Overlay (and P2P) Networks

Part II

- Navigability of Complex Networks
- Mathematics and the Internet: A Source of Enormous Confusion and Great Potential
- Internet Indirection Architecture
- SIP / P2P-SIP

Samu Varjonen

Ashwin Rao

Navigability of Complex Networks

Boguna, Marian et al. **"Navigability of complex networks."** Nature Physics 5, no. 1 (2009): 74-80.

Clustering of Nodes

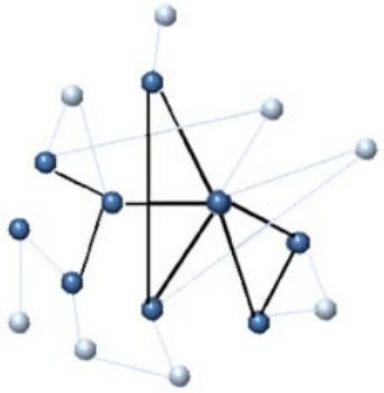

$k_T = 2$

k_T: extract subgraph $G(k_T)$ with nodes having degree $k > k_T$

Serrano, M. Angeles, Dmitri Krioukov, and Marián Boguná. **"Self-similarity of complex networks and hidden metric spaces."** Physical review letters 100, no. 7 (2008): 078701.

Clustering of Nodes

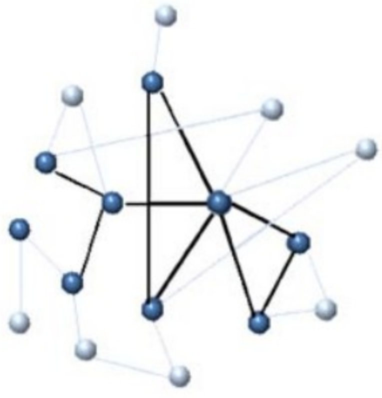

$k_T = 2$

k_T: extract subgraph $G(k_T)$ with nodes having degree $k > k_T$

$< k_i(k_T) >$: average degree of this subgraph $G(k_T)$

$\bar{c}(k_T)$: clustering coefficient of subgraph $G(k_T)$

Serrano, M. Angeles, Dmitri Krioukov, and Marián Boguná. **"Self-similarity of complex networks and hidden metric spaces."** Physical review letters 100, no. 7 (2008): 078701.

Clustering of Nodes

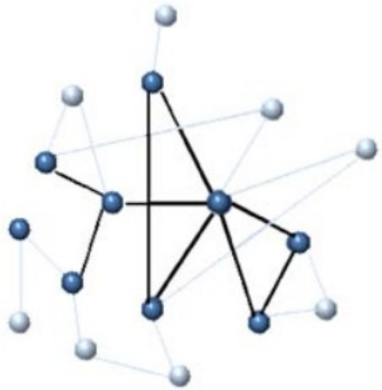

$k_T = 2$

k_T: extract subgraph $G(k_T)$ with nodes having degree $k > k_T$

$< k_i(k_T) >$: average degree of this subgraph $G(k_T)$

$\bar{c}(k_T)$: clustering coefficient of subgraph $G(k_T)$

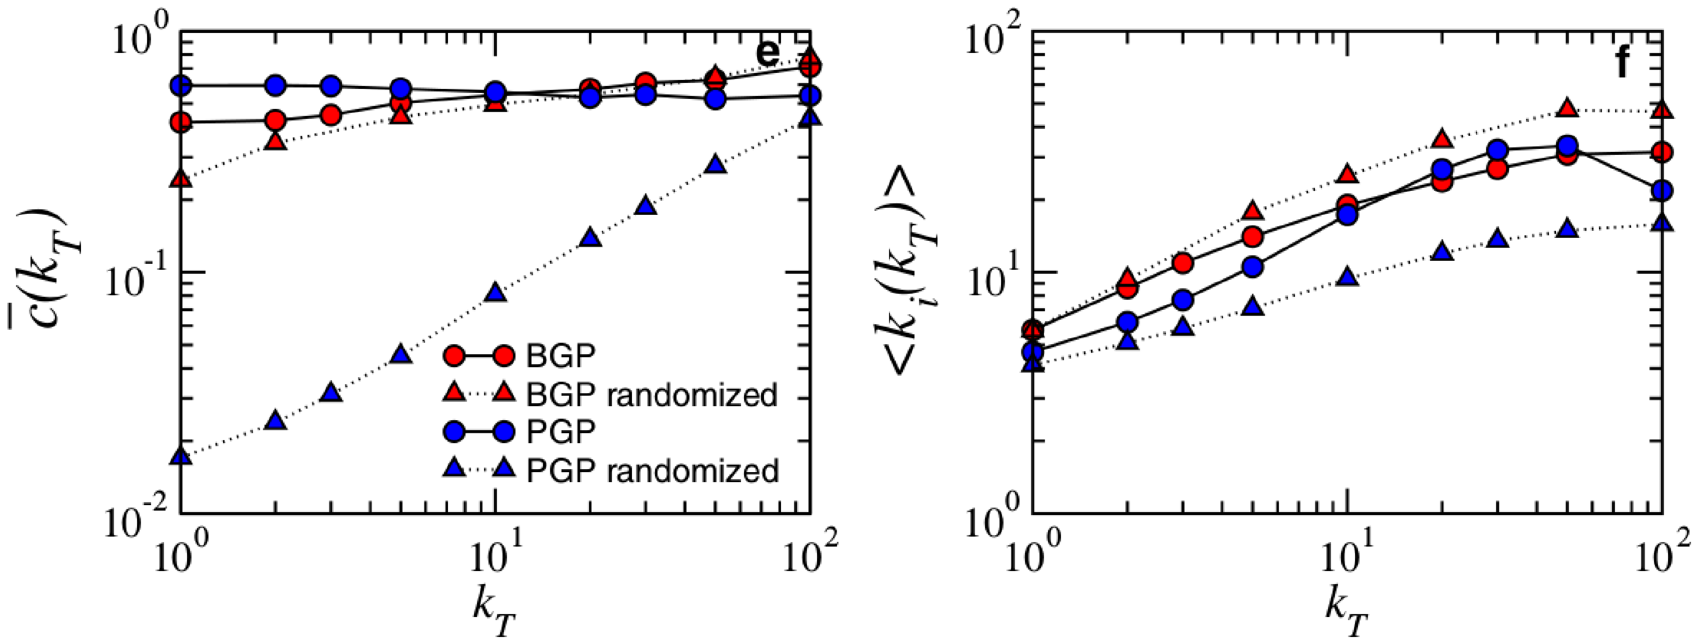

Serrano, M. Angeles, Dmitri Krioukov, and Marián Boguná. **"Self-similarity of complex networks and hidden metric spaces."** Physical review letters 100, no. 7 (2008): 078701.

Clustering of Nodes

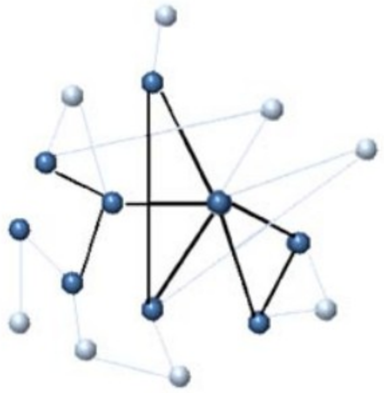

k_T: extract subgraph $G(k_T)$ with nodes having degree $k > k_T$

$< k_i(k_T) >$: average degree of this subgraph $G(k_T)$

$\bar{c}(k_T)$: clustering coefficient of subgraph $G(k_T)$

$k_T = 2$

Constant clustering coeff.

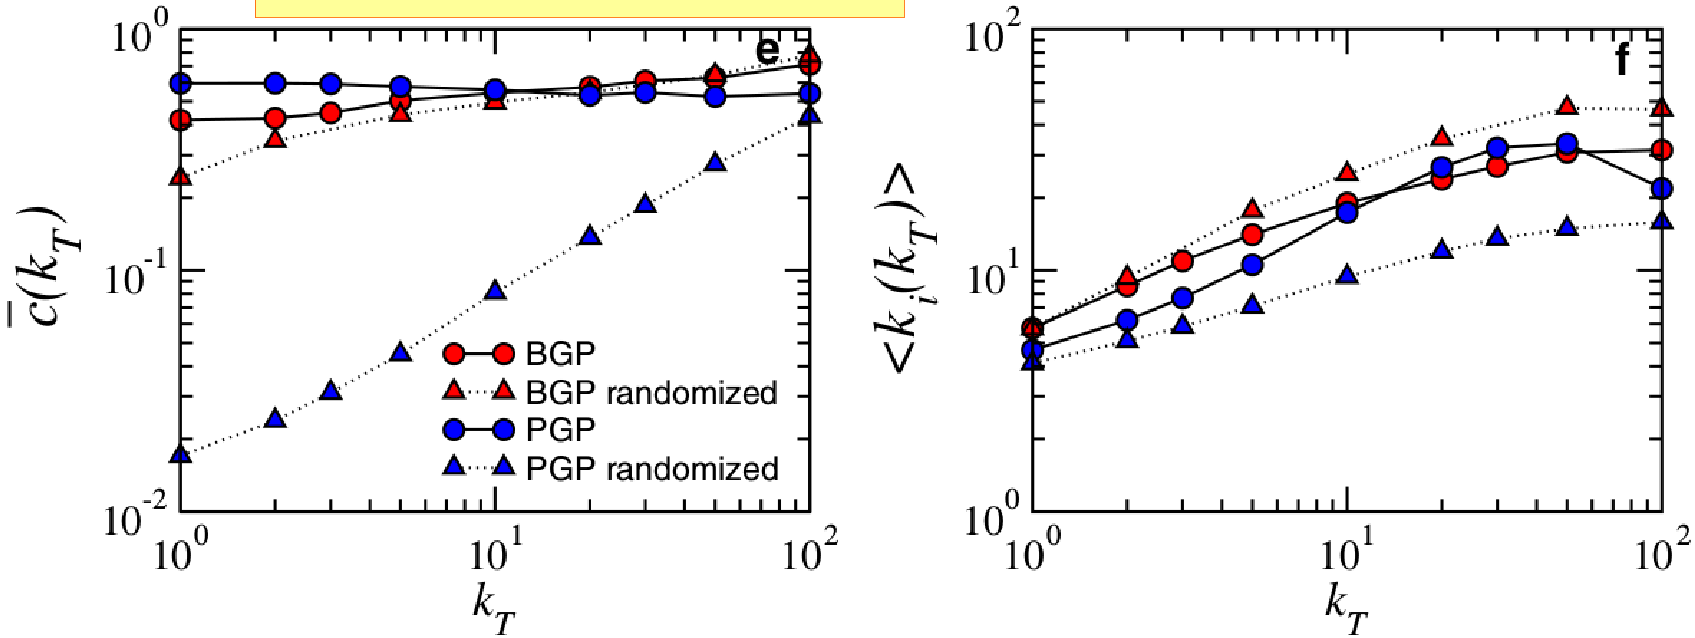

Serrano, M. Angeles, Dmitri Krioukov, and Marián Boguná. **"Self-similarity of complex networks and hidden metric spaces."** Physical review letters 100, no. 7 (2008): 078701.

Clustering of Nodes

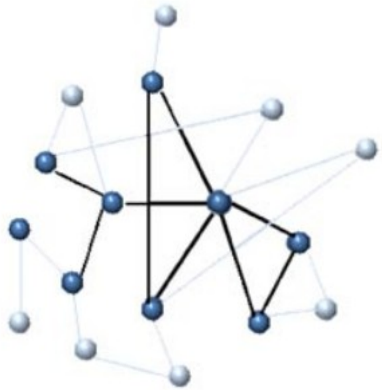

k_T: extract subgraph $G(k_T)$ with nodes having degree $k > k_T$

$< k_i(k_T) >$: average degree of this subgraph $G(k_T)$

$\bar{c}(k_T)$: clustering coefficient of subgraph $G(k_T)$

$k_T = 2$

Constant clustering coeff.

Average Degree Increases

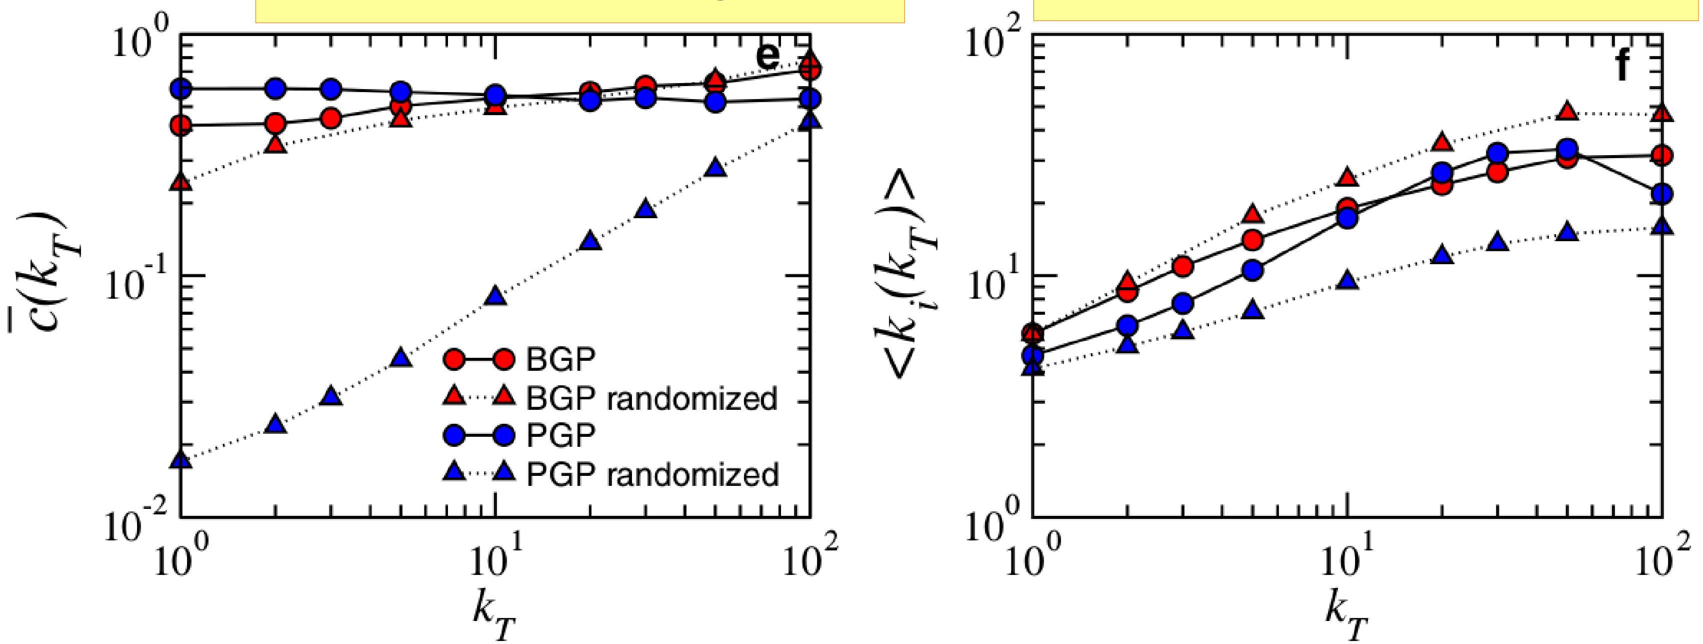

Serrano, M. Angeles, Dmitri Krioukov, and Marián Boguná. **"Self-similarity of complex networks and hidden metric spaces."** Physical review letters 100, no. 7 (2008): 078701.

Clustering of Nodes

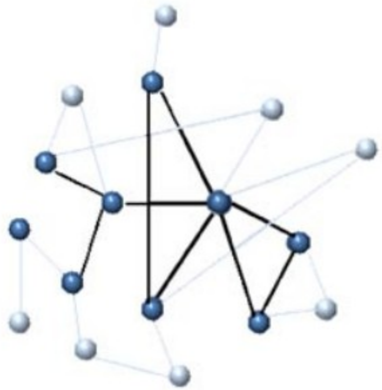

k_T: extract subgraph $G(k_T)$ with nodes having degree $k > k_T$

$< k_i(k_T) >$: average degree of this subgraph $G(k_T)$

$\bar{c}(k_T)$: clustering coefficient of subgraph $G(k_T)$

$k_T = 2$

Constant clustering coeff.

