Peer-to-Peer Networks

Chapter 4: Graphs and Methods

Thorsten Strufe

Note: these slides have been prepared with influence by material of Prof. Michael Welzl, Prof. Pietro Michiardi, and Dr. Stefan Schmid

Chapter Outline

- P2P Overlays as Graphs (This chapter is a reminder)
- Graphs
- Metrics in and Properties of Graphs
- Algorithms on Graphs
- A tiny introduction to game theory

Some questions…

- How scalable is Gnutella?
- How robust is Kazaa?
- Why does FreeNet work?
- What would an ideal (unstructured) P2P system look like?
- What do the overlay networks of existing P2P systems look like?

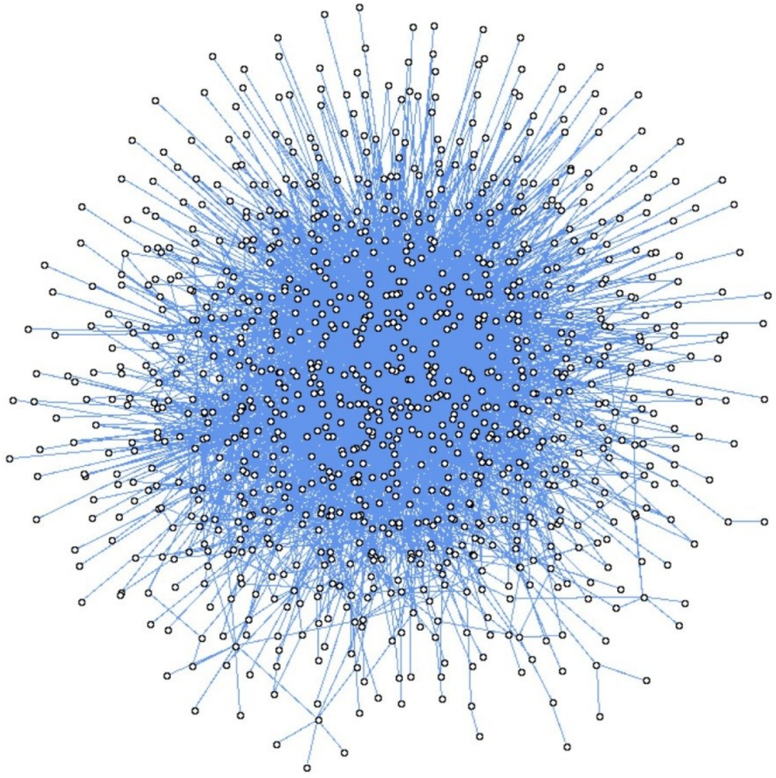

Gnutella snapshot, 2000

Scalability of Gnutella: quick answer

- Bandwidth Generated in Bytes (Message 83 bytes) *[SIC]
  - Searching for a 18 byte string

	T=2	*T=3*	*T=4*	*T=5*	*T=6*	*T=7*	*T=8*
N=2	332	498	664	830	996	1,162	1,328
N=3	747	1,743	3,735	7,719	15,687	31,623	63,495
N=4	1,328	4,316	13,280	40,172	120,848	362,876	1,088,960
N=5	2,075	8,715	35,275	141,515	566,475	2,266,315	9,065,675
N=6	2,988	15,438	77,688	388,938	1,945,188	9,726,438	48,632,688
N=7	4,067	24,983	150,479	903,455	5,421,311	35,528,447	192,171,263
N=8	5,312	37,848	262,600	1,859,864	13,019,712	91,138,648	637,971,200

  - N = number of connections
  - T = number of hops

** [SIC]: Error already in source, orders of magnitude are important, here ;)*
Source: Jordan Ritter: Why Gnutella Can't Scale. No, Really.

Graphs

- Rigorous analysis of P2P systems: based on graph theory
  - Refresher of graph theory needed
- First: graph families and models
  - Random graphs
  - Small world graphs
  - Scale-free graphs
- Then: graph theory and P2P
  - How are the graph properties reflected in real systems?
    - Users (peers) are represented by vertices in the graph
    - Edges represent connections in the overlay (routing table entries)
- Concept of self-organization
  - Network structures emerge from simple rules
  - E.g. also in social networks, www, actors playing together in movies

What Is a Graph?

- Definition of a graph:

  Graph $G = (V, E)$ consists of two finite sets, set V of vertices (nodes) and set E of edges (arcs) for which the following applies:

  1. If $e \in E$, then exists $(v, u) \in V \times V$, such that $v \in e$ and $u \in e$
  2. If $e \in E$ and above (v, u) exists, and further for $(x, y) \in V \times V$ applies $x \in e$ and $y \in e$, then $\{v, u\} = \{x, y\}$

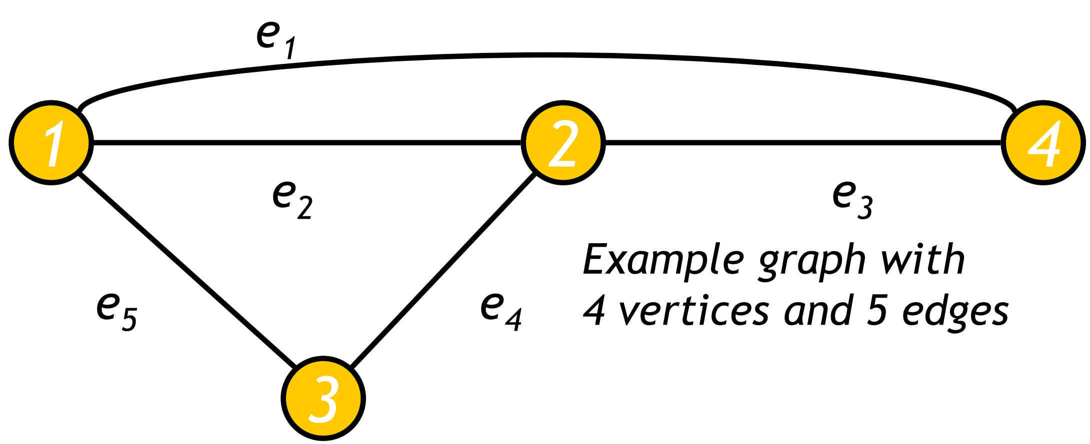

Example graph with 4 vertices and 5 edges

Side note:

Edges can have (multiple) "weights" $w : E \rightarrow \mathbb{R}$

Recall Graphs

- Graph G = (V,E),
  - Edges can have (multiple) “weights” $w : E \rightarrow R$
  - Edges are ***directed*** or ***undirected***
  - The graph hence is ***directed*** or ***undirected***
  - A graph can be ***connected*** (strong and weak connectivity)
- Undirected, acyclic graphs are trees
  - Directed, acyclic graphs are DAGs

- Graphs have an ***order, size, distance, diameter,*** and ***density***

How are Graphs Implemented?

