, V. Cerf, D. Clark, R. Kahn, L. Kleinrock, D. Lynch, J. Postel, L. Roberts, and S. Wolff. A brief history of the Internet. *Commun ACM*, 40(2):102–108, 1997.

[29] D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, and A. Tomkins. Geographic routing in social networks. *PNAS*, 102:11623–11628, 2005.

[30] H. Lim, J. C. Hou, and C.-H. Choi. Constructing Internet coordinate system based on delay measurement. *IEEE ACM T Network*, 13(3), 2005.

[31] 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 Surveys and Tutorials*, 7(2):72–93, 2005.

[32] P. Mahadevan, D. Krioukov, K. Fall, and A. Vahdat. Systematic topology analysis and generation using degree correlations. In *SIGCOMM*, 2006.

[33] J. Matoušek. *Lectures on Discrete Geometry*, chapter 15, Embedding Finite Metric Spaces into Normed Spaces. Springer, New York, 2002.

[34] F. Menczer. Growing and navigating the small world Web by local content. *PNAS*, 99:14014–14019, 2002.

[35] D. Meyer, L. Zhang, and K. Fall, editors. *Report from the IAB Workshop on Routing and Addressing*. The Internet Architecture Board, 2007.

[36] M. Motiwala, N. Feamster, and S. Vempala. Path splicing: Reliable connectivity with rapid recovery. In *HotNets*, 2007.

[37] L. Muchnik, R. Itzhack, S. Solomon, and Y. Louzoun. Self-emergence of knowledge trees: Extraction of the Wikipedia hierarchies. *Phys Rev E*, 76:016106, 2007.

[38] M. Nei and S. Kumar. *Molecular Evolution and Phylogenetics*. Oxford University Press, Oxford, 2000.

[39] M. E. J. Newman. The structure and function of complex networks. *SIAM Rev*, 45(2):167–256, 2003.

[40] M. E. J. Newman. Power laws, Pareto distributions and Zipf's law. *Contemp Phys*, 46(5):323–351, 2005.

[41] T. S. E. Ng and H. Zhang. Predicting Internet network distance with coordinates-based approaches. In *INFOCOM*, 2002.

[42] A. Rao, S. Ratnasamy, C. Papadimitriou, S. Shenker, and I. Stoica. Geographic routing without location information. In *MobiCom*, 2003.

[43] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content-addressable network. In *SIGCOMM*, 2001.

[44] S. Redner. How popular is your paper? An empirical study of the citation distribution. *Eur Phys J B*, 4:131–134, 1998.

[45] University of Oregon RouteViews Project. `http://www.routeviews.org/`.

[46] O. Sandberg. Distributed routing in small-world networks. In *ALENEX*, 2006.

[47] Y. Shavitt and T. Tankel. On the curvature of the Internet and its usage for overlay construction and distance estimation. In *INFOCOM*, 2004.

[48] Y. Shavitt and T. Tankel. Hyperbolic embedding of Internet graph for distance estimation and overlay construction. *To appear in IEEE ACM T Network*, 2008.

[49] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup protocol for Internet applications. *IEEE ACM T Network*, 11(1), 2003.

[50] L. Subramanian, M. Caesar, C. T. Ee, M. Handley, M. Mao, S. Shenker, and I. Stoica. HLP: A next generation inter-domain routing protocol. In *SIGCOMM*, 2005.

[51] L. Tang and M. Crovella. Virtual landmarks for the Internet. In *IMC*, 2003.

[52] M. Thorup and U. Zwick. Compact routing schemes. In *SPAA*, 2001.

[53] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32:425–443, 1967.

[54] D. J. Watts, P. S. Dodds, and M. E. J. Newman. Identity and search in social networks. *Science*, 296:1302–1305, 2002.