Average Degree Increases

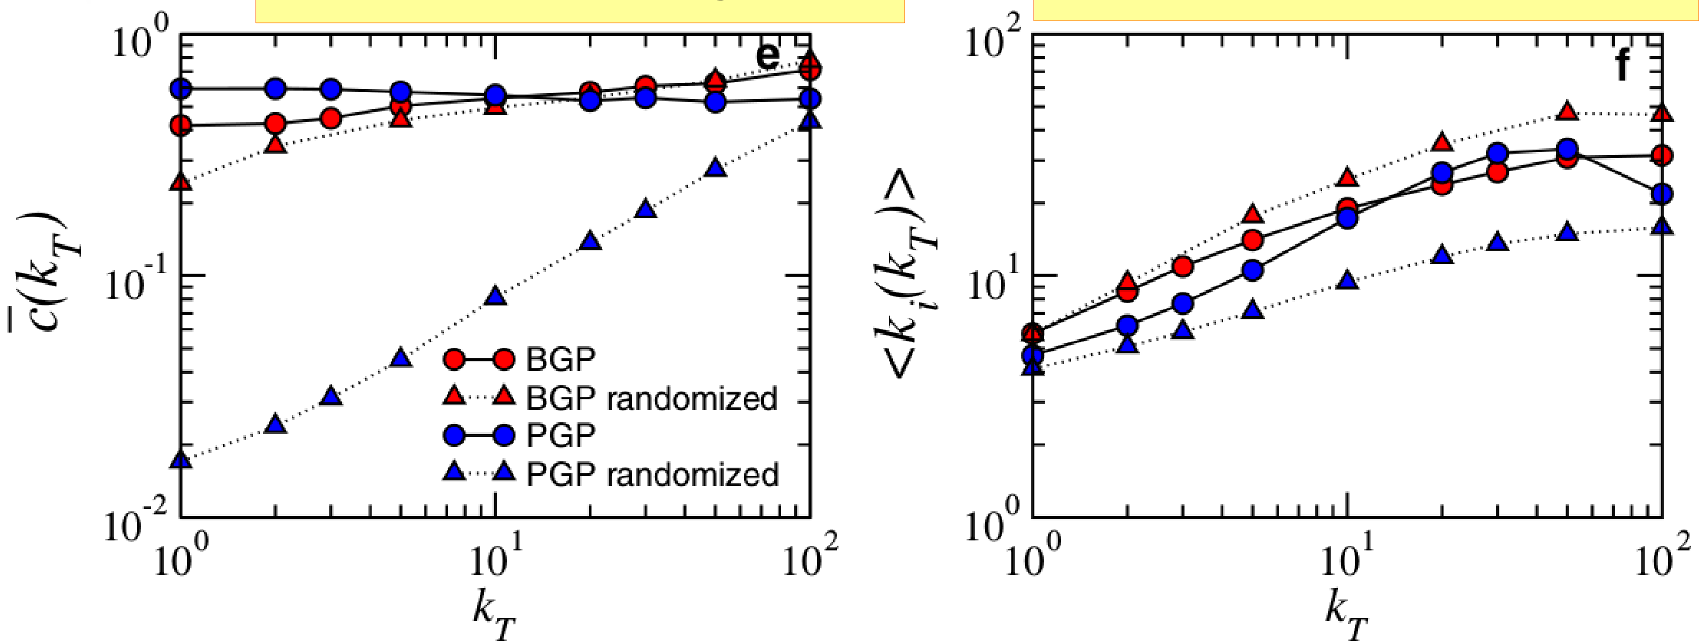

How to Generate Scale-Free Graphs with Strong Clustering

Serrano, M. Angeles, Dmitri Krioukov, and Marián Boguná. **"Self-similarity of complex networks and hidden metric spaces."** Physical review letters 100, no. 7 (2008): 078701.

Generating Scale-Free Graphs with Strong Clustering

Generating Scale-Free Graphs with Strong Clustering

Take all nodes and distribute them within an underlying circle

Generating Scale-Free Graphs with Strong Clustering

Take all nodes and distribute them within an underlying circle

Assign each node an expected degree k where P(k) $\sim k^{-\gamma}$

Generating Scale-Free Graphs with Strong Clustering

Take all nodes and distribute them within an underlying circle

Assign each node an expected degree k where P(k) $\sim k^{-\gamma}$

Connect each pair of nodes with a connection probability r(d; k,k$'$)

d is the distance between these two nodes in the circle

$\mathrm{d}_c = kk'$ is also called the characteristic distance

Generating Scale-Free Graphs with Strong Clustering

Take all nodes and distribute them within an underlying circle

Assign each node an expected degree k where $\mathrm{P(k)} \sim k^{-\gamma}$

Connect each pair of nodes with a connection probability $\mathrm{r(d;\ k,k')}$

d is the distance between these two nodes in the circle

$\mathrm{d}_c = kk'$ is also called the characteristic distance

$$\mathrm{r\ (d;\ k,\ k')} = \left(1 + \frac{d}{d_c}\right)^{-\alpha}$$

Generating Scale-Free Graphs with Strong Clustering

Take all nodes and distribute them wit

Assign each node an expected degree k

Connect each pair of nodes with a conn

d is the distance between these two no

$d_c = kk'$ is also called the characteristi

$$r\,(d;\, k,\, k') = \left(1 + \frac{d}{d_c}\right)^{-\alpha}$$

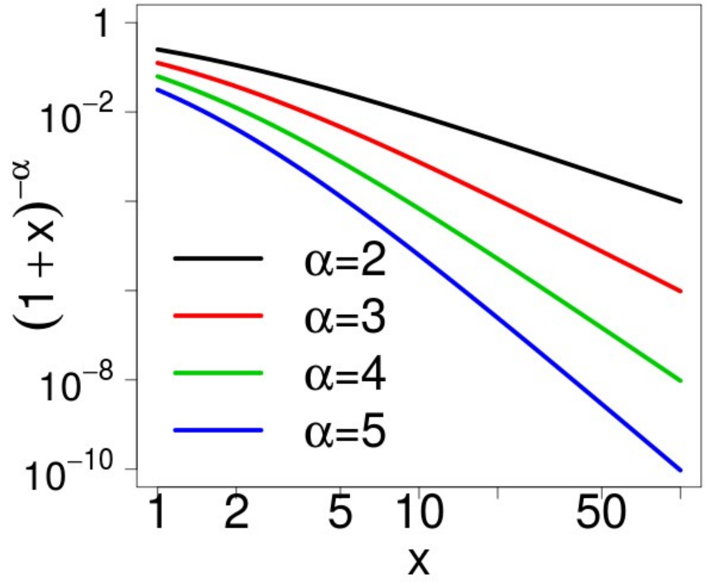

Hubs will be connected with a high probability because of large d_c

Generating Scale-Free Graphs with Strong Clustering

Take all nodes and distribute them wit

Assign each node an expected degree k

Connect each pair of nodes with a conn

d is the distance between these two noc

$d_c = kk'$ is also called the characteristi

$$r\,(d;\, k,\, k') = \left(1 + \frac{d}{d_c}\right)^{-\alpha}$$

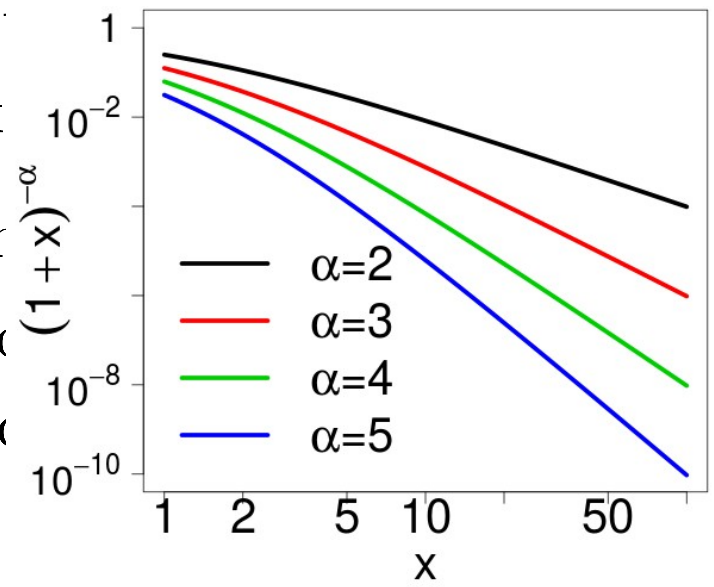

Hubs will be connected with a high probability because of large d_c

Low degree nodes connected only if (hidden distance) d is small

Generating Scale-Free Graphs with Strong Clustering

Take all nodes and distribute them wit

Assign each node an expected degree k

Connect each pair of nodes with a conn

d is the distance between these two noc

$d_c = kk'$ is also called the characteristic

$$r\left(d; k, k'\right) = \left(1 + \frac{d}{d_c}\right)^{-\alpha}$$

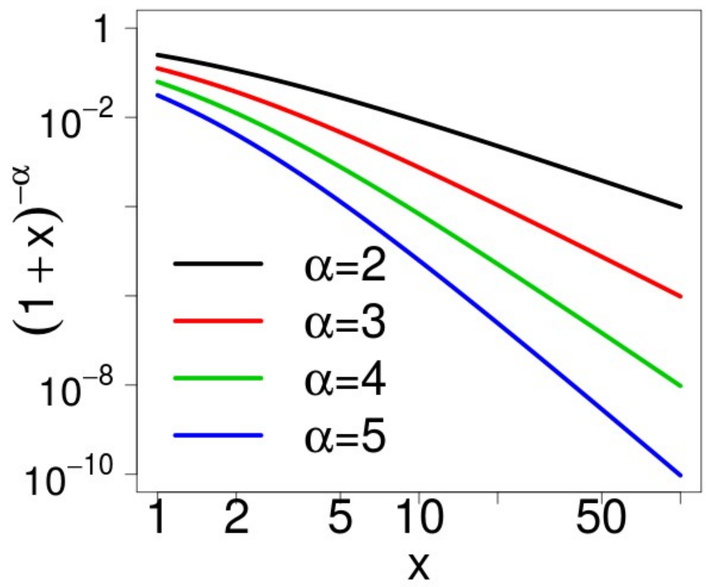

Hubs will be connected with a high probability because of large d_c

Low degree nodes connected only if (hidden distance) d is small

Hubs connected to low degree nodes at moderate hidden distance

Generating Scale-Free Graphs with Strong Clustering

Take all nodes and distribute them wit

Assign each node an expected degree k

Connect each pair of nodes with a conn

d is the distance between these two no

$d_c = kk'$ is also called the characteristi

$$r\,(d;\,k,\,k') = \left(1 + \frac{d}{d_c}\right)^{-\alpha}$$

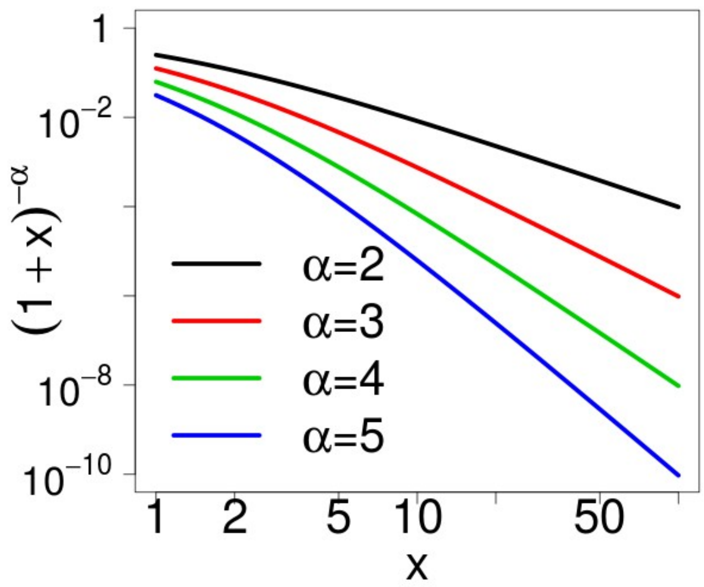

Hubs will be connected with a high probability because of large d_c

Low degree nodes connected only if (hidden distance) d is small

Hubs connected to low degree nodes at moderate hidden distance

α importance of hidden distance

Generating Scale-Free Graphs with Strong Clustering

Take all nodes and distribute them within an underlying circle

Assign each node an expected degree k where P(k) $\sim k^{-\gamma}$

Connect each pair of nodes with a connection probability r(d; k,k′)

d is the distance between these two nodes in the circle

$\mathrm{d}_c = kk'$ is also called the characteristic distance

$$\mathrm{r}\ (\mathrm{d};\ \mathrm{k},\ \mathrm{k}') = \left(1 + \frac{d}{d_c}\right)^{-\alpha}$$

Hubs will be connected with a high probability because of large d_c

Low degree nodes connected only if (hidden distance) d is small

Hubs connected to low degree nodes at moderate hidden distance

α importance of hidden distance

Path Length
(Greedy Routing)

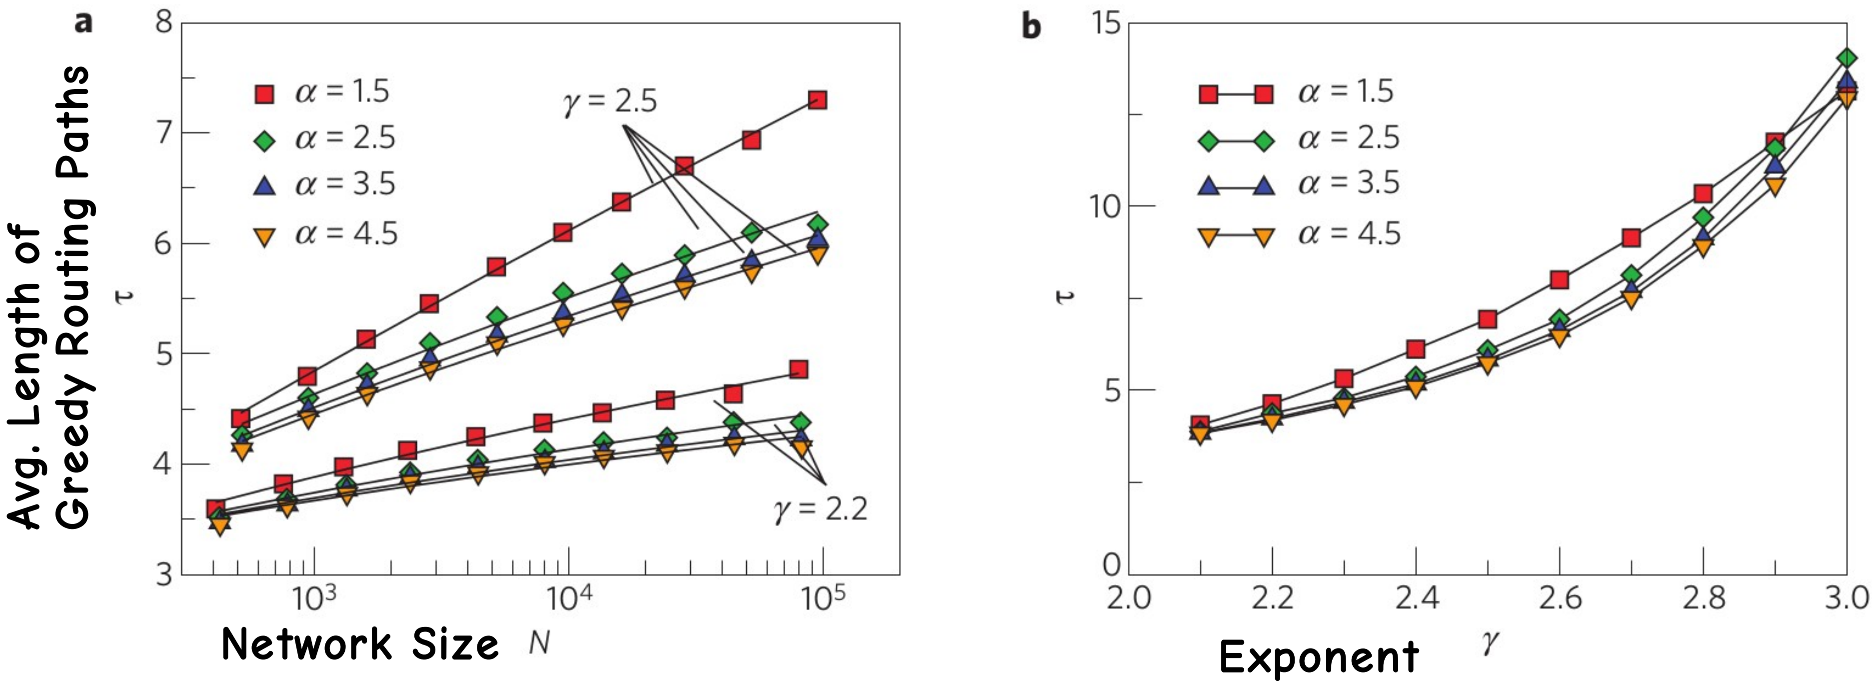

Boguna, Marian et al. **"Navigability of complex networks."** Nature Physics 5, no. 1 (2009): 74-80.