- Adjacency/Incidence Matrix

	1	**2**	**3**
1	0	1	0
2	1	0	1
3	0	1	0

- Adjacency/Incidence List

(1,2)
(2,1),(2,3)
(3,2)

1:2
2:1,3
3:2

- (Plus specialized others..)

VERY good book is: Sedgewick: Algorithms in C, part 3 (Graph Algorithms)

Properties of Graphs

- An edge $e \in E$ is directed if the start and end vertices in condition 2 above are identical: $v = x$ and $y = u$
- An edge $e \in E$ is undirected if $v = x$ and $y = u$ as well as $v = y$ and $u = x$ are possible
- A graph G is directed (undirected) if the above property holds for all edges
- A *loop* is an edge with identical endpoints
- Graph $G_1 = (V_1, E_1)$ is a subgraph of $G = (V, E)$, if $V_1 \subseteq V$ and $E_1 \subseteq E$ (such that conditions 1 and 2 are met)

Important Types of Graphs

- Vertices $v, u \in V$ are connected if there is a path from v to u: $(v, v_2), (v_2, v_3), \ldots, (v_{k-1}, u) \in E$

- Graph G is connected if all $v, u \in V$ are connected
  - Strong connectivity of directed graphs means, that paths between each node pair exist
  - Weak connectivity: edges between all node pairs exist, but not paths…

- An undirected, connected, acyclic graph is called a tree
  - Side note: Undirected, acyclic graphs which are not connected are called forest

- Directed, connected, acyclic graph is called DAG
  - DAG = Directed Acyclic Graph (connectivity is assumed)

- An induced graph $G(V_C) = (V_C, E_C)$ is a graph $V_C \subseteq V$ and with edges $E_C = \{e = (i, j) \mid i, j \in V_C\}$ (all edges from G connecting the nodes in G_C)

- An induced graph that is connected is called a component

Overlays?

- A CLIQUE is a graph that is fully connected

  $(u,v) \in E \mid$ for all $u \in V$ and $v \in V$, $u \neq v$

- A P2P Overlay (V_o,E_o) (in general) is a subgraph such that $V_o=V$ and $E_o \subseteq E$ (edges are selected edges from a CLIQUE graph)

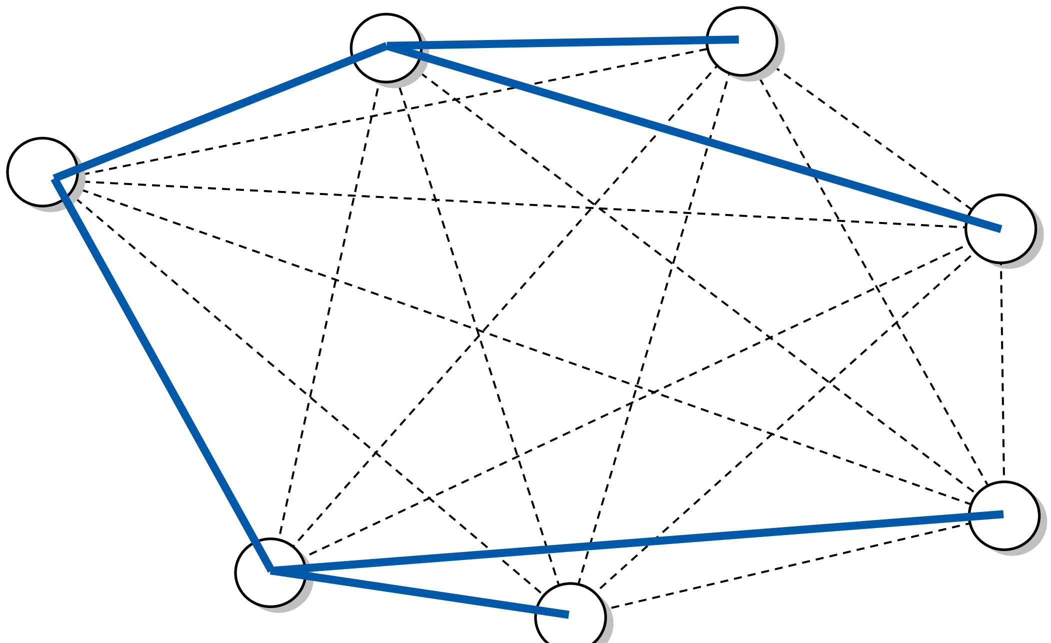

- ***Why?*** Considering the nodes to be on the Internet, they all can create connections between each other…

Important Graph Metrics

- Order: the number of vertices in a graph
- Size of the graph is the number of edges |E|

- Distance: *d(v, u)* between vertices *v* and *u* is the length of the shortest path between *v* and *u*

- Diameter: *d(G)* of graph *G* is the maximum of *d(v, u)* for all $v, u \in V$

- The density of a graph is the ratio of the number of edges and the number of possible edges.

Graph Metrics: Vertex Degree

- In graph $G = (V, E)$, the degree of vertex $v \in V$ is the total number of edges $(v, u) \in E$ and $(u, v) \in E$
  - Degree is the number of edges which touch a vertex

- For directed graph, we distinguish between in-degree and out-degree
  - In-degree is number of edges coming to a vertex
  - Out-degree is number of edges going away from a vertex

- The degree of a vertex can be obtained as:
  - Sum of the elements in its row in the incidence matrix
  - Length of its vertex incidence list

- The ***degree distribution*** is the distribution over all node degrees

  (given as a frequency distribution or (often) complementary cumulative distribution function CCDF (Komplement der Verteilungsfunktion))

Graph Metrics: Degree Distribution (Examples)

