Greedy Routing on Hidden Metric Spaces as a Foundation of Scalable Routing Architectures

Dmitri Krioukov,
kc claffy, and Kevin Fall
CAIDA/UCSD, and
Intel Research, Berkeley

Problem

- High-level
  - Routing is one of the most fundamental and yet most complex functions that networks perform
  - Scalability of routing limits the scalability of future network designs and architectures that FIND seeks to explore
    - Routing in “harsh conditions” (ad-hoc, sensor, etc., networks) scales poorly due to convergence requirements (even simulations cannot get beyond a few thousand nodes)
    - Routing even in “normal conditions” (today’s Internet!) also bears ominous scalability concerns, primarily with convergence, but also with routing table growth (IAB Routing Scalability Report, RFC4984)
  - Until we solve the fundamental routing scalability problems, we cannot hope that our future network architectures will truly scale
- Low-level
  - Routing to a destination requires knowing where it is
  - Maintaining this knowledge in current Internet requires routers to continually propagate topology update messages
  - [Korman and Peleg, 2006] show that the number of such messages cannot scale better than linearly with network size in the worst case
  - [Krioukov et al, 2007] show that this worst case holds for observed Internet topologies
  - Scalable routing for Internet-like topologies is impossible, until we learn to route without topology updates, i.e., without precise knowledge of the network topology

Main idea: greedy routing on hidden metric spaces as a possible solution

- How can we route if we do not know network topology? That's heresy!
- [Milgram, 1966]'s experiments demonstrated that humans can efficiently route through the global acquaintance network topology knowing only some meta-information about their immediate neighbors and final destination
- It is true for other complex networks also!
- Social, biological, and communication networks, including the Internet, share the following peculiarities:
  - *scale-free*: node degree distribution follows power law $P(k) \sim k^{-\gamma}$ with $2 \leq \gamma \leq 3$
  - *strong clustering*: many triangular subgraphs
- Recent progress on compact routing in theoretical computer science achieves astonishing scalability of hierarchical routing on static scale-free graphs
- The above insights and discoveries led us to conjecture that the scale-free, strongly clustered structure of complex networks maximizes the efficiency of greedy routing guided by hidden metric spaces that underlie these networks
- Reconstruction of the hidden metric space underlying the real Internet would lead to an addressing architecture that allows for efficient and infinitely scalable routing without topology updates

