Routing in the Internet and Navigability of Scale-Free Networks

Dmitri Krioukov
CAIDA/UCSD
dima@caida.org
F. Papadopoulos, M. Boguñá, A. Vahdat

~~Science or engineering?~~

- ~~Network science vs. network engineering~~
- ~~Computer science vs. computer engineering~~
- ~~Study existing networks vs. designing new ones~~
- We cannot really *design* truly large-scale systems (e.g., Internet)
  - We can design their building blocks (e.g., IP)
  - But we cannot fully control their large-scale behavior
  - Since they acquire some elements of self-organization, or self-evolving (self-*) behavior beyond our control
- Let us study existing large-scale networks and try to use what we learn in designing new ones
  - Discover "nature-designed" efficient mechanisms that we can reuse (or respect) in our future designs

Internet

- Microscopic view ("design")
  - IP/TCP, routing protocols
  - Routers
  - Per-ISP router-level topologies
- Macroscopic view ("non-design")
  - Global AS-level topology is a cumulative result of local, decentralized, and rather complex interactions between AS pairs
  - Surprisingly, in 1999, it was found to look completely differently than engineers had thought
    - It is not a grid, tree, or classical random graph
    - It shares all the main features of topologies of other complex networks
      - scale-free (power-law) node degree distributions ($P(k) \sim k^{-\gamma}$, $\gamma \in [2,3]$)
      - strong clustering (large numbers of 3-cycles)
- The big problem is that "design" has now to deal with "non-design"
  - Routing protocols have to find and promptly update paths to all destinations in the Internet

IPv4 INTERNET

TOPOLOGY MAP

AS-level INTERNET GRAPH

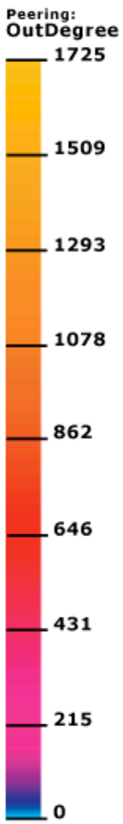

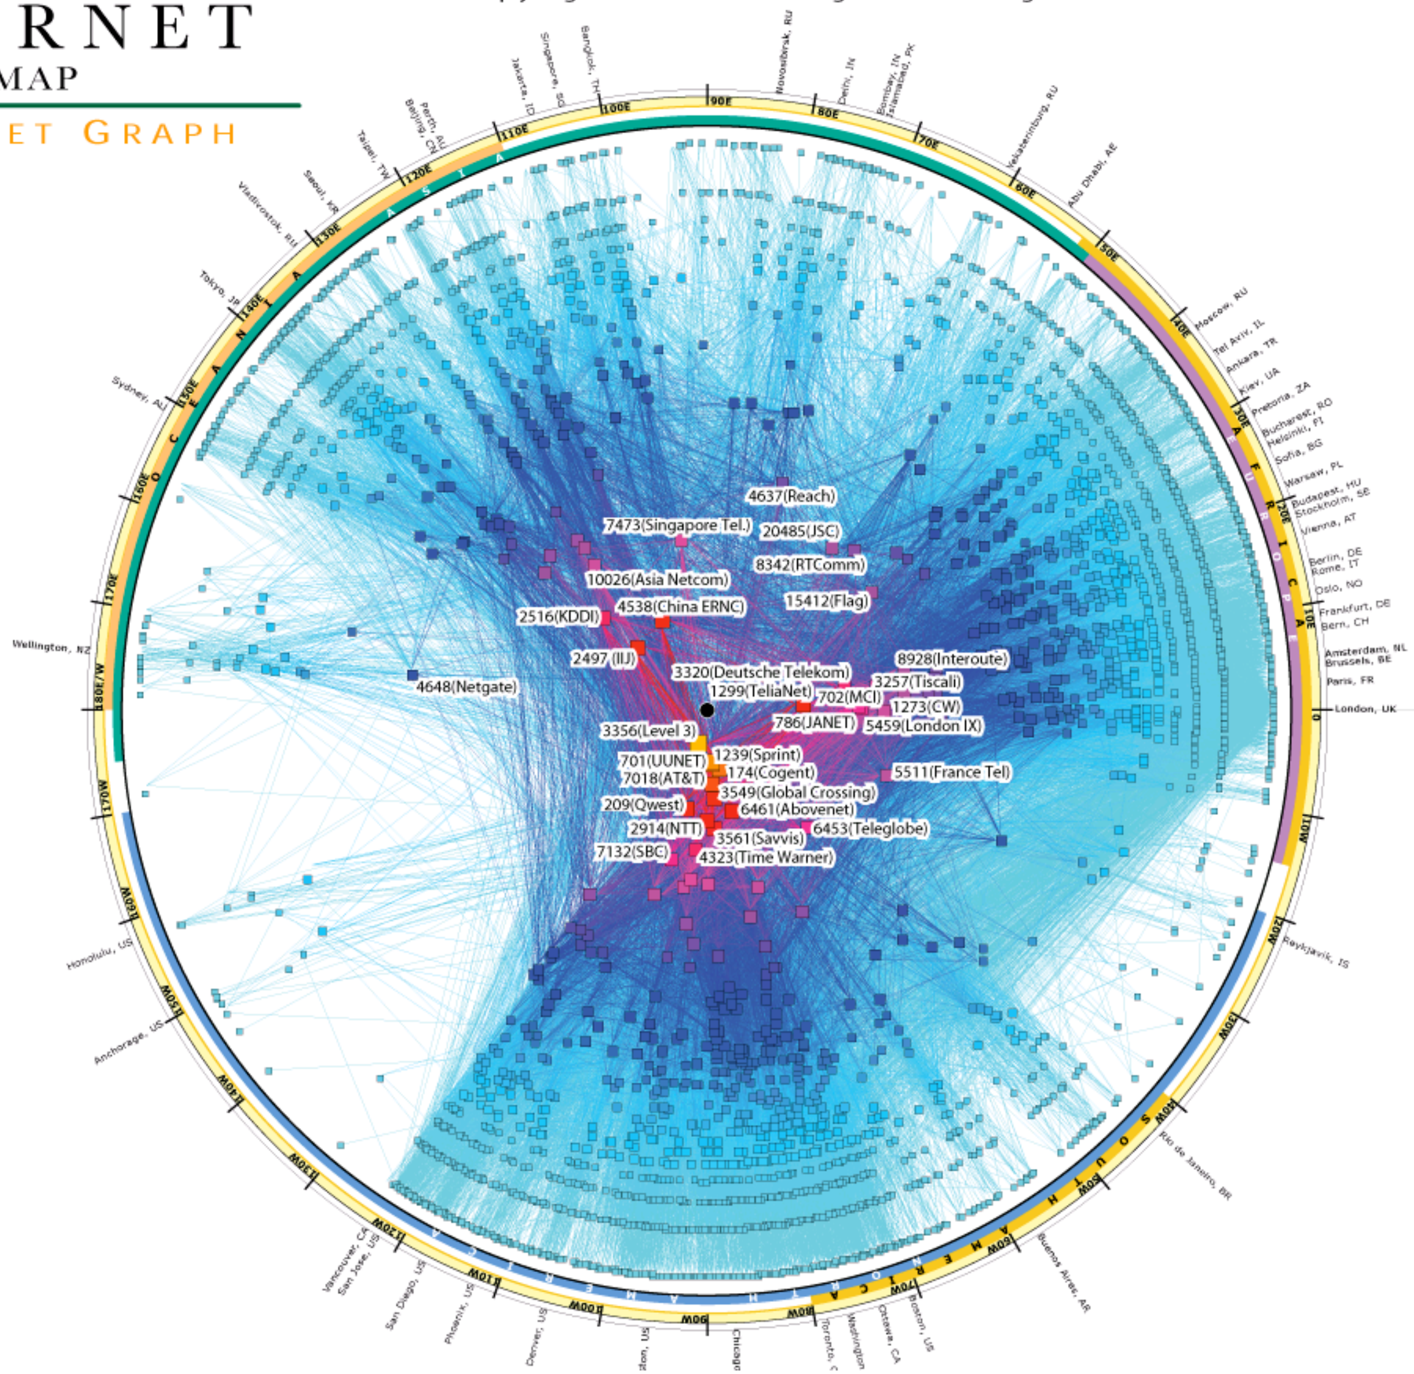

Routing practice

- Global (DFZ) routing tables
  - 300,000 prefix entries (and growing)
  - 30,000 ASs (and growing)
- Routing overhead/convergence
  - BGP updates
    - 2 per second on average
    - 7000 per second peak rate
  - Convergence after a single event can take up to tens of minutes
- Problems with design?
  - Yes and no

Routing theory

- There can be no routing algorithm with the number of messages per topology change scaling better than linearly with the network size in the worst case
- Small-world networks are this worst case

- *Is there any workaround?*
- *If topology updates/convergence is so expensive, then may be we can route without them, i.e., without global knowledge of the network topology?*
- *Let us look at the existing systems*

Milgram's experiments

- Settings: random people were asked to forward a letter to a random individual by passing it to their friends who they thought would maximize the probability of letter moving "closer" to the destination
- Results: surprisingly many letters (*30%*) reached the destination by making only *~6* hops on average
- Conclusion:
  - People do not know the global topology of the human acquaintance network
  - But they can still find (short) paths through it

Navigability of complex networks

- In many (if not all) existing complex networks, nodes communicate without any global knowledge of network topologies
- How is this possible???

Hidden metric space explanation

- All nodes exist in a metric space
- Distance in this space abstract node similarities
- Network consists of links that exist with probability that decreases with the hidden distance
- More similar/close nodes are more likely to be connected
- The result is that all nodes exist in “two places at once”:
  - a network
  - a hidden metric space
- So that there are two distances between each pair of nodes
  - the length of shortest path between them in the network
  - hidden distance

Greedy routing (Kleinberg)

- To reach a destination, each node forwards information to the one of its neighbors that is closest to the destination in the hidden space