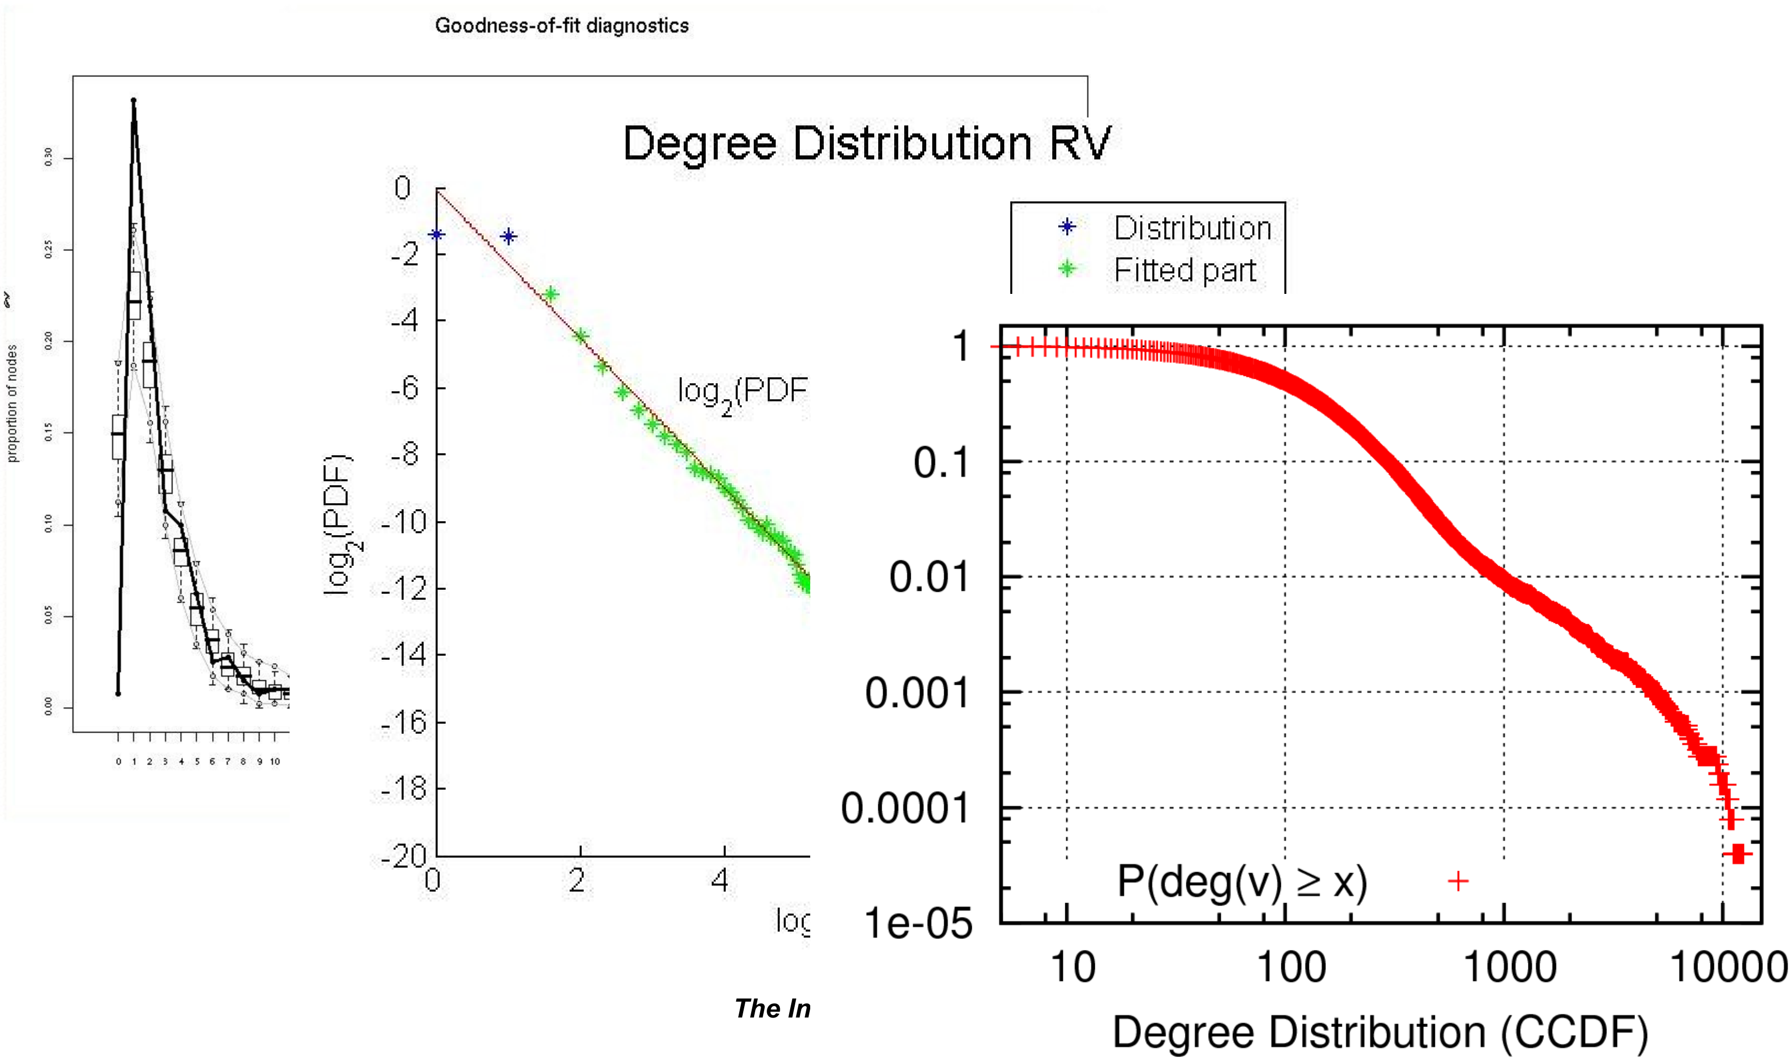

The In

Online Social Network (xing crawl) [strufe10popularity]

Further Graph Metrics

- All pairs shortest paths (APSP): *d(v, u)* | all $v,u \in V$

- Hop Plot: Distance distribution over all distances
  Hist(APSP(G))

- Average/characteristic path length: Sum of the distances over all pairs of nodes divided by the number of pairs

For defined routings on graphs:

Characterisic Routing Length: average length of paths *found* (potentially stochastic…)

Sanity Check:

- APSP:

1,1,1,2,2

1,1,1,2,2

1,1,1,1,1

1,1,1,1,2

1,1,1,2,2

1,1,2,2,2

- HopPlot:

