Routing in Generalized Geometric Inhomogeneous Random Graphs
(Extended Abstract)⋆

Andrés SEVILLA[1⋆⋆][0000−0001−8867−5528] and Antonio FERNÁNDEZ ANTA[2][0000−0001−6501−2377]

[1] Dpto. de Sistemas Informáticos, Technical University of Madrid, Spain
andres.sevilla@upm.es
[2] IMDEA Networks Institute, Madrid, Spain
antonio.fernandez@imdea.org

Abstract. In this paper we study a new random graph model that we denote (κ,π)-KG and new greedy routing algorithms (of deterministic and probabilistic nature). The (κ,π)-KG graphs have power-law degree distribution and small-world properties. (κ,π)-KG roots on the Geometric Inhomogeneous Random Graph (GIRG) model, and hence they both preserve the properties of the hyperbolic graphs and avoid the problems of using hyperbolic cosines. In order to construct (κ,π)-KG graphs, we introduce two parameters κ and π in the process of building a (κ,π)-KG graph. With these parameters we can generate Kleinberg and power-law networks as especial cases of (κ,π)-KG. Also, we propose two new greedy routing algorithms to reduce the fail ratio and maintaining a good routing performance. The first algorithm is deterministic and the second is, in essence, a weighted random walk. We use simulation techniques to test our network model, and evaluate the new routing algorithms on the two graph models (GIRG and (κ,π)-KG). In our simulations, we evaluate the number of hops to reach a destination from a source and the routing fail ratio, and measure the impact of the parameters (κ and π) on the performance of the new routing algorithms. We observe that our graph model (κ,π)-KG is more flexible than GIRG, and the new routing algorithms have better performance than the routing algorithms previously proposed.

1 Introduction

In the latest years, geometric approaches have been used to build scale-free networks for modeling complex/large networks [1, 7, 9]. In particular, hyperbolic geometric models have been used to construct these networks [9]. In these models, each node is assigned a virtual coordinate in the hyperbolic space, and nodes are linked using a probability distribution based on a distance function. For instance, Boguñá et al. [1] proposed a

⋆ This submission is a short paper. This work was partially funded by the Spanish grant TIN2017-88749-R (DiscoEdge), the Region of Madrid EdgeData-CM program (P2018/TCS-4499), and the NSF of China grant 61520106005.
⋆⋆ Corresponding author.

method to embed Internet in a hyperbolic disc of radius R, where the node density grows exponentially with the distance from the disk center. The nodes are linked with a probability according to the product of their expected degrees (drawn from a power-law distribution) and re-scaled by their distance. Using this embedding, the authors claim a 97% of path success with a geographic greedy routing algorithm[3].

Graphs embedded in hyperbolic spaces are also frequently used to propose new routing algorithms. For instance, Kleinberg [8] uses the hyperbolic coordinates to implement a geographic routing algorithm for ad-hoc wireless networks with guarantee of successfully reach a destination from a source. On their hand, Cassagnes et al. [4] propose a dynamic P2P overlay embedded in a hyperbolic space in which a node computes its coordinates without global knowledge of the graph topology.

The GIRG Model. Brigmann et al. [2] proposed an alternative model of hyperbolic graphs called Geometric Inhomogeneous Random Graphs (GIRG). It is inspired in the Chung-Lu [5, 6] random graphs and basically is a model for scale-free networks with an underlying geometry. The GIRG model assigns to each node a weight which is used by a probabilistic process to link the nodes.

More formally, a GIRG is a graph $G = (V, E)$ where each node $v \in V$ has an uniform random position (x_v) in a geometric space T^δ and a weight $w_v \in R^+$. The weights are drawn from a power-law distribution. The links E are also random, so that two nodes are linked with a probability that increases with the node weight (power-law factor) and decreases with the distance between nodes (Kleinberg [7] factor). Concretely, two nodes u and v are linked with probability

$$P_{uv} = \left(\frac{1}{d_{uv}}\right)^{\delta\alpha} \left(\frac{w_u w_v}{W}\right)^{\alpha}, \tag{1}$$

where d_{uv} is the distance between node u and v, $\alpha > 1$ is a decay parameter, δ is the dimension of the geometric space, and W is the aggregated weight of all nodes. In particular, the nodes are placed in a δ-dimensional torus ($T^\delta = R^\delta / Z^\delta$), and as distance metric the ∞-norm on T^δ is used (the authors claim that other metrics can be used). The GIRG model avoids the use of cosines and preserves the properties of a hyperbolic random graphs.