Hidden space visualized

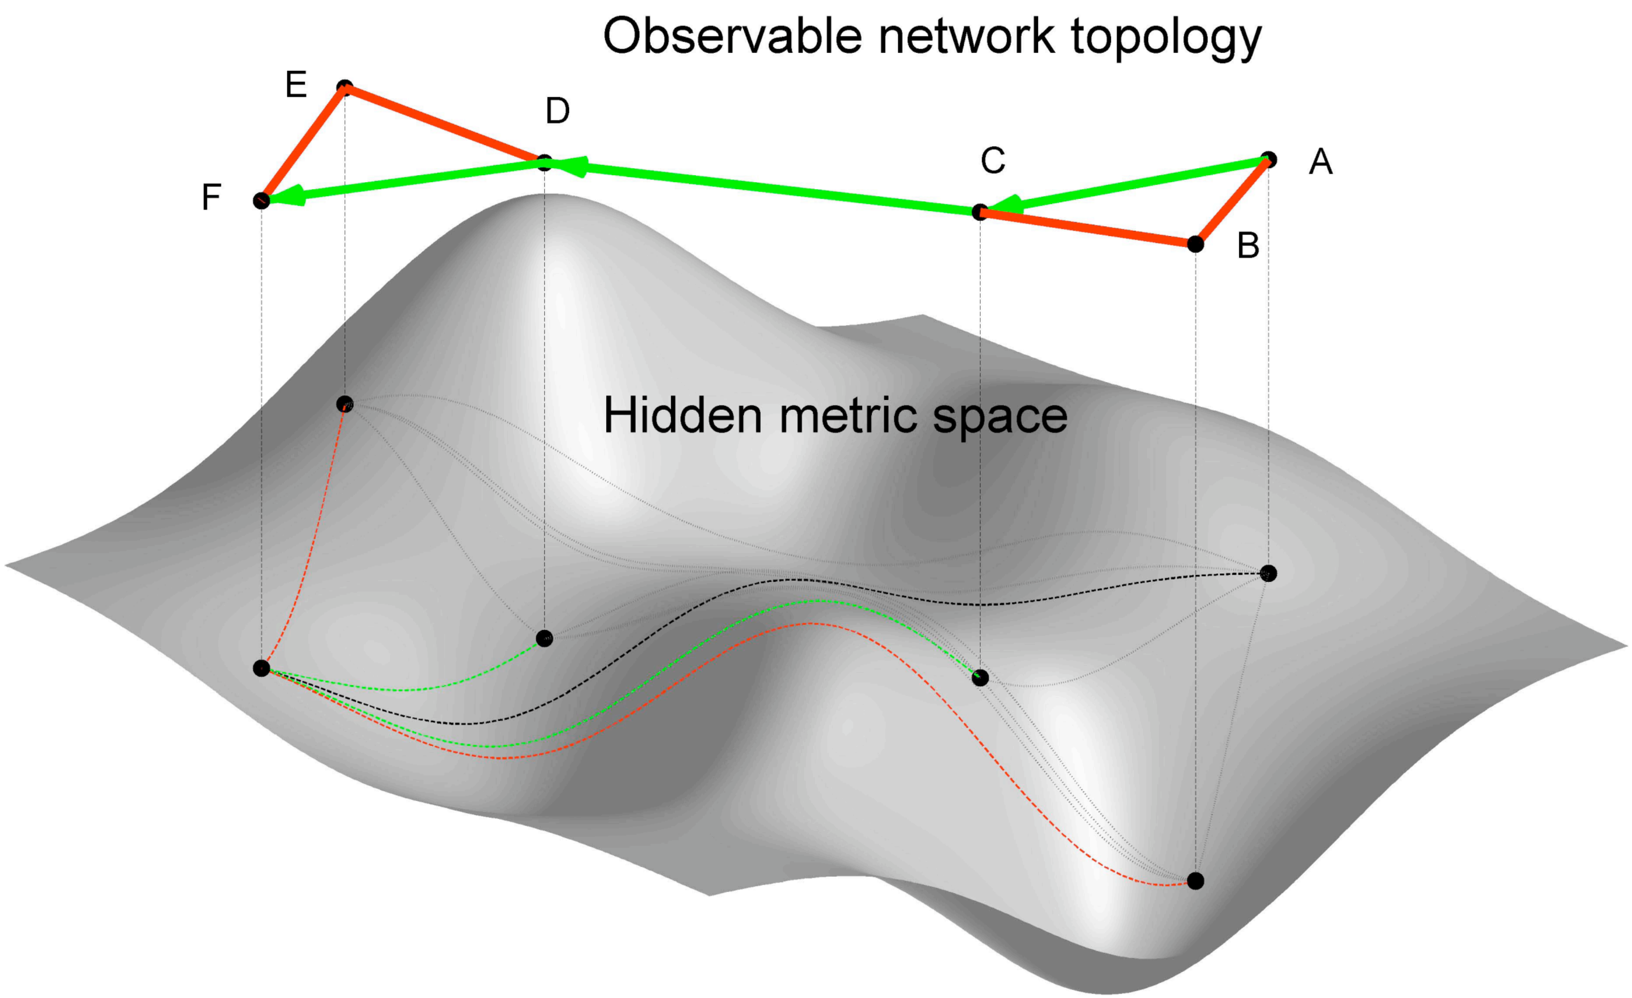

Questions raised by the approach

- What is the hidden space?
- What are the node positions in it?
- What is the connection probability?
- How efficient is the greedy routing process?
  - How often greedy-routing paths get stuck at nodes that do not have any neighbors closer to the destination than themselves
  - How closely greedy-routing paths follow the shortest paths in the network
- What topologies are navigable, i.e., congruent w.r.t. greedy routing, i.e., make it efficient?

Hidden spaces are metric spaces

- Using the simplest metric space (a circle), we show that
  - the triangle inequality in hidden spaces
    - transitivity of being similar/close

  explains
  - strong clustering in real networks
    - transitivity of being connected
- It also explains their self-similarity

Self-similarity of complex networks (existing knowledge)

- Self-similarity w.r.t. rescaling (of distances, time, etc.)
  - Fluctuations at phase transitions
  - Fractals
- Fractal dimension
  - Box covering procedures
- Complex networks
  - Degree distributions are self-similar
  - Some networks are self-similar w.r.t. box covering
  - But no distance rescaling since these networks are small worlds

Self-similarity of complex networks (new perspective)

- If complex networks are embedded in hidden metric spaces, then distance rescaling in "large-world" hidden spaces is equivalent to degree renormalization
- W.r.t. this degree renormalization, all of the following are self-similar in real networks and modeled networks with metric spaces underneath
  - degree distribution
  - degree correlations
  - clustering
- Only degree distributions are self-similar in maximally random networks with degree distribution of real or modeled networks
- Evidence that metric spaces do underlie real networks which are self-similar w.r.t. hidden distance rescaling

Navigability mechanisms

- More navigable networks are networks with
  - more heterogeneous node degree distributions
    - more hubs
  - stronger clustering
    - stronger influence of hidden distances on links
    - stronger congruency between hidden geometries and observed topologies
    - stronger congruency between greedy and shortest paths
- Greedy routing paths follow navigable path pattern

Navigable path pattern

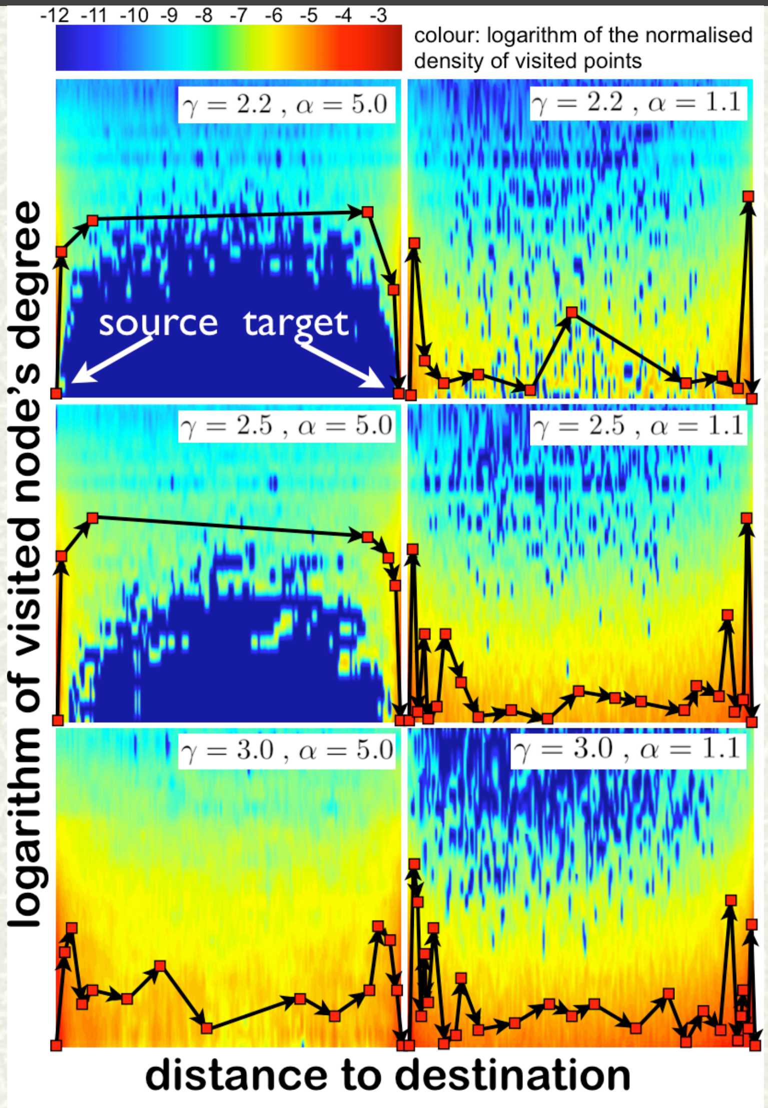

Real networks are navigable

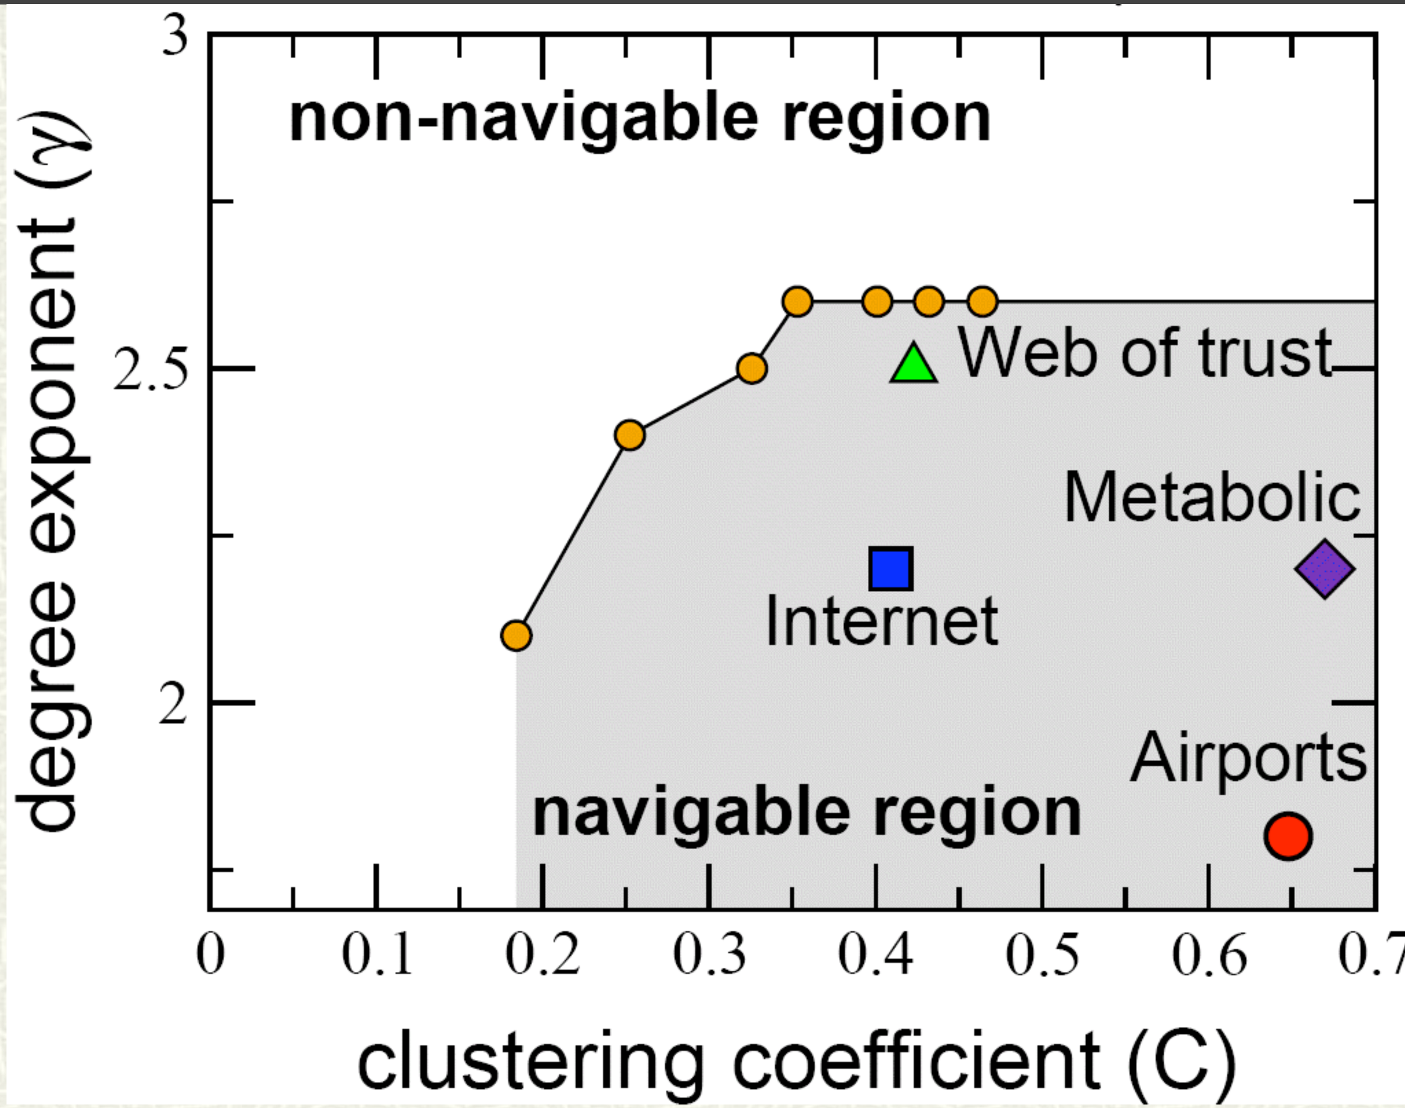

Hidden geometries

- What hidden geometries are maximally congruent with the navigability mechanisms of the observed complex network topologies?