1: 20

2: 10

- CPL: *40/30 = 1.333*
- CRL (CHORD): *2.033*

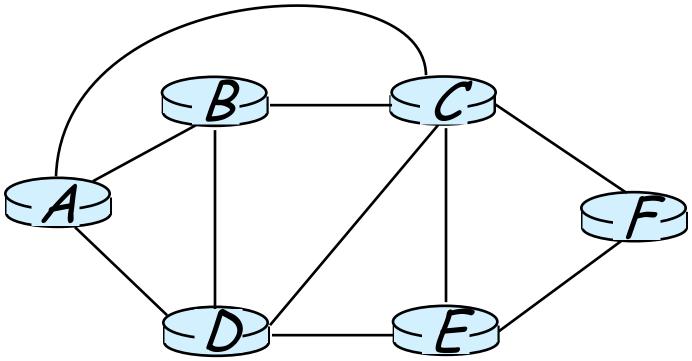

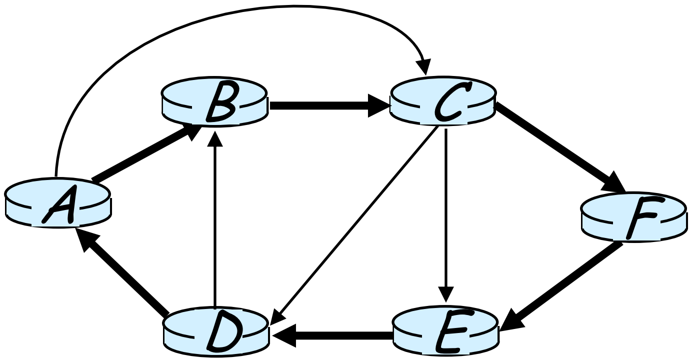

Important Graph Metrics (3)

- Edge connectivity: is the minimum number of edges that have to be removed to separate the graph into at least two components

- Vertex connectivity: the minimum number of nodes..

> - *Does each network have **ONE** maximum flow?*
> - *What is the edge/vertex connectivity if we have a node with degree=1?*
>   - *Is this a sensible metric? How to „heal“ this?*
>
> *==> balanced cut, size of the remaining giant connected component, fraction of disconnected nodes*

- How to calculate them?
- Which of both is higher?
- In which cases are they the same?

- maxflow, Menger‘s Theorem…

Important Graph Metrics (2)

- Clustering coefficient: number of edges between neighbors divided by maximum number of edges between them
  - k neighbors: k(k-1)/2 possible edges between them

$$C(i) = \frac{2E(N(i))}{d(i)(d(i)-1)}$$

E(N(i)) = number of edges between neighbors of i

d(i) = degree of i

- ***What if***: a node has only one neighbor? ☺

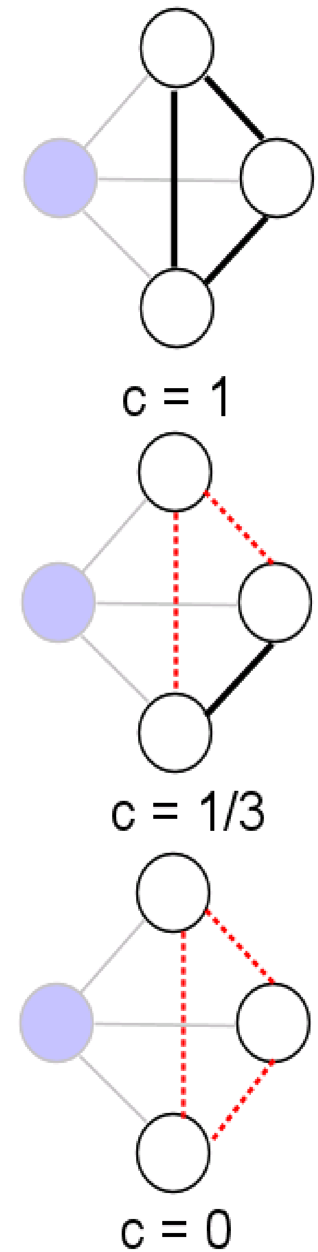

Source: Wikipedia

P2P Self Organization and Routing

- So what can a peer actually do?
  - *(Initially: bootstrapping, ID selection)*

- Select neighbors
  - Randomly
  - According to some rule

- Select next hop (when delegation is needed..)
  - Randomly
    - Single next hop
    - Multiple next hops (request replication… flooding)
  - According to some rule

- *(Change ID, but that‘s already advanced.. ☺)*

So what!?

Location Overlay $\mathcal{L} = (V_L, E_L)$ *(of Vertices **V** and Edges **E**)*

- Reliability
  - High success probability

→ ***hit ratio*** = $\frac{|Hit|}{|Requests|}$

→ ***response time*** =

- Low response times

- Resource usage
  - Low message complexity

→ ***message complexity*** = $\frac{\sum_{m_e \in \mathcal{M}} m_e \cdot d(e)}{|Requests|}$

- ***Great! We can do maths, now!*** ☺ (Plus: we can define metrics..)

Classes of Graphs

- Regular graphs
- Random graphs
- Graphs with Small-World characteristic
- Scale-free graphs
- …Graphs with plenty more characteristics
  - (dis-) assortativity
  - Rich-club connectivity
  - …

Regular Graphs

- Regular graphs have traditionally been used to model networks

- Regular graphs:
  - Node degree is constant
  - Different topologies possible

- But the model does not reflect reality of nature very well…

Random Graphs

- Random graphs are first widely studied graph family
  - Many P2P networks choose neighbors more or less randomly

- Two different notations generally used:
  - Erdös and Renyi
  - Gilbert

- Gilbert's definition: Graph $G_{n,p}$ (with n nodes) is a graph where the probability of an edge $e = (v, w)$ is p

Construction algorithm:

- For each possible edge, draw a random number
- If the number is smaller than p, then the edge exists
- p can be function of n or constant

Basic Results for Random Graphs

Giant Connected Component

Let $c > 0$ be a constant and $p = c/n$.

If $c < 1$ every component of $G_{n,p}$ has order $O(\log N)$ with high probability.

*If $c > 1$ then there is one component of size **$n*(f(c) + O(1))$** where $f(c) > 0$, with high probability. All other components have size $O(\log N)$*

- ***English***: Giant connected component emerges with high probability when average degree is about 1

Node degree distribution

- If we take a random node, how high is the probability *P(k)* that it has degree *k*?
- Node degree is Poisson distributed
  - Parameter c = expected number of occurrences

$$P(k) = \frac{c^k e^{-c}}{k!}$$

Clustering coefficient

- Clustering coefficient of a random graph is asymptotically equal to *p* with high probability

Random Graphs: Summary

- Random graphs have two advantages over regular graphs
  1. Many interesting properties analytically solvable
  2. Much more realistic (for many applications, e.g., social networks)

- Note: Does not mean social/p2p networks are random graphs; just that the properties of social/p2p networks are well-described by random graphs

- Question: How to model networks with local clusters and small diameter?

