

Navigability of complex networks

Marián Boguñá^{1*}, Dmitri Krioukov² and K. C. Claffy²

Routing information through networks is a universal phenomenon in both natural and man-made complex systems. When each node has full knowledge of the global network connectivity, finding short communication paths is merely a matter of distributed computation. However, in many real networks, nodes communicate efficiently even without such global intelligence. Here, we show that the peculiar structural characteristics of many complex networks support efficient communication without global knowledge. We also describe a general mechanism that explains this connection between network structure and function. This mechanism relies on the presence of a metric space hidden behind an observable network. Our findings suggest that real networks in nature have underlying metric spaces that remain undiscovered. Their discovery should have practical applications in a wide range of areas where networks are used to model complex systems.

Networks are ubiquitous in all domains of science and technology, and permeate many aspects of daily human life^{1–4}, especially on the rise of the information technology society^{5,6}. Our growing dependence on them has inspired a burst of activity in the new field of network science, keeping researchers motivated to solve the difficult challenges that networks offer. Among these, the relation between network structure and function is perhaps the most important and fundamental. Transport is one of the most common functions of networked systems. Examples can be found in many domains: transport of energy in metabolic networks, of mass in food webs, of people in transportation systems, of information in cell signalling processes or of bytes across the Internet.

In many of these examples, routing—or signalling of information propagation paths through a complex network maze—has a determinant role in the transport properties of the system, in particular in systems such as the Internet or airport networks that have transport as their primary function. The observed efficiency of this routing process in real networks poses an intriguing question: how is this efficiency achieved? When each element of the system has a full view of the global network topology, finding short routes to target destinations is a well-understood computational process. However, in many networks observed in nature, including those in society and biology (signalling pathways, neural networks and so on), nodes efficiently find intended communication targets even though they do not possess any global view of the system. For example, neural networks would not function so well if they could not route specific signals to appropriate organs or muscles in the body, although no neurone has a full view of global inter-neurone connectivity in the brain.

Here, we identify a general mechanism that explains routing conductivity, or navigability of real networks based on the concept of similarity between nodes^{7–12}. Specifically, intrinsic characteristics of nodes define a measure of similarity between them, which we abstract as a hidden distance. Taken together, hidden distances define a hidden metric space for a given network. Our recent work shows that these spaces explain the observed structural peculiarities of several real networks, in particular social and technological ones¹³. Here, we show that this underlying metric structure can be used to guide the routing process, leading to efficient communication without global information in arbitrarily

large networks. Our analysis reveals that, remarkably, real networks satisfy the topological conditions that maximize their navigability within this framework. Therefore, hidden metric spaces offer explanations of two open problems in complex networks science: the communication efficiency networks so often exhibit and their unique structural characteristics.

Node similarity and hidden metric spaces

Our work is inspired by the seminal work of sociologist Stanley Milgram on the small-world problem. The small-world paradigm refers to the existence of short chains of acquaintances among individuals in societies¹⁴. At Milgram's time, direct proof of such a paradigm was impossible owing to the lack of large databases of social contacts, so Milgram conceived an experiment to analyse the small-world phenomenon in human social networks. Randomly chosen individuals in the United States were asked to route a letter to an unknown recipient using only friends or acquaintances that, according to their judgement, seemed most likely to know the intended recipient. The outcome of the experiment revealed that, without any global network knowledge, letters reached the target recipient using, on average, 5.2 intermediate people, demonstrating that social acquaintance networks were indeed small worlds.

The small-world property can be easily induced by adding a small number of random connections to a 'large world' network¹⁵. More striking is the fact that social networks are navigable without global information. Indeed, the only information that people used to make their routing decisions in Milgram's experiment was a set of descriptive attributes of the destined recipient, such as place of living and occupation. People then determined who among their contacts was 'socially closest' to the target. The success of the experiment indicates that social distances among individuals—even though they may be difficult to define mathematically—have a role in shaping the network architecture and that, at the same time, these distances can be used to navigate the network. However, it is not clear how this coupling between the structure and function of the network leads to efficiency of the search process, or what the minimum structural requirements are to facilitate such efficiency¹⁶.

Here, we show how network navigability depends on the structural parameters characterizing the two most prominent and common properties of real complex networks: (1) scale-free (power-law) node degree distributions characterizing the

¹Departament de Física Fonamental, Universitat de Barcelona, Martí i Franquès 1, 08028 Barcelona, Spain, ²Cooperative Association for Internet Data Analysis (CAIDA), University of California, San Diego (UCSD), 9500 Gilman Drive, La Jolla, California 92093, USA. *e-mail: marian.boguna@ub.edu.

heterogeneity in the number of connections that different nodes have and (2) clustering, a measure of the number of triangles in the network topology. We assume the existence of a hidden metric space, an underlying geometric frame that contains all nodes of the network, shapes its topology and guides routing decisions, as illustrated in Fig. 1. Nodes are connected in the observable topology, but a full view of their global connectivity is not available at any node. Nodes are also positioned in the hidden metric space and identified by their coordinates in it. Distances between nodes in this space abstract their similarity^{7–12}. These distances influence both the observable topology and routing function: (1) the smaller the distance between two nodes in the hidden space, that is, the more similar the two nodes, the more likely they are connected in the observable topology; (2) nodes also use hidden distances to select, as the next hop, the neighbour closest to the destination in the hidden space. Kleinberg introduced the term greedy routing to describe this forwarding process¹⁶. Greedy routing and its modifications have been studied extensively in recent computer science literature^{17–29} (see also Kleinberg’s review³⁰ and references therein). However, most of these works do not study greedy routing on scale-free topologies, which are known as the common signature of many large-scale self-evolving complex networks^{1–3}.