Path Length
(Greedy Routing)

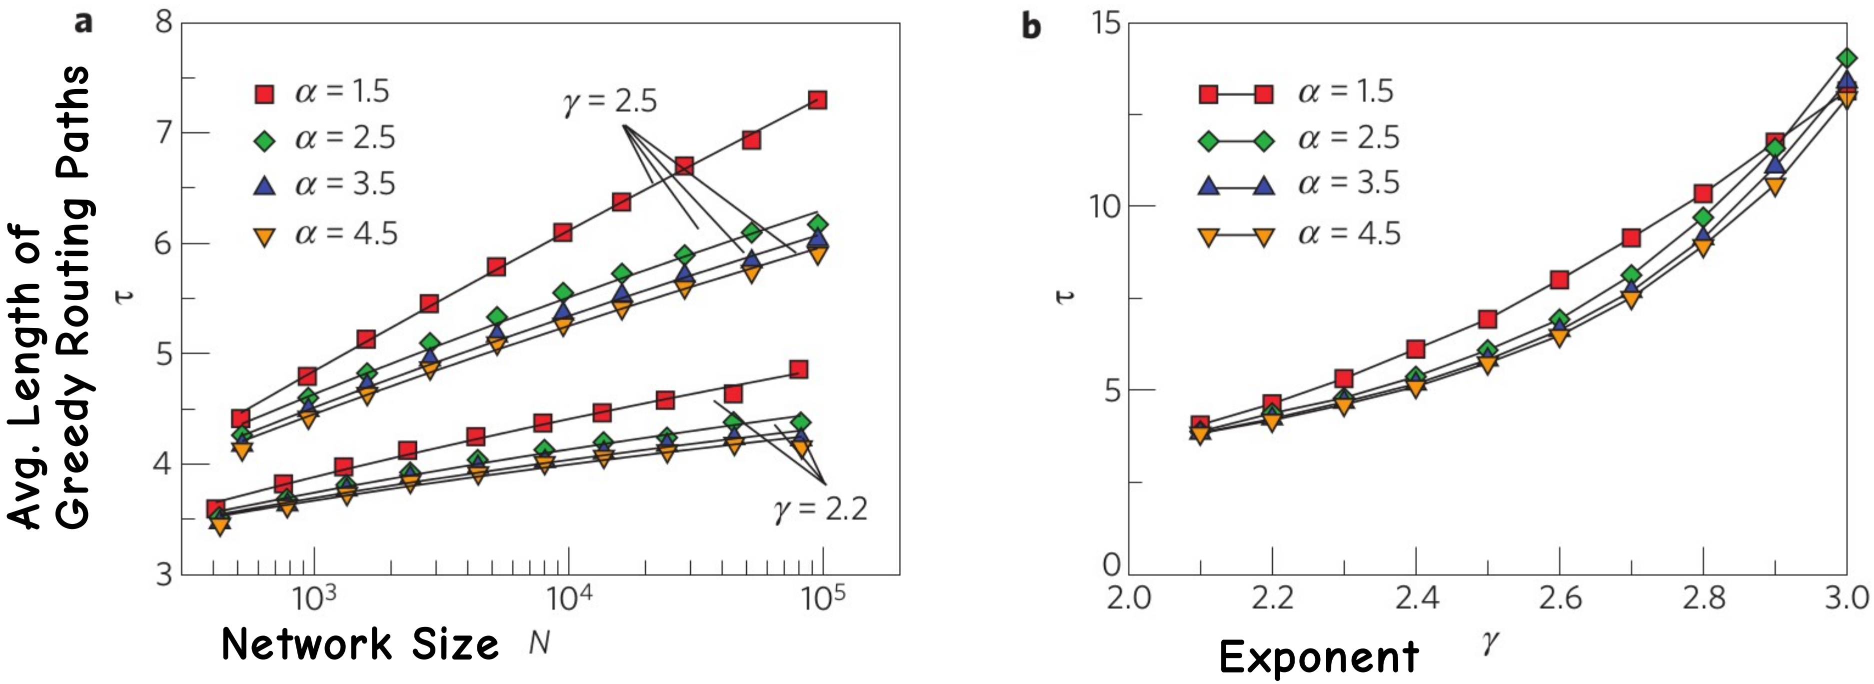

Path length grows polylogarithmically with the network size

Boguna, Marian et al. **"Navigability of complex networks."** Nature Physics 5, no. 1 (2009): 74-80.

Path Length
(Greedy Routing)

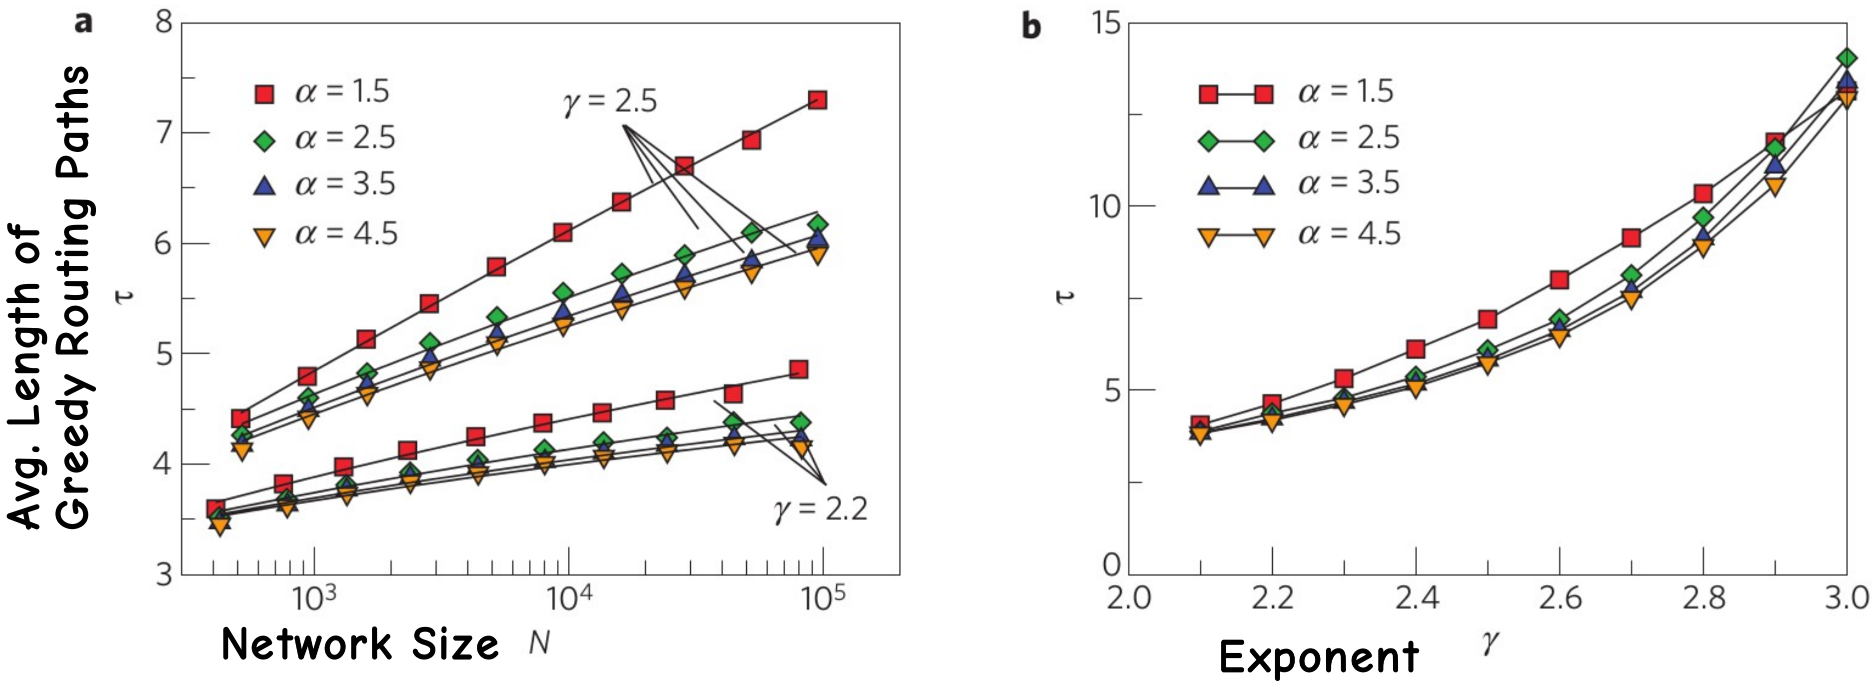

Path length grows polylogarithmically with the network size

Paths shorter for smaller exponents and stronger clustering

Boguna, Marian et al. **"Navigability of complex networks."** Nature Physics 5, no. 1 (2009): 74-80.

Greedy Routing

Greedy Routing

- Hidden Space as the coordinate space
  - ***Hidden space is circle in this example***

Greedy Routing

- Hidden Space as the coordinate space
  - ***Hidden space is circle in this example***
- Greedy Routing: Send to neighbor who is closer to the destination (in hidden space)

Greedy Routing

- Hidden Space as the coordinate space
  - ***Hidden space is circle in this example***
- Greedy Routing: Send to neighbor who is closer to the destination (in hidden space)
- Unsuccessful Paths: None of your neighbors are closer to the destination in the hidden space

Success Probability
(Greedy Routing)

Boguna, Marian et al.
"Navigability of complex networks."
Nature Physics 5, no. 1 (2009): 74-80.

Success Probability
(Greedy Routing)

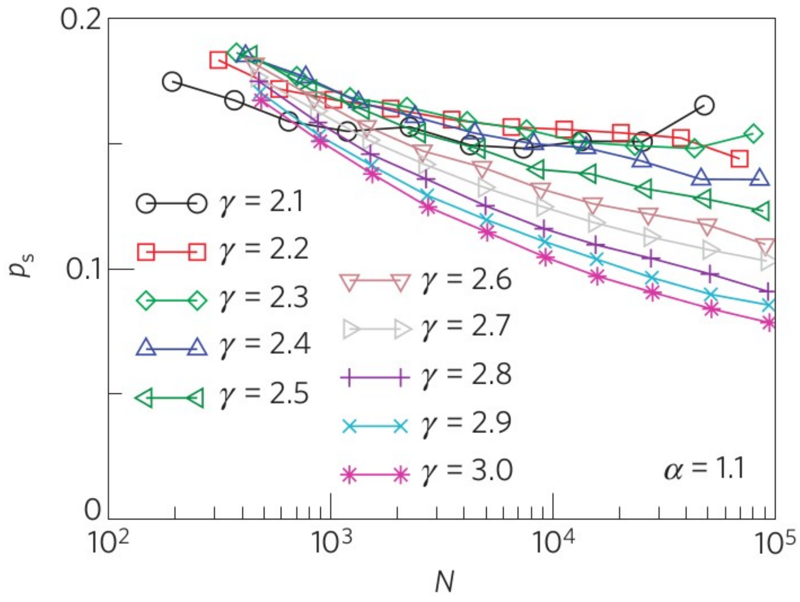

Boguna, Marian et al.
"Navigability of complex networks."
Nature Physics 5, no. 1 (2009): 74-80.

Success Probability
(Greedy Routing)

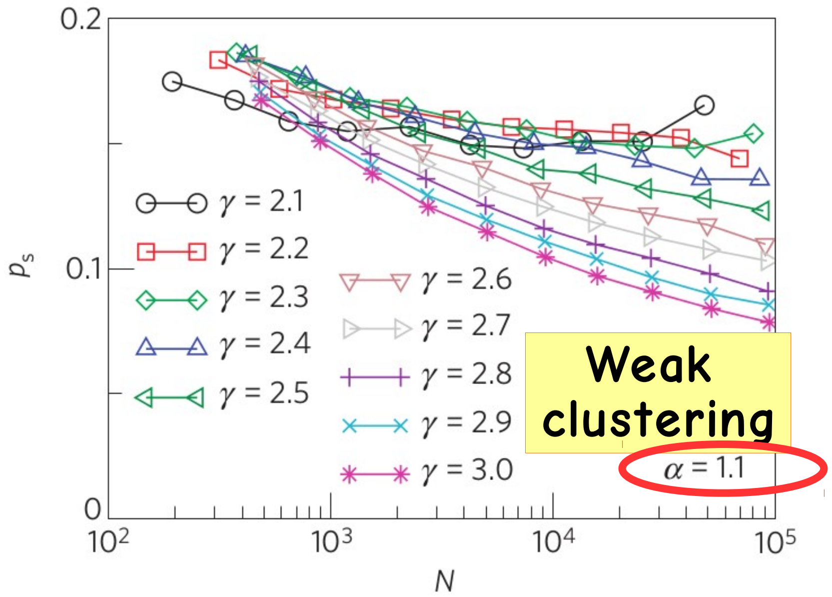

Boguna, Marian et al.
"Navigability of complex networks."
Nature Physics 5, no. 1 (2009): 74-80.

Success Probability
(Greedy Routing)

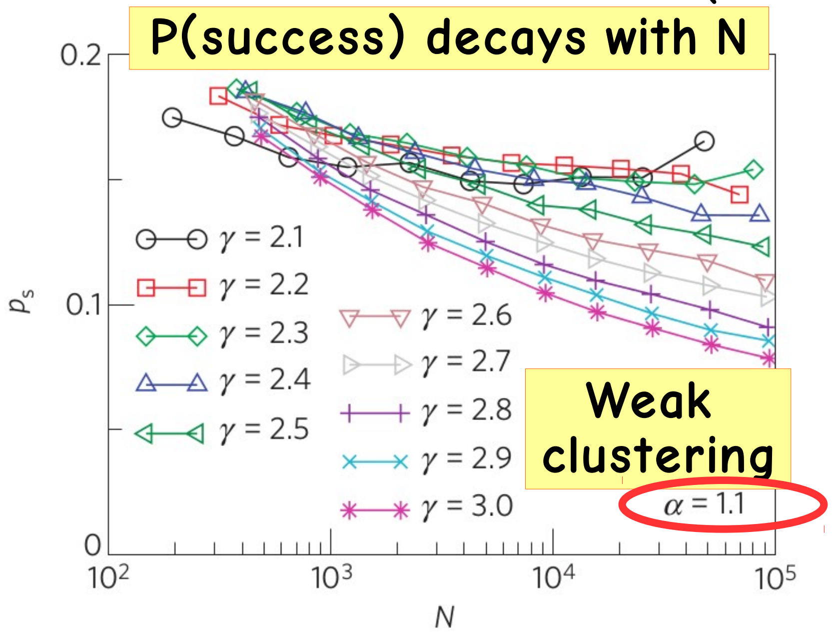

Boguna, Marian et al.
"Navigability of complex networks."
Nature Physics 5, no. 1 (2009): 74-80.

Success Probability
(Greedy Routing)

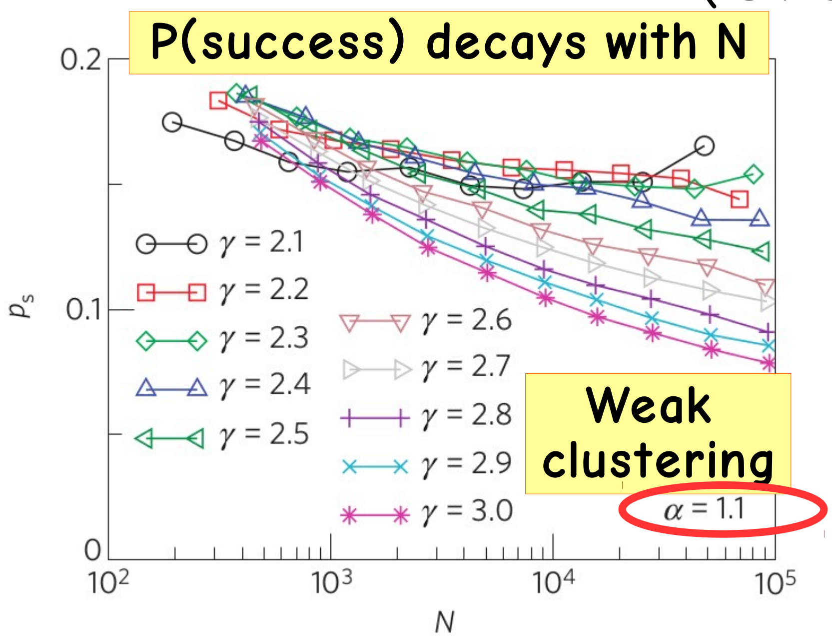

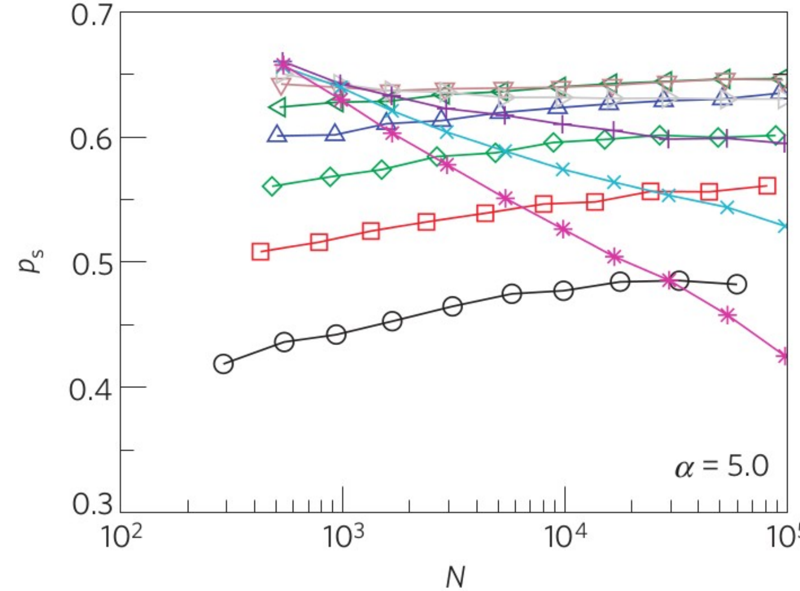

Boguna, Marian et al.
"Navigability of complex networks."
Nature Physics 5, no. 1 (2009): 74-80.