- Answer: Small-world networks

Six Degrees of Separation

- Famous experiment from 1960's (S. Milgram)

- Send a letter to random people in Kansas and Nebraska and ask people to forward letter to a person in Boston
  - Person identified by name, profession, and city

- Rule: Give letter only to people you know by first name and ask them to pass it on according to same rule
  - Some letters reached their goal

- Letter needed six steps on average to reach the person

- Graph theoretically: Social networks have dense local structure, but (apparently) small diameter
  - Generally referred to as "small world effect"
  - Usually, small number of persons act as "hubs"

Milgram's Small World Experiment

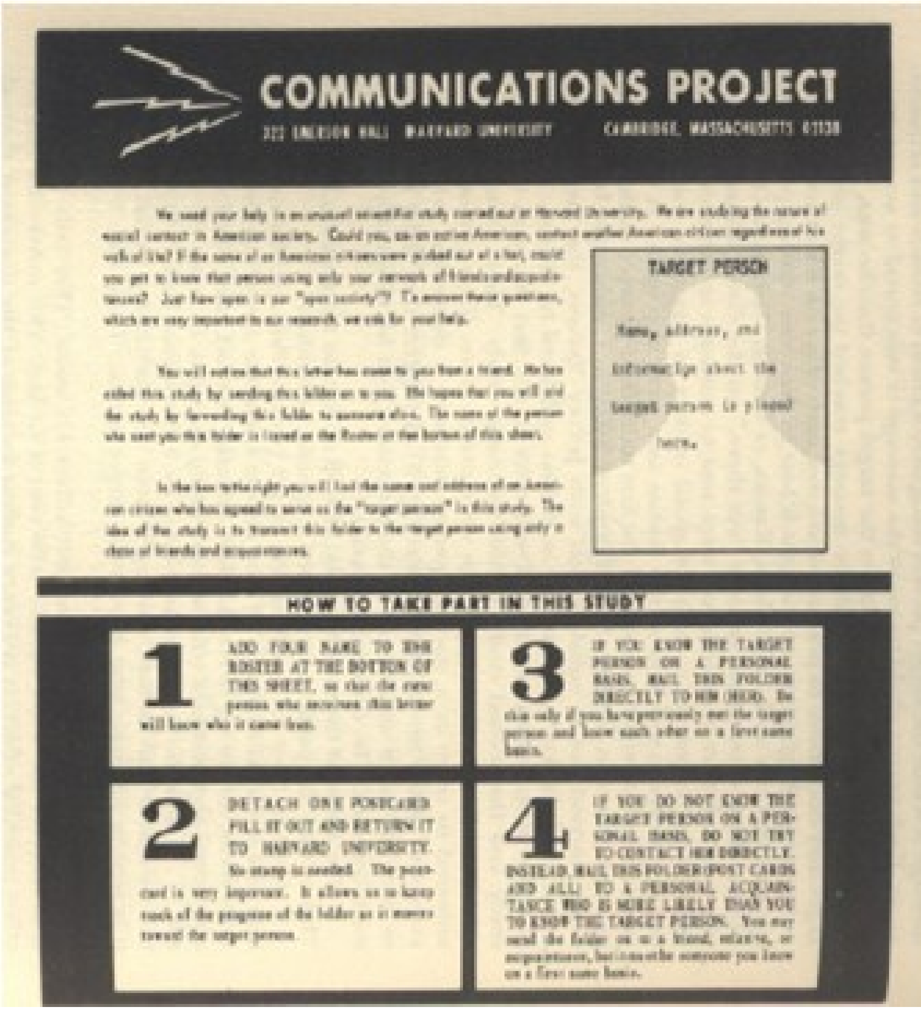

COMMUNICATIONS PROJECT

HARVARD UNIVERSITY — CAMBRIDGE, MASSACHUSETTS

TARGET PERSON

Name, address, and information about the target person is placed here.

HOW TO TAKE PART IN THIS STUDY

1 ADD YOUR NAME TO THE ROSTER AT THE BOTTOM OF THIS SHEET, so that the next person who receives this letter will know who it came from.

2 DETACH ONE POSTCARD. FILL IT OUT AND RETURN IT TO HARVARD UNIVERSITY. No stamp is needed. The postcard is very important. It allows us to keep track of the progress of the folder as it moves toward the target person.

3 IF YOU KNOW THE TARGET PERSON ON A PERSONAL BASIS, MAIL THIS FOLDER DIRECTLY TO HIM (HER). Do this only if you have previously met the target person and know each other on a first name basis.

4 IF YOU DO NOT KNOW THE TARGET PERSON ON A PERSONAL BASIS, DO NOT TRY TO CONTACT HIM DIRECTLY. INSTEAD, MAIL THIS FOLDER (POST CARDS AND ALL) TO A PERSONAL ACQUAINTANCE WHO IS MORE LIKELY THAN YOU TO KNOW THE TARGET PERSON. You may send the folder on to a friend, relative, or acquaintance, but it must be someone you know on a first name basis.

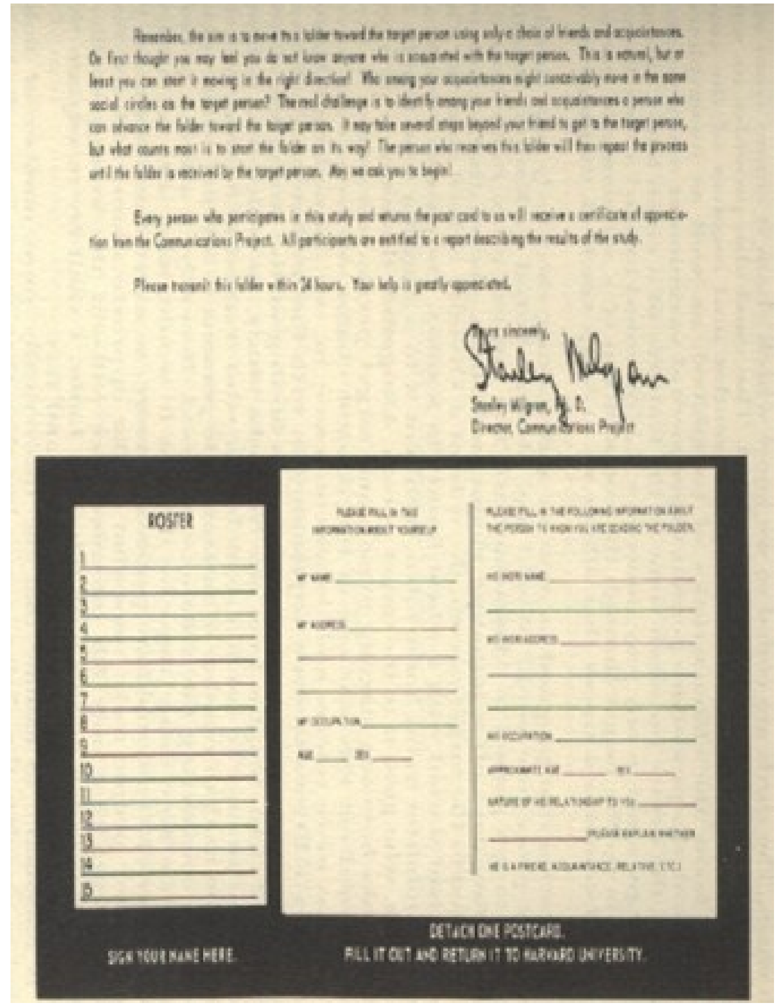

Every person who participates in this study and returns the post card to us will receive a certificate of appreciation from the Communications Project. All participants are entitled to a report describing the results of the study.