How hidden metric spaces shape the network structure and function

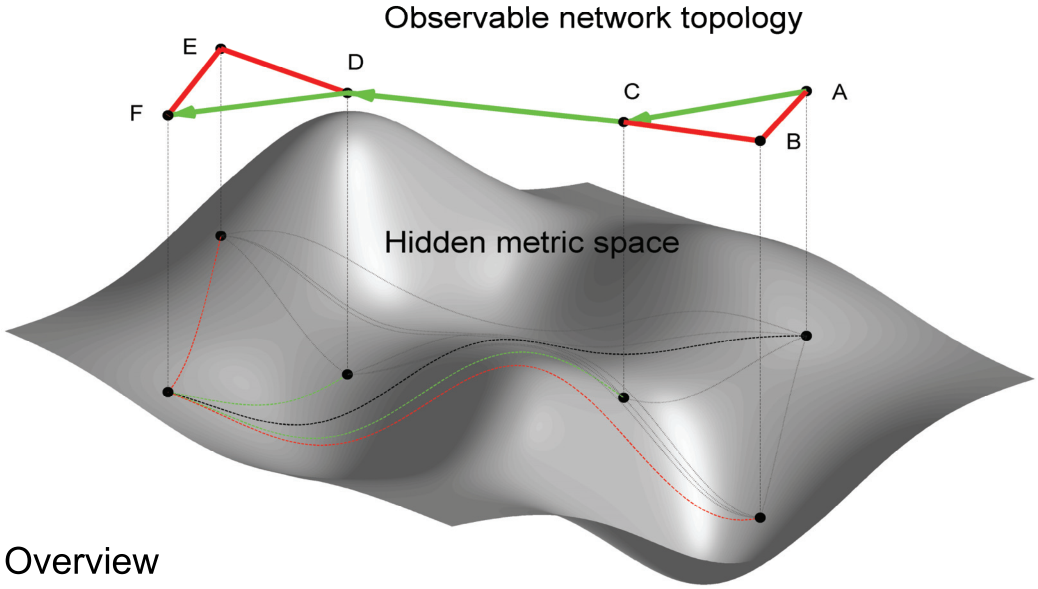

- Overview
  - The smaller the distance between two nodes in the hidden metric space (HMS), the more likely they are connected in the observable network.
  - If node *A* is close to node *B*, and *B* is close to *C* in the hidden space, then *A* and *C* are necessarily close, and all three nodes are likely connected in the observable network. The geometry of the hidden space thus provides a natural explanation for clustering observed in many complex networks.
  - The HMS also guides the greedy routing process: If node *A* wants to reach node *F*, it looks for which of its neighbors is closest to target F *in the hidden space*. Hidden distance *CF* is smaller than *BF*, therefore *A* forwards to *C*; node *C* performs similar calculations and selects its neighbor *D* as the next hop. Node *D* is directly connected to *F*; the resulting path is $A \to C \to D \to F$
- Concrete model
  - To demonstrate the effect of HMSs on navigability of complex networks, we construct a network model in which the HMS is the simplest possible metric space – a circle
  - Our model has only two parameters:
    - γ, the exponent of the power-law node degree distribution in the network; and
    - α, which is, roughly, the strength of clustering in the network
  - All nodes are uniformly distributed on the circle and
  - Each pair of nodes in the network:
    - have expected degrees k and k', drawn from probability distribution $P(k) \sim k^{-\gamma}$
    - are HMS-located at polar angles θ and θ' (hidden distance $d(\theta,\theta')$) on the circle
    - are connected with the probability that decreases with the hidden distance d between them as $1/d^{\alpha}$ and increases with their expected degrees as $(kk')^{\alpha}$

Self-similarity of complex networks as the first evidence that hidden metric spaces do exist

- We discovered that degree distributions, correlations, *and clustering* in real complex networks, including the Internet, are self-similar w.r.t. a simple degree renormalization illustrated on the right
- In randomized versions of these networks degree distributions and correlations are still self-similar, *but clustering is not*

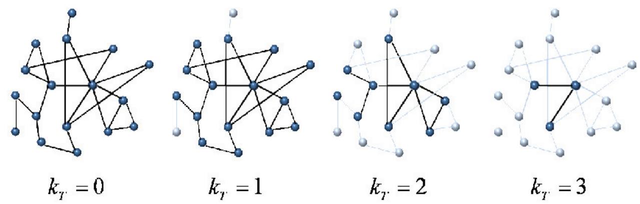

- k_T - renormalization degree threshold
- k_i - internal node degree in k_T-subgraphs
- c - clustering

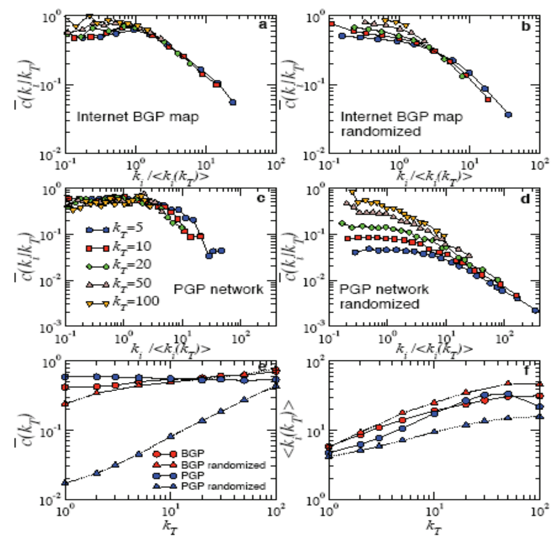

- Networks constructed by our HMS-model exhibit qualitatively similar effects, which suggests that
- **Hidden metric spaces are a plausible explanation of self-similarity of complex networks w.r.t. degree renormalization, which is equivalent to *hidden distance rescaling***