Hidden metric spaces are hyperbolic

- Network nodes can often be hierarchically classified
  - Community structure (social and biological networks)
  - Customer-provider hierarchies (Internet)
  - Hierarchies of overlapping balls/sets (all networks)
    - The mapping between balls in R^d $B(x,r)$ and points $\alpha = (x,r)$ in H^{d+1} satisfies
      - If $|\alpha\text{-}\alpha'| \leq C$, then there exist $k(C)$ s.t. $k^{-1} \leq r/r' \leq k$ and $|x\text{-}x'| \leq k\ r$
      - If $|x\text{-}x'| \leq k\ r$ and $k^{-1} \leq r/r' \leq k$, then there exist $C(k)$ s.t. $|\alpha\text{-}\alpha'| \leq C$
- Hierarchies are (approximately) trees
- Trees embed isometrically in hyperbolic spaces

Hyperbolic geometry

- Geometry in which through a point not belonging to a line passes not one but infinitely many lines parallel to the given line

Poincaré disc model

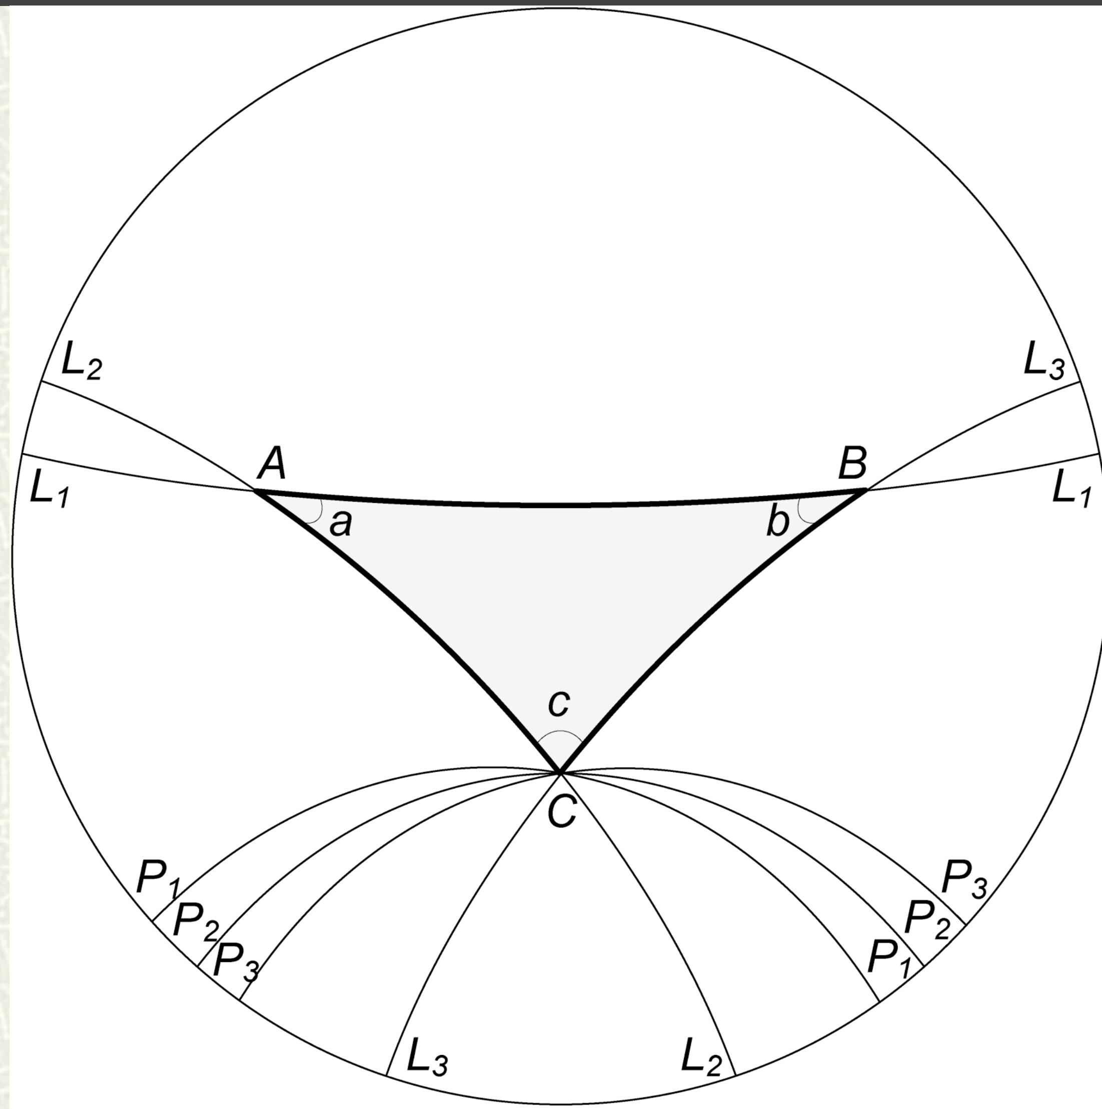

Tessellation and tree embedding

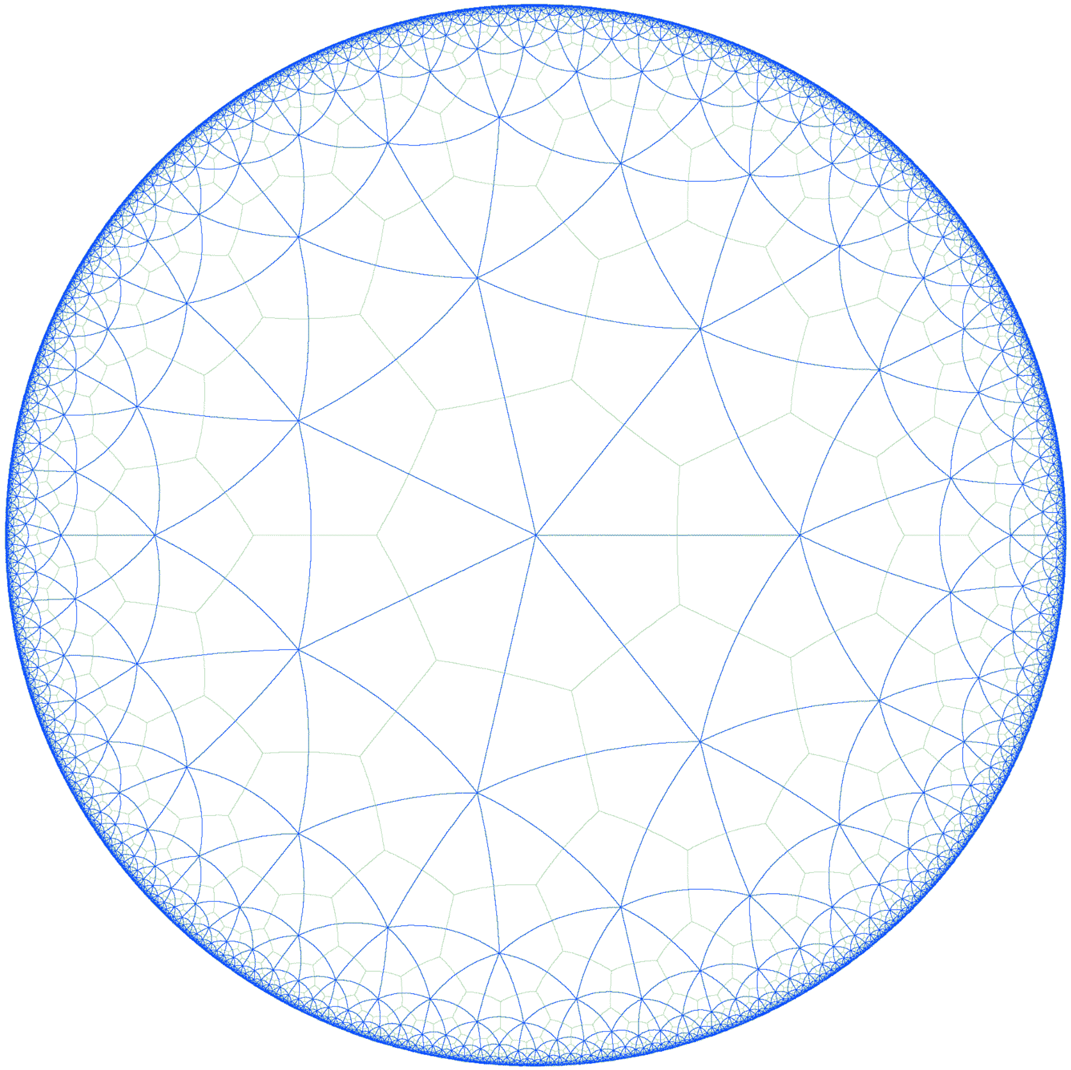

Tessellation art

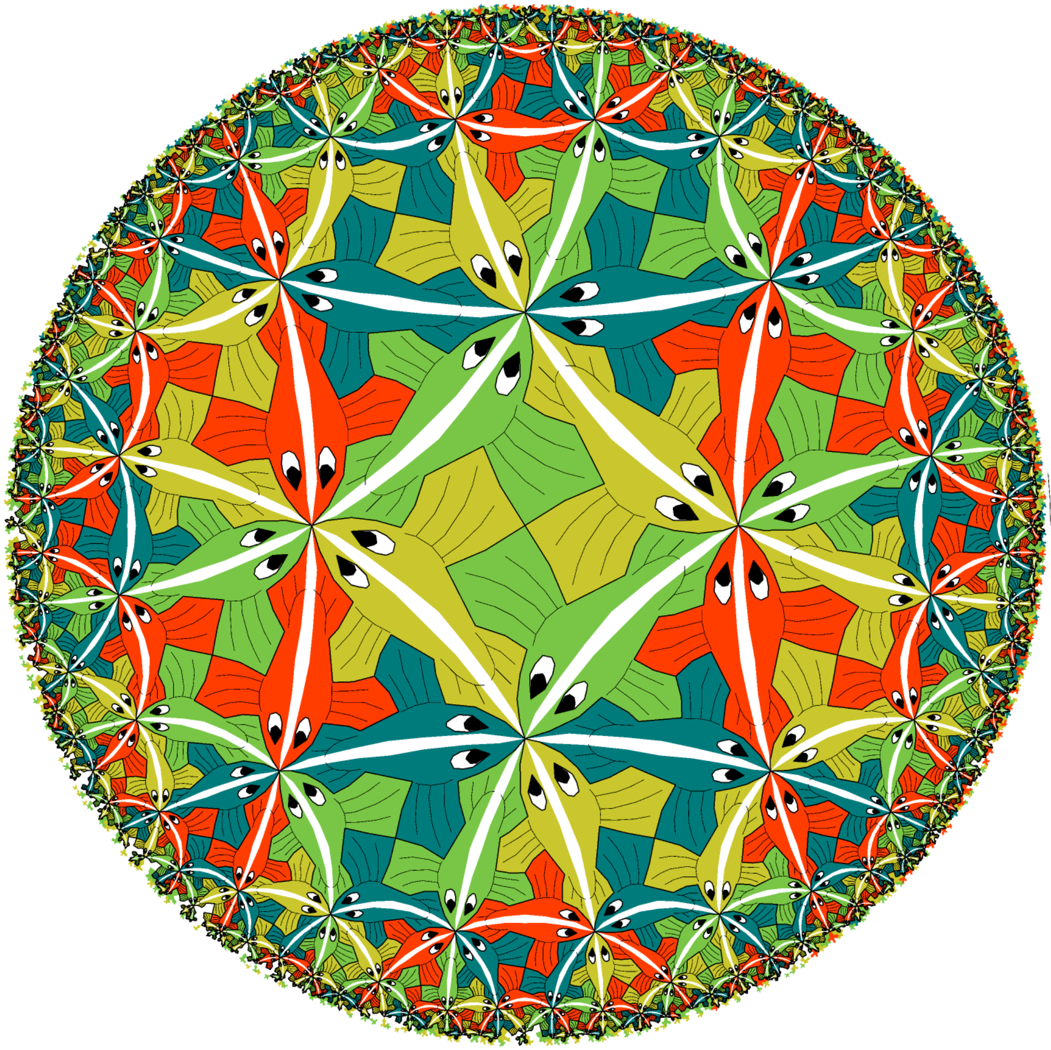

Geometry properties

Property	Euclid.	Spherical	Hyperbolic
Curvature	0	1	-1
Parallel lines	1	0	∞
Triangles are	normal	thick	thin
Shape of triangles	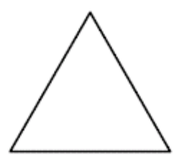	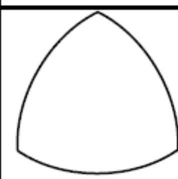	
Sum of angles	π	$> \pi$	$< \pi$
Circle length	$2\pi R$	$2\pi \sin R$	$2\pi \sinh R$
Disc area	$2\pi R^2/2$	$2\pi(1 - \cos R)$	$2\pi(\cosh R - 1)$

Main hyperbolic property

- The volume of balls and surface of spheres grow with their radius r as

$$e^{\alpha r}$$

  where $\alpha = (-K)^{1/2}(d-1)$, K is the curvature and d is the dimension of the hyperbolic space
- The numbers of nodes in a tree within or at r hops from the root grow as

$$b^r$$

  where b is the tree branching factor
- The metric structures of hyperbolic spaces and trees are essentially the same ($\alpha = \ln b$)

