

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

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

Marián Boguñá
Universitat de Barcelona
marian.boguna@ub.edu

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

1. 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 [6] and emerging wireless and sensor networks [7].

In many other networks in nature however, nodes can efficiently communicate, even though they do not exchange any information about the current global state of the network topology. Milgram's 1969 experiments [9] showed a classic demonstration of this surprising effect: humans can find paths to destinations through their social acquaintance network, even though no human has global knowledge of its structure. Much later, Jon Kleinberg offered the first popular explanation [3]. 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 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. Therefore, 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 messages 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 [8] 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 [1].

Our work follows Kleinberg's formalism. 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. The main results of our work is that if we model hidden spaces as non-Euclidean *hyperbolic spaces*, then their negative curvature leads to: (i) natural emergence of scale-free topologies constructed over such hidden spaces; and (ii) extremely efficient greedy forwarding on these topologies, achieving almost 100% reachability and optimal (i.e., shortest) path lengths, even under dynamic network conditions.

2. SCALE-FREE NETWORKS AND HYPERBOLIC METRIC SPACES

The main metric property of hyperbolic geometry 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 $\sim e^R$. 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 surface of a sphere or the volume of a ball of radius R , measured as the number of nodes lying at or within R hops from the root, grow as b^R . Informally, hyperbolic spaces can therefore be thought of as "continuous versions" of trees.

To see why this exponential expansion of hidden space is intrinsic to complex 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 sub-subgroups, etc. The relationships between such groups and sub-groups can be approximated by tree-like structures, in which the distance between two nodes estimates how similar they are [10]. 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 [2], makes the space hyperbolic.

We now put these 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 simple 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 . The simplest way to place

N nodes on the hyperbolic plane is to distribute them uniformly over a disc of radius R . It turns out that selecting $N = \kappa e^{R/2}$ gives us control over the average node degree \bar{k} , which we can tune using parameter κ . The hyperbolically uniform node density implies that we assign the angular coordinates $\theta \in [0, 2\pi]$ to nodes with the uniform density $\rho(\theta) = 1/(2\pi)$, while the density for the radial coordinate $r \in [0, R]$ is exponential $\rho(r) = \sinh r / (\cosh R - 1) \approx e^{r-R} \sim e^r$ (vs. $\rho(r) \sim r$ in the Euclidean plane). The simplest connection probability we could think of is the step function $p(x) = \Theta(R - x)$, 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 $x \leq R$, where x is given by the hyperbolic law of cosines: $\cosh x = \cosh r \cosh r' - \sinh r \sinh r' \cos \Delta\theta$, with $\Delta\theta = \min(|\theta - \theta'|, 2\pi - |\theta - \theta'|)$. The described model produces graphs with the power-law node degree distribution $P(k) \approx n(k)/N \sim k^{-3}$.

We emphasize that we have done nothing to enforce this power law. It appears as a consequence of the negative curvature of the underlying space. To understand why, we have to calculate the average degree $k(r)$ of nodes located at distance r from the disc center. These analytic calculations are complicated, and the exact expression is rather long, so that we omit it here for brevity. What matters is that $k(r)$ decreases exponentially, $k(r) \sim e^{-r/2}$. Therefore, the inverse function is logarithmic, $r(k) \sim -2 \ln k$, and the node degree distribution in the network is approximately a power law, $P(k) \approx \rho[r(k)] |r'(k)| \sim k^{-3}$.

We can generalize the described model. For example, we can distribute nodes non-uniformly on the disc, $\rho(r) \approx \alpha e^{\alpha(r-R)} \sim e^{\alpha r}$, with $\alpha = 1$ corresponding to the hyperbolically uniform node distribution. It turns out—we omit the calculations—that the average node degree $k(r)$ decreases as $\sim e^{-r/2}$ if $\alpha \geq 1/2$, and as $\sim e^{-\alpha r}$ otherwise. As above, the node degree distribution $P(k)$ is a power law: $P(k) \sim k^{-\gamma}$, with $\gamma = 2\alpha + 1$ 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.

Our modeled 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 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. However, the step function connection probability does not allow us to tune clustering. To solve this problem we introduce the following modified connection probability $p(x) = 1/(1 + e^{\beta(x-R)/2})$, where x is the hyperbolic distance between nodes and $\beta > 1$ is a parameter. If $\beta \rightarrow \infty$, $p(x)$ becomes the step function, and clustering is maximized, while for finite values of β , the step function “smooths out” reducing the clustering strength. At $\beta \rightarrow 1$ clustering goes to zero.

As an example, in [5] we demonstrate that we can create networks whose topological properties closely match those of the AS-level Internet.