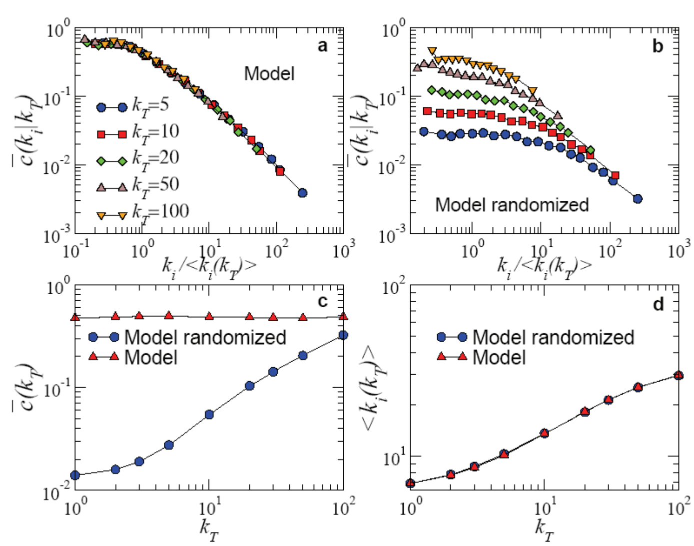

Navigability of complex networks as the first evidence that greedy routing is efficient on the topologies of real networks

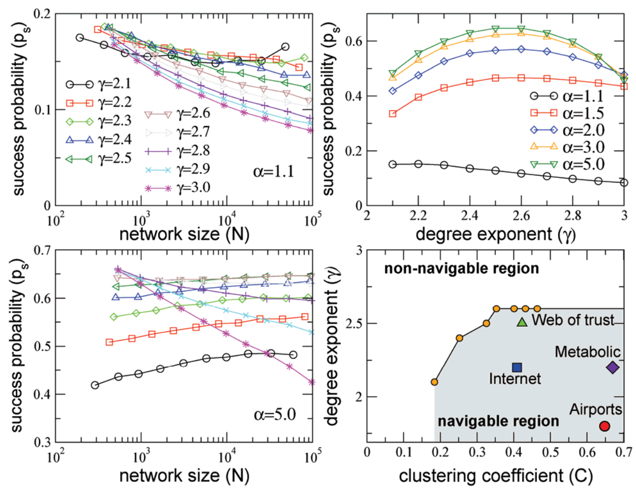

- We discovered that the topological properties of real complex networks correspond to the parameter values in our model (*2.1≤γ≤2.6, α≥1.5*) that yield the most efficiently navigable networks in the greedy routing framework; this finding confirms that
  - our main conjecture is plausible
  - greedy routing architectures built on top of hidden metric spaces are efficient
  - real networks may explicitly or implicitly evolve to increase their navigability
- Explanation of navigability of real networks
  - specific values of *γ* (*2.1≤γ≤2.6*) correspond to exactly the densities of hubs (high-degree nodes) needed to interconnect all hidden distance scales, necessary for routing efficiency
  - strong clustering (large *α*) ensures that orientation of links is congruent with the node degree hierarchy, i.e., most links connect the network core to the periphery
  - **navigable networks thus have a peculiar hierarchical structure congruent with the optimal layout of efficient greedy routing paths**

• navigable region – combination of γ and C for which p_s increases with network size