With the graph model, Brigmann et al. [3] proposed a new greedy routing algorithm with constant probability of success, that we denote GIRG-Φ. The GIRG greedy routing algorithm [3] is used to send a message from a source node to a destination node. In addition to the position of the destination (which is part of the message), it only uses local information at the current node holding the message (i.e., the weights and coordinates of its neighbors). The routing algorithm works in rounds. In each round, an objective function Φ (Eq. 2) is evaluated to obtain the objective value for the current node and each of its neighbors. If no neighbor has larger objective value than the current node, the routing fails, and the message is dropped. Otherwise the message is sent to the neighbor with largest objective value.

[3] A geographic algorithm routes a message to the neighbor closest to the destination

The objective function proposed in [3] is

$$\Phi(v,t) = \frac{w_v}{N \cdot w_{\min} \cdot d_{ut}^{\delta}}, \tag{2}$$

where t is the destination node, $w_{\min} = \min_{v \in V}\{w_v\}$ is the minimum weight drawn from the weight power-law distribution, N is the number of nodes of the graph, and δ is the dimension of the geometric space ($\delta = 2$ by default).

Contributions. In this paper we present a new model, denoted (κ, π)-KG, to build graphs with power-law and small-world properties, inspired on the Kleinberg [7] and GIRG [3] models. As in the former, our model uses an underlying complete torus T^{δ}, in which (unlike GIRG) all torus links are preserved (short links). One long link is added per node (like in Kleinberg model), using a probability expression similar GIRG's (Eq. 1). In this new probability expression, we have introduced two parameters, κ and π, in order to tune and evaluate separately the role and impact of Kleinberg and power-law factors. For example, this modification allows giving more influence to the Kleinberg (resp., power-law) factor in order to study its influence in the topology, the degree distribution, or the performance of routing algorithms.

Then, we propose greedy routing algorithms that could be used in both GIRG and (κ, π)-KG graphs. The routing algorithms proposed are inspired in GIRG's. They are fully distributed, and try to find a tradeoff between a low stretch factor (the maximum ratio between the length of the paths obtained by the routing algorithm and the distance between the source and the destination) and a small fail ratio.

The first algorithm is deterministic, guided by an objective function that has to be increased in each step until reaching the destination. The algorithm is based on GIRG's [3], but adapted so that the route never fails if the message gets close enough to the destination node. This reduces significantly the fail ratio without increasing the stretch. Additionally, we propose a second probabilistic routing algorithm. This algorithm uses the objective function as a weight for a probabilistic decision, in a similar way to a traditional random walk. This algorithm never fails, and shows by simulation route lengths very close to the lengths obtained with the deterministic routing algorithm.

2 (κ, π)-KG Model

In the (κ, π)-KG model, the nodes of a graph $G = (V, E)$ are those of a complete torus T^{δ}, as Kleinberg [7] used in his model. For simplicity, we assume $\delta = 2$ in the rest of the presentation. Hence, each node $u \in V$ has a pair of coordinates (x_u, y_u) and four local neighbors[4], so that there is always a route for all pair of nodes (which is not guaranteed in GIRG). The weight $w_u \geq 0$ of a node $u \in V$ is drawn from a power-law probability distribution with parameter $\beta \in (2, 3)$, with maximum and minimum weights w_{max} and $w_{min} > 0$, respectively. Then, each node $u \in V$ chooses an extra neighbor $v \neq u$ independently with probability

$$P_{uv} = \left(\frac{1}{d_{uv}}\right)^{\delta\kappa} \left(\frac{w_v}{W}\right)^{\pi}. \tag{3}$$

[4] All links are considered bidirectional.