Hidden space in our model

- Hyperbolic disc of radius R, where $N = \kappa\, e^{R/2}$, N is the number of nodes in the network and κ controls its average degree
  - Average degree is fixed (by κ) to the same value (~*6*, like in many real networks) for all modeled networks

Node distribution

- Number of nodes $n(r)$ located at distance r from the disc center is

$$n(r) \sim e^{\alpha r}$$

where $\alpha = 1$ corresponds to the uniform node distribution in the hyperbolic plane of curvature -1

Connection probability

- Connected each two nodes if the distance between them is less than or equal to R

Average node degree at distance *r* from the disc center

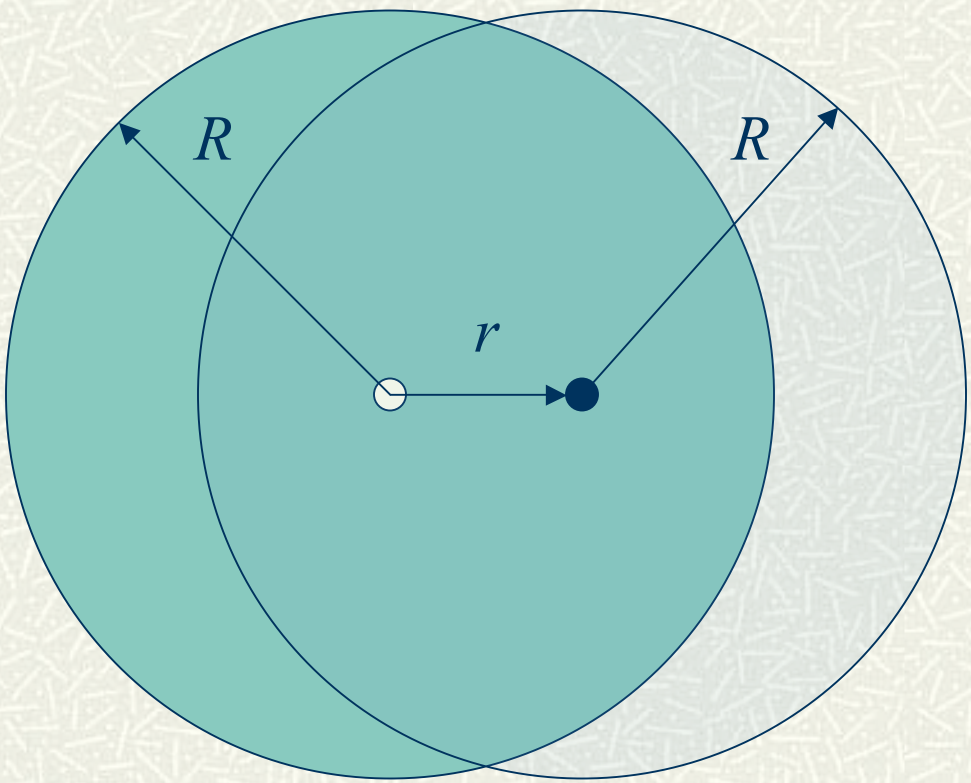

Average node degree at distance *r* from the disc center

- For $\alpha = 1$, we obtain a terse but exact expression

$$
\begin{aligned}
k(r) = \delta \Bigg\{ &(\cosh R - 2)\pi + 2 \left(\cosh\chi \arccos \frac{\cosh r \cosh\chi - \cosh R}{\sinh r \sinh\chi} \right. \\
&+ \left. \cosh R \arctan \frac{\cosh\chi - \cosh R \cosh r}{\sqrt{(\cosh r - \cosh(R-\chi))(\cosh(R+\chi) - \cosh r)}} \right) \\
&- \arctan \frac{8(\cosh r - \cosh R \cosh\chi)\sqrt{(\cosh r - \cosh(R-\chi))(\cosh(R+\chi) - \cosh r)}}{16(\cosh r - \cosh R\cosh\chi)\cosh r - 8\cosh^2 r + 4(\cosh^2 R + \cosh^2 R \cosh^2\chi + \cosh^2\chi) - 1} \Bigg\}. \qquad (42)
\end{aligned}
$$

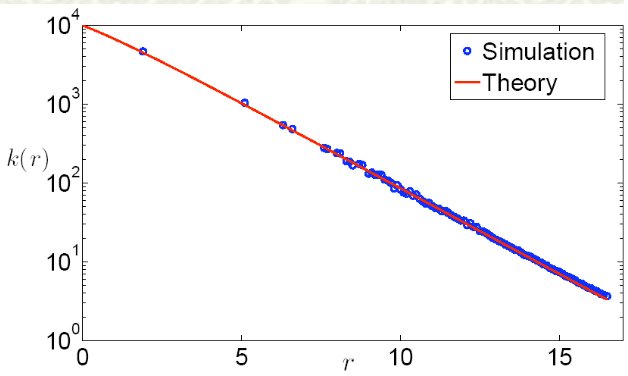

- For other α:

$$k(r) \sim e^{-\beta r}$$

where

$\beta = \alpha$ if $\alpha \leq ½$

$\beta = ½$ otherwise

Node degree distribution

- Is given by the combination of exponentials to yield a power law

  $$P(k) \sim k^{-\gamma}$$

  where

  $\gamma = 1 + \alpha/\beta =$

  2 if $\alpha \leq ½$; or

  $2\alpha + 1$ otherwise
- The uniform node distribution in the plane $(\alpha = 1)$ yields $\gamma = 3$