Please transmit this folder within 24 hours. Your help is greatly appreciated.

Yours sincerely,

Stanley Milgram, Ph. D.
Director, Communications Project

ROSTER

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15

SIGN YOUR NAME HERE.

DETACH ONE POSTCARD.
FILL IT OUT AND RETURN IT TO HARVARD UNIVERSITY.

Small-World Networks

- Developed/discovered by Watts and Strogatz (1998)
  - Over 30 years after Milgram's experiment!
- Watts and Strogatz looked at three networks
  - Film collaboration between actors, US power grid, Neural network of worm *Caenorhabditis elegans ("C. elegans")*
- Measured characteristics:
  - Clustering coefficient as a measure for 'regularity', or 'locality' of the network
    - If it is high, edges are rather connecting neighbors than nodes far apart
  - The average path length between vertices
- Results:
  - Grid-like networks:
    - High clustering coefficient $\Rightarrow$ high average path length (edges are not 'random', but rather 'local')
  - Most real-world (natural) networks have a **high clustering coefficient** (0.3-0.4), but nevertheless a **low average path length**

Small-World Networks and Random Graphs

- Results
  - Compared to a random graph with same number of nodes
  - Diameters similar, slightly higher for real graph
  - Clustering coefficient orders of magnitude higher
- Definition of small-worlds network
  - Dense local clustering structure and small diameter comparable to that of a same-sized random graph

	$D_\oslash(\text{real})$	$D_\oslash(\text{random})$	$C(\text{real})$	$C(\text{random})$
Film collaboration	3.65	2.99	0.79	0.00027
Power grid	18.7	12.4	0.08	0.005
C. elegans	2.65	2.25	0.28	0.05

Constructing Small-World Graphs

- Put all *n* nodes on a ring, number them consecutively from 1 to *n*
- Connect each node with its *k* clockwise neighbors
- Traverse ring in clockwise order
- For every edge
  - Draw random number *r*
  - If $r < p$, then re-wire edge by selecting a random target node from the set of all nodes (no duplicates)
  - Otherwise keep old edge
- Different values of *p* give different graphs
  - If *p* is close to 0, then original structure mostly preserved
  - If *p* is close to 1, then new graph is random
  - Interesting things happen when *p* is somewhere in-between

