

Navigating Ultrasmall Worlds in Ultrashort Time

Marián Boguñá¹ and Dmitri Krioukov²

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

²*Cooperative Association for Internet Data Analysis (CAIDA), University of California, San Diego (UCSD),
9500 Gilman Drive, La Jolla, California 92093, USA*

(Received 18 September 2008; published 3 February 2009)

Random scale-free networks are ultrasmall worlds. The average length of the shortest paths in networks of size N scales as $\ln \ln N$. Here we show that these ultrasmall worlds can be navigated in ultrashort time. Greedy routing on scale-free networks embedded in metric spaces finds paths with the average length scaling also as $\ln \ln N$. Greedy routing uses only local information to navigate a network. Nevertheless, it finds asymptotically the shortest paths, a direct computation of which requires global topology knowledge. Our findings imply that the peculiar structure of complex networks ensures that the lack of global topological awareness has asymptotically no impact on the length of communication paths. These results have important consequences for communication systems such as the Internet, where maintaining knowledge of current topology is a major scalability bottleneck.

DOI: 10.1103/PhysRevLett.102.058701

PACS numbers: 89.75.Hc

Random scale-free networks are ultrasmall worlds [1–3]. The average and maximum lengths of the shortest paths in scale-free networks with power-law degree distributions $P(k) \sim k^{-\gamma}$, $\gamma \in [2, 3]$, scale with network size N as $\ln \ln N$ [1,2,4]. However, finding such shortest paths requires global topology knowledge, which is not available to nodes in many real networks. It may seem surprising at first that, having no global topological awareness, nodes can find any paths to destinations at all. In Ref. [6], we address this apparent paradox by showing that the observed topological characteristics of complex networks maximize their navigability, measured by the efficiency of the greedy routing process.

Greedy routing (GR) [7–11] relies on the hidden metric space abstraction [12]. In this abstraction a network is embedded in a metric space, with distances in this space representing intrinsic node similarities. To route information to a given destination, a node forwards the information to its network neighbor closest to the destination in this space. This general mechanism underlies processes ranging from search in social networks [13] to protein folding [14]. The existence of hidden metric spaces under real networks in general is a conjecture, but we found empirical evidence of their existence for some real networks, such as the Internet or some social networks [12]. In other cases, the metric space may be visible. In the airport network, for example, this space is geographic [6,15].

In Ref. [15], numerical experiments show that scale-free networks are navigable in a wide region of parameters. Specifically, GR and its modifications are found to perform generally well, in terms of the length and number of successful paths, on scale-free networks embedded in a plane. The GR efficiency is attributed to network heterogeneity. In Ref. [6], the analytic results and simulations show that not only heterogeneity but also clustering affect strongly the GR efficiency. The thermodynamic limit is

considered, and a network is called navigable if, in this limit, GR can find paths for a macroscopic fraction of source-destination pairs. Navigable networks are shown to have sufficiently strong clustering and heterogeneity of node degrees, i.e., $\gamma \approx 2$.

Here we show analytically and in simulations that the average hop length of paths that GR produces in these navigable networks scales with network size as $\bar{\tau} \sim \ln \ln N / |\ln(\gamma - 2)|$. Given that the average length of shortest paths in these networks, as shown in Refs. [1,2], also scales as $\mathcal{D} \sim \ln \ln N / |\ln(\gamma - 2)|$, we conclude that the GR paths are asymptotically shortest.

To obtain this result, we use the generic class of models introduced in Ref. [12]. These models generate scale-free networks embedded in metric spaces as follows. Given a target network size N , first assign to all nodes their coordinates in the metric space and an additional hidden variable κ representing their expected degrees. To generate scale-free networks, the variable κ is power-law distributed according to $\rho(\kappa) \propto \kappa^{-\gamma}$, $\kappa \in [\kappa_0, \infty)$, where κ_0 is the minimum expected degree. The metric space can be any homogeneous and isotropic D -dimensional space. Nodes are distributed in it with a uniform density δ that is set to $\delta = 1$ without loss of generality. Then each pair of vertices i and j is connected by an edge with probability $r(x)$, $x \equiv d_{ij}/(\mu\kappa_i\kappa_j)^{1/D}$, where d_{ij} is the distance between the two vertices in the metric space and κ_i and κ_j are their expected degrees.