Success Probability
(Greedy Routing)

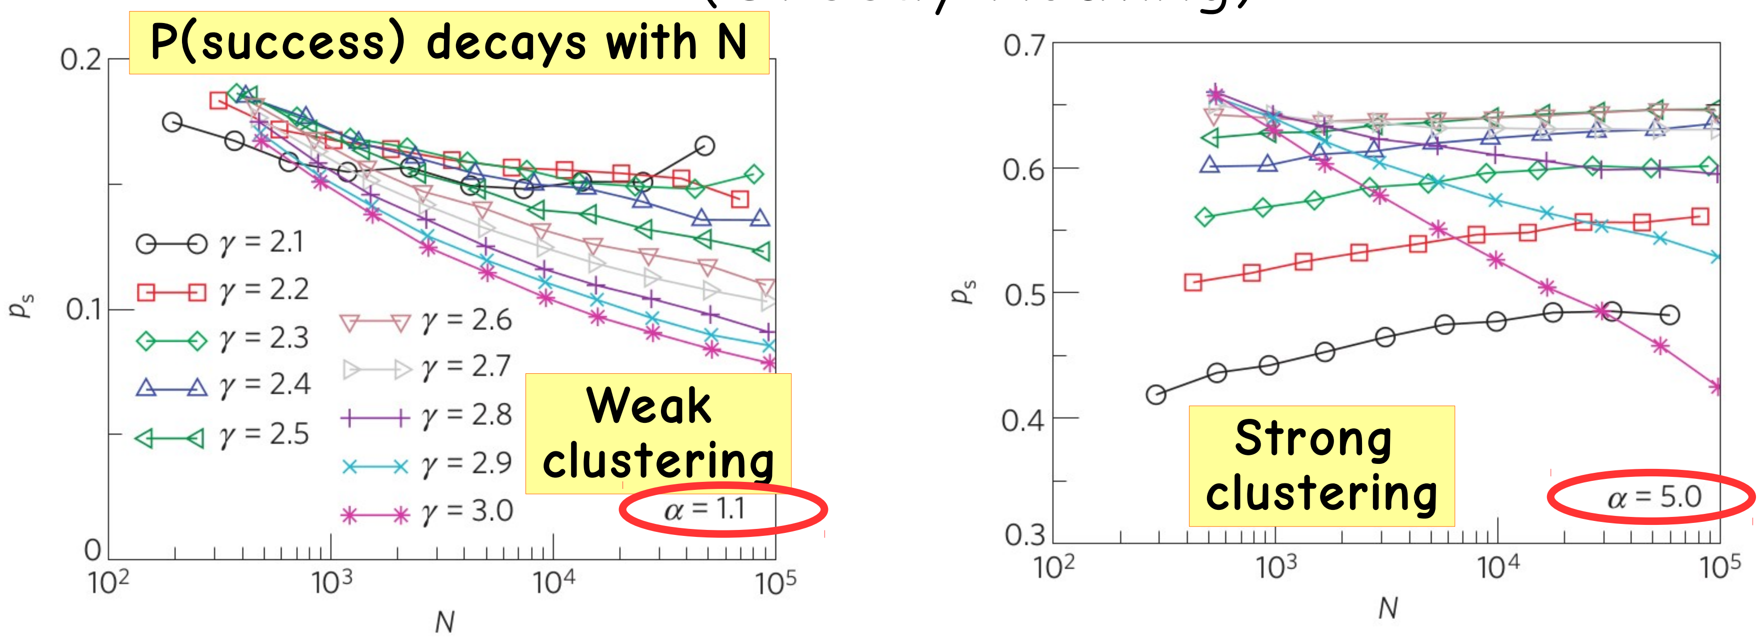

Boguna, Marian et al.
"Navigability of complex networks."
Nature Physics 5, no. 1 (2009): 74-80.

Success Probability
(Greedy Routing)

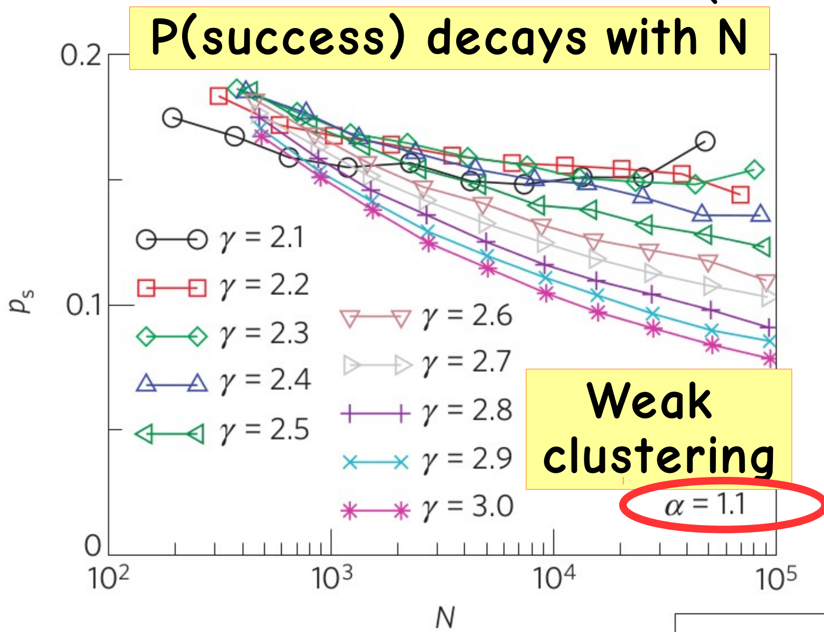

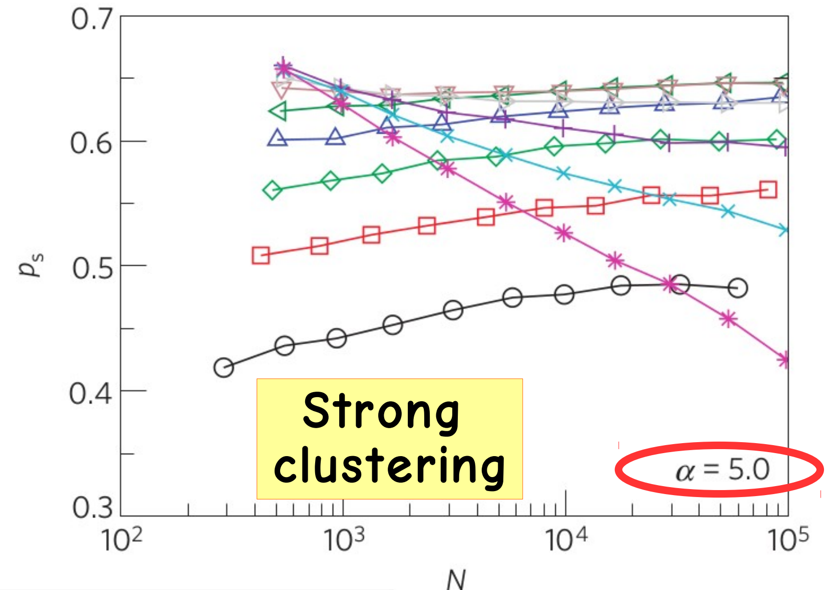

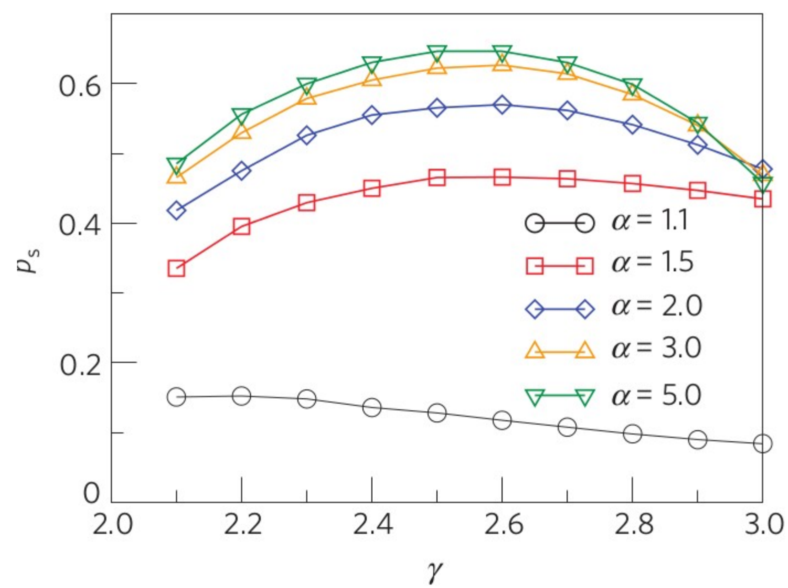

Boguna, Marian et al.
"Navigability of complex networks."
Nature Physics 5, no. 1 (2009): 74-80.

Success Probability
(Greedy Routing)

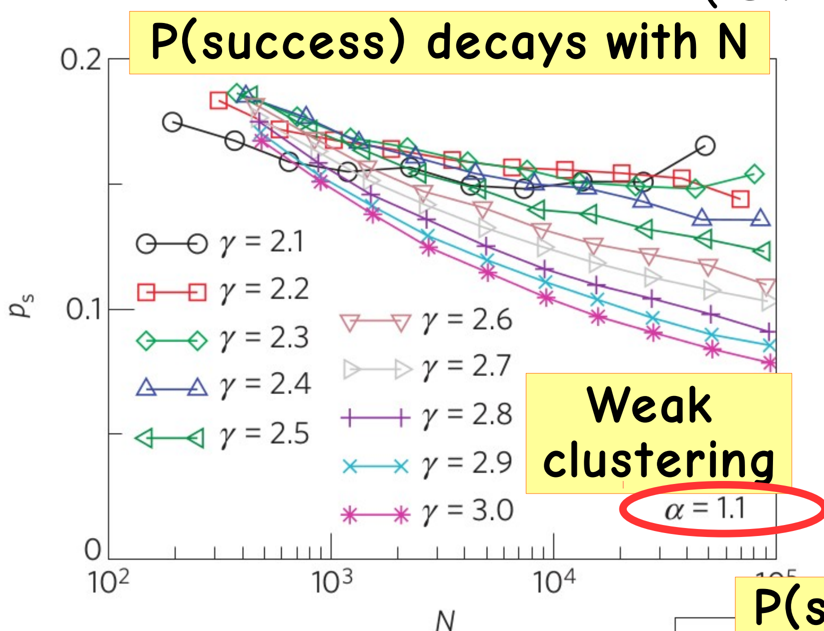

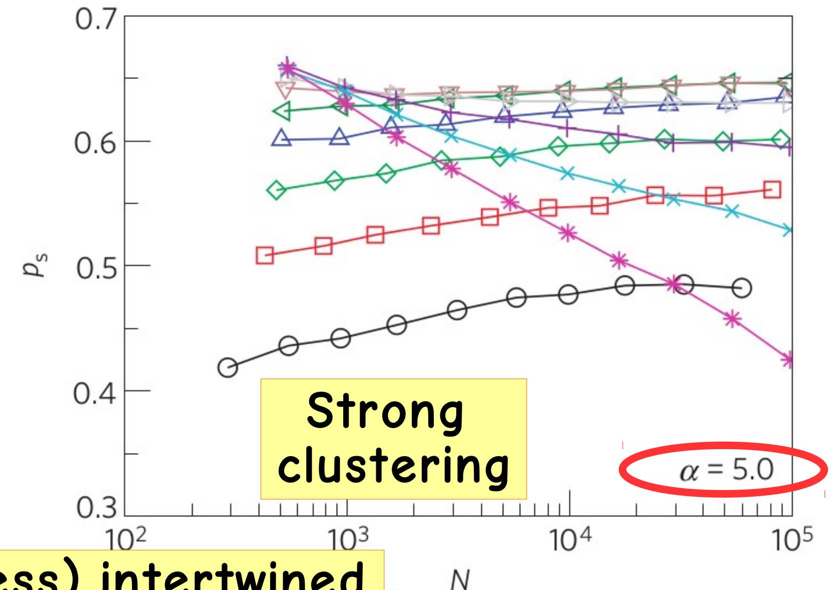

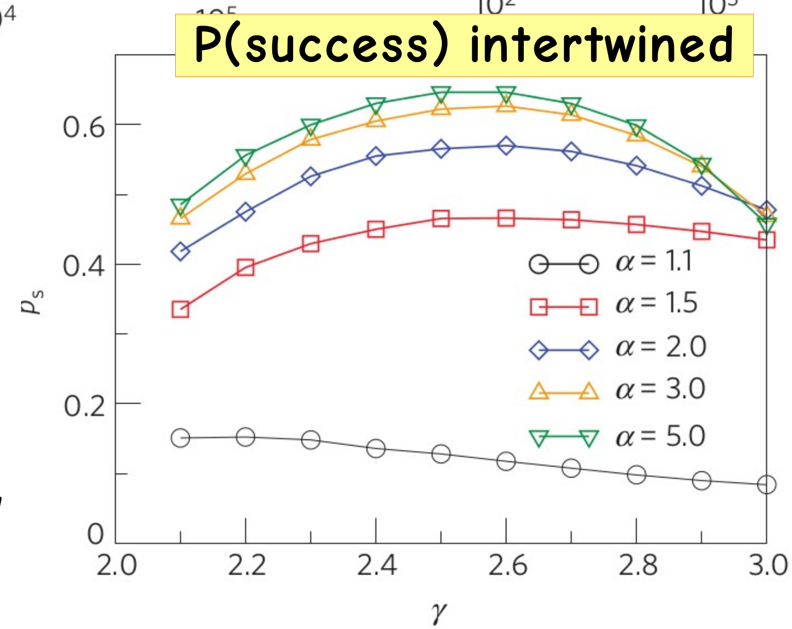

Boguna, Marian et al.
"Navigability of complex networks."
Nature Physics 5, no. 1 (2009): 74-80.

Navigation in Scale Free Networks

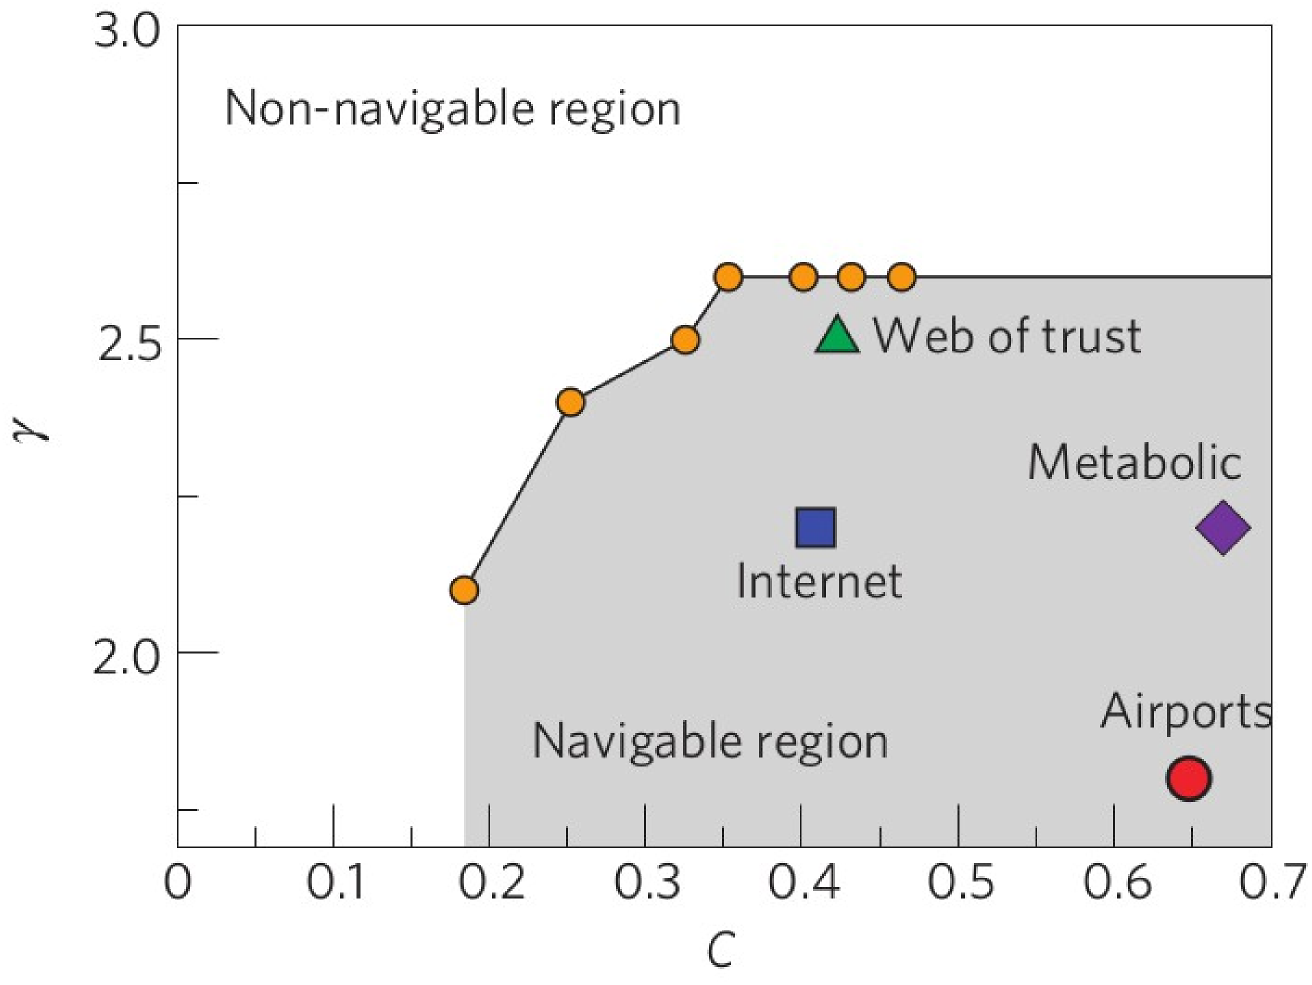

Navigation in Scale Free Networks

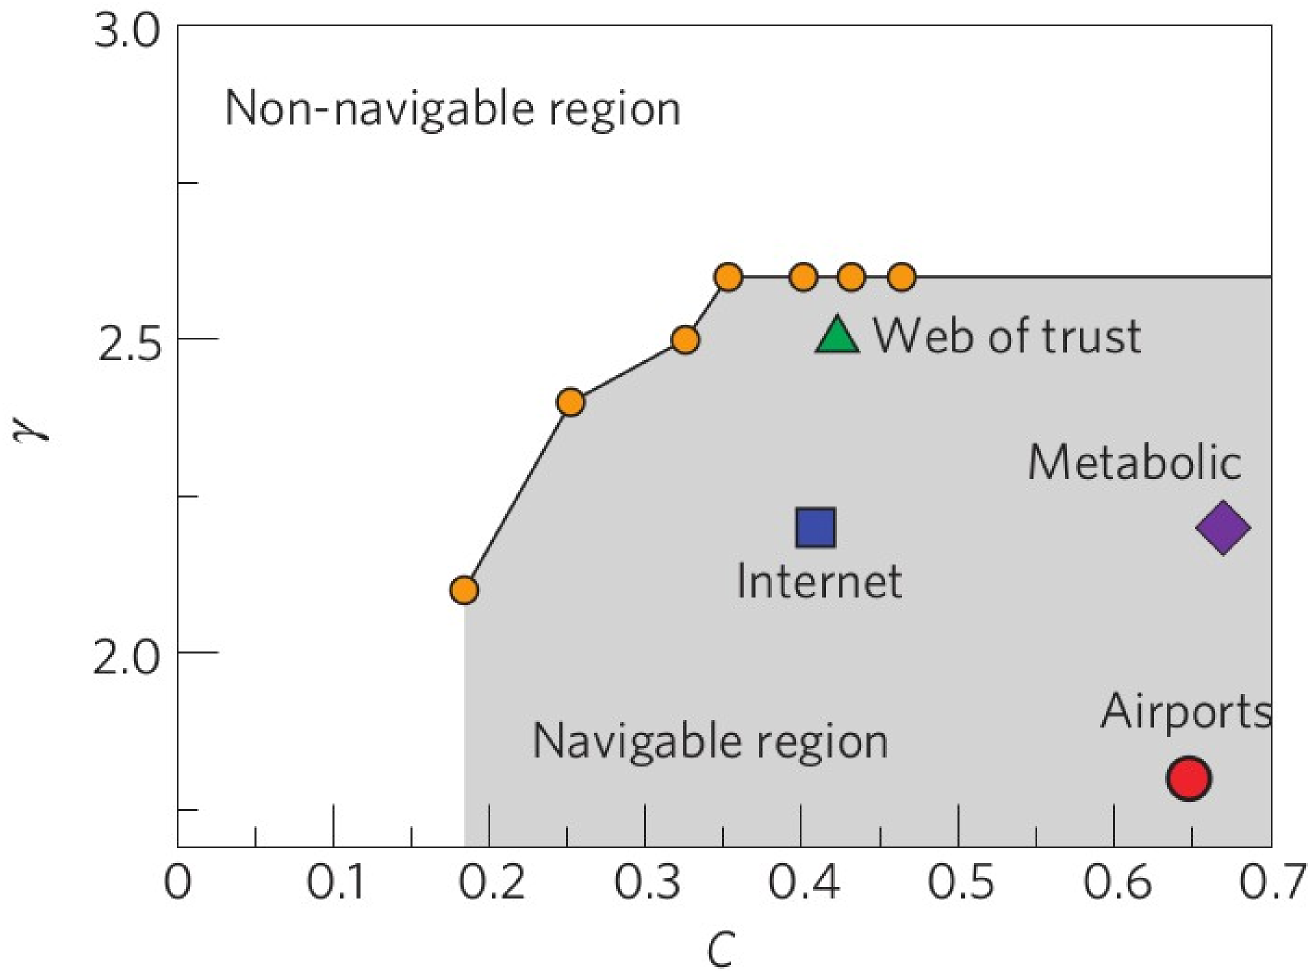

Clustering Coefficient $C = f(\gamma, \alpha)$

Navigation in Scale Free Networks

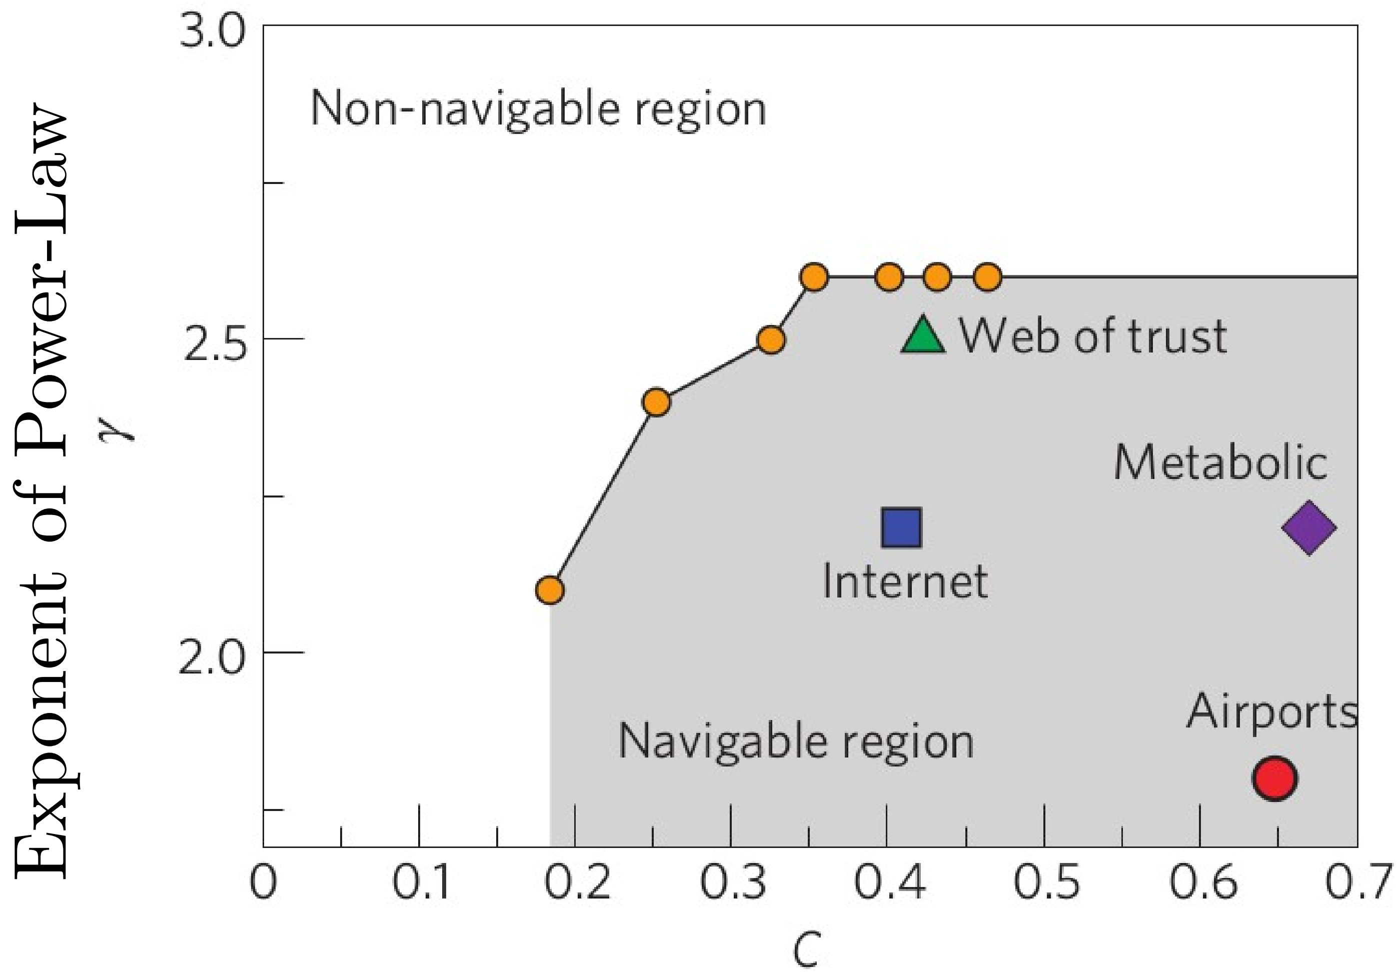

Clustering Coefficient $C = f(\gamma, \alpha)$

Navigation in Scale Free Networks

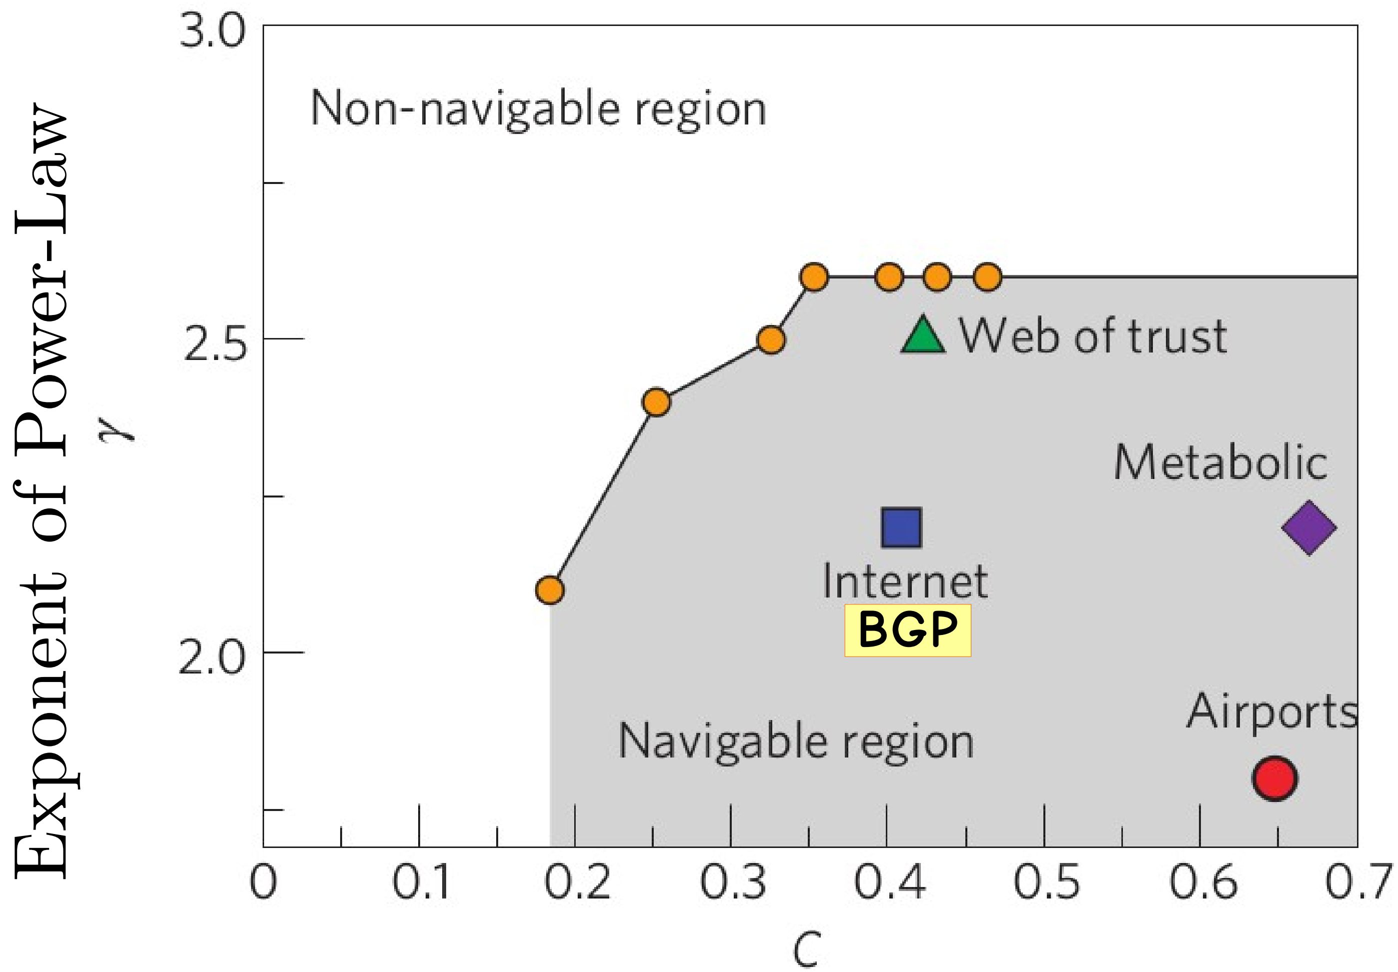

Clustering Coefficient $C = f(\gamma, \alpha)$

Navigation in Scale Free Networks

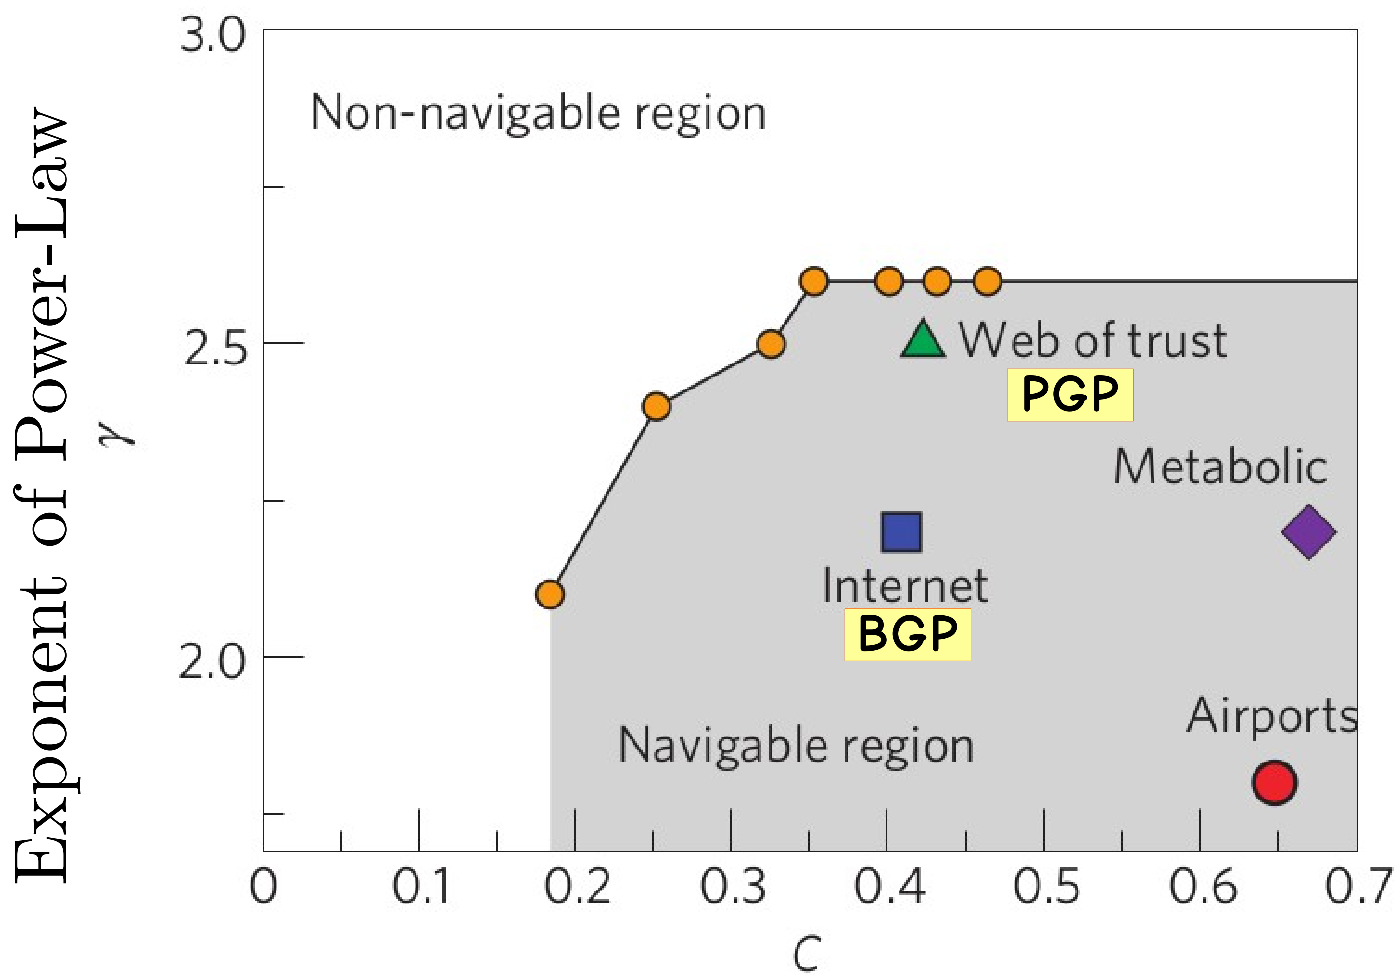

Clustering Coefficient $C = f(\gamma, \alpha)$

Navigation in Scale Free Networks

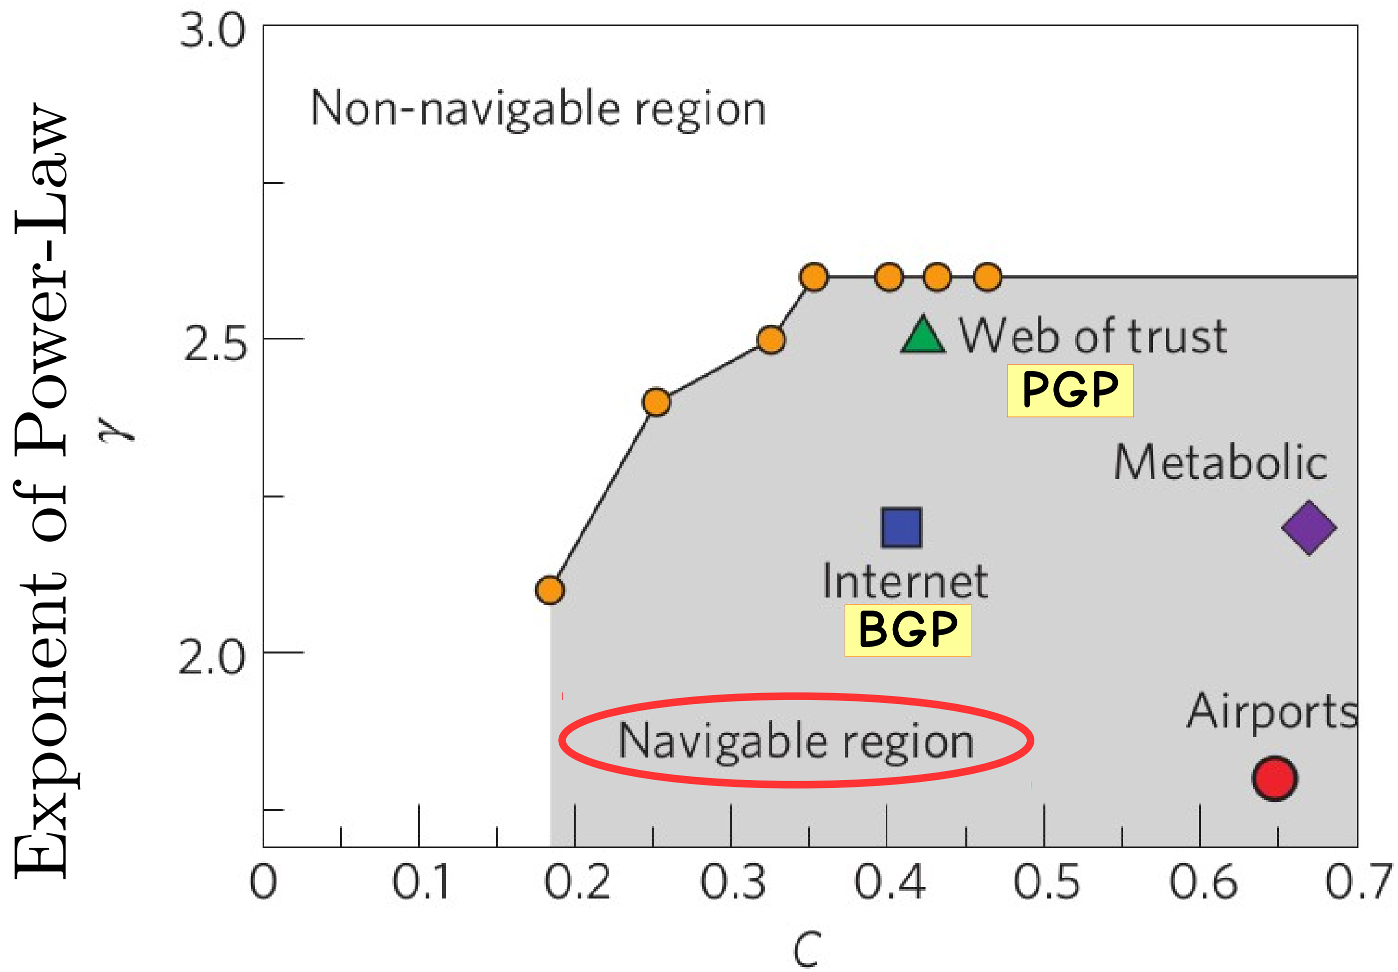

Clustering Coefficient $C = f(\gamma, \alpha)$

Navigation in Scale Free Networks

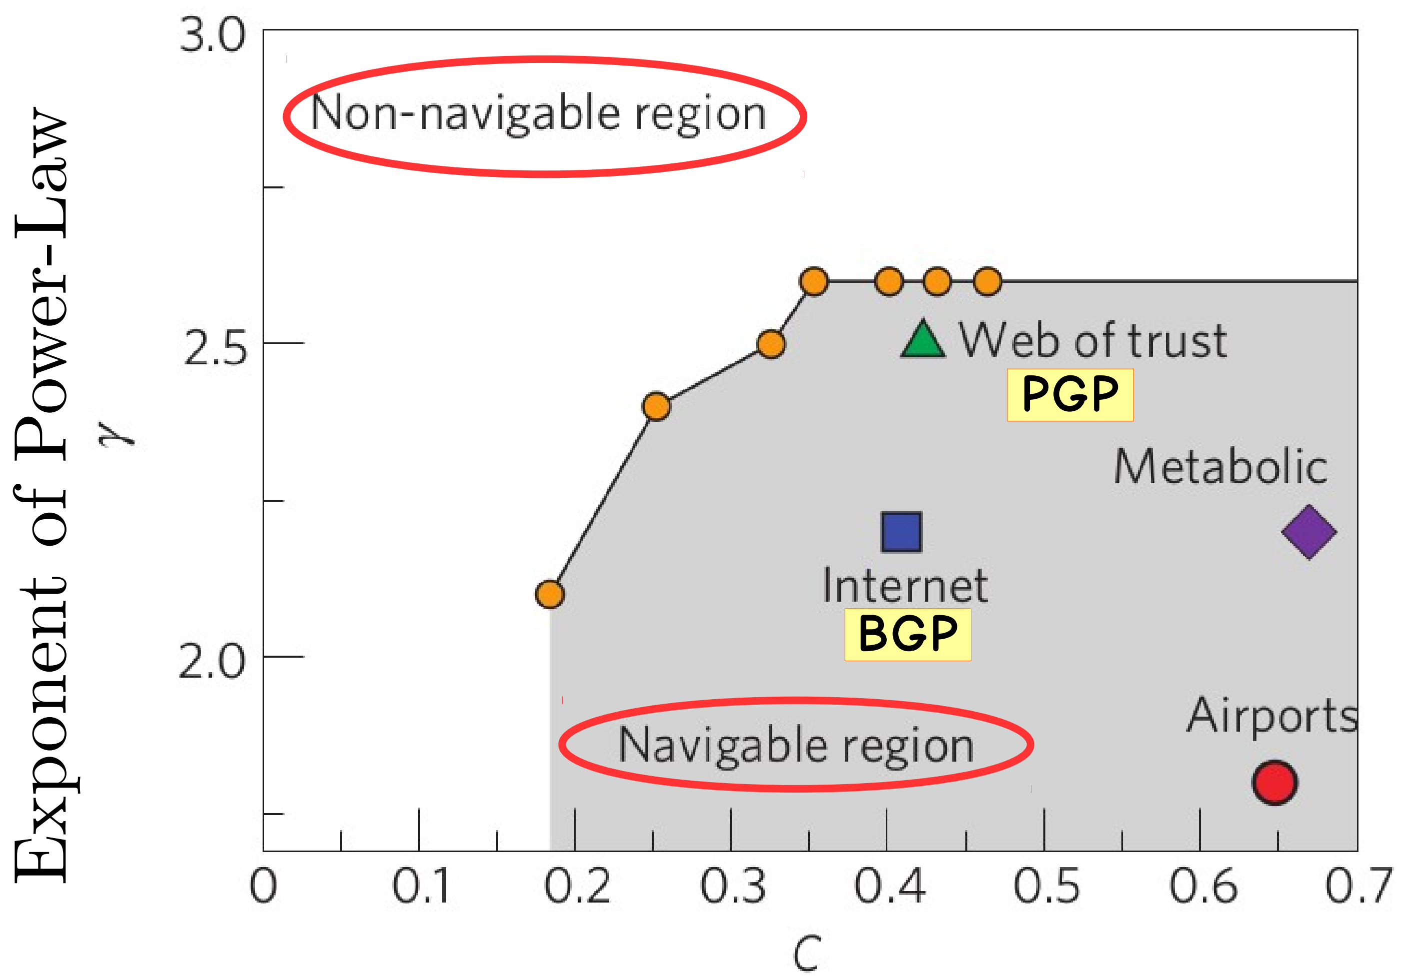

Clustering Coefficient $C = f(\gamma, \alpha)$

Implications of Result

- Internet Routing

Implications of Result

- Internet Routing
  - Routers currently exchange signals to keep cohenrent view of network
  - Network size increasing with time
  - Hidden metric space eliminates the need for control signals exchanged to notify changes in network

Implications of Result

- Internet Routing
  - Routers currently exchange signals to keep cohenrent view of network
  - Network size increasing with time
  - Hidden metric space eliminates the need for control signals exchanged to notify changes in network
- How to proceed to discover the hidden metric space

Implications of Result

- Internet Routing
  - Routers currently exchange signals to keep cohenrent view of network
  - Network size increasing with time
  - Hidden metric space eliminates the need for control signals exchanged to notify changes in network
- How to proceed to discover the hidden metric space
- Does Shortest Path imply Shortest Time to destination?
  - What happens in case of congestion at hubs?

Mathematics and the Internet: A Source of Enormous Confusion and Great Potential

W Willinger et al. “**Mathematics and the internet: A source of enormous confusion and great potential.**” In Notices of the AMS. 2009.

Scale-Free Model for AS-Graph

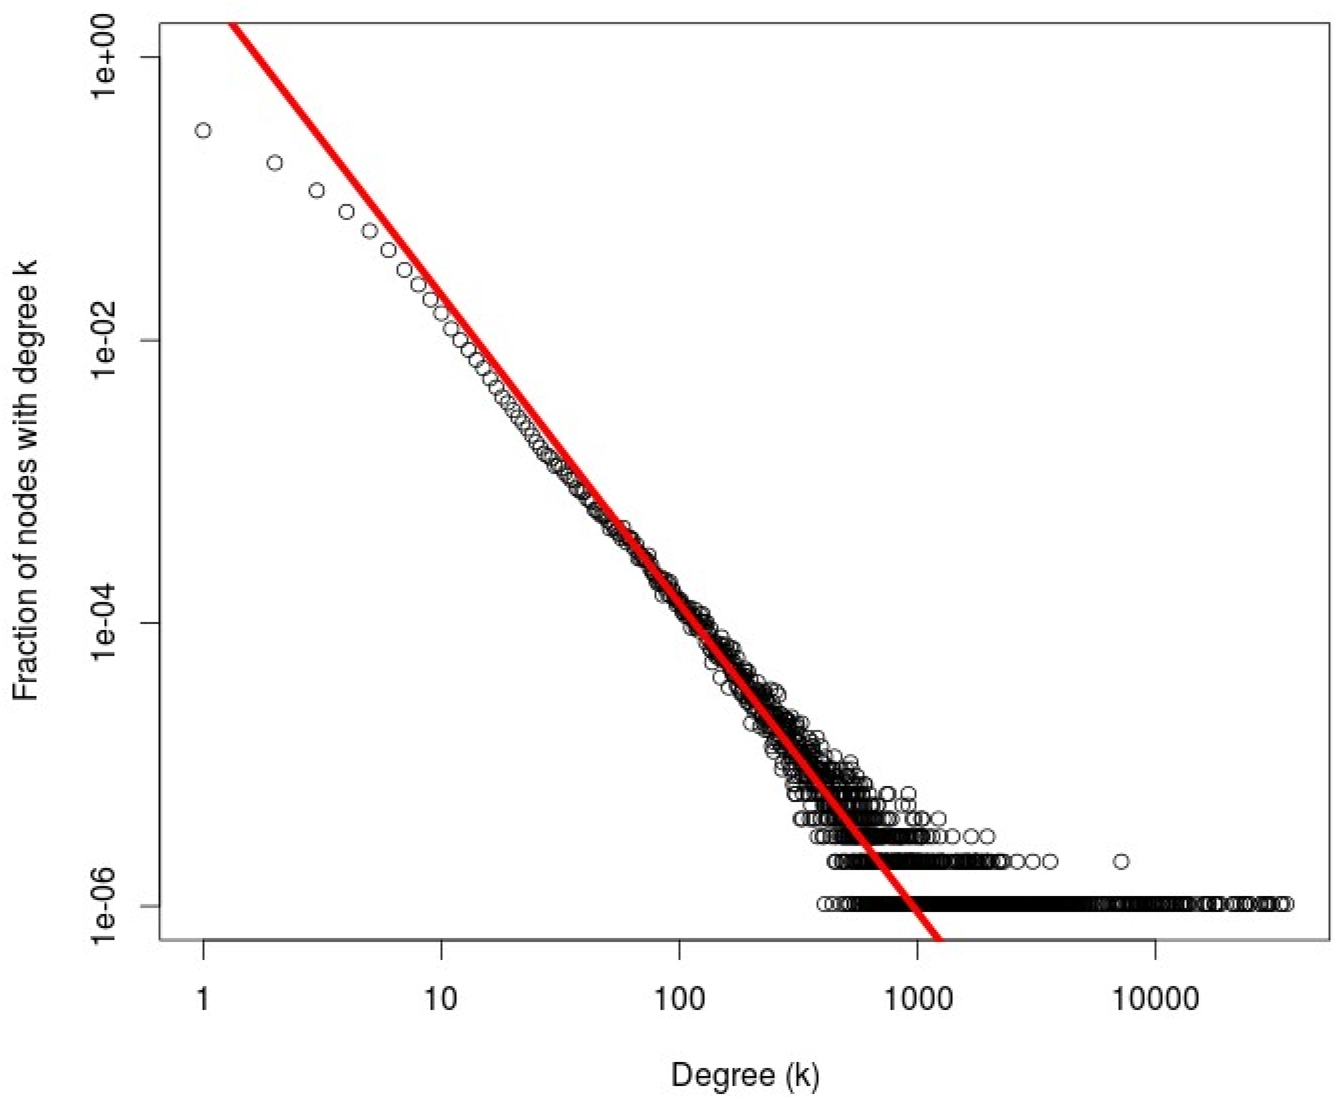

AS Topology of skitter dataset parsed by SNAP team
http://snap.stanford.edu/data/as-skitter.html

Scale-Free Model for AS-Graph

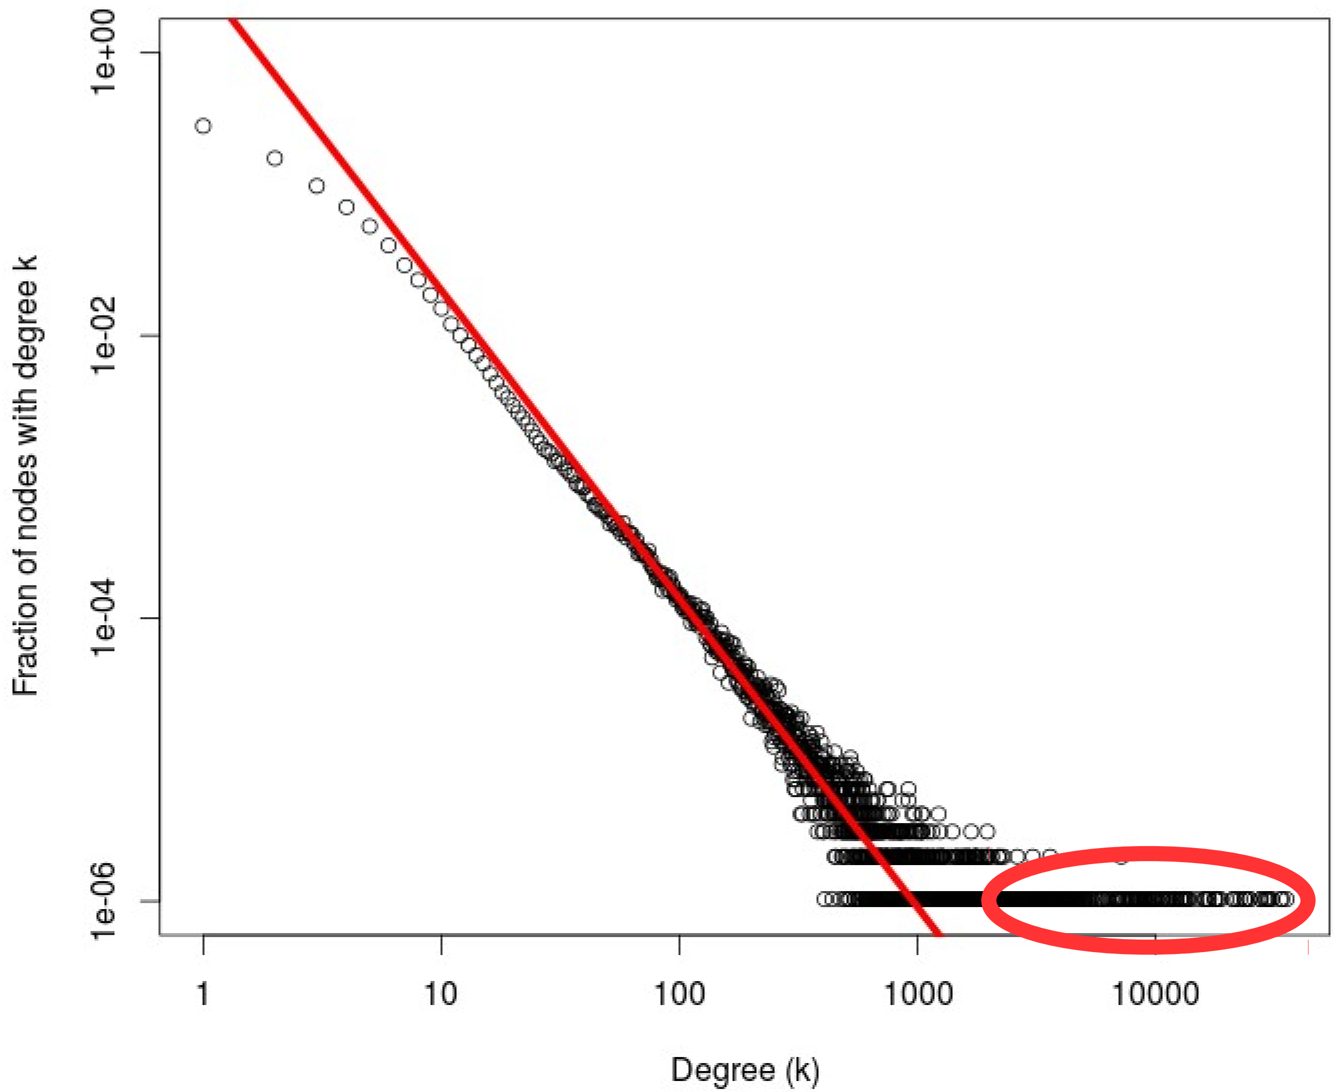

AS Topology of skitter dataset parsed by SNAP team
http://snap.stanford.edu/data/as-skitter.html

Is the Scale-Free Internet A Myth?

Is the Scale-Free Internet A Myth?

- What we have seen till now wrt to Preferential Attachment
  - Preferential attachment results in Hubs
  - Hubs vulnerable to coordinated attacks
  - Why is the Internet still up and running

Is the Scale-Free Internet A Myth?

- What we have seen till now wrt to Preferential Attachment
  - Preferential attachment results in Hubs
  - Hubs vulnerable to coordinated attacks
  - Why is the Internet still up and running
- Is the Scale-Free modeling paradigm consistent with the engineered nature of the Internet and the design constraints imposed by existing technology?
  - Is the simplistic toy model too generic?
  - Do the available measurements, their analysis, and their modeling efforts support the claims made by "Error and Attack Tolerance" paper?

Importance of Measurements

- Tool for measurement study for AS-measurements

Importance of Measurements

- Tool for measurement study for AS-measurements
  - Traceroute

Importance of Measurements

- Tool for measurement study for AS-measurements
  - Traceroute
- Biases of traceroute

Importance of Measurements

- Tool for measurement study for AS-measurements
  - Traceroute
- Biases of traceroute
  - Uses IPv4 Protocol
    - What about non-IPv4 protocols like MPLS?
    - Entry points to non-IPv4 regions can aggregate to Hubs

Importance of Measurements

- Tool for measurement study for AS-measurements
  - Traceroute
- Biases of traceroute
  - Uses IPv4 Protocol
    - What about non-IPv4 protocols like MPLS?
    - Entry points to non-IPv4 regions can aggregate to Hubs
  - Only reports the interfaces traversed by the packet
    - Routers can have multiple interfaces and appear on different routes with different IP addresses

Leverage Domain Knowledge

Leverage Domain Knowledge

- Device Constraints
  - Finite number of interfaces on routers
  - Finite capacity of routers

Leverage Domain Knowledge

- Device Constraints
  - Finite number of interfaces on routers
  - Finite capacity of routers
- Placement of High Degree Nodes
  - Edge vs Core

Leverage Domain Knowledge

- Device Constraints
  - Finite number of interfaces on routers
  - Finite capacity of routers
- Placement of High Degree Nodes
  - Edge vs Core

How would you deploy the network if you are a network engineer?

- Leverage domain knowledge to identify driving forces behind the design of high engineered systems such as the Internet

Leverage Domain Knowledge

- Device Constraints
  - Finite number of interfaces on routers
  - Finite capacity of routers
- Placement of High Degree Nodes
  - Edge vs Core

What about Overlay Networks?

How would you deploy the network if you are a network engineer?

- Leverage domain knowledge to identify driving forces behind the design of high engineered systems such as the Internet

Scale Free for Gnutella?

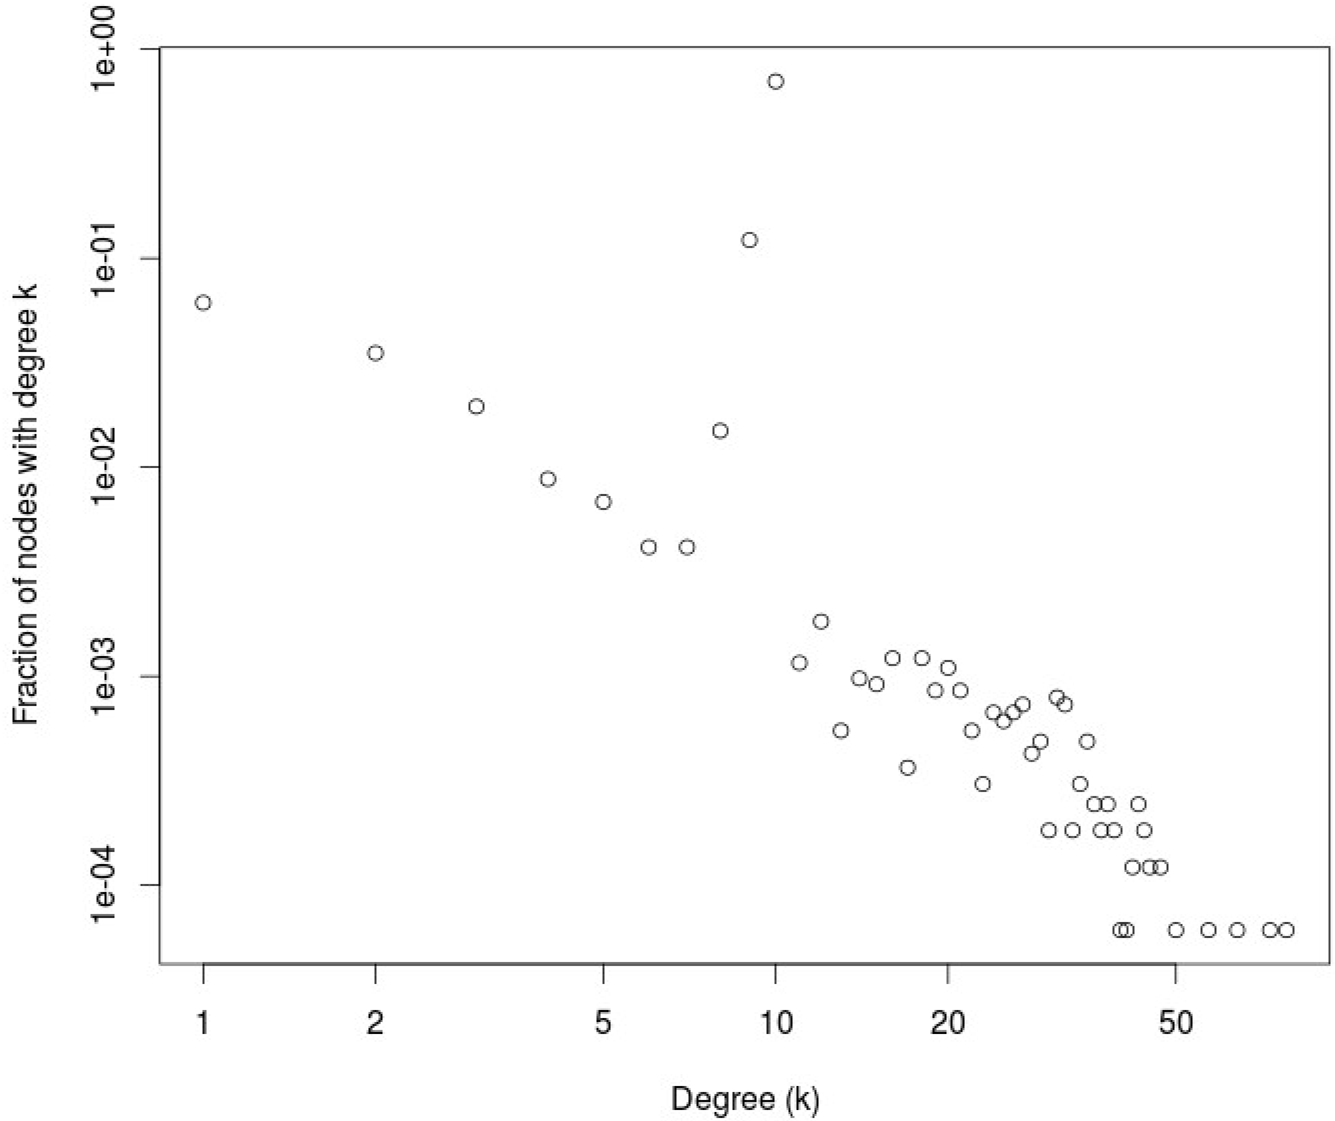

Gnutella August 2002 dataset parsed by SNAP team
http://snap.stanford.edu/data/p2p-Gnutella31.html

Summary
(Modeling Overlay Networks)

Recap of Modeling Overlay Networks

Recap of Modeling Overlay Networks

- Milgram's Experiment

Recap of Modeling Overlay Networks

- Milgram's Experiment
- Duncan Watts Random Rewiring Model

Recap of Modeling Overlay Networks

- Milgram's Experiment
- Duncan Watts Random Rewiring Model
- Scale-Free Networks

Recap of Modeling Overlay Networks

- Milgram's Experiment
- Duncan Watts Random Rewiring Model
- Scale-Free Networks
  - Preferential attachment
  - Evolving Copying Model (Copying Generative Model)
  - Scale-Free with Strong Clustering

Recap of Modeling Overlay Networks

- Milgram's Experiment
- Duncan Watts Random Rewiring Model
- Scale-Free Networks
  - Preferential attachment
  - Evolving Copying Model (Copying Generative Model)
  - Scale-Free with Strong Clustering
- Error and Fault Tolerance of Complex Networks

Recap of Modeling Overlay Networks

- Milgram's Experiment
- Duncan Watts Random Rewiring Model
- Scale-Free Networks
  - Preferential attachment
  - Evolving Copying Model (Copying Generative Model)
  - Scale-Free with Strong Clustering
- Error and Fault Tolerance of Complex Networks
- Navigation (Greedy Routing)

Recap of Modeling Overlay Networks

- Milgram's Experiment
- Duncan Watts Random Rewiring Model
- Scale-Free Networks
  - Preferential attachment
  - Evolving Copying Model (Copying Generative Model)
  - Scale-Free with Strong Clustering
- Error and Fault Tolerance of Complex Networks
- Navigation (Greedy Routing)
  - In Small World (Kleinberg's Small World)
  - In Complex Networks (Scale-Free with Strong Clustering)

Recap of Modeling Overlay Networks

- Milgram's Experiment
- Duncan Watts Random Rewiring Model
- Scale-Free Networks
  - Preferential attachment
  - Evolving Copying Model (Copying Generative Model)
  - Scale-Free with Strong Clustering
- Error and Fault Tolerance of Complex Networks
- Navigation (Greedy Routing)
  - In Small World (Kleinberg's Small World)
  - In Complex Networks (Scale-Free with Strong Clustering)
- Mathematics and the Internet: A Source of Enormous Confusion and Great Potential

Commonly used metrics

- Clustering Coefficient
- Diameter
- Degree Distribution

Methodology

1) Make observations (conduct measurement studies)
2) Build model to explain observations
   - Choose the right level of granularity (zoom level)
   - Strip the problem to a simple form
   - Attempt to formulate the problem and model the system
3) Validate model
   - Reproduce observations/measurements
   - Explain observations
4) Revisit step 2 (and 1) to improve understanding

Important Articles

- Milgram, Stanley. "**The small world problem.**" Psychology today 2.1 (1967): 60-67
- Watts, Duncan and Strogatz, Steven. "**Collective dynamics of 'small-world' networks.**" Nature 393.6684 (1998): 440-442.
- Barabási, Albert-László, and Albert, Réka. "**Emergence of scaling in random networks.**" Science 286, no. 5439 (1999): 509-512.
- Kleinberg, Jon. "**The small-world phenomenon: An algorithmic perspective.**" In ACM Symposium on Theory of computing, pp. 163-170. 2000.
- Ravi Kumar et al. "**Stochastic models for the web graph.**" In Annual Symposium on Foundations of Computer Science, 2000.
- Albert, Réka, and Barabási, Albert-László. "**Statistical mechanics of complex networks.**" Reviews of modern physics 74.1 (2002): 47.
- Newman, Mark. **"The structure and function of complex networks."** SIAM review 45, no. 2 (2003): 167-256.
- Mitzenmacher, M. (2004). "**A brief history of generative models for power law and lognormal distributions**." Internet mathematics, 1(2), 226-251.
- Mark Newman. **"Power laws, Pareto distributions and Zipf's law."** Contemporary physics 46, no. 5 (2005): 323-351
- Jure Leskovec et al. "**Graphs over time: densification laws, shrinking diameters and possible explanations.**" In ACM SIGKDD, pp. 177-187. 2005.
- Boguna, Marian et al. **"Navigability of complex networks."** Nature Physics 5, no. 1 (2009): 74-80.
- W Willinger et al. "**Mathematics and the internet: A source of enormous confusion and great potential.**" In Notices of the AMS. 2009.

Internet Indirection Architecture (I3)

http://i3.cs.berkeley.edu/

Ion Stoica et al. 2002. **"Internet indirection infrastructure."** In SIGCOMM '02.

Recap of Chord

Recap of Chord

- Distributed lookup protocol
  - Given a key, map the key to a node

Recap of Chord

- Distributed lookup protocol
  - Given a key, map the key to a node
- Assign a unique m-bit key (identifier) to a node
  - Consistent hashing (of IP address) generates keys

Recap of Chord

- Distributed lookup protocol
  - Given a key, map the key to a node
- Assign a unique m-bit key (identifier) to a node
  - Consistent hashing (of IP address) generates keys
- Identifiers (nodes) ordered in a circle module as 2^m
  - Every node has predecessor and successor

Recap of Chord

- Distributed lookup protocol
  - Given a key, map the key to a node
- Assign a unique m-bit key (identifier) to a node
  - Consistent hashing (of IP address) generates keys
- Identifiers (nodes) ordered in a circle module as 2^m
  - Every node has predecessor and successor
- A key k is assigned to a node whose identifier is equal to or follows k the identifier space.
  - k stored on successor(k)

Recap of Chord

- Distributed lookup protocol
  - Given a key, map the key to a node
- Assign a unique m-bit key (identifier) to a node
  - Consistent hashing (of IP address) generates keys
- Identifiers (nodes) ordered in a circle module as 2^m
  - Every node has predecessor and successor
- A key k is assigned to a node whose identifier is equal to or follows k the identifier space.
  - k stored on successor(k)
- Routing table (at most m entries on each node)
  - i^{th} entry → first node succeeds node by at least 2^{i-1}

Chord Properties

Chord Properties

- Each node responsible for ≈ K/N keys
  - K → total #keys, N → total #nodes
  - When a node joins or leaves the network only O(K/N) keys will be relocated
    - Relocation is local to the node

Chord Properties

- Each node responsible for ≈ K/N keys
  - K → total #keys, N → total #nodes
  - When a node joins or leaves the network only O(K/N) keys will be relocated
    - Relocation is local to the node
- Lookups take O(log N) messages

Chord Properties

- Each node responsible for ≈ K/N keys
  - K → total #keys, N → total #nodes
  - When a node joins or leaves the network only O(K/N) keys will be relocated
    - Relocation is local to the node
- Lookups take O(log N) messages
- $O(\log^2 N)$ messages required to re-establish routing invariants after join/leave
  - Each node's successor is correctly maintained
  - For every key (k), the node responsible for k is successor(k)

Packet's Perspective of Internet Services

- **Unicast**: One fixed source to one fixed destination
- **Broadcast:** One source to all destinations
- **Multicast**: One fixed source to multiple destinations who are part of a group
- **Anycast**: One source to exactly one destination who is a member of a group

Packet's Perspective of Internet Services

- **Unicast**: One fixed source to one fixed destination
- **Broadcast:** One source to all destinations
- **Multicast**: One fixed source to multiple destinations who are part of a group
- **Anycast**: One source to exactly one destination who is a member of a group

Internet Services using Unicast, Broadcast, Multicast, and Anycast are built over the **point-to-point abstraction**

Packet's Perspective of Internet Services

- **Unicast**: One fixed source to one fixed destination
- **Broadcast:** One source to all destinations
- **Multicast**: One fixed source to multiple destinations who are part of a group
- **Anycast**: One source to exactly one destination who is a member of a group

Internet Services using Unicast, Broadcast, Multicast, and Anycast are built over the **point-to-point abstraction**

Can we use another abstraction?

Rendezvous-based Communication

Rendezvous-based Communication

- Source sends packets to a ***logical identifier***. Receivers express interest in packets sent to an identifier

Rendezvous-based Communication

- Source sends packets to a ***logical identifier***. Receivers express interest in packets sent to an identifier
- Packet is a pair **(id, data)**
  - id → host/object/session/...

Rendezvous-based Communication

- Source sends packets to a ***logical identifier***. Receivers express interest in packets sent to an identifier
- Packet is a pair **(id, data)**
  - id → host/object/session/...

Rendezvous-based Communication

- Source sends packets to a ***logical identifier***. Receivers express interest in packets sent to an identifier
- Packet is a pair **(id, data)**
  - id → host/object/session/…

- Receivers use triggers **(id, addr)** to express interest
  - Forward packet with identifier **(id)** to receiver with IP address **(addr)**

Rendezvous-based Communication

- Source sends packets to a ***logical identifier***. Receivers express interest in packets sent to an identifier
- Packet is a pair **(id, data)**
  - id → host/object/session/…

- Receivers use triggers **(id, addr)** to express interest
  - Forward packet with identifier **(id)** to receiver with IP address **(addr)**
- Packet sent to receivers if the
  - the interest (**id_T**) from receiver is a longest prefix match
  - the match is longer than matching threshold k (k < m)

Rendezvous-based Communication

- Source sends packets to a ***logical identifier***. Receivers express interest in packets sent to an identifier
- Packet is a pair **(id, data)**
  - id → host/object/session/…

- Receivers use triggers **(id, addr)** to express interest
  - Forward packet with identifier **(id)** to receiver with IP address **(addr)**
- Packet sent to receivers if the
  - the interest ($\mathbf{id_T}$) from receiver is a longest prefix match
  - the match is longer than matching threshold k ($k < m$)

Abstraction decouples the act of sending from the act of receiving

API to Implement Indirection

- SendPacket (p)
- InsertTrigger(t)
- RemoveTrigger(t)

API to Implement Indirection

- SendPacket (p)
- InsertTrigger(t)
- RemoveTrigger(t)

API Implemented in an i3 Overlay Network

API to Implement Indirection

- SendPacket (p)
- InsertTrigger(t)
- RemoveTrigger(t)

API Implemented in an i3 Overlay Network

- Overlay Consists of i3 Servers
    - Store Triggers
    - Forward packets using IP between i3 nodes and end-hosts
        - ***Packets are not stored at the Servers***
    - Implemented using Chord (or any other DHT)

Benefits of i3

- Support for mobility

Benefits of i3

- Support for mobility
  - On moving to new address (addr'), receiver sends new trigger (id, addr')
  - Receivers periodically refresh triggers

Benefits of i3

- Support for mobility
  - On moving to new address (addr'), receiver sends new trigger (id, addr')
  - Receivers periodically refresh triggers
- Multicast

Benefits of i3

- Support for mobility
  - On moving to new address (addr'), receiver sends new trigger (id, addr')
  - Receivers periodically refresh triggers
- Multicast
  - Source is agnostic to the set of receivers
  - Receiver agnostic to the set of sources
  - Trigger chains can be used to minimize triggers

Benefits of i3

- Support for mobility
  - On moving to new address (addr'), receiver sends new trigger (id, addr')
  - Receivers periodically refresh triggers
- Multicast
  - Source is agnostic to the set of receivers
  - Receiver agnostic to the set of sources
  - Trigger chains can be used to minimize triggers
- Anycast

Benefits of i3

- Support for mobility
  - On moving to new address (addr'), receiver sends new trigger (id, addr')
  - Receivers periodically refresh triggers
- Multicast
  - Source is agnostic to the set of receivers
  - Receiver agnostic to the set of sources
  - Trigger chains can be used to minimize triggers
- Anycast
  - Id contains a common prefix component and a **suffix**

Benefits of i3

- Support for mobility
  - On moving to new address (addr'), receiver sends new trigger (id, addr')
  - Receivers periodically refresh triggers
- Multicast
  - Source is agnostic to the set of receivers
  - Receiver agnostic to the set of sources
  - Trigger chains can be used to minimize triggers
- Anycast
  - Id contains a common prefix component and a **suffix**
- Anonymity

Benefits of i3

- Support for mobility
  - On moving to new address (addr'), receiver sends new trigger (id, addr')
  - Receivers periodically refresh triggers
- Multicast
  - Source is agnostic to the set of receivers
  - Receiver agnostic to the set of sources
  - Trigger chains can be used to minimize triggers
- Anycast
  - Id contains a common prefix component and a **suffix**
- Anonymity
- Service Composition

Benefits of i3

- Support for mobility
  - On moving to new address (addr'), receiver sends new trigger (id, addr')
  - Receivers periodically refresh triggers
- Multicast
  - Source is agnostic to the set of receivers
  - Receiver agnostic to the set of sources
  - Trigger chains can be used to minimize triggers
- Anycast
  - Id contains a common prefix component and a **suffix**
- Anonymity
- Service Composition
  - Stacked identifiers

P2P SIP

SIP

- Session Initiation Protocol
  - An Application-layer control (signaling) protocol for creating, modifying and terminating sessions with one or more participants
  - Sessions include Internet multimedia conferences, Internet telephone calls and multimedia distribution
  - Members can communicate via multicast or mesh of unicast relations, or a combination of the two
  - Text based, model similar to HTTP

J. Rosenberg et al. **"RFC 3261: SIP: session initiation protocol."** (2003).

Network Elements

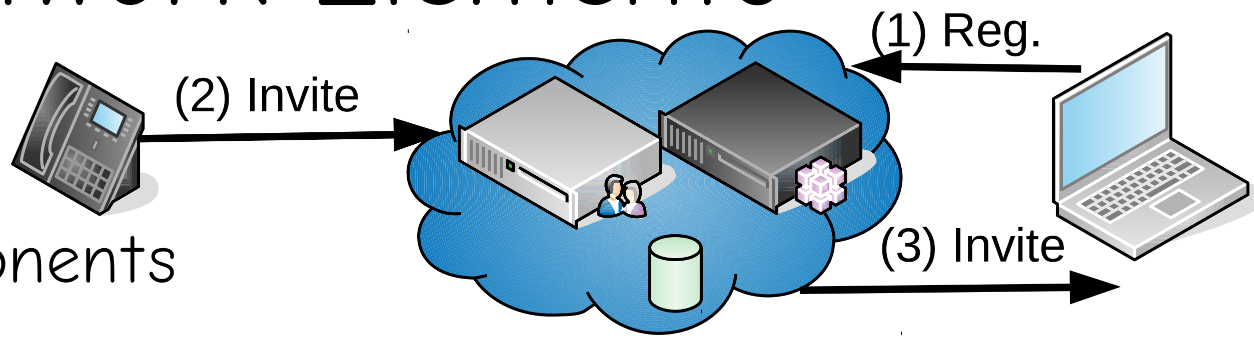

- User-agent
  - End-point components

http://sipsense.com/

Network Elements

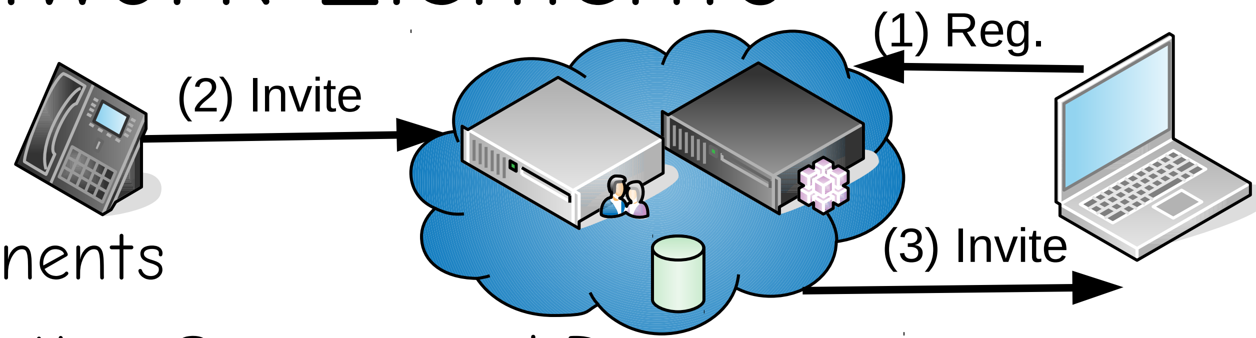

- User-agent
  - End-point components
- SIP Registrar, Location Server, and Proxy
  - Helps users resolve the IP address of each other
- Feature Servers
  - Value added services (call forwarding, recording, etc.)
- Session Border Controller
  - Protect a SIP sub-network from attacks
- Gateways: Signalling Gateway and Media Gateway
  - Transcode Media
  - Support interaction with non-SIP clients

http://sipsense.com/

P2P SIP

- Can we leverage P2P technologies to implement SIP?
  - Do away with central servers

Kundan Singh et al. **"Peer-to-peer internet telephony using SIP."** In NOSSDAV 2005

P2P SIP

- Can we leverage P2P technologies to implement SIP?
  - Do away with central servers

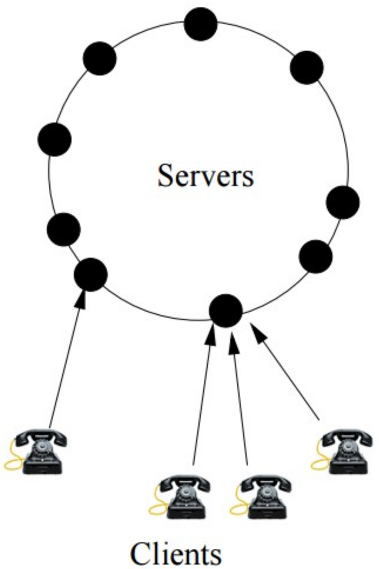

Only Servers in DHT
Unmodified Clients

Kundan Singh et al. **"Peer-to-peer internet telephony using SIP."** In NOSSDAV 2005

P2P SIP

- Can we leverage P2P technologies to implement SIP?
  - Do away with central servers

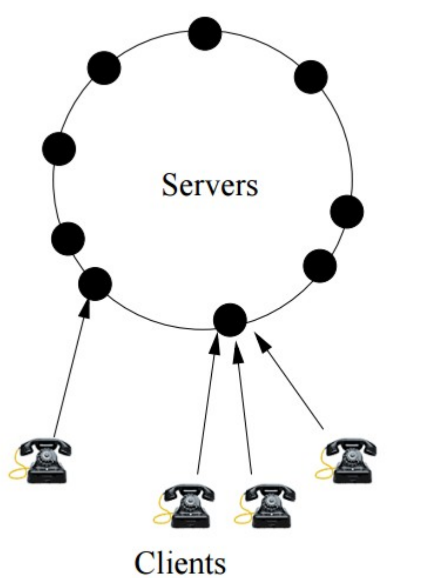

Only Servers in DHT
Unmodified Clients

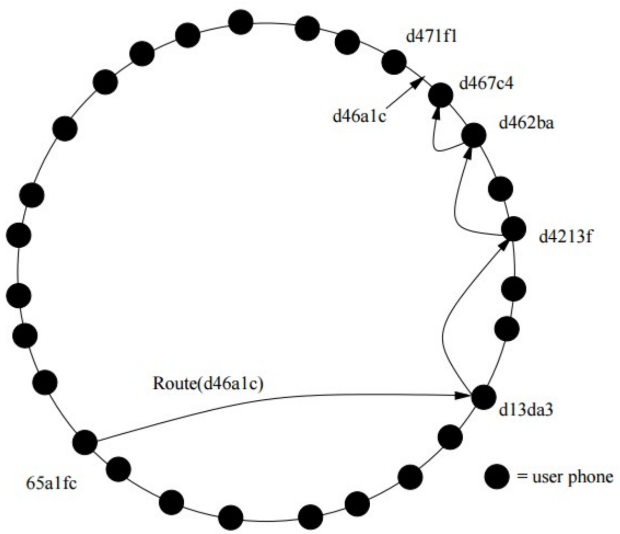

All Clients in DHT
Requires modification of clients

Kundan Singh et al. **"Peer-to-peer internet telephony using SIP."** In NOSSDAV 2005

P2P SIP

- Can we leverage P2P technologies to implement SIP?
  - Do away with central servers

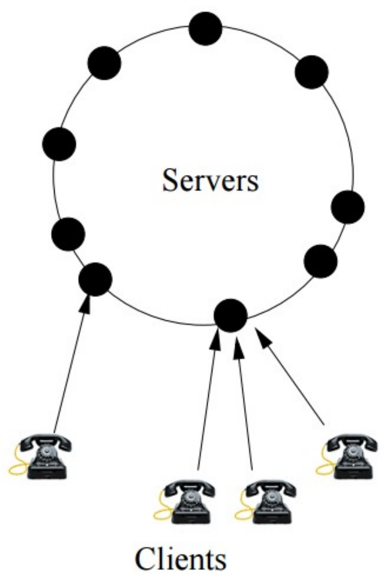

Only Servers in DHT
Unmodified Clients

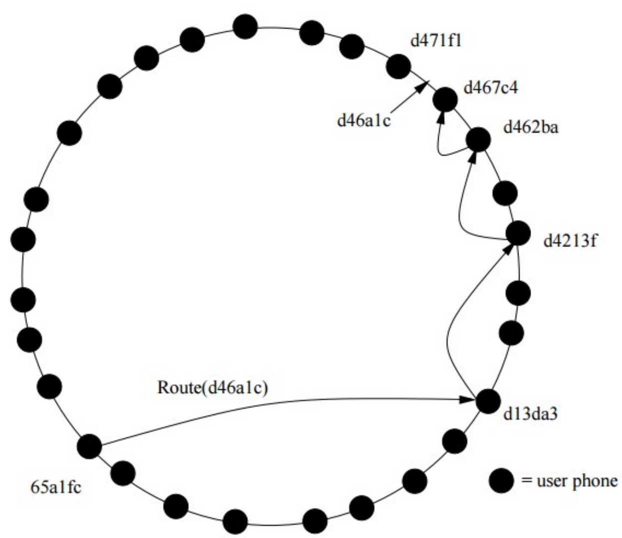

All Clients in DHT
Requires modification of clients

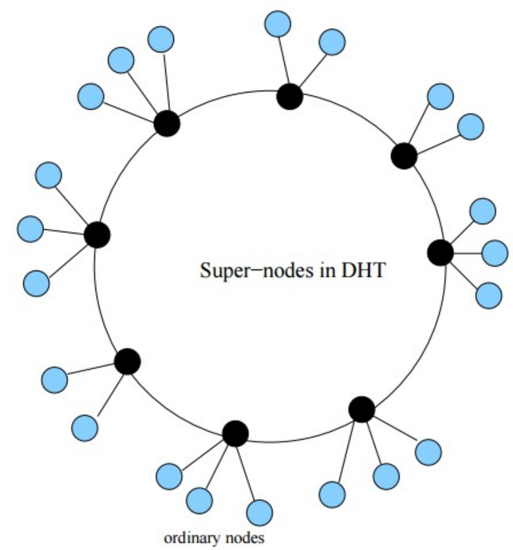

Super-nodes in DHT

Kundan Singh et al. **"Peer-to-peer internet telephony using SIP."** In NOSSDAV 2005