• Internet – AS-level topology from BGP RouteViews

• Web of trust – PGP trust network

• Metabolic – network of metabolic reactions of *E. coli*

• Airports – global airport network

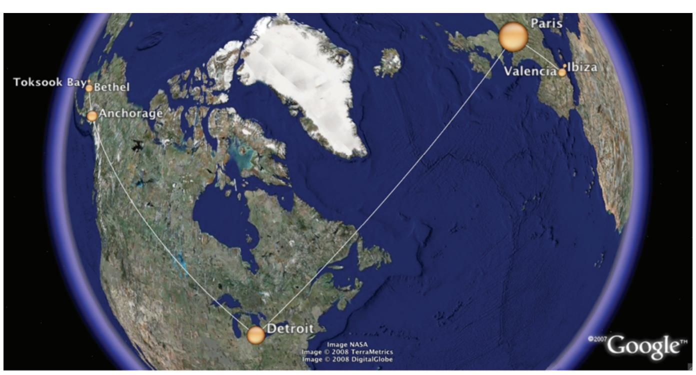

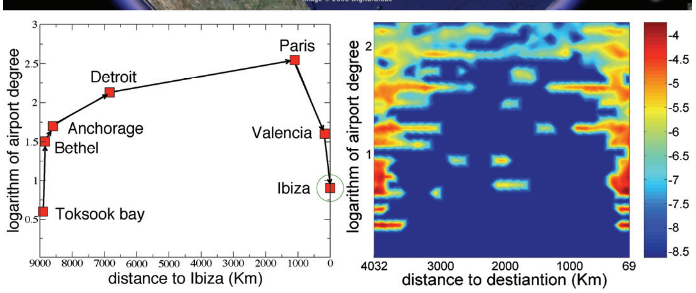

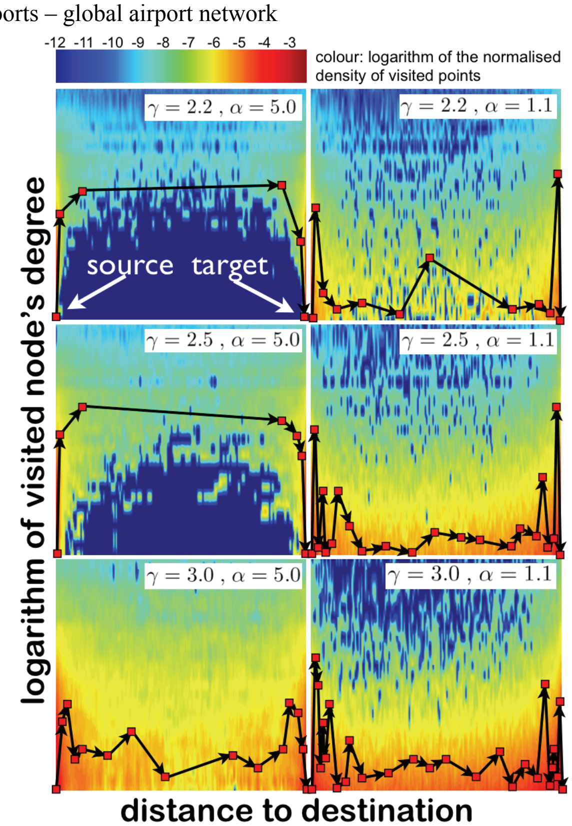

k-core decomposition of real and modeled networks

- Node coreness is a measure of how deep within the network core the node is located (the *k*-core is the maximal subgraph induced by nodes that have *k* or more connections to other nodes in the subgraph; and a node's coreness is *k* such that the *k*-core contains the node but the *k+1*-core does not)
- All nodes are color-coded based on their coreness and size-coded based on their degrees; higher-coreness nodes are also closer to circle centers
- The first row depicts two real networks: the AS-level Internet as seen by the Border Gateway Protocol (BGP) and the Pretty Good Privacy (PGP) social network of mutual trust relationships
- The rest of the rows show modeled networks for different values of power-law exponent γ in cases with strong ($\alpha=5.0$) and weak ($\alpha=1.1$) clustering
- The network size for all real and modeled cases is approximately 10^4
- There is a remarkable similarity between real networks and modeled networks with $\gamma=2.2$ and $\alpha=5.0$, i.e., with the most navigable parameter values

Internet AS topology (BGP tables)