As can be observed, this probability has two factors (as in Eq. 1): the Kleinberg factor is a function of the distance d_{uv} between nodes[5], while the power-law factor depends on the node weights (w_u and w_v). (κ, π)-KG introduces the parameters $\pi \geq 0$ and $\kappa \geq 0$ in the probability expression (3) in order to modulate the two factors independently. Given different values to these parameters it is possible to build networks with different properties and test on them the behavior of greedy routing algorithms.

3 Routing Algorithms

We present in this section two routing algorithms for GIRG and (κ, π)-KG graphs.

Deterministic Routing. We propose first a new greedy routing algorithm denoted τ-Det, where τ is a parameter of the algorithm. τ-Det works as GIRGS's greedy routing with a new objective function

$$\Phi_1(u,t) = w_u / d_{ut}^k, \tag{4}$$

where $k = \tau \ln(w_{\max}/w_{\min})$ and $d_{ut} \in \mathbb{N}$ is the distance from u to the destination t. Function Φ_1 uses the parameter τ (via k) to ensure greedy routing success. It makes sure a message reaches the destination when it is at a distance no higher than a threshold τ, as the following lemma shows.

Lemma 1 *Let $u, t \in V$ be nodes such that $d_{ut} \leq \tau$, then for any $v \in V$ such that $d_{vt} < d_{ut}$ it holds that $\Phi_1(v,t) > \Phi_1(u,t)$.*

Proof. We want to prove that $\Phi_1(v,t) = w_v/d_{vt}^k > \Phi_1(u,t) = w_u/d_{ut}^k$. In the extreme case, $w_v = w_{\min}$ and $w_u = w_{\max}$. Hence, it is enough to prove that $w_{\min}/w_{\max} > (d_{vt}/d_{ut})^k$. We have that $d_{vt} \leq d_{ut} - 1$, $d_{ut} \leq \tau$, and $\frac{x-1}{x}$ is strictly increasing for $x > 0$. Hence, it is enough to prove that $w_{\min}/w_{\max} > \left(\frac{\tau-1}{\tau}\right)^k = (1 - 1/\tau)^k$. Now, since $1 - x < e^{-x}$ for $x \neq 0$, it is enough that $w_{\min}/w_{\max} \geq e^{k/\tau}$, which holds since $k = \tau \ln(w_{\max}/w_{\min})$. □

The value of k in Eq. 4 is a function of the maximum and minimum weights, and the threshold value. The successful ratio of the greedy routing increases with τ (and k), making this ratio tunable. Unfortunately, we have observe that the route length increases with τ as well.

Probabilistic Routing. We present now a second distributed routing algorithm for GIRG and (κ, π)-KG denoted τ-RW. With the new algorithm a message follows a random walk. When the message is at a node $u \in V$, the next hop is selected among its neighbors with probabilities proportional to their value of the objective function Φ_1 (Eq.4).

In can be proven that with this new algorithm all messages eventually reach their destinations, since the graph is connected. This is the main difference with the deterministic algorithm (and GIRG's algorithm), with which the routing process can fail. We will study the route length with this new algorithm by simulation, and compare it with the route length with the deterministic algorithms. We observe that the removal of failed routing comes as a very low cost in route length.

[5] The distance can be Euclidean, Manhattan, or based on the $\infty - norm$.

		GIRG ($\alpha = 1.5$)	($\kappa = 1, \pi = 2$)-KG
Deterministic	Geo	12.82 / 0.00	8.92 / 0.00
	GIRG-Φ	12.03 / 0.90	**5.88** / 0.50
	τ-Det ($\tau = 1$)	**10.36** / 0.39	6.51 / 0.19
	τ-Det ($\tau = 2$)	10.71 / 0.26	7.24 / 0.16
	τ-Det ($\tau = 3$)	10.97 / 0.19	7.59 / 0.13
	τ-Det ($\tau = 4$)	11.15 / **0.13**	7.79 / **0.11**
Random Walk	Geo	131.79 / 0.00	60.32 / 0.00
	GIRG-Φ	71.68 / 0.00	17.19 / 0.00
	τ-RW ($\tau = 1$)	11.33 / 0.00	**6.68** / 0.00
	τ-RW ($\tau = 2$)	**11.02** / 0.00	7.24 / 0.00
	τ-RW ($\tau = 3$)	11.11 / 0.00	7.56 / 0.00
	τ-RW ($\tau = 4$)	11.25 / 0.00	7.77 / 0.00

Table 1: Comparing the best average hops and fail ratio on GIRG and (κ, π)-KG graphs using deterministic and probabilitic routing algorithms. In a table entry *h/f* means *h* hops and *f* fail ratio. Values in bold represent the best result of the routing algorithms.