3. GREEDY FORWARDING

We now evaluate the performance of greedy forwarding (GF) 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 desti-

nation in the hyperbolic space. We present simulation results for two 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 first focus on static networks, where the network topology does not change, and then emulate the network topology dynamics by randomly removing one or more links from the topology. We fix the target number of nodes in the network to $N = 10^4$ and its average degree to $\bar{k} = 6.5$, which roughly matches the Internet’s AS topology. We use the step function connection probability. For each generated network, we extract the Giant Connected Component (GCC), and perform GF between 10^4 random source-destination pairs.

Static networks. Figures 1(a) and 1(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. Remarkably, 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. 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 2.

Dynamic networks. We next study the GF performance in dynamic networks 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 in the GCC. We repeat the procedure for 1000 random links, and report the average value for p_s^l . Figures 1(c) and 1(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. Due to space limitations we do not present stretch results. 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”

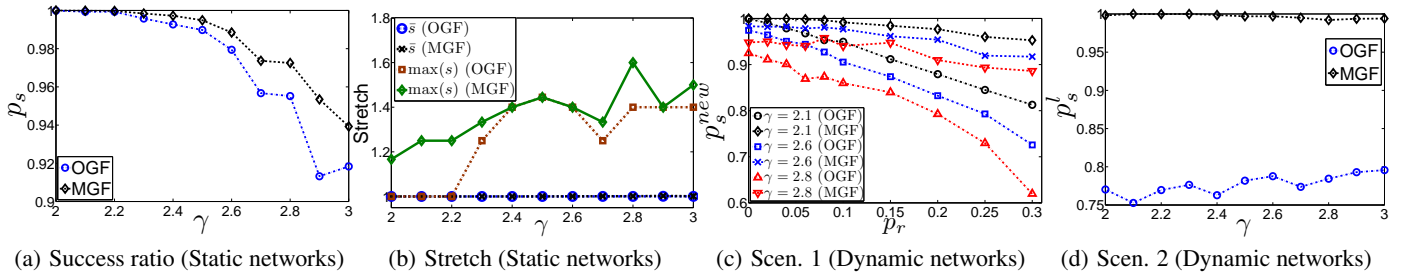


Figure 1: Performance of greedy forwarding (GF) in static and dynamic networks.




Figure 2: 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). The geodesics and greedy paths follow the same pattern, i.e., they are approximately congruent.

to find them.

As a final note, although the success ratios in scale-free networks with small γ 's are extremely close to 1, they are not exactly 1, thanks to randomness of graph construction in our models. However, we do have simple techniques that can boost the success to 1, discussed in [5].

4. FUTURE WORK

We can classify potential applications of our findings in two categories. The first category concerns synthetic networks, such as P2P overlays that provide a large flexibility for tasks, such as application-level routing, information sharing, and so on. 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 and maintaining networks that grow and shrink over hyperbolic spaces. The network model considered in this paper generates a whole network at once, and therefore it is not

immediately applicable for overlay-like applications.

More challenging are the applications for real networks—the Internet in the first place. Can we embed the real Internet topology into a hyperbolic space and greedily forward through this embedding with similar efficiency? This task is quite challenging as we no longer have freedom to build a graph to be congruent with the underlying space by construction; the graph is given to us. 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? The most related to this earlier work is the one by R. Kleinberg who shows in [4] how any graph can be embedded in the hyperbolic plane such that greedy forwarding's success ratio is 100%. However, the embedding requires full global knowledge of the graph topology and may need to be recomputed on link/node failures, things we want to avoid.

5. REFERENCES

- [1] S. N. Dorogovtsev and J. F. F. Mendes. Evolution of networks. *Advances in Physics*, 51:1079–1187, 2002.
- [2] M. Gromov. *Metric Structures for Riemannian and Non-Riemannian Spaces*. Birkhäuser, Boston, 2007.
- [3] J. Kleinberg. Navigation in a small world. *Nature*, 406:845, 2000.
- [4] R. Kleinberg. Geographic routing using hyperbolic space. In *INFOCOM*, 2007.
- [5] D. Krioukov, F. Papadopoulos, M. Boguna, and A. Vahdat. Efficient navigation in scale-free networks embedded in hyperbolic metric spaces. 2008. <http://arxiv.org/abs/0805.1266>.
- [6] D. Meyer, L. Zhang, and K. Fall, editors. *RFC4984*. The Internet Architecture Board, 2007.
- [7] V. Naumov and T. Gross. Scalability of routing methods in ad hoc networks. *Performance Evaluation*, 62:193–209, 2005.
- [8] M. E. J. Newman. The structure and function of complex networks. *SIAM Rev*, 45(2):167–256, 2003.
- [9] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32:425–443, 1969.
- [10] D. J. Watts, P. S. Dodds, and M. E. J. Newman. Identity and search in social networks. *Science*, 296:1302–1305, 2002.