Node degree distribution: theory vs. simulations

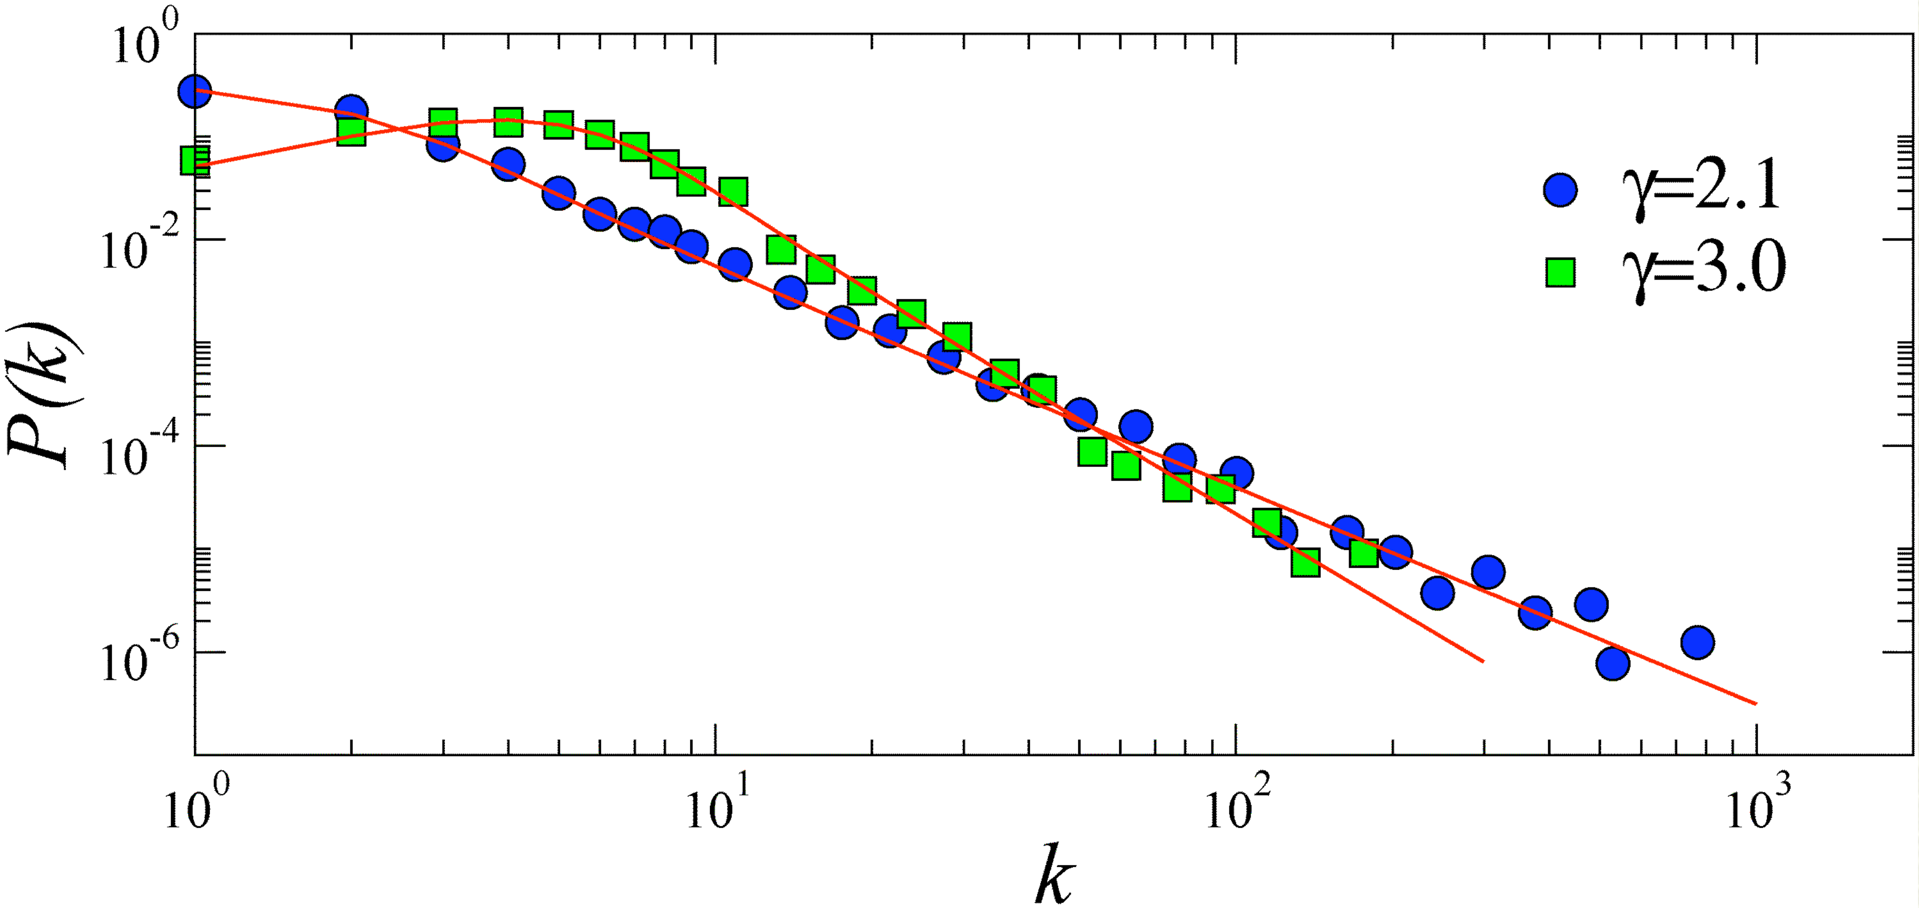

Node degree distribution: model vs. Internet

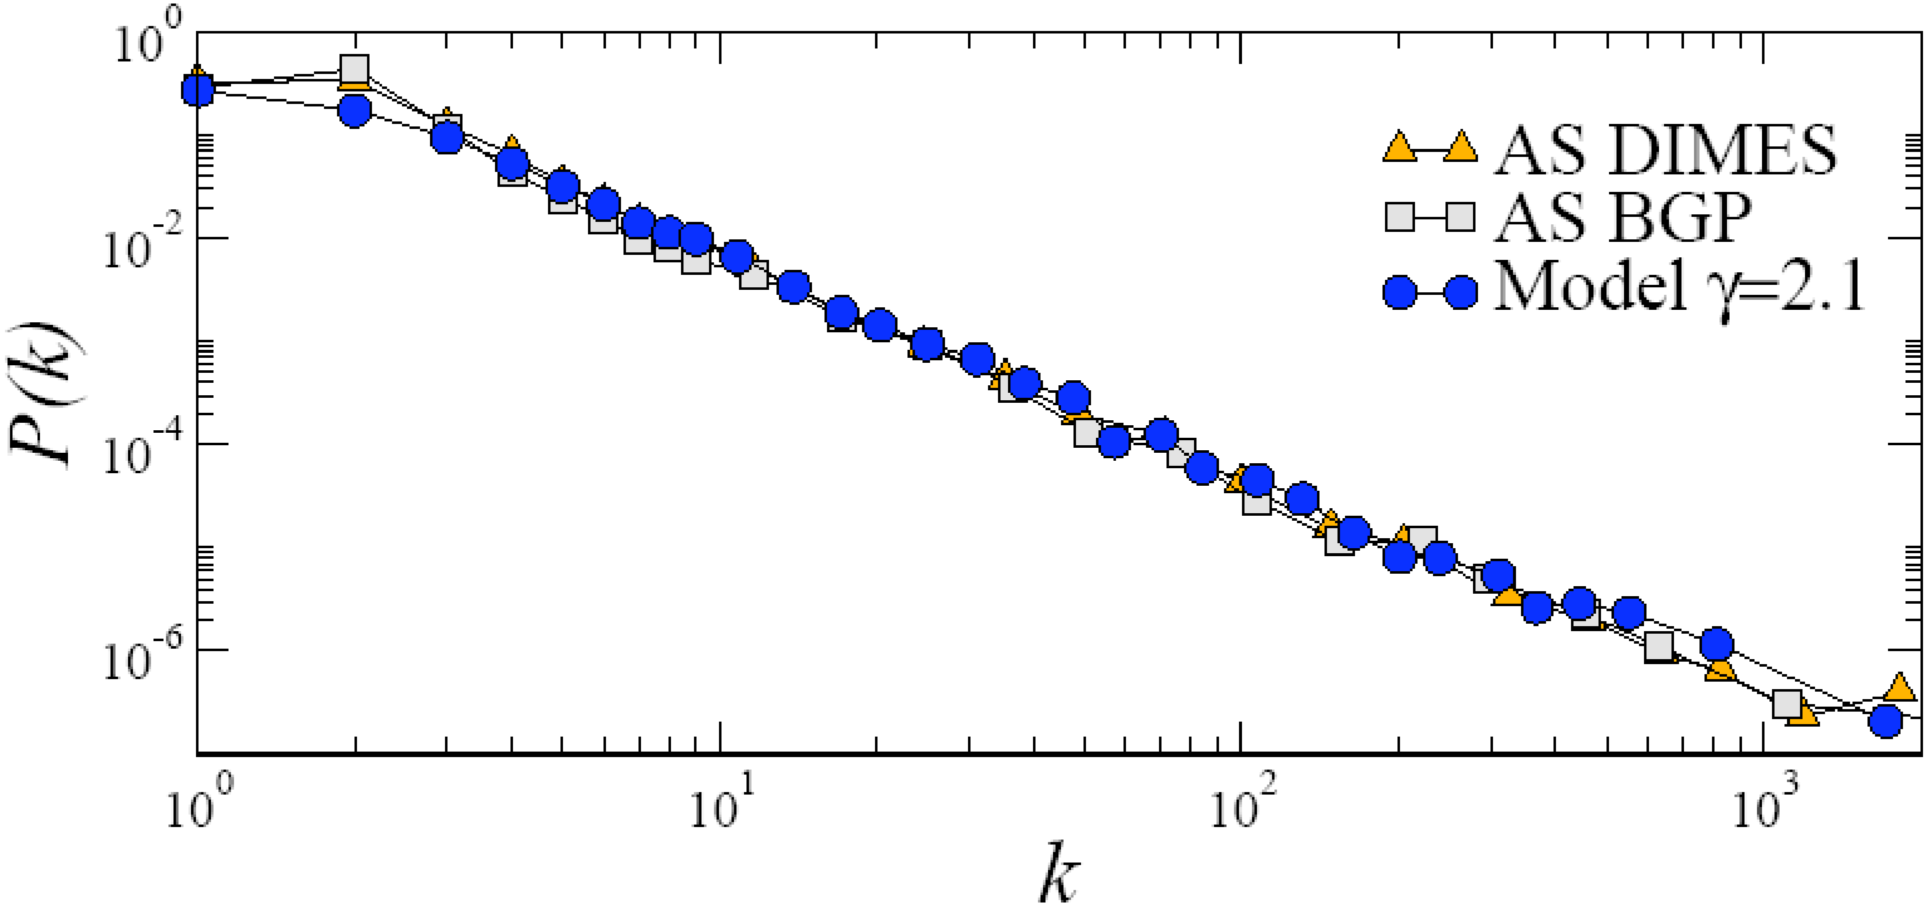

Degree correlations: model vs. Internet

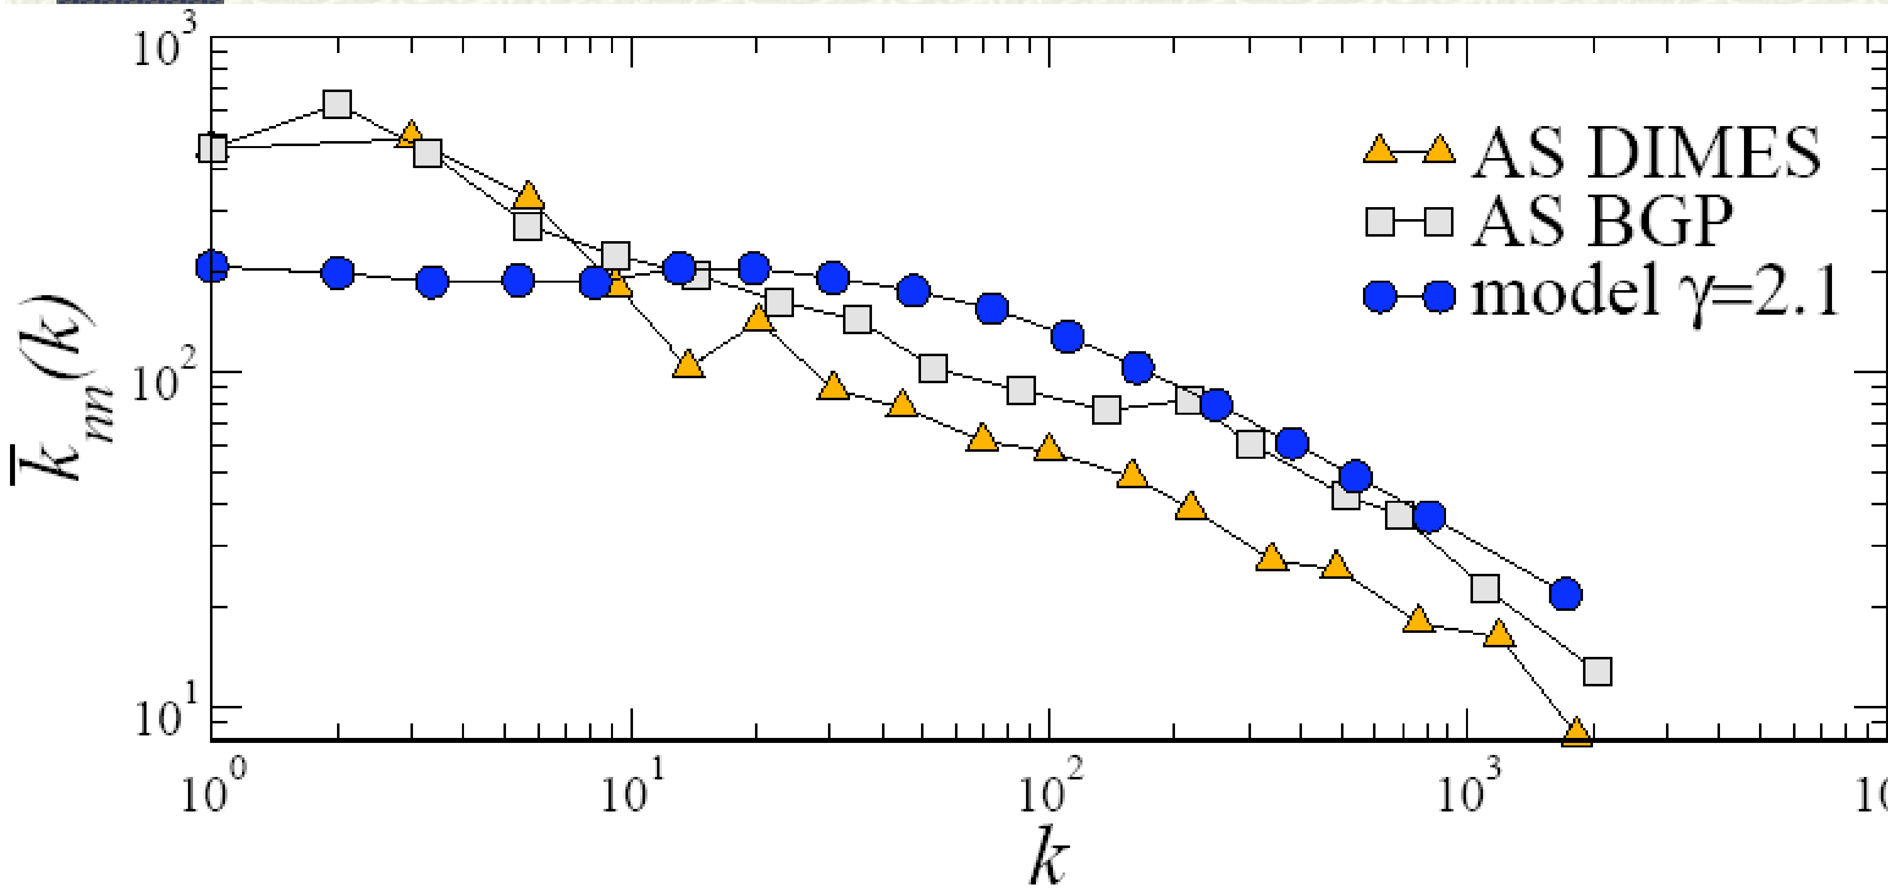

Clustering: model vs. Internet

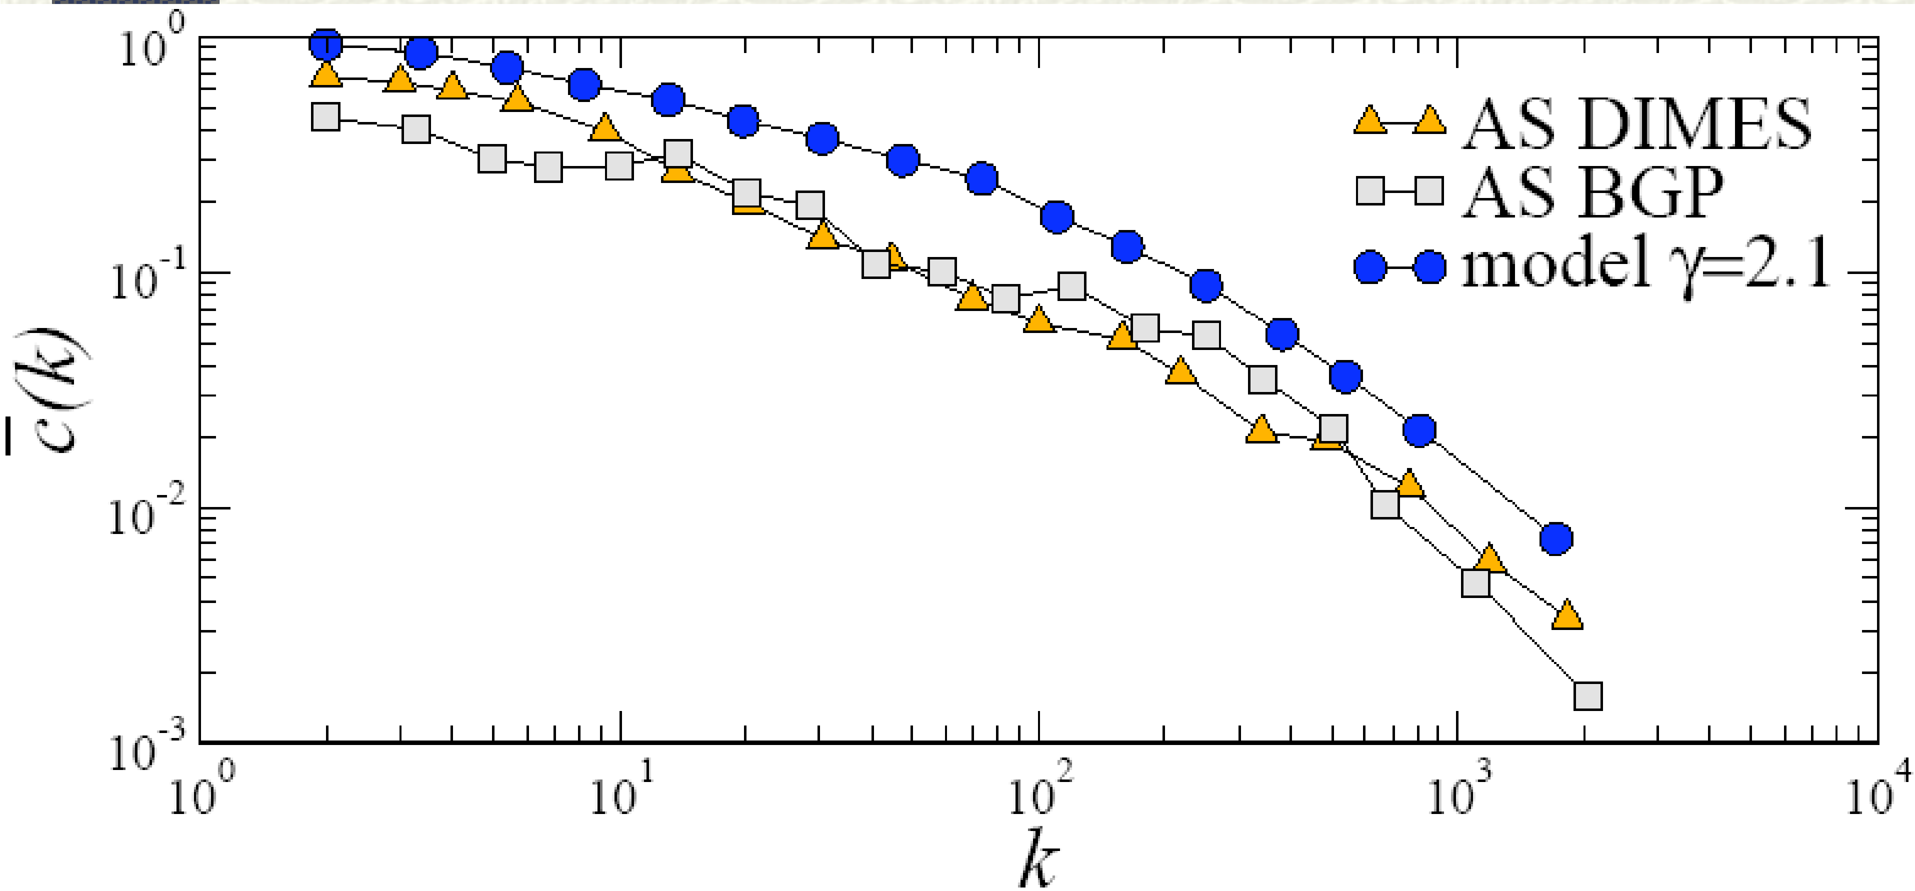

Visualization of a modeled network

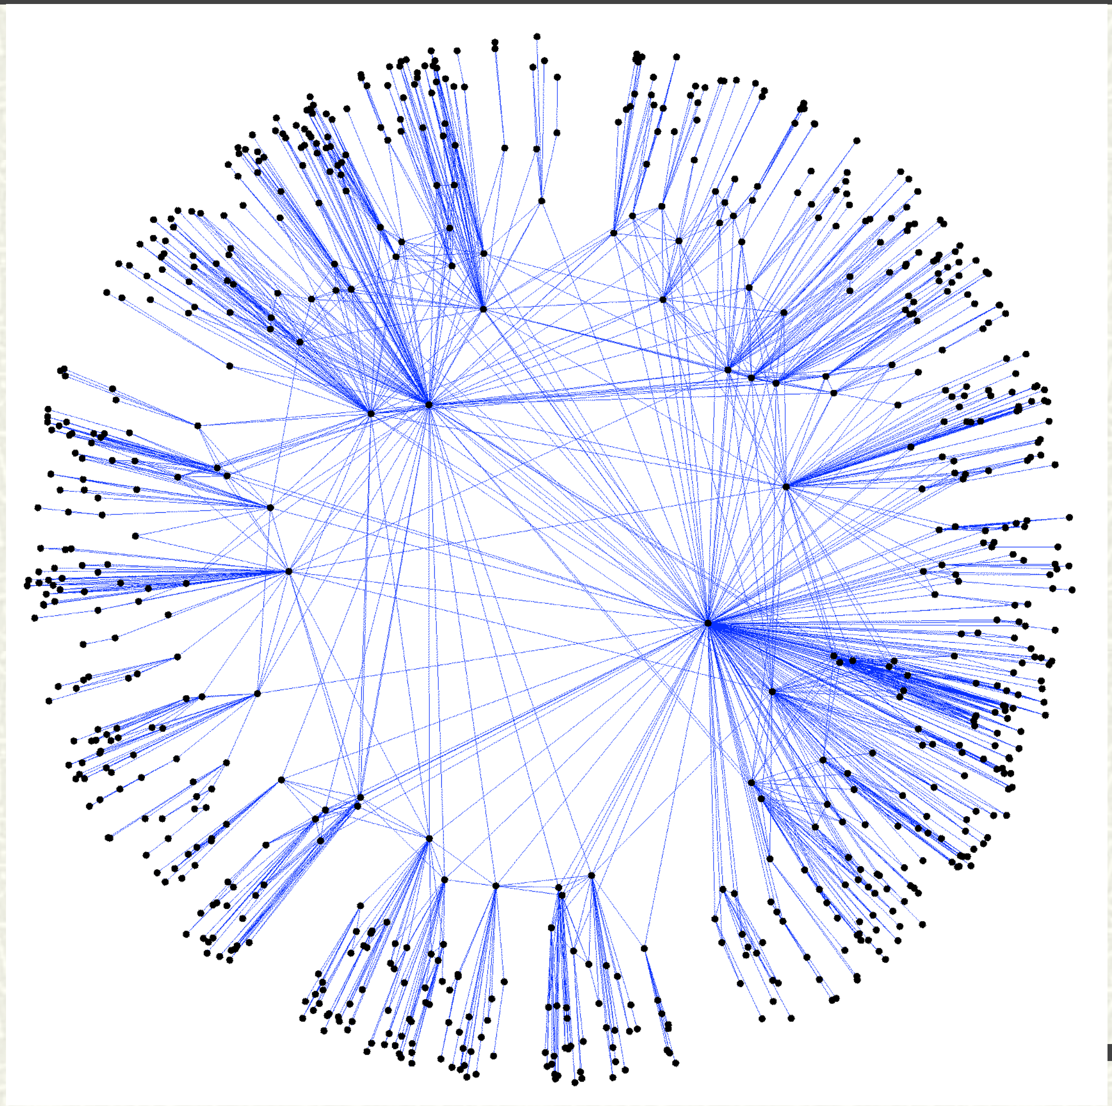

Successful greedy paths

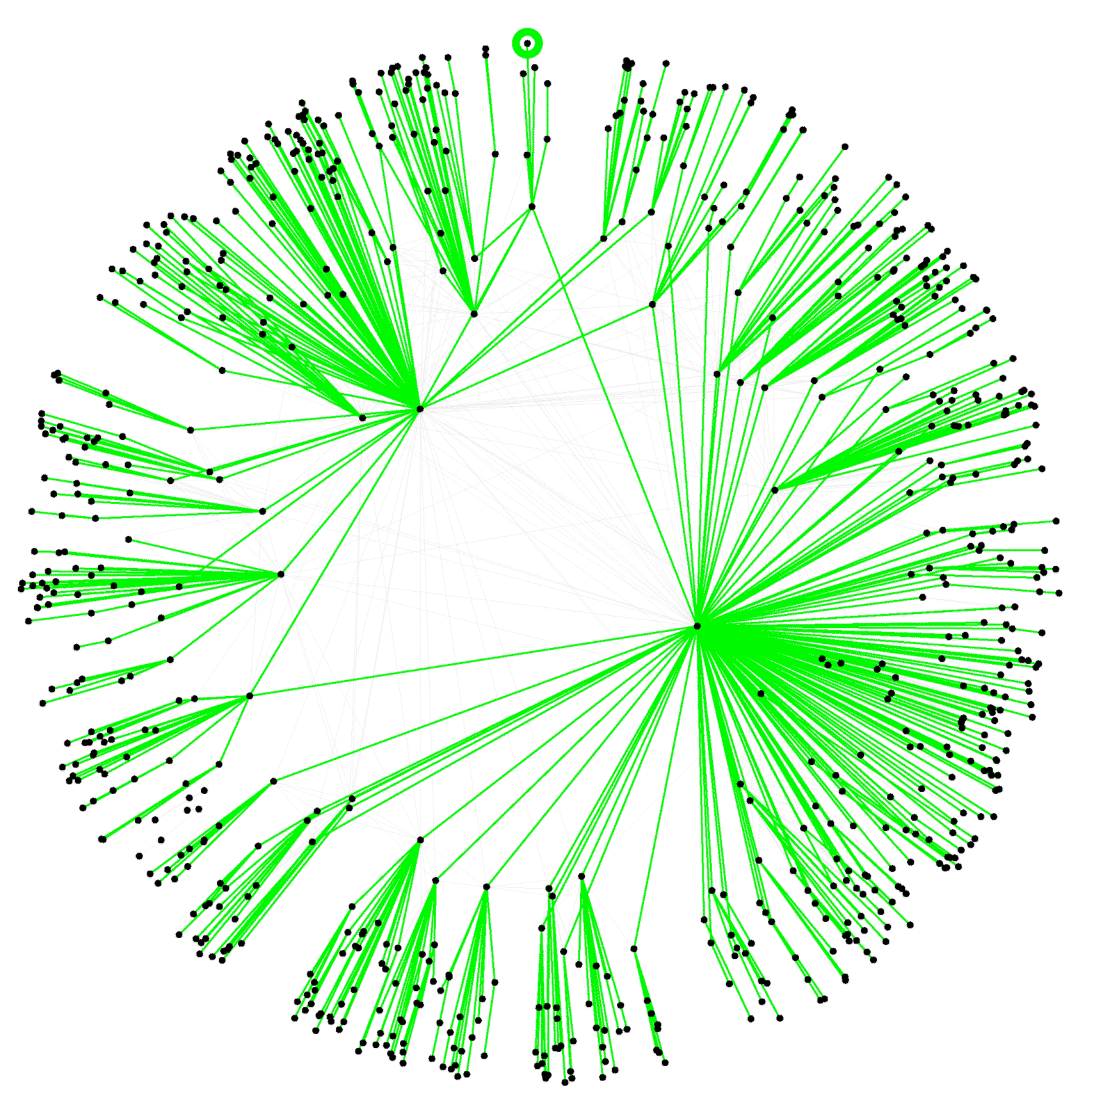

Unsuccessful greedy paths

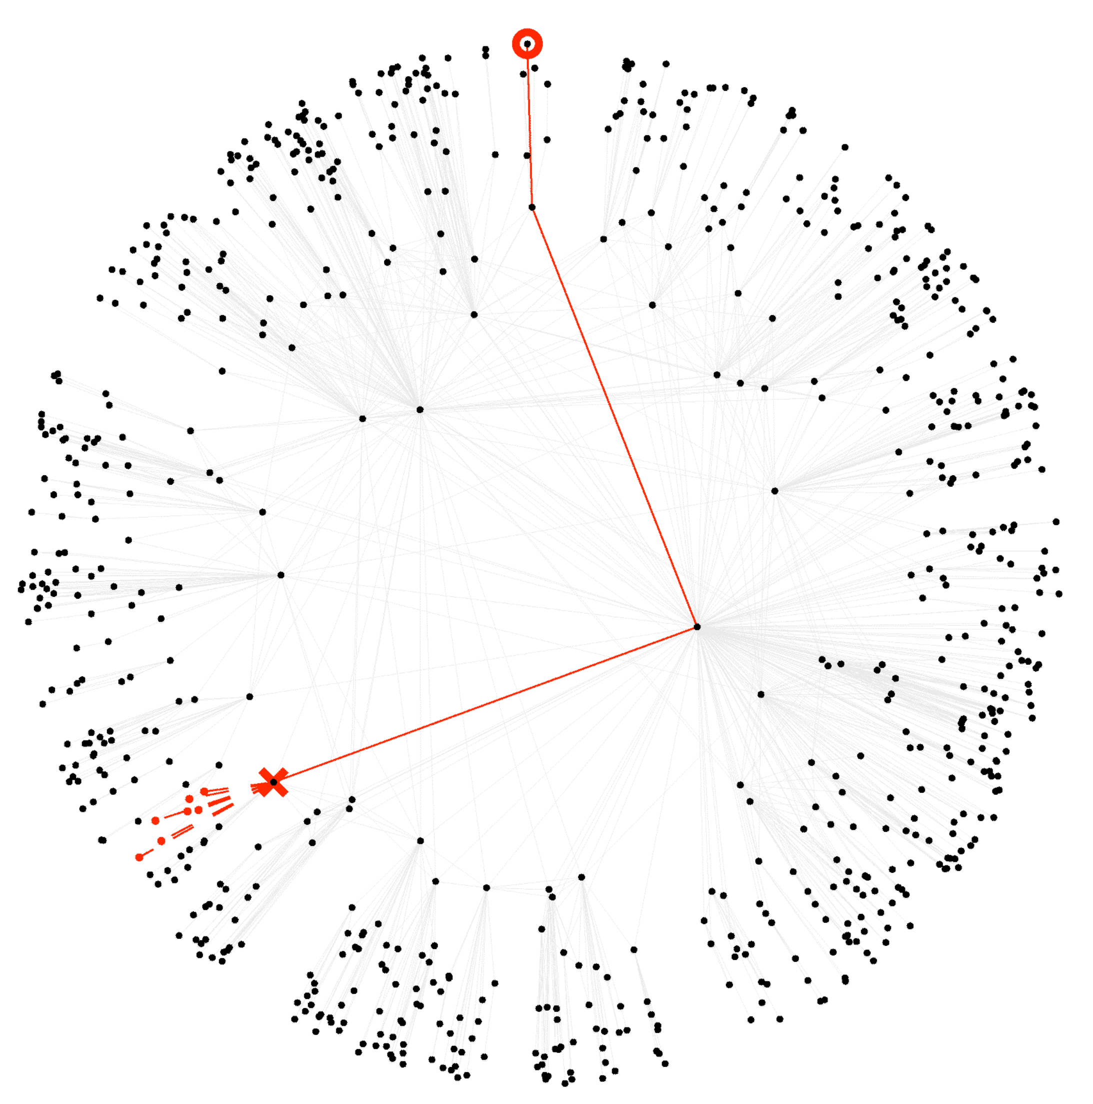