4 Experimental Evaluation

We have developed a tool for creating multiple GIRG and (κ, π)-KG graphs, and simulating the routing of messages in them with the described algorithms. In all the graphs created, we have simulated the following routing algorithms.

- Deterministic algorithms: GIRG-Φ, τ-Det, and geographic.
- Probabilistic algorithms: GIRG-Φ, τ-RW, and geographic random walks[6].

We have introduced a little change in the algorithm GIRG-Φ. As describe above, in the original algorithm when the objective function (Eq. 2) reaches a local minimum, the message is dropped. In our version, geographic routing is used to continue the routing to the destination node.

Simulation Parameters. We run the simulations on graphs obtained from a 2-dimensional torus 101×101. The node weights are drawn from a power-law distribution $p(w) \sim w^{-}\beta$ with $\beta = 2.1$, $w_{\min} = 1$ and $w_{\max} = 10^5$. The distance threshold τ of the Φ_1 objective function will take the following value $3, 10, 30, 100$ (100 is largest distance for the Manhattan and the ∞-norm metrics, and it is larger than the Euclidean maximum distance). Note that this parameter is only relevant for the τ-Det and τ-RW routing algorithms. Each simulation includes the routing of 50.000 messages. Every message is routed from a source to a destination chosen uniformly at random.

Simulation Results. Table 1 shows the average results using Manhattan distance of the routing algorithms on GIRG (with $\alpha = 1.5$) and (κ, π)-KG (with $\kappa = 1$ and $\pi = 2$)

[6] The next hop is chosen with probability inversely proportional to the distance to the destination.

graphs. These graphs yield the best routing algorithm results among the collection of values of α, κ and π explored (not shown due to space limit). Comparing the graphs, (κ, π)-KG presents better performance than GIRG both in hops and fail ratio. Regarding routing algorithms, while deterministic routing algorithms have the best performance in hops, they have a positive fail ratio. Among them, the algorithm τ-Det outperforms the other deterministic algorithms, both in hops and fail ratio (excluding Geo). Probabilistic algorithms never fail at the cost of longer routes. However, the τ-Det and τ-RW algorithms show almost identical results in number of hops.

References

1. Marián Boguná, Fragkiskos Papadopoulos, and Dmitri Krioukov. Sustaining the internet with hyperbolic mapping. *Nature Communications*, 1(62):1–8, 2010.
2. Karl Bringmann, Ralph Keusch, and Johannes Lengler. Geometric inhomogeneous random graphs. *Theoretical Computer Science*, 760:35 – 54, 2019.
3. Karl Bringmann, Ralph Keusch, Johannes Lengler, Yannic Maus, and Anisur Rahaman Molla. Greedy routing and the algorithmic small-world phenomenon. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*, PODC '17, pages 371–380, New York, NY, USA, 2017. ACM.
4. C. Cassagnes, T. Tiendrebeogo, D. Bromberg, and D. Magoni. Overlay addressing and routing system based on hyperbolic geometry. In *2011 IEEE Symposium on Computers and Communications (ISCC)*, pages 294–301, June 2011.
5. Fan Chung and Linyuan Lu. The average distances in random graphs with given expected degrees. *Proceedings of the National Academy of Sciences*, 99(25):15879–15882, 2002.
6. Fan Chung and Linyuan Lu. Connected components in random graphs with given expected degree sequences. *Annals of Combinatorics*, 6(2):125–145, Nov 2002.
7. J. M. Kleinberg. Navigation in a small world. *Nature*, 406(6798), August 2000.
8. R. Kleinberg. Geographic routing using hyperbolic space. In *IEEE INFOCOM 2007 - 26th IEEE International Conference on Computer Communications*, pages 1902–1909, May 2007.
9. Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá. Hyperbolic geometry of complex networks. *Phys. Rev. E*, 82:036106, Sep 2010.