Regular, Small-World, Random

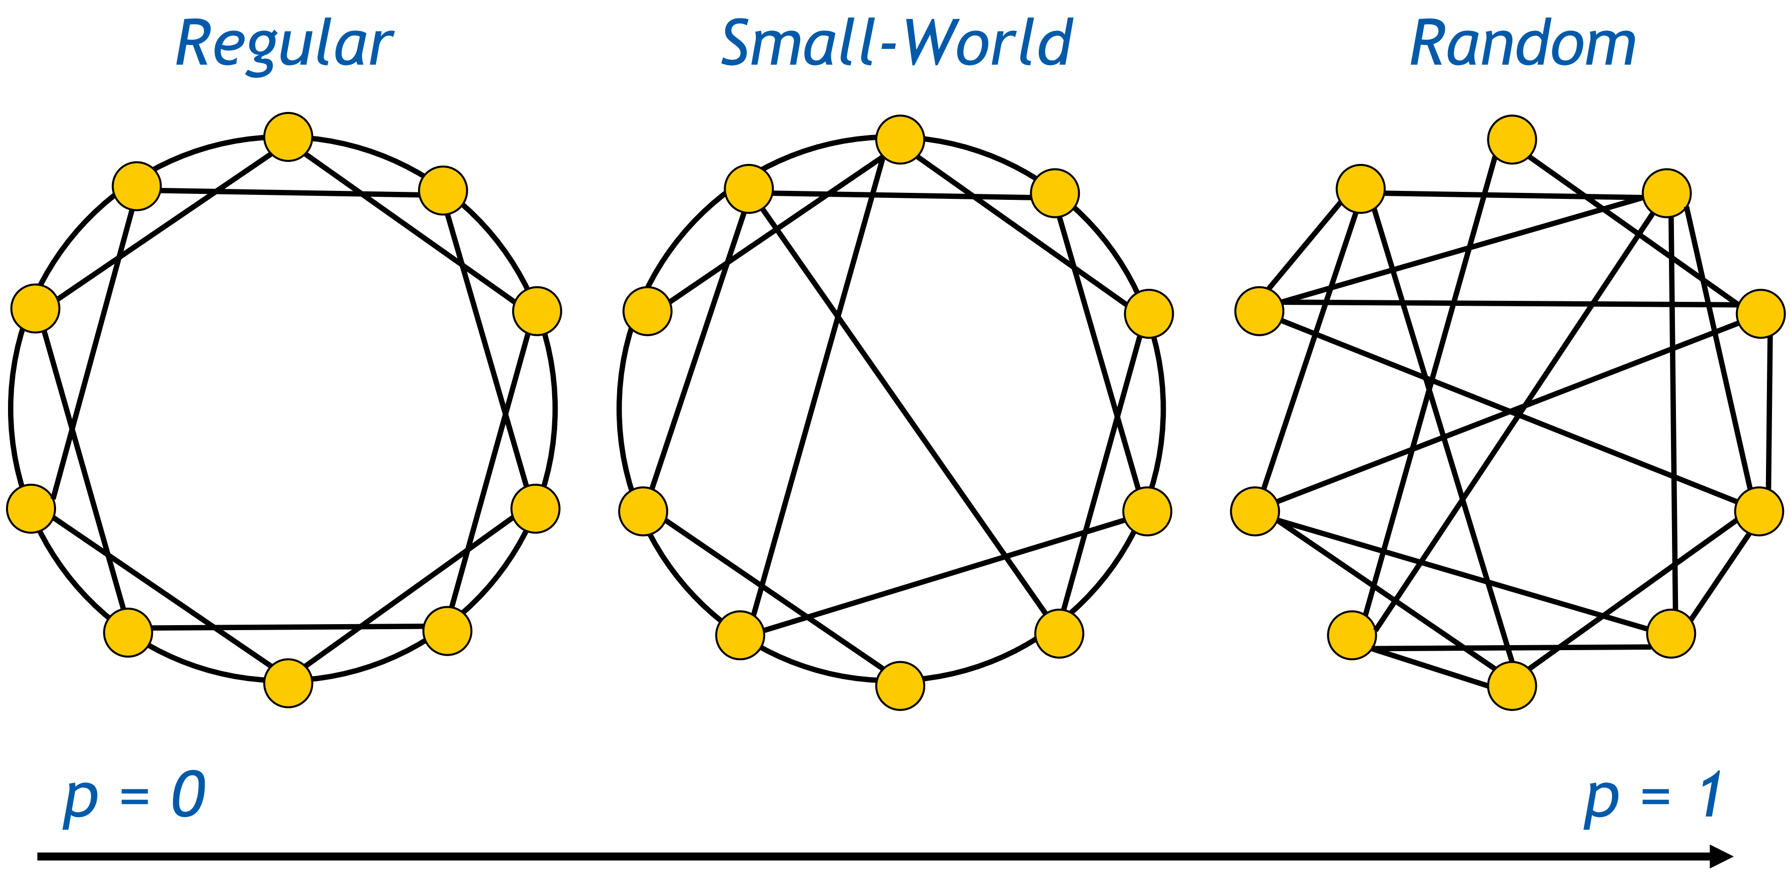

Kleinberg's Small-World Navigability Model

- Small-world model and power law explain why short paths exist

- Missing piece in the puzzle: why can we find these paths?
  - Each node has only local information
  - Even if a short cut exists, how do people know about it?
  - Milgram's experiment:
    - Some additional information (profession, address, hobbies etc.) is used to decide which neighbor is "closest" to recipient
    - results showed that first steps were the largest

- Kleinberg's Small-World Model
  - Set of points in an *n* x *n* grid
  - Distance is the number of "steps" separating points
    - $d(i, j) = |x_i - x_j| + |y_i - y_j|$

Kleinberg's Topologies

- Take *d*-dimensional grid in which all nodes are connected to all neighbors along each axis
- *Additionally* connect nodes in higher distance with probability decreasing with growing distance

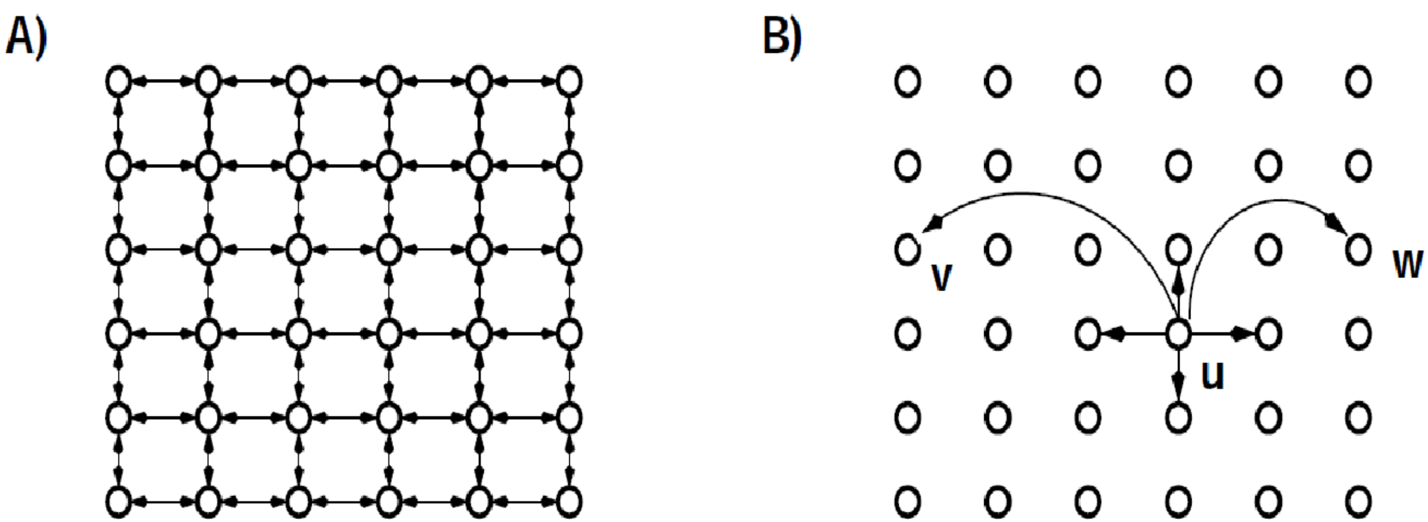

Creating Kleinberg's Topologies

- Construct graph as follows:
  - Every node *i* is connected to node *j* within distance *d'*
  - For every node *i*, additional *q* edges are added. Probability that node *j* is selected is proportional to $d(i, j)^{-r}$, for constant *r* (harmonic distribution *p*)

$$p_{sw}(v_i) = \frac{1}{x^r} + c$$

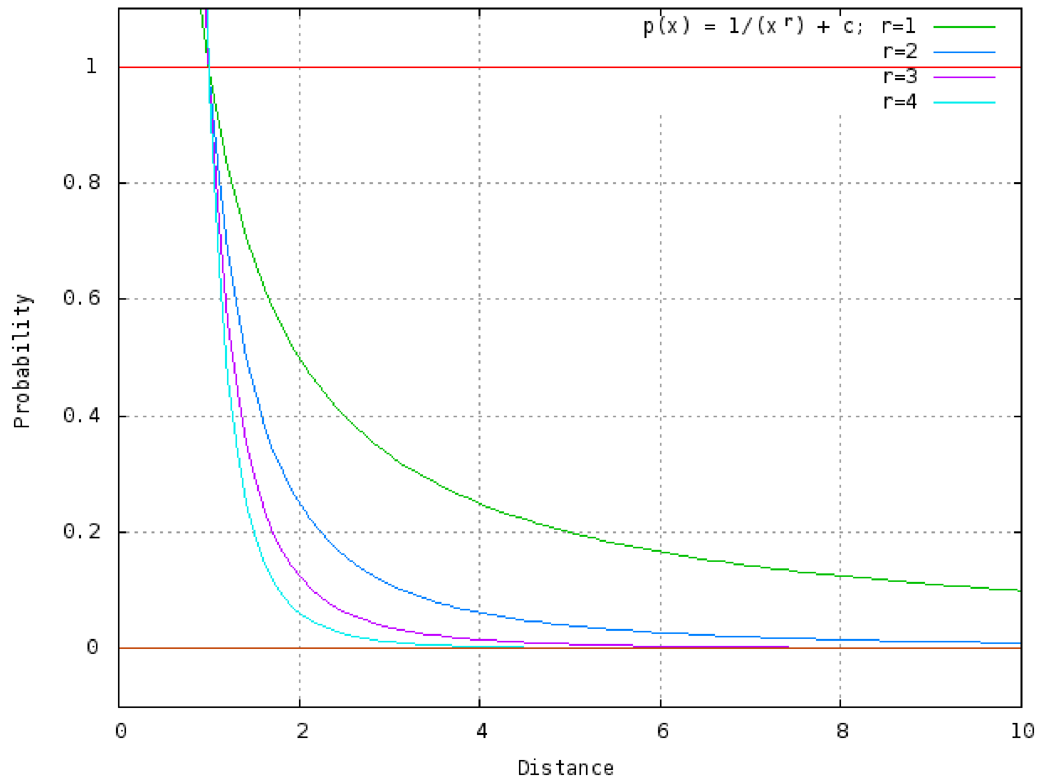

Navigation in Kleinberg's Model

- Simple greedy routing: nodes only know local links and target position, always use the link that brings message closest to target
  - If r=2, expected lookup time is $O(\log^2 n)$
  - If $r \neq 2$, expected lookup time is $O(n^{\varepsilon})$, where ε depends on r