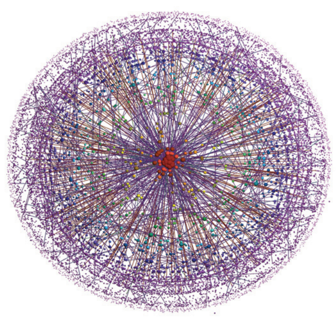

Social trust network (PGP)

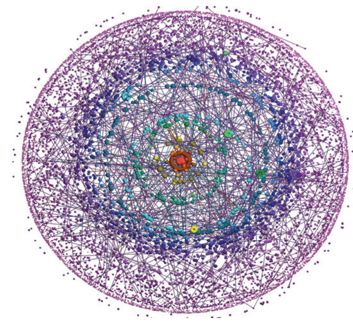

Model with $\gamma=2.2$ and $\alpha=5.0$

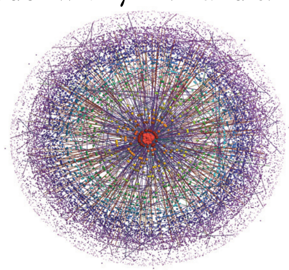

Model with $\gamma=2.2$ and $\alpha=1.1$

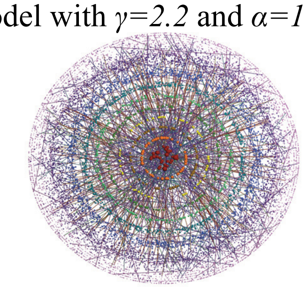

Model with $\gamma=2.5$ and $\alpha=5.0$

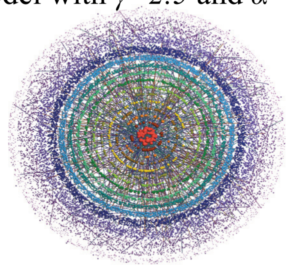

Model with $\gamma=2.5$ and $\alpha=1.1$

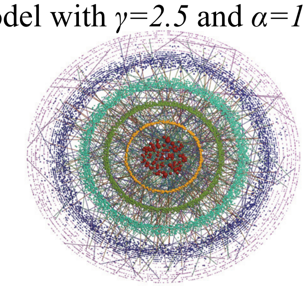

Model with $\gamma=3.0$ and $\alpha=5.0$

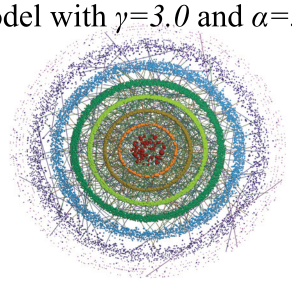

Model with $\gamma=3.0$ and $\alpha=1.1$

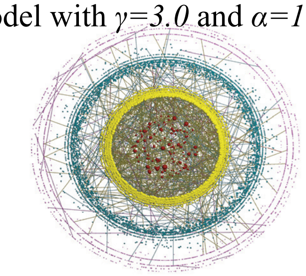

Negative curvature as the main property of hidden metric spaces

- We discovered that
  - Scale-free networks are congruent w.r.t. hidden hyperbolic geometries
  - This congruency is robust w.r.t. network dynamics/evolution
- The high-level reason why HMSs are negatively curved
  - Nodes in complex networks can often be hierarchically classified
  - Hierarchies are (approximately) trees
  - Trees embed “almost” isometrically in hyperbolic spaces
- Our specific findings are
  - Hidden hyperbolic metric spaces explain, simultaneously, the scale-free degree distributions and strong clustering in complex networks
  - Greedy routing mechanisms in these settings may offer virtually infinitely scalable routing algorithms for future communication networks