A proper choice of the parameter μ , which depends on a specific form of the connection probability $r(x)$, guarantees that the average degree of vertices with hidden variable κ is $\bar{k}(\kappa) = \kappa$, so that κ can indeed be identified with the degree. The exponent γ in $\rho(\kappa)$ is then the power-law exponent of the degree distribution in the resulting networks [12]. These properties of the model hold for any dimension D of the metric space and for any form of the

connection probability $r(x)$, as long as the integral $\int_0^\infty x^{D-1}r(x)dx$ is bounded. We thus have a very versatile class of models since we can independently fix the average degree and the exponent γ without specifying the function $r(x)$, which can then be used to control clustering in the network. For example, in Ref. [12] we use $r(x) = (1+x)^{-\alpha}$, with $\alpha > D$ and $\mu = \Gamma(D/2)\Gamma(\alpha)/[2\pi^{D/2}\langle k \rangle\Gamma(\alpha-D)\Gamma(D)]$. This form of $r(x)$ leads to the following two extremes. In the limit $\alpha \rightarrow D$, clustering vanishes. The network loses its metric properties and becomes equivalent to a random graph, where the probability that two nodes are connected depends only on their expected degrees and not on the metric distance between them. In the opposite extreme $\alpha \rightarrow \infty$, clustering converges to a finite value, and the topology of the network is strongly influenced by the metric properties of the underlying space. This latter extreme yields networks lying in the navigable region.

We next give an intuitive explanation, illustrated in Fig. 1, for why GR is efficient in these navigable networks with strong coupling between network topology and underlying geometry. Suppose the GR process starts at some low-degree node and intends to reach a destination located far away in the metric space. Ideally, the process should proceed to hubs, high-degree nodes, that likely cover long distances by their numerous connections. However, GR is degree-agnostic; it checks only underlying distances. Therefore, this ideal scenario with propagation through the hubs can be implemented only if the node's neighbor closest to the destination is also its highest-degree neighbor. But this condition is the more likely satisfied the faster $r(x)$ decreases, because the faster $r(x)$ decreases, the stronger the dependency between a node's degree and the characteristic scale of distances that the node covers by its connections. This dependency is simple: the higher the

degree of a node, the larger its characteristic distance scale. (In the non-navigable limit $\alpha \rightarrow D$, this dependency disappears.) Consequently, if the next node along a path has a higher degree, then the node after the next one has an even higher degree, and the metric distance between these nodes also increases. On the other hand, the faster $r(x)$ decreases (e.g., the larger α), the stronger the clustering. We thus see that, in the navigable case with strong clustering, GR first travels over a sequence of nodes with increasing degrees and increasing internode distances. At some point, after the current distance to the destination becomes comparable to the internode distance, this pattern changes, and the process completes in a finite number of hops.

We now put this intuition on quantitative grounds. We first compute the probability that a node of expected degree κ has a neighbor with expected degree κ' at a distance d from it: $P(\kappa', d|\kappa)$. Using results from Refs. [12,16], it is easy to show that this probability is

$$P(\kappa', d|\kappa) = \frac{\rho(\kappa')}{\mu\langle k \rangle\kappa} d^{D-1} r\left(\frac{d}{(\mu\kappa\kappa')^{1/D}}\right). \quad (1)$$

The marginal distribution with respect to κ is

$$P(\kappa'|\kappa) = \frac{\kappa'\rho(\kappa')}{\langle k \rangle} \quad (2)$$

for any function $r(x)$ and any dimension D . We next compute the correlation between variables κ' and d conditioned on κ . Using Bayes' rule and Eq. (2), we write

$$P(d|\kappa, \kappa') = \frac{P(\kappa', d|\kappa)}{P(\kappa'|\kappa)} = \frac{d^{D-1}}{\mu\kappa\kappa'} r\left(\frac{d}{(\mu\kappa\kappa')^{1/D}}\right). \quad (3)$$

The average metric distance between two connected vertices with expected degrees κ and κ' is then

$$\bar{d}(\kappa, \kappa') = (\mu\kappa\kappa')^{1/D} \int_0^{x_c} x^D r(x) dx, \quad (4)$$

where $x_c = d_c(N)(\mu\kappa\kappa')^{-1/D}$ and $d_c(N)$ is the maximum distance between nodes in the metric space: $d_c(N) \sim N^{1/D}$.

If $r(x) = (1+x)^{-\alpha}$, with $\alpha > D+1$, then the integral in Eq. (4) is bounded, and we observe positive correlations between degrees and distances: the higher the node degree, the longer the characteristic distances that it covers by its connections, which is exactly the property guaranteeing GR efficiency. If $D < \alpha < 1+D$, the integral in Eq. (4) diverges, and we obtain

$$\bar{d}(\kappa, \kappa') \sim (\mu\kappa\kappa')^{(\alpha/D)-1} [d_c(N)]^{D+1-\alpha}. \quad (5)$$

In the limit $\alpha \rightarrow D$, $\bar{d}(\kappa, \kappa')$ loses any dependence on κ and κ' and becomes a large value diverging in the thermodynamic limit. As a consequence, degrees and distances are no longer correlated. The furthest neighbor no longer tends to have the highest degree. These arguments explain why the network cannot be navigated if it loses its metric properties and clustering vanishes.




FIG. 1 (color online). Illustration of the efficient greedy routing mechanism. The figure shows the current vertex v with its local neighborhood. The size of each vertex is proportional to its degree, and the plane represents the underlying metric space. Vertex v' is the neighbor of v , which is closest to the target and also one of its furthest and highest-degree neighbors. At the next hop, greedy routing proceeds from v' to v'' , reaching an even higher-degree vertex, traveling an even longer distance, and getting much closer to the target.

We now shift our attention entirely to navigable networks with $\alpha > D + 1$ and compute the lengths of greedy paths in them. As the first step, we calculate the maximum expected degree $\kappa_{c,nn}(\kappa)$ among all neighbors of a node of expected degree κ . In a finite-size network, the variable κ is bounded by a natural cutoff $\kappa_c \sim N^{1/(\gamma-1)}$ [5]. This cutoff is calculated as the value of $\kappa = \kappa_c$ such that we expect to find only one node with $\kappa > \kappa_c$ out of a sample of N vertices: $N \int_{\kappa_c}^{\infty} \rho(\kappa) d\kappa \sim 1$. Following the same reasoning, the value of $\kappa_{c,nn}(\kappa)$ can be evaluated as

$$\kappa \int_{\kappa_{c,nn}(\kappa)}^{\kappa_c} P(\kappa'|\kappa) d\kappa' \sim 1, \quad (6)$$

which leads to

$$\kappa_{c,nn}(\kappa) \sim \begin{cases} \kappa^{1/(\gamma-2)}, & \kappa < \kappa_c^{\gamma-2} \\ \kappa_c, & \kappa > \kappa_c^{\gamma-2} \end{cases} \quad (7)$$

This result, together with Eq. (4), yields the following expression for the average distance to the next node along a GR path from a node of degree κ :

$$\bar{d}_{nn}(\kappa) \sim \kappa^{(\gamma-1)/D(\gamma-2)} \quad \text{for } \kappa < \kappa_c^{\gamma-2}. \quad (8)$$

Equations (7) and (8) turn out to be central to our analysis. First, we see from Eq. (7) that only if $1/(\gamma - 2) > 1$, i.e., if $\gamma < 3$, is the degree of the next node along a GR path, on average, higher than the degree of the current node. This property explains why only scale-free networks with $\gamma < 3$ are navigable.

Equation (8) also shows that the expected distance between the next node and the current node of degree $\kappa \sim N^{(\gamma-2)/(\gamma-1)}$ is $\bar{d}_{nn} \sim N^{1/D}$, which is of the order of the maximum distance between all nodes in the metric space. In other words, we can cross the entire network in a single hop, landing at a node located at a finite and size-independent distance from the target. Putting these observations together, we conclude that the time to reach a target from a low-degree source located far away ($\sim N^{1/D}$) from the target is roughly the number of hops that it takes to reach a node of expected degree $\kappa \sim N^{(\gamma-2)/(\gamma-1)}$, which is a size-dependent contribution, plus the number of hops needed to cover a finite distance from this node to the target, which is a size-independent contribution.

Following these observations, we iterate Eq. (7)

$$\kappa_{\tau+1} \propto \kappa_{\tau}^{1/(\gamma-2)}, \quad \tau = 0, 1, \dots, \quad (9)$$

to find the value of τ such that $\kappa_{\tau} \sim N^{(\gamma-2)/(\gamma-1)}$. The solution is

$$\bar{\tau} = A + \frac{\ln[\ln N + B]}{|\ln(\gamma - 2)|}, \quad (10)$$

where A and B are functions of γ and $\langle k \rangle$.

This result is remarkable in many respects. First, in the large-size limit we obtain $\bar{\tau} \sim \ln \ln N$, meaning that greedy paths are ultrashort. Second, the prefactor in front of the

logarithm is just a function of γ , surprisingly independent of the average degree. Finally, this prefactor is equal to the prefactor of the average shortest-path lengths in scale-free networks [1,2]. It was also shown in Refs. [1,2] that fluctuations around the average shortest-path lengths are constant. Therefore, in the thermodynamic limit the shortest-path length distribution becomes a delta function. This fact, together with the equality between the average shortest- and greedy-path lengths, implies that for $N \gg 1$ the distribution of greedy-path lengths also converges to the same delta function. Consequently, in large networks, all greedy paths are shortest paths.

To check the accuracy of our theory, we perform extensive numerical simulations for the model with $D = 1$ (a circle) and $\alpha = \infty$, which is equivalent to taking $r(x) = e^{-x}$ and $\mu = 1/(2\langle k \rangle)$. We also fix the minimum expected degree to $\kappa_0 = 2$. We note that parameters κ_0 and δ are dummies and can be set to arbitrary values; the only independent parameters in the model are the average degree $\langle k \rangle$, the degree exponent γ , and clustering strength α . Fixing $\kappa_0 = 2$ helps to generate networks that are fully connected almost surely. If κ_0 is fixed to a constant, then the average degree depends on γ as $\langle k \rangle = (\gamma - 1)\kappa_0/(\gamma - 2)$. Varying $\langle k \rangle$ is desirable as it allows us to directly check with simulations if there is indeed no dependency on $\langle k \rangle$ of the prefactor in Eq. (10). We also verified that networks with a fixed average degree yield the same results.

Once a network with these parameters is generated, we simulate the GR process by choosing at random a source and destination and forwarding at each node to the node's neighbor closest to the destination on the circle. The number of source-destination pairs is 10^6 , and the results are averaged over a number of network realizations ranging between 400 and 4000. This process is performed for different power-law exponents γ and network sizes N . The average GR path length $\bar{\tau}$ is then computed as a function of N for different γ .

The top plot in Fig. 2 shows the results of these simulations. We then fit empirical $\bar{\tau}(N)$ to a function of the form $A_1 + A_2 \ln[\ln N + A_3]$, where the constants A_1 , A_2 , and A_3

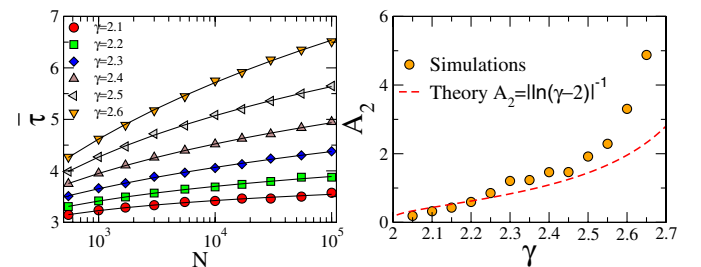


FIG. 2 (color online). Left: Average length of GR paths in generated networks with different values of γ as a function of the system size N . Solid lines are fits of the form $A_1 + A_2 \ln[\ln N + A_3]$. Right: Parameter A_2 obtained from the fit of $\bar{\tau}(N)$ compared to the theoretical prediction $A_2 = |\ln(\gamma - 2)|^{-1}$.

are free parameters estimated using the least square fit to the data. The bottom plot in Fig. 2 shows the empirical estimate of the coefficient A_2 compared with the theoretical prediction $A_2 = |\ln(\gamma - 2)|^{-1}$. The agreement is very good for the values of γ close to 2 and deteriorates as γ approaches 3. This deterioration is a consequence of the mean field approximation that assumes that at each hop the degree increases, which is true only on average. In fact, for γ approaching 3, there is a increasingly non-negligible probability of making a hop toward a smaller-degree node [6]. We have also checked [17] that the fluctuations of GR path lengths around their average $\sqrt{\bar{\tau}^2 - \bar{\tau}^2}$ stay constant with increasing N or even slightly decrease for small γ . These observations confirm that in the thermodynamic limit the GR path length distribution converges to a delta function.

In summary, we have shown that greedy routing finds asymptotically the shortest paths in scale-free networks with strong clustering and power-law node degree distribution exponents $\gamma < 3$. Given that topologies of many real networks do have these properties [18–20], our findings imply, surprisingly, that, even without any global knowledge of network topology, nodes in complex networks can propagate information along the shortest routes. In other words, topologies of many real networks have a peculiar structure that guarantees that the lack of global topological awareness imposes asymptotically no impact on the structure of information flows in the network: with or without the global topology knowledge, information can flow along the shortest routes. There are other, regular networks, such as lattices, that also possess these properties, but they require specific embeddings into specific spaces. Greedy routes on scale-free networks, on the other hand, are shortest regardless of the specifics of a hidden metric space or connection probability.

Complex networks thus have the structure that allows them to perform, in the most efficient way, one of their most basic and common functions: to propagate or signal information to specific targets through a complex network maze whose global connectivity is unknown to any node. It remains an open question if real networks evolve to become navigable [21,22] or which networks do have hidden metric spaces underneath and which do not. Even if such spaces exist, it may be quite challenging to identify their exact structure. At the same time, our findings have optimistic practical implications as they open up a possibility to find shortest-path routing strategies for the Internet that would not require any global topology knowledge. The requirement for routers to have and constantly update this knowledge is a major scalability bottleneck in the Internet today [23].

We thank K. C. Claffy and M. Ángeles Serrano for useful suggestions. This work was supported by FIS2007-66485-C02-02, Generalitat de Catalunya Grant No. SGR00889, the Ramón y Cajal program of the Span-

ish Ministry of Science, NSF CNS-0434996 and CNS-0722070, DHS N66001-08-C-2029, and Cisco Systems.

-
- [1] R. Cohen and S. Havlin, Phys. Rev. Lett. **90**, 058701 (2003).
 - [2] S. N. Dorogovtsev, J. F. F. Mendes, and A. N. Samukhin, Nucl. Phys. **B653**, 307 (2003).
 - [3] F. Chung and L. Lu, Proc. Natl. Acad. Sci. U.S.A. **99**, 15 879 (2002).
 - [4] The diameter in Ref. [3] scales as $\ln N$ because as opposed to Refs. [1,2], to avoid degree correlations for tractability purposes, the authors impose the maximum-degree cutoff at $\sim \sqrt{N}$, which is smaller than the natural maximum-degree cutoff $\sim N^{1/(\gamma-1)}$ in scale-free networks with $\gamma < 3$ [5]. The average shortest-path length is still $\sim \ln \ln N$ in Ref. [3].
 - [5] S. N. Dorogovtsev and J. F. F. Mendes, Adv. Phys. **51**, 1079 (2002).
 - [6] M. Boguñá, D. Krioukov, and K. C. Claffy, Nature Phys. **5**, 74 (2009).
 - [7] J. M. Kleinberg, Nature (London) **406**, 845 (2000).
 - [8] C. Martel and V. Nguyen, in *Proceedings of the Twenty-Third Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing* (Association for Computing Machinery, New York, 2004).
 - [9] Ö. Şimşek and D. Jensen, in *Proceedings of the 2005 International Joint Conferences on Artificial Intelligence* (Morgan Kaufmann, San Francisco, 2005).
 - [10] P. Fraigniaud, C. Gavoille, A. Kosowski, E. Lebhar, and Z. Lotker, in *Proceedings of the 19th ACM Symposium on Parallelism in Algorithms and Architectures* (Association for Computing Machinery, New York, 2007).
 - [11] J. Kleinberg, in *Proceedings of the International Congress of Mathematicians, Madrid, 2006* (European Mathematical Society, Zurich, 2006).
 - [12] M. A. Serrano, D. Krioukov, and M. Boguñá, Phys. Rev. Lett. **100**, 078701 (2008).
 - [13] J. Travers and S. Milgram, Sociometry **32**, 425 (1969).
 - [14] E. Ravasz, S. Gnanakaran, and Z. Toroczkai, arXiv:0705.0912.
 - [15] H. Thadakamalla, R. Albert, and S. Kumara, New J. Phys. **9**, 190 (2007).
 - [16] M. Boguñá and R. Pastor-Satorras, Phys. Rev. E **68**, 036112 (2003).
 - [17] M. Boguñá and D. Krioukov, arXiv:0809.2995v1.
 - [18] R. Albert and A.-L. Barabási, Rev. Mod. Phys. **74**, 47 (2002).
 - [19] M. E. J. Newman, SIAM Rev. **45**, 167 (2003).
 - [20] S. N. Dorogovtsev and J. F. F. Mendes, *Evolution of Networks: From Biological Nets to the Internet and WWW* (Oxford University Press, Oxford, 2003).
 - [21] A. Clauset and C. Moore, arXiv:cond-mat/0309415.
 - [22] A. Chaintreau, P. Fraigniaud, and E. Lebhar, in *Proceedings of the 35th International Colloquium on Automata, Languages and Programming*, Lect. Notes Comput. Sci. Vols. 5125–5126 (Springer-Verlag, Berlin, 2008).
 - [23] *RFC 4984*, edited by D. Meyer, L. Zhang, and K. Fall (The Internet Architecture Board, Amsterdam, 2007).