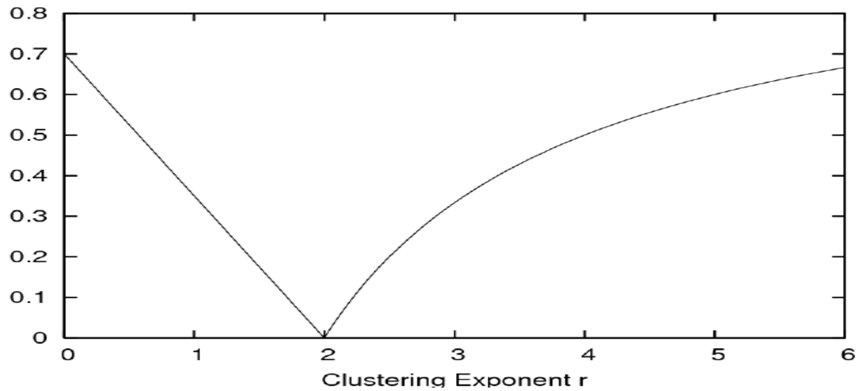

- Kleinberg has shown: Number of messages needed is proportional to

  $O(\log^2 n)$ iff r=s (s = number of dimensions)
  - Idea behind proof: for any r > s there are too few random edges to make paths short
  - For r < s there are too many random edges ⇒ too many choices for passing message
    - The message will make a (long) random walk through the network

Summary of Kleinberg

- Kleinberg small worlds thus provide a way of building a peer-to-peer overlay network, in which a very simple, greedy and local routing protocol is applicable
  - Practical algorithm: Forward message to contact who is closest to target
  - Assumes some way of associating nodes with points in grid (know about “closest”)
  - Compare with CAN

Problems with Small-World Graphs

Small-world graphs explain why:

- Highly clustered graphs can have short average path lengths ("short cuts")

Small-world graphs do *NOT* explain why:

- This property emerges in real networks
  - Real networks are practically never ring-like

Further problem with small-world graphs:

- Nearly all nodes have same degree
- Not true for random graphs
- What about real networks?

Internet

- Faloutsos et al. study from 99: Internet topology examined in 1998
  - AS-level topology, during 1998 Internet grew 45%

- Motivation:
  - What does the Internet look like?
  - Are there any topological properties that don’t change over time?
  - How to generate Internet-like graphs for simulations?

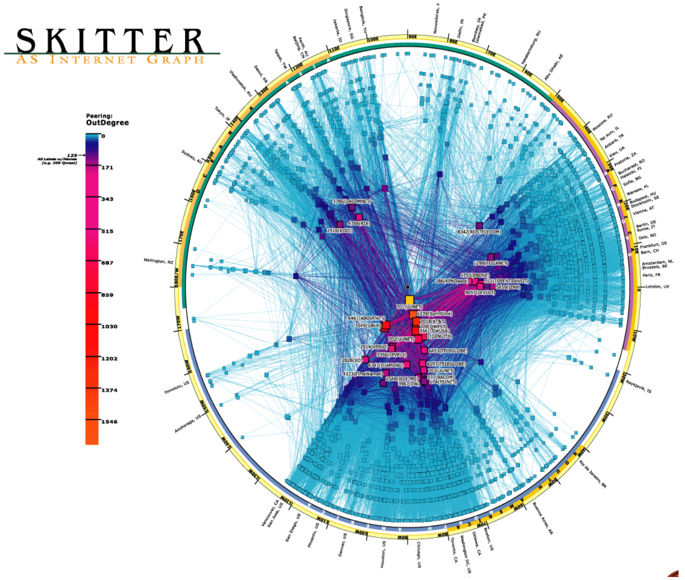

- 4 key properties found, each follows a power-law; Sort nodes according to their (out)degree
  1. *Outdegree of a node is proportional to its rank to the power of a constant*
  2. *Number of nodes with same outdegree is proportional to the outdegree to the power of a constant*
  3. *Eigenvalues of a graph are proportional to the order to the power of a constant*
  4. *Total number of pairs of nodes within a distance d is proportional to d to the power of a constant*

World Wide Web

- Links between documents in the World Wide Web
  - 800 Mio. documents investigated (S. Lawrence, 1999)

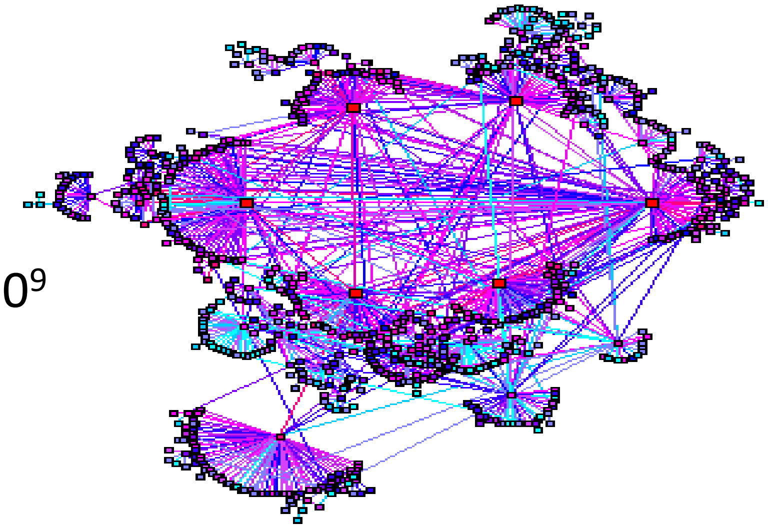

- What was expected so far?
  - Number of links per web page: $\langle k \rangle \sim 6$
  - Number of pages in the WWW: $N_{WWW} \sim 10^9$