We use the class of network models developed in recent work¹³. They generate networks with topologies similar to those of real networks—small-world, scale-free and with strong clustering—and, simultaneously, with hidden metric spaces lying underneath. The simplest model in this class uses a one-dimensional circle as the underlying metric space, in which nodes are uniformly distributed. The model first assigns to each node its expected degree k , drawn from a power-law degree distribution $P(k) \sim k^{-\gamma}$, with $\gamma > 2$, and then connects each pair of nodes with connection probability $r(d; k, k')$ that depends both on the distance d between the two nodes in the circle and their assigned degrees k and k' ,

$$r(d; k, k') \equiv r(d/d_c) = (1 + d/d_c)^{-\alpha}, \quad (1)$$

where $\alpha > 1$ and $d_c \sim kk'$,

which means that the probability of link connection between two nodes in the network decreases with the hidden distance between them (as $\sim d^{-\alpha}$) and increases with their degrees (as $\sim (kk')^\alpha$).

These two properties have a clear interpretation. The connection cost increases with hidden distance, thus discouraging long-range links. However, in making connections, rich (well-connected, high-degree) nodes care less about distances (connection costs) than poor nodes. Furthermore, the characteristic distance scale d_c provides a coupling between node degrees and hidden distances, and ensures the following three topological characteristics that we commonly see in real networks. First, pairs of richly connected, high-degree nodes—hubs—are connected with high probability regardless of the hidden distance between them because their characteristic distance d_c is so large that any actual distance d between them will be short in comparison: regardless of d , connection probability r in equation (1) is close to 1 if d_c is large. Second, pairs of low-degree nodes will not be connected unless the hidden distance d between them is short enough to compare to the small value of their characteristic distance d_c . Third, following similar arguments, pairs composed of hubs and low-degree nodes are connected only if they are located at moderate hidden distances.

The parameter α in equation (1) determines the importance of hidden distances for node connections. The larger α , the more preferred are connections between nodes close in the hidden space. Consequently, the triangle inequality in the metric space leads to stronger clustering in the network, see Fig. 1. Clustering has a clear interpretation in our approach as a reflection of the network’s

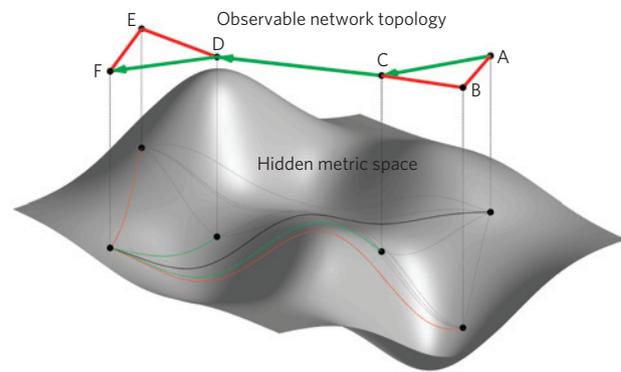


Figure 1 | How hidden metric spaces influence the structure and function of complex networks. The smaller the distance between two nodes in the hidden metric space, the more likely they are connected in the observable network topology. If node A is close to node B, and B is close to C, then A and C are necessarily close because of the triangle inequality in the metric space. Therefore, triangle ABC exists in the network topology with high probability, which explains the strong clustering observed in real complex networks. The hidden space also guides the greedy-routing process: if node A wants to reach node F, it checks the hidden distances between F and its two neighbours B and C. Distance CF (green dashed line) is smaller than BF (red dashed line); therefore, A forwards information to C. Node C then carries out similar calculations and selects its neighbour D as the next hop on the path to F. Node D is directly connected to F. The result is path $A \rightarrow C \rightarrow D \rightarrow F$ shown by green edges in the observable topology.

metric strength: the more powerful is the influence of the network’s underlying metric space on the observable topology, the more strongly it is clustered.

Although our toy model is not designed to exactly match any specific real network, it generates graphs that are surprisingly similar to some real networks, such as the Internet at the autonomous system level or the USA airport network. See Supplementary Information for details.

Navigability of modelled networks

We use the model to generate scale-free networks with different values of power-law degree distribution exponent γ and clustering strength α , covering the observed values in most documented complex networks^{1–3}. We then simulate greedy routing for a large sample of paths on all generated networks, and compare the following two navigability parameters: (1) the average hop length τ from source to destination of successful greedy-routing paths and (2) the success ratio p_s , defined as the percentage of successful paths. Unsuccessful paths are paths that get stuck at nodes without neighbours closer to the destination in the hidden space than themselves. These nodes usually have small degrees. See the Methods section for simulation details.

Figure 2 shows the impact of the network’s degree distribution and clustering on the average length τ of greedy-routing paths. We observe a straightforward dependency: paths are shorter for smaller exponents γ and stronger clustering (larger α values). The dependency of the success ratio (the fraction of successful paths) p_s on the two topology parameters γ and α is more intertwined. Figure 3 shows that the effect of one parameter, γ , on the success ratio depends on the other parameter, the level of clustering. If clustering is weak (low α), the percentage of successful paths decays with network size N regardless of the value of γ (Fig. 3a). However, with strong clustering (large α), the percentage of successful paths increases with N and attains a maximum for large networks if $\gamma \lesssim 2.6$, whereas it degrades for large networks if $\gamma > 2.6$ (Fig. 3b). Figure 3c shows this effect for networks of the same size ($N = 10^5$) with different γ and α . The value of $\gamma = 2.6 \pm 0.1$ maximizes the

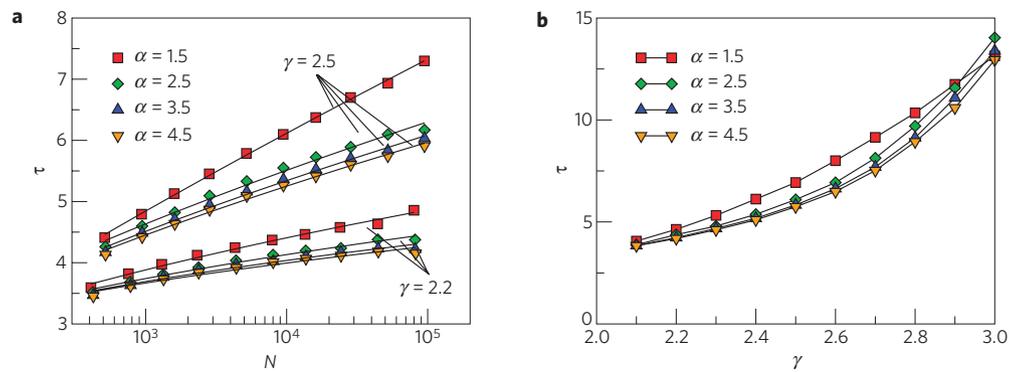


Figure 2 | Average length of greedy-routing paths. **a**, The average hop length of successful paths, τ , as a function of the network size N for different values of γ and α . Results for values of $\gamma > 2.5$ look similar but with longer paths and are omitted for clarity. In all cases, the path length grows polylogarithmically with the network size: the observed values of τ are fitted well by $\tau(N) = A[\log N]^\nu$ (solid lines), where A and ν are some constants. **b**, τ as a function of γ and α for networks of fixed size $N \approx 10^5$. The effect of the two parameters on average path length is straightforward: paths are shorter for smaller exponents γ and stronger clustering (larger α values).

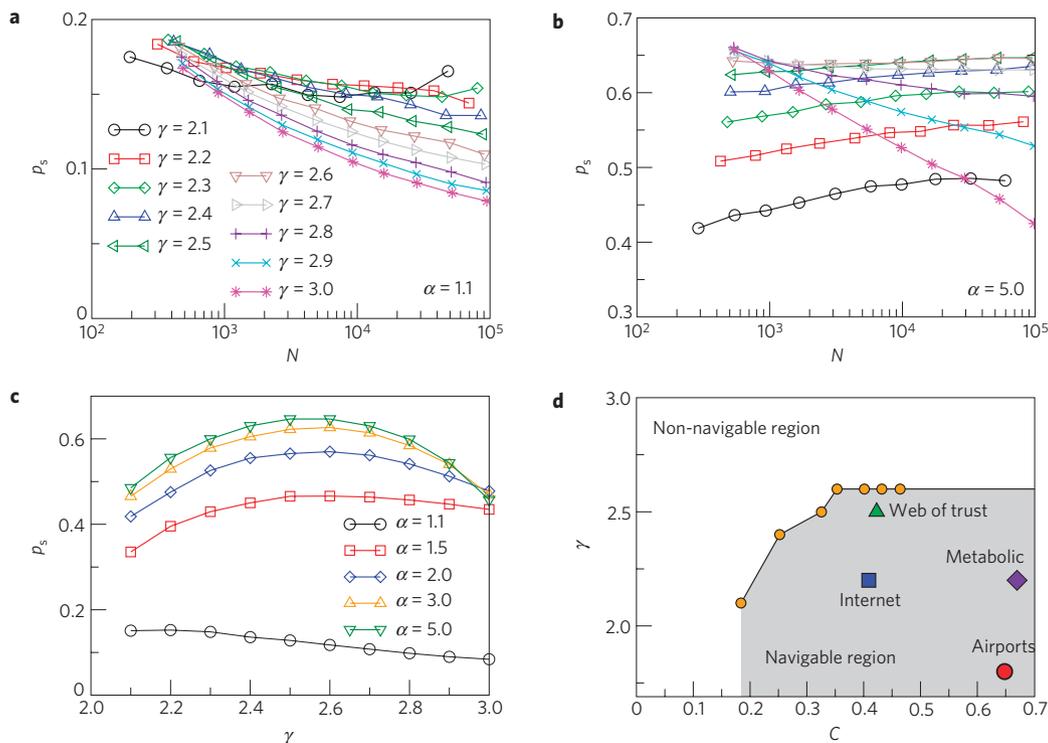


Figure 3 | Success probability of greedy routing. **a, b**, Success probability p_s as a function of network size N for different values of γ with weak (**a**) and strong (**b**) clustering. **c**, p_s as a function of γ and α for networks of fixed size $N \approx 10^5$. **d**, Parameter α mapped to clustering coefficient C (ref. 15) by computing C for each network with given γ and α . For each value of C , there is a critical value of $\gamma = \gamma_c(C)$ such that the success ratio in networks with this C and $\gamma > \gamma_c(C)$ decreases with the network size ($p_s(N) \rightarrow 0$ as $N \rightarrow \infty$), whereas $p_s(N)$ reaches a constant value for large N in networks with $\gamma < \gamma_c(C)$. The solid line in the plot shows these critical values $\gamma_c(C)$, separating the low- γ , high- C navigable region, in which greedy routing remains efficient in the large-graph limit, from the high- γ , low- C non-navigable region, where the efficiency of greedy routing degrades for large networks. The plot labels measured values of γ and C for several real complex networks. Internet is the global Internet topology of autonomous systems as seen by the Border Gateway Protocol³⁹; Web of trust is the Pretty Good Privacy social network of mutual trust relationships⁴⁰; Metabolic is the network of metabolic reactions of *E. coli*⁴¹; and Airports is the network of the public air transportation system⁴².

number of successful paths once clustering is above a threshold, $\alpha \geq 1.5$. These observations mean that for a fixed clustering strength, there is a critical value of the exponent γ (Fig. 3d) below which networks remain navigable as their size increases, but above which their navigability deteriorates with their size.

In summary, strong clustering improves both navigability metrics. We also find a delicate trade-off between values of γ close to 2 minimizing path lengths, and higher values—not

exceeding $\gamma \approx 2.6$ —maximizing the percentage of successful paths. We explain these findings in the next section, but we note here that qualitatively, this navigable parameter region contains most complex networks observed in reality^{1–3}, as confirmed in Fig. 3d, where we juxtapose a few paradigmatic examples of communication, social, biological and transportation networks versus the identified navigable region of clustering and degree distribution exponent. Interestingly, power grids, which propagate

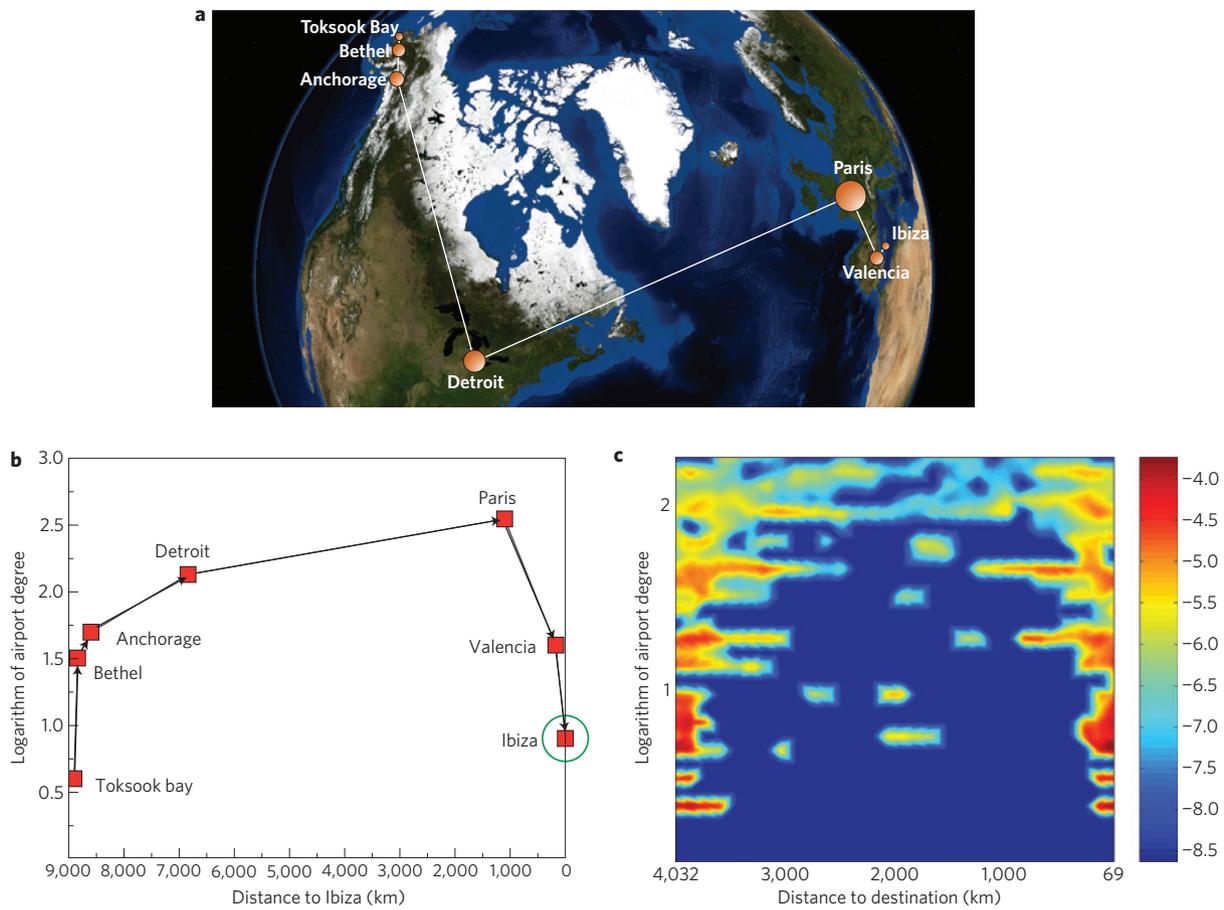


Figure 4 | Greedy routing in the airport network. **a**, The structure of the single greedy-routing path from Toksook Bay to Ibiza (image from NASA). At each intermediate airport, the next hop is the airport closest to Ibiza geographically. Sizes of symbols representing the airports are proportional to the logarithm of their degrees. **b**, The changing distance to Ibiza versus the degree of the visited airports. **c**, The structure of greedy-routing paths between a collection of airports in the USA (ref. 43). We include an airport pair in the collection if the distance between the airports is between 3,900 and 4,100 kilometres. The number of airport pairs in this collection is 7,620. We use colour to indicate how often paths in the collection go through an airport of a given degree located at a given geographical distance from the destination: blue/red indicates exponentially less/more visits to those airports, or more specifically, the colour is the logarithm of the normalized density of visited airports.

electricity rather than route information, are neither scale-free nor clustered^{15,31}.

Air travel by greedy routing as an explanation

We illustrate the greedy-routing function, and the structure of networks conducive to such routing, with an example of passenger air travel. Suppose we want to travel from Toksook Bay, Alaska, to Ibiza, Spain, by the public air transportation network. Nodes in this network are airports, and two airports are connected if there is at least one flight between them. We travel according to the greedy-routing strategy using geography as the underlying metric space. At each airport we choose the next-hop airport geographically closest to the destination. Under these settings, our journey goes first to Bethel, then to Anchorage, to Detroit, over the Atlantic to Paris, then to Valencia and finally to Ibiza, see Fig. 4. The sequence and sizes of airport hops reveal the structure of our greedy-routing path. The path proceeds from a small airport to a local hub at a small distance, from there to a larger hub at a larger distance, and so on until we reach Paris. At that point, when the distance to the destination becomes sufficiently small, greedy routing leads us closer to our final destination by choosing not another hub, but a less-connected neighbouring airport.

We observe that the navigation process has two, somewhat symmetric phases. The first phase is a coarse-grained search,

travelling longer and longer distances per hop towards hubs, thus ‘zooming out’ from the starting point. The second phase corresponds to a fine-grained search, ‘zooming in’ onto the destination. The turning point between the two phases appears naturally: once we are in a hub near the destination, the probability that it is connected to a bigger hub closer to the destination sharply decreases, but at this point we do not need hubs anyway, and greedy routing directs us to smaller airports at shorter distances next to the destination.

This zoom-out/zoom-in mechanism works efficiently only if the coupling between the airport network topology and the underlying geography satisfies the following two conditions: the sufficient hubs condition and the sufficient clustering condition. The first condition ensures that a network has enough hub airports (high-degree nodes) to provide an increasing sequence during the zoom-out phase. This condition is fulfilled by the real airport network and by other scale-free networks with small values of degree distribution exponent γ , because the smaller the γ , the larger the proportion of hubs in the network.

However, the presence of many hubs does not ensure that greedy routing will use them. Unlike humans, who can use their knowledge of airport size to selectively travel via hub airports, greedy routing uses only one constraint at each hop: minimize distance to the destination. Therefore, the network topology must satisfy the

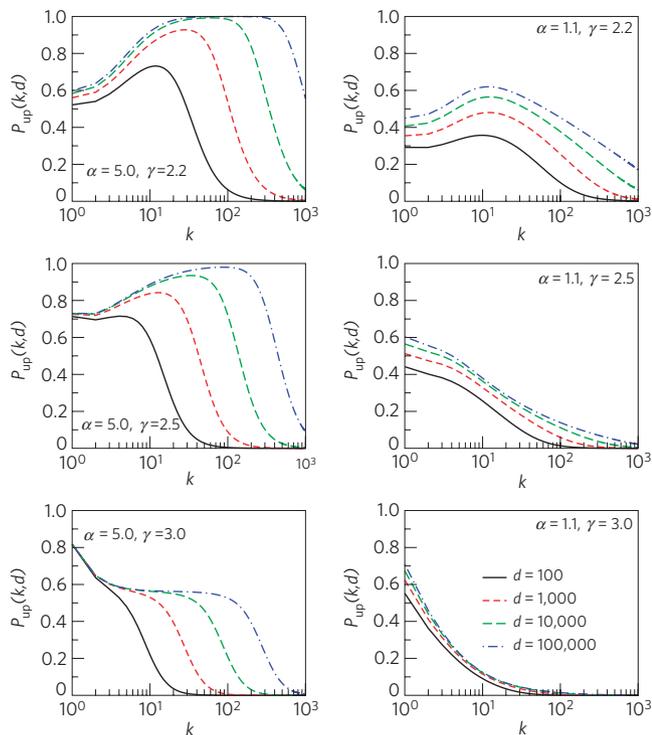


Figure 5 | Probability that greedy routing travels to higher-degree nodes.

The probability $P_{up}(k, d)$ that the greedy-routing next hop after a node of degree k located at distance d from a destination has higher degree $k' \geq k$ and is closer to the destination. The distance legend in the bottom right plot applies to all of the plots. The results are for the large-graph limit $N \rightarrow \infty$.

second condition, which ensures that Bethel is larger than Toksook Bay, Anchorage larger than Bethel and so on. More generally, this condition is that the next greedy hop from a remote low-degree node probably has a higher degree, so that greedy paths typically head first towards the highly connected network core. But the network metric strength is exactly the required property: preference for connections between nodes nearby in the hidden space means that low-degree nodes are less likely to have connectivity to distant low-degree nodes; only high-degree nodes can have long-range connection that greedy routing will effectively select. The stronger this coupling between the metric space and topology (the higher α in equation (1)), the stronger the clustering in the network.

To illustrate, imagine an airport network without sufficient clustering, one where the airport closest to our destination (Ibiza) among all airports connected to our current node (Toksook Bay, Alaska) is not Bethel, which is bigger than Toksook Bay, but Nightmute, Alaska, a nearby airport of comparable size to Toksook Bay. As greedy routing first leads us to Nightmute, then to another small nearby airport, and then to another, we can no longer get to Ibiza in few hops. Worse, travelling via these numerous small airports, we could reach one with no connecting flights heading closer to Ibiza. Our greedy routing would be stuck at this airport with an unsuccessful path.

These factors explain why the most navigable topologies correspond to scale-free networks with small exponents of the degree distribution, that is, a large number of hubs, and with strong clustering, that is, strong coupling between the hidden geometry and the observed topology.

The structure of greedy-routing paths

We observe the discussed zoom-out/zoom-in mechanism in analytical calculations and numerical simulations. Specifically, we

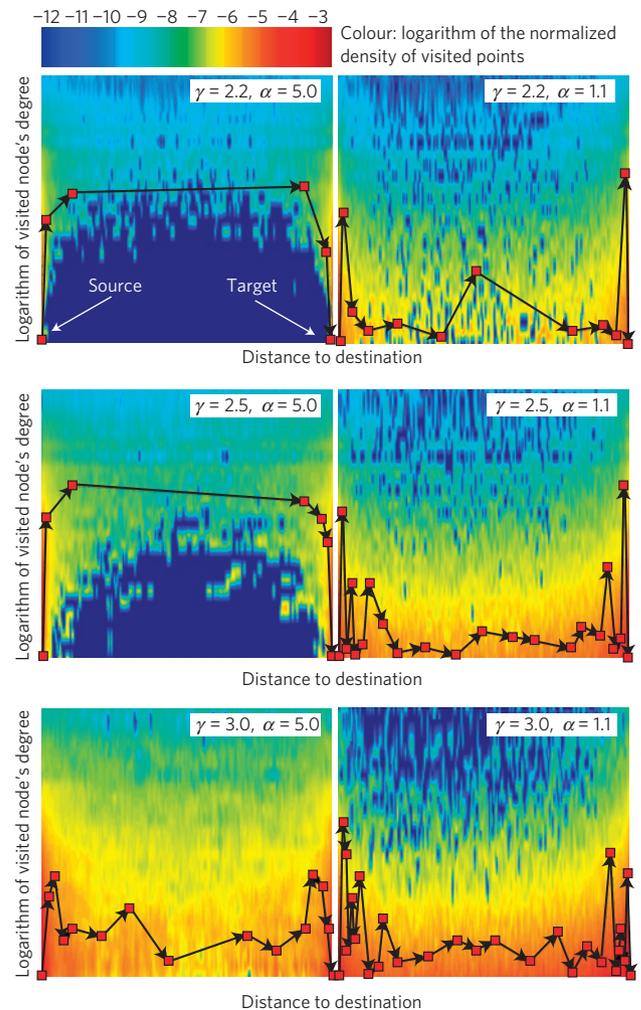


Figure 6 | The structure of greedy-routing paths. We visualize the results of our simulation of greedy routing in modelled networks with different values of γ and α observed in real complex networks. The hidden distance between the starting point and the destination is always approximately 10^4 , and the network size N and number of attempted paths is always 10^5 for each (γ, α) combination, but the number of successful paths and path hop lengths vary, see Figs 2, 3. All paths start and end at low-degree nodes located, respectively, in the bottom left and bottom right corners of the diagrams (see top left plot). For each (γ, α) , we show a single typical path in black and, as in Fig. 4, use colour to indicate how often paths included a node of a given degree located at a given distance from the destination. The simulations confirm that only when γ is small and α is large does the average path structure follow the zoom-out/zoom-in pattern that characterizes successful greedy routing in real networks, for example, in the airport network in Fig. 4.

calculate (in Supplementary Information) the probability that the next hop from a node of degree k located at hidden distance d from the destination has a larger degree $k' > k$, in which case the path moves towards the high-degree network core, see Fig. 5. In the most navigable case, with small degree-distribution exponent and strong clustering, the probability of increasing the node degree along the path is high at low-degree nodes, and sharply decreases to zero after reaching a node of a critical degree value, which increases with distance d . This observation implies that greedy-routing paths first propagate up to higher-degree nodes in the network core and then exit the core towards low-degree destinations in the periphery. In contrast, with low clustering, paths are less likely to find higher-degree nodes regardless of the distance to the destination. This

path structure violates the zoom-out/zoom-in pattern required for efficient navigation.

Figure 6 shows the structure of greedy-routing paths in simulations, further confirming our analysis. We again see that for small degree-distribution exponents and strong clustering (upper left and middle left), the routing process quickly finds a way to the high-degree core, makes a few hops there and then descends to a low-degree destination. In the other, non-navigable cases, the process can almost never get to the core of high-degree nodes. Instead, it wanders in the low-degree periphery, increasing the probability of getting lost at low-degree nodes.

Discussion

Our main motivation for this work comes from long-standing scalability problems with the Internet routing architecture³². To route information packets to a given destination, Internet routers must communicate to maintain a coherent view of the global Internet topology. The constantly increasing size and dynamics of the Internet thus leads to immense and quickly growing communication and information processing overhead, a major bottleneck in routing scalability³³ causing concerns among Internet experts that the existing Internet routing architecture may not sustain even another decade³². Discovery of the Internet's hidden metric space would remove this bottleneck, eliminating the need for the inherently unscalable communication of topology changes. Instead routers would be able to just forward packets greedily to the destination based on hidden distances.

In a similar manner, reconstruction of hidden metric spaces underlying other real networks may prove practically useful. For example, in social or communication networks (for example, the Web, overlay or online social networks) hidden spaces would yield efficient strategies for searching specific individuals or content based only on local knowledge. The metric spaces hidden under some biological networks (such as neural, gene regulatory networks, signalling or even protein folding^{34,35} pathways) can become a powerful tool in studying the structure of information or signal flows in these networks, enabling investigation of such processes without detailed global knowledge of the network structure or organization.

The natural question we thus face is how to proceed towards discovery of the explicit structure of hidden metric spaces underlying real networks. We do not expect spaces underlying different networks to be exactly the same. For example, the similarity spaces of Web pages⁹ and Wikipedia editors¹¹ probably differ. However, the main contribution of this work establishes the general mechanisms behind navigability of scale-free, strongly clustered topologies that characterize many different real networks. The next step is to find the common properties of hidden spaces that render them congruent with these mechanisms. Specifically, we are interested in what geometries of hidden spaces lead to such congruency³⁶.

In general, we believe that the present and future work on hidden metric spaces and network navigability will deepen our understanding of the fundamental laws describing relationships between structure and function of complex networks.

Methods

A model with the circle as a hidden metric space. In our model, we place all nodes on a circle by assigning them a random variable θ , that is, their polar angle, distributed uniformly in $[0, 2\pi)$. The circle radius R grows linearly with the total number of nodes N , $2\pi R = N$, to keep the average density of nodes on the circle fixed to 1. We next assign to each node its expected degree κ drawn from some distribution $\rho(\kappa)$. The connection probability between two nodes with hidden coordinates (θ, κ) and (θ', κ') takes the form

$$r(\theta, \kappa; \theta', \kappa') = \left(1 + \frac{d(\theta, \theta')}{\mu \kappa \kappa'}\right)^{-\alpha}, \quad \mu = \frac{(\alpha - 1)}{2\langle k \rangle},$$

where $d(\theta, \theta')$ is the geodesic distance between the two nodes on the circle and $\langle k \rangle$ is the average degree. It can be shown that the average degree of nodes with hidden variable κ , $\langle k(\kappa) \rangle$, is proportional to κ (ref. 37). This proportionality guarantees that the shape of the node degree distribution $P(k)$ in generated networks is approximately the same as the shape of $\rho(\kappa)$. The choice of $\rho(\kappa) = (\gamma - 1)\kappa_0^{\gamma-1}\kappa^{-\gamma}$, $\kappa > \kappa_0 \equiv (\gamma - 2)\langle k \rangle / (\gamma - 1)$, $\gamma > 2$, generates random networks with a power-law degree distribution of the form $P(k) \sim k^{-\gamma}$, where γ is a model parameter that regulates the heterogeneity of the degree distribution in the network. This parameter abstracts the heterogeneity of node degrees in real networks, where degree distributions may not perfectly follow power laws, or may exhibit various forms of high-degree cutoffs^{38,39}. The specific effects are less important than the overall measure of heterogeneity. We note that instead of a circle in our model we could use any isotropic space of any dimension¹³.

Numerical simulations. Our model has three independent parameters: exponent γ of power-law degree distributions, clustering strength α and average degree $\langle k \rangle$. We fix $\langle k \rangle$ to 6, which is roughly equal to the average degree of some real networks of interest^{39,40}, and vary $\gamma \in [2.1, 3]$ and $\alpha \in [1.1, 5]$, covering their observed ranges in documented complex networks¹⁻³. For each (γ, α) pair, we produce networks of different sizes $N \in [10^3, 10^5]$ generating, for each (γ, α, N) , a number of different network instances—from 40 for large N to 4,000 for small N . In each network instance G , we randomly select 10^6 source–destination pairs (a, b) and execute the greedy-routing process for them starting at a and selecting, at each hop h , the next hop as the h 's neighbour in G closest to b in the circle. If for a given (a, b) , this process visits the same node twice, then the corresponding path leads to a loop and is unsuccessful. We then average the measured values of path hop lengths τ and percentage of successful paths p_s across all pairs (a, b) and networks G for the same (γ, α, N) . Note that we are not concerned with the absolute values of the success ratio p_s . Instead we use it as a measure of navigability to compare networks with different (γ, α, N) . For this purpose, we could use the success ratio of any (improved) modification of standard greedy routing.

Shortest path versus shortest time. All results derived in the present paper are about finding short paths across a network topology. The total physical time from source to destination is implicitly assumed to be proportional to the number of hops. In real transportation systems, such as the Internet or the airport network, the finite capacity of nodes implies that the end-to-end path latency may be longer when intermediate nodes are congested. Although our results most cleanly apply to uncongested systems, there are obvious modifications, such as choosing the second- or third-nearest rather than the nearest neighbour, that could still find nearly shortest paths while reducing and balancing load on the system.

Received 12 March 2008; accepted 26 September 2008; published online 16 November 2008

References

- Albert, R. & Barabási, A.-L. Statistical mechanics of complex networks. *Rev. Mod. Phys.* **74**, 47–97 (2002).
- Newman, M. E. J. The structure and function of complex networks. *SIAM Rev.* **45**, 167–256 (2003).
- Dorogovtsev, S. N. & Mendes, J. F. F. *Evolution of Networks: From Biological Nets to the Internet and WWW* (Oxford Univ. Press, 2003).
- Boccaletti, S., Latora, V., Moreno, Y., Chavez, M. & Hwang, D.-U. Complex networks: Structure and dynamics. *Phys. Rep.* **424**, 175–308 (2006).
- Castells, M. *The Rise of the Network Society* (Blackwell, 1996).
- Pastor-Satorras, R. & Vespignani, A. *Evolution and Structure of the Internet. A Statistical Physics Approach* (Cambridge Univ. Press, 2004).
- Watts, D. J., Dodds, P. S. & Newman, M. E. J. Identity and search in social networks. *Science* **296**, 1302–1305 (2002).
- Girvan, M. & Newman, M. E. J. Community structure in social and biological networks. *Proc. Natl Acad. Sci. USA* **99**, 7821–7826 (2002).
- Menczer, F. Growing and navigating the small world Web by local content. *Proc. Natl Acad. Sci. USA* **99**, 14014–14019 (2002).
- Leicht, E. A., Holme, P. & Newman, M. E. J. Vertex similarity in networks. *Phys. Rev. E* **73**, 026120 (2006).
- Crandall, D., Cosley, D., Huttenlocher, D., Kleinberg, J. & Suri, S. *Proc. ACM Int. Conf. on Knowledge Discovery and Data Mining* (ACM, 2008).
- Clauset, A., Moore, C. & Newman, M. E. J. Hierarchical structure and the prediction of missing links in networks. *Nature* **453**, 98–101 (2008).
- Ángeles Serrano, M., Krioukov, D. & Boguñá, M. Self-similarity of complex networks and hidden metric spaces. *Phys. Rev. Lett.* **100**, 078701 (2008).
- Travers, J. & Milgram, S. An experimental study of the small world problem. *Sociometry* **32**, 425–443 (1969).
- Watts, D. J. & Strogatz, S. H. Collective dynamics of small-world networks. *Nature* **393**, 440–442 (1998).
- Kleinberg, J. Navigation in a small world. *Nature* **406**, 845 (2000).
- Kleinberg, J. *STOC '00: Proc. Thirty-Second Annual ACM Symp. on Theory of Computing* 163–170 (ACM, 2000).

18. Kleinberg, J. M. in *Small-World Phenomena and the Dynamics of Information* (eds Dietterich, T. G., Becker, S. & Ghahramani, Z.) 431–438 (MIT Press, 2001).
19. Manku, G. S., Naor, M. & Wieder, U. *Proc. 36th ACM Symp. on Theory of Computing (STOC)* 54–63 (ACM, 2004).
20. Martel, C. *23rd ACM Symp. Principles of Distributed Computing (PODC)* 179–188 (ACM, 2004).
21. Nguyen, V. & Martel, C. *SODA '05: Proc. Sixteenth Annual ACM-SIAM Symp. on Discrete Algorithms* 311–320 (Society for Industrial and Applied Mathematics, 2005).
22. Nguyen, V. & Martel, C. *Proc. Fourth Workshop on Analytic Algorithmics and Combinatorics* 213–227 (SIAM, 2008).
23. Simsek, Ö. & Jensen, D. in *Proc. Nineteenth Int. Joint Conf. in Artificial Intelligence* (eds Kaelbling, L. P. & Saffioti, A.) 304–310 (Professional Book Center, 2005).
24. Lebhar, E. & Schabanel, N. in *Proc. 31st Int. Colloquium on Automata, Languages and Programming* Vol. 3142 (eds Díaz, J., Karhumäki, J., Lepistö, A. & Sannella, D.) 894–905 (Lecture Notes in Computer Science, Springer, 2004).
25. Fraigniaud, P. in *Proc of the 13th Annual European Symposium* Vol. 3669 (eds Brodal, G. S. & Leonardi, S.) 791–802 (Lecture Notes in Computer Science, Springer, 2005).
26. Fraigniaud, P., Gavoille, C. & Paul, C. Eclecticism shrinks even small worlds. *Distrib. Comput.* **18**, 279–291 (2006).
27. Fraigniaud, P., Gavoille, C., Kosowski, A., Lebhar, E. & Lotker, Z. in *Proc. Nineteenth Annual ACM Symp. on Parallel Algorithms and Architectures* (eds Gibbons, P. B. & Scheideler, C.) 1–7 (ACM, 2007).
28. Fraigniaud, P. & Gavoille, C. in *Proc. Twentieth Annual Symposium on Parallelism in Algorithms and Architectures* (eds auf der Heide, F. M. & Shavit, N.) 62–69 (ACM, 2008) SPAA.
29. Chaintreau, A., Fraigniaud, P. & Lebhar, E. in *Automata, Languages and Programming* Vol. 5125 (eds Aceto, L. *et al.*) 133–144 (Lecture Notes in Computer Science, Springer, 2008).
30. Kleinberg, J. Complex networks and decentralized search algorithms. *Proc. Int. Congr. Math. (ICM)* **3**, 1019–1044 (2006).
31. Amaral, L. A. N., Scala, A., Barthélemy, M. & Stanley, H. E. Classes of small-world networks. *Proc. Natl Acad. Sci.* **97**, 11149–11152 (2000).
32. Meyer, D., Zhang, L. & Fall, K. (eds). *RFC4984* (The Internet Architecture Board, 2007).
33. Krioukov, D., Claffy, K. C., Fall, K. & Brady, A. On compact routing for the Internet. *Comput. Commun. Rev.* **37**, 41–52 (2007).
34. Ravasz, E., Gnanakaran, S. & Toroczkai, Z. Network structure of protein folding pathways. Preprint at <<http://arxiv.org/abs/0705.0912>> (2007).
35. Toroczkai, Z., Kozma, B., Bassler, K. E., Hengartner, N. W. & Korniss, G. Gradient networks. *J. Phys. A: Math. Theor.* **41**, 155103 (2008).
36. Krioukov, D., Papadopoulos, F., Boguñá, M. & Vahdat, A. Efficient navigation in scale-free networks embedded in hyperbolic metric spaces. Preprint at <<http://arxiv.org/abs/0805.1266>> (2008).
37. Boguñá, M. & Pastor-Satorras, R. Class of correlated random networks with hidden variables. *Phys. Rev. E* **68**, 036112 (2003).
38. Giot, M. *et al.* A protein interaction map of *Drosophila melanogaster*. *Science* **302**, 1727–1736 (2003).
39. Mahadevan, P. *et al.* The Internet AS-level topology: Three data sources and one definitive metric. *Comput. Commun. Rev.* **36**, 17–26 (2006).
40. Boguñá, M., Pastor-Satorras, R., Díaz-Guilera, A. & Arenas, A. Models of social networks based on social distance attachment. *Phys. Rev. E* **70**, 056122 (2004).
41. Jeong, H., Tombor, B., Albert, R., Oltvai, Z. N. & Barabási, A.-L. The large-scale organization of metabolic networks. *Nature* **407**, 651–654 (2000).
42. Barrat, A., Barthélemy, M., Pastor-Satorras, R. & Vespignani, A. The architecture of complex weighted networks. *Proc. Natl Acad. Sci. USA* **101**, 3747–3752 (2004).
43. <<http://www.transtats.bts.gov/>>

Acknowledgements

We thank M. Á. Serrano for useful comments and discussions. This work was supported in part by DGES grant FIS2007-66485-C02-02, Generalitat de Catalunya grant No. SGR00889, the Ramón y Cajal program of the Spanish Ministry of Science, by NSF CNS-0434996 and CNS-0722070, by DHS N66001-08-C-2029 and by Cisco Systems.

Additional information

Supplementary Information accompanies this paper on www.nature.com/naturephysics. Reprints and permissions information is available online at <http://npg.nature.com/reprintsandpermissions>. Correspondence and requests for materials should be addressed to M.B.