Percentage of successful paths

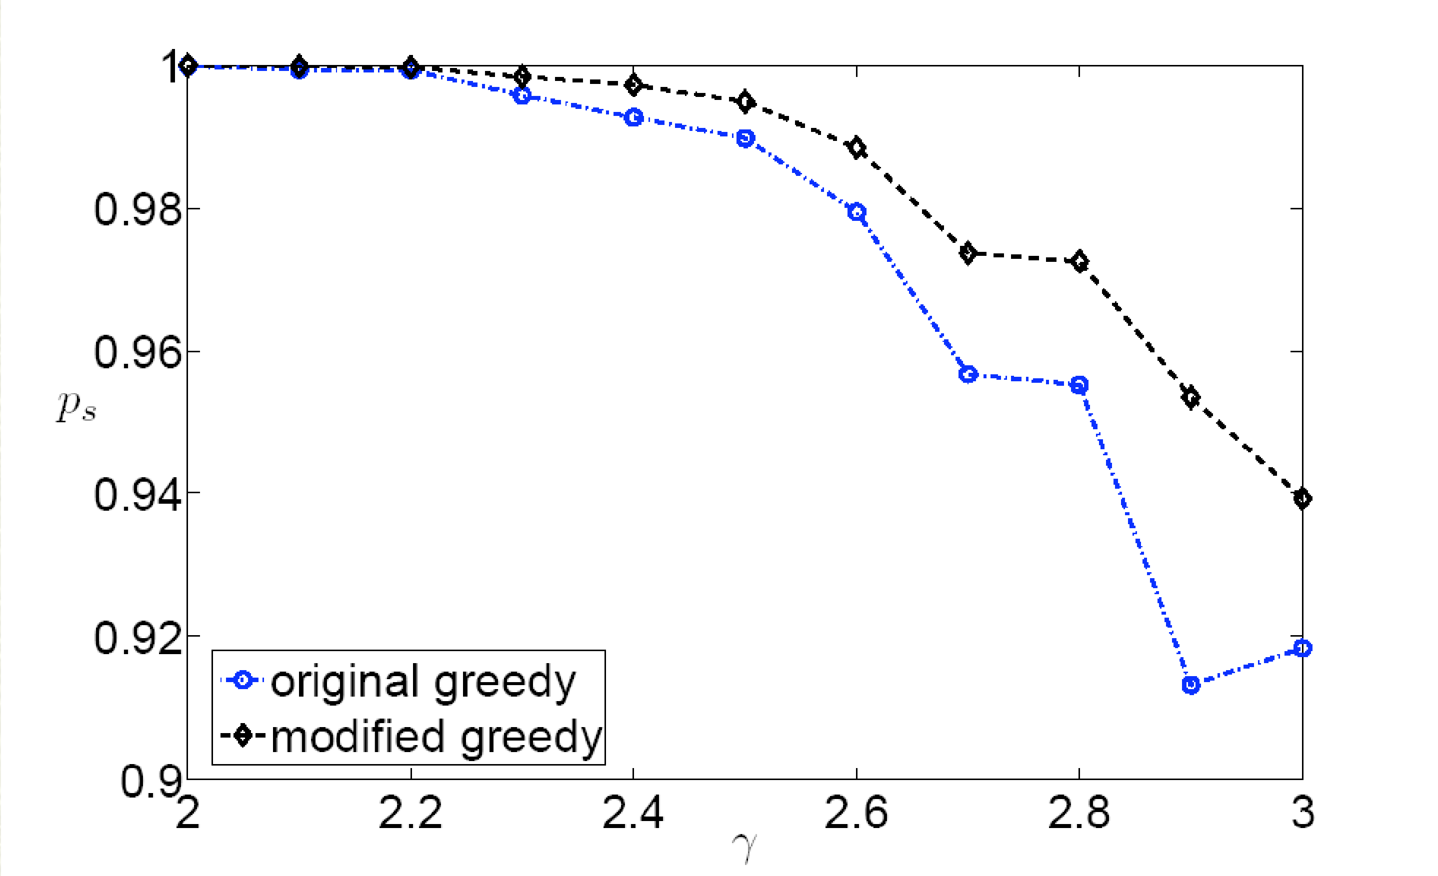

Average length of successful paths

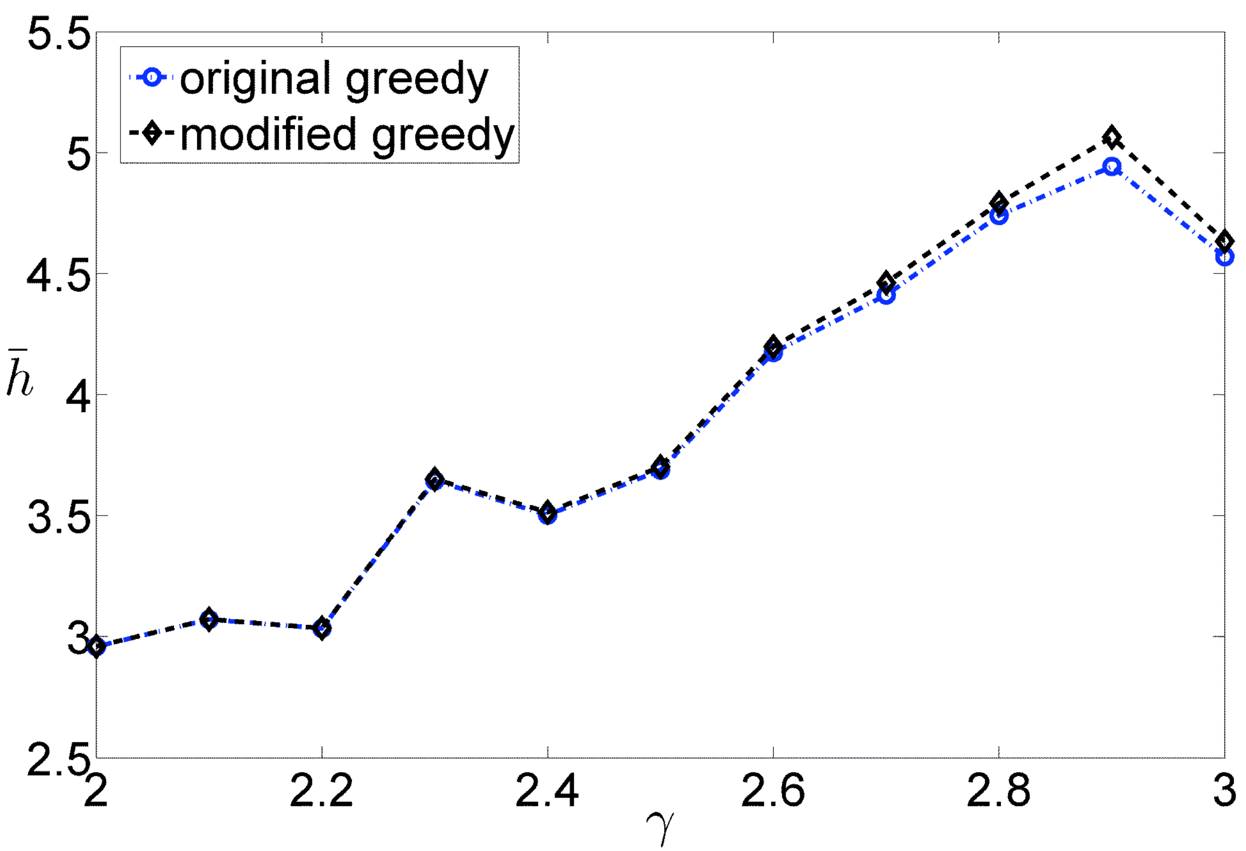

Multiplicative average stretch

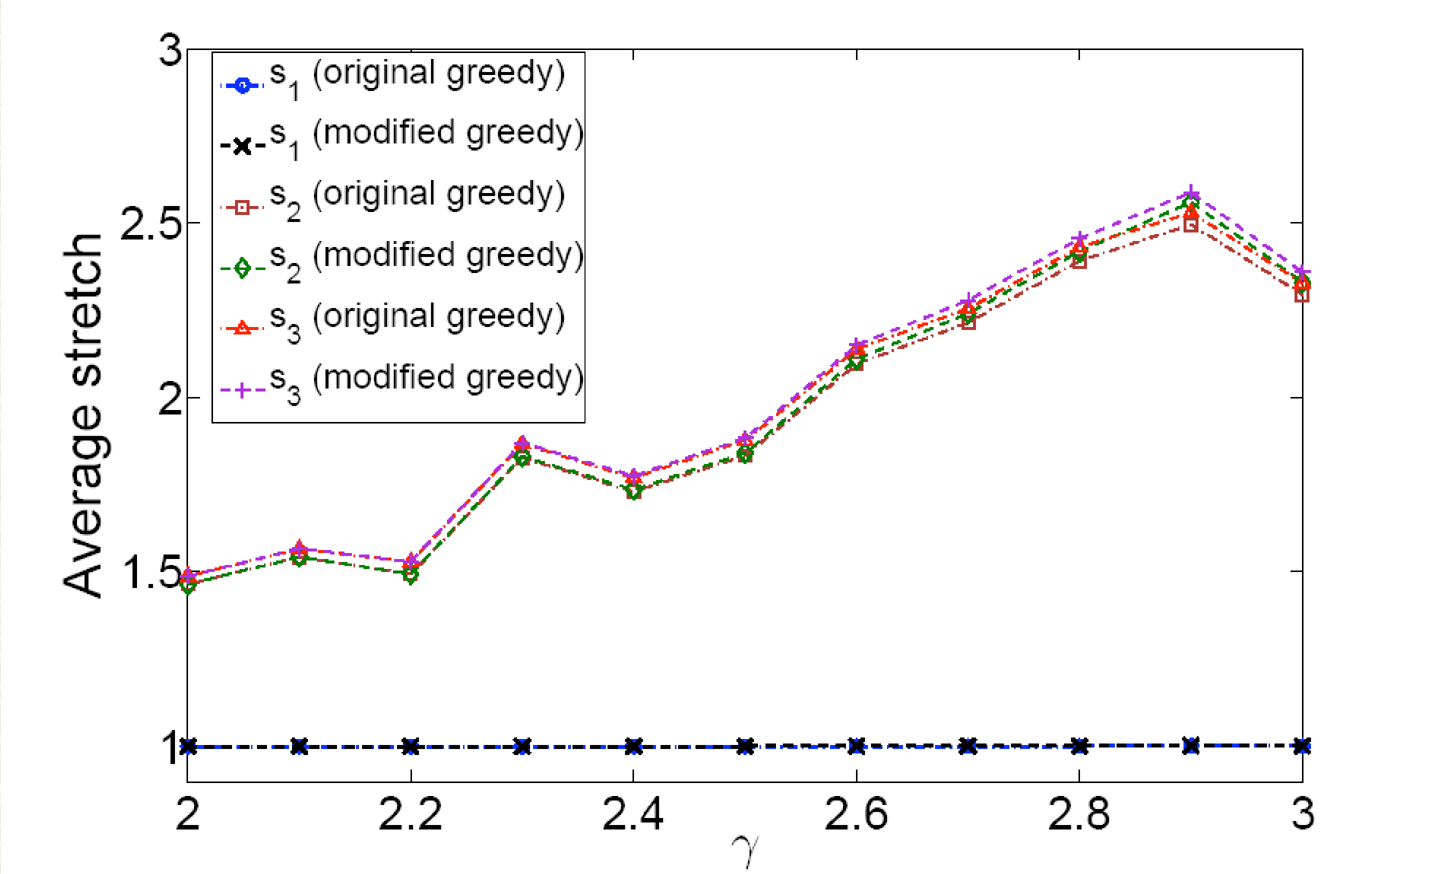

Multiplicative maximum stretch

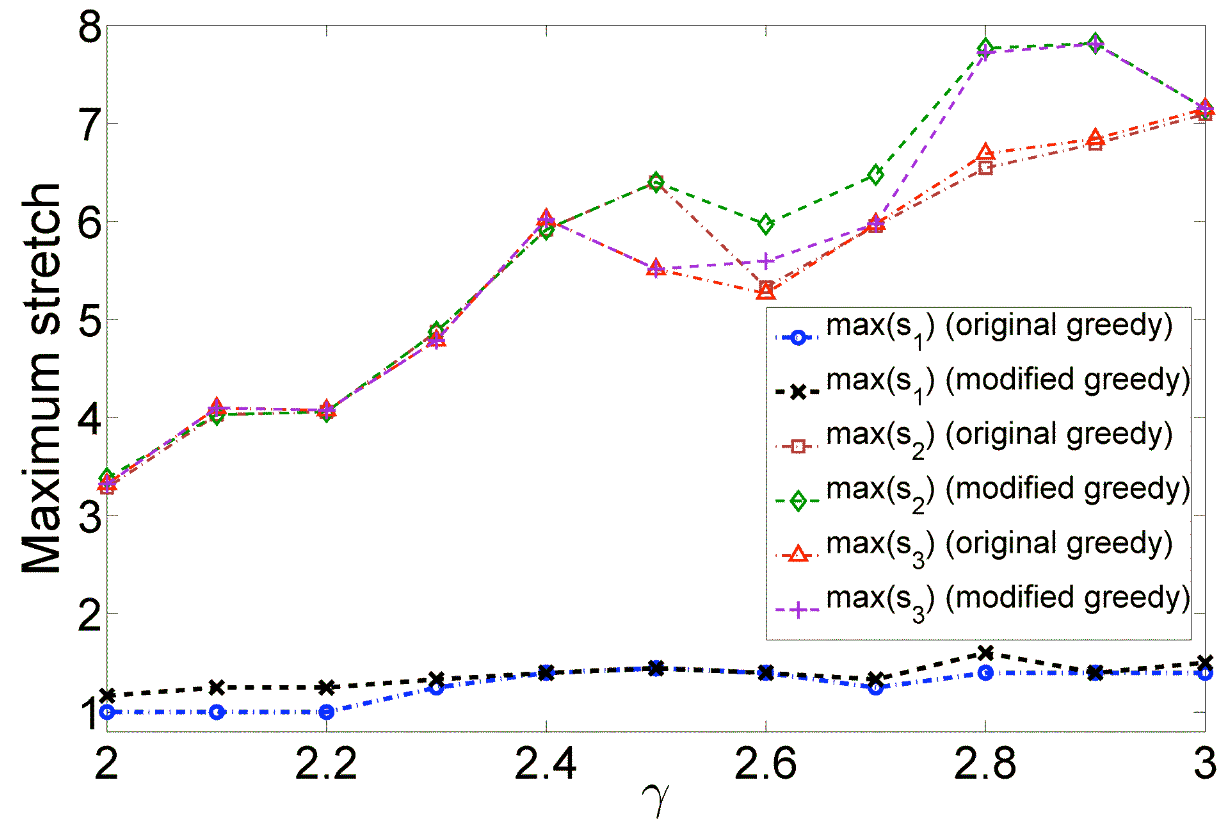

Robustness of greedy routing w.r.t. network dynamics

- As network topology changes, the greedy routing efficiency characteristics deteriorate very slowly
- For example, 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%*

In summary

- Scale-free networks are congruent w.r.t. hidden hyperbolic geometries
  - Greedy paths follow shortest paths that approximately follow shortest hidden paths, i.e., geodesics in the hyperbolic space
- This congruency is robust w.r.t. network dynamics/evolution
  - There are many shortest paths between the same source and destination that satisfy the above properties
  - If some of them go away, others remain available, and greedy routing still finds them

Conclusion

- Hidden hyperbolic metric spaces explain, simultaneously, the two main topological characteristics of complex networks
  - scale-free degree distributions
  - strong clustering
- Greedy routing mechanism in these settings may offer virtually infinitely scalable routing algorithms for future communication networks
  - Zero communication costs (no routing updates!)
  - Constant routing table sizes (coordinates in the space)
  - No stretch (all paths are shortest, stretch=*1*)

Problems to solve

- Find the exact structure of hidden metric spaces underlying real networks
- Find the coordinates of nodes in them