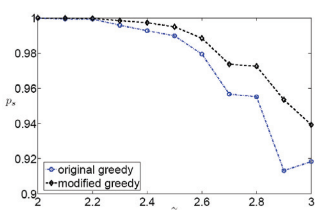

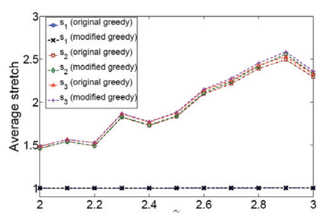

- Notations
  - p_s – percentage of successful paths
  - s_1 – stretch of greedy paths w.r.t. shortest paths in the graph
  - $s_{2,3}$ – stretch of shortest and greedy paths w.r.t. lengths of hidden hyperbolic geodesics
  - γ – exponent of power-law degree distributions
  - *Original greedy* gets stuck if no neighbor closer to the destination
  - *Modified greedy* still forwards to the neighbor closest to the destination among all the neighbors and gets stuck if the packet loops
- As network topology changes, the greedy routing efficiency characteristics deteriorate very slowly
  - For $\gamma \leq 2.5$, removal of up to *10%* of the links from the topology degrades the percentage of successful path by less than *1%*

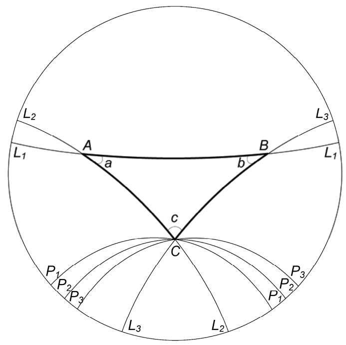

Poincaré disc model of the hyperbolic plane: L_i and P_i are infinite straight lines

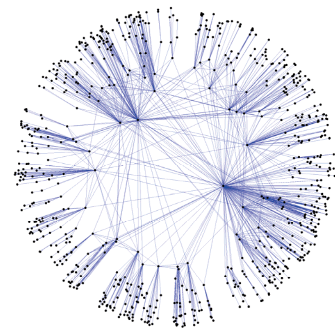

Visualization of a modeled network

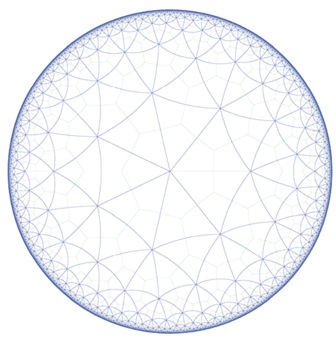

Tessellation and tree embedding

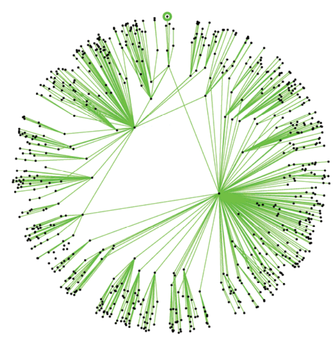

Successful paths

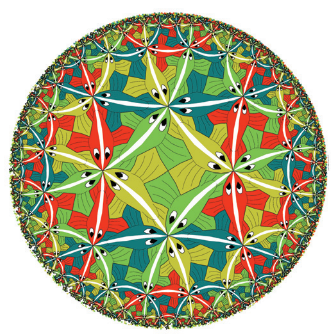

Tessellation art

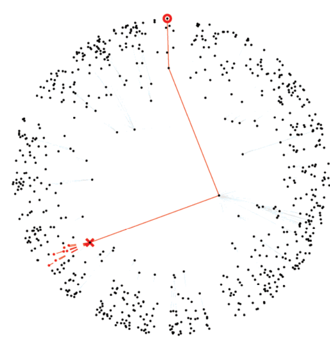

Unsuccessful paths

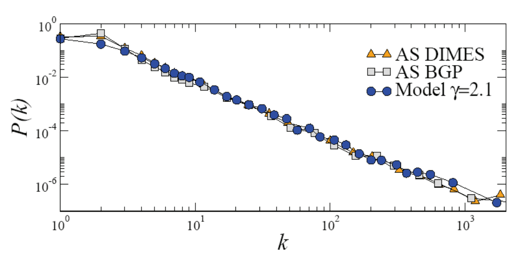

Node degree distributions in modeled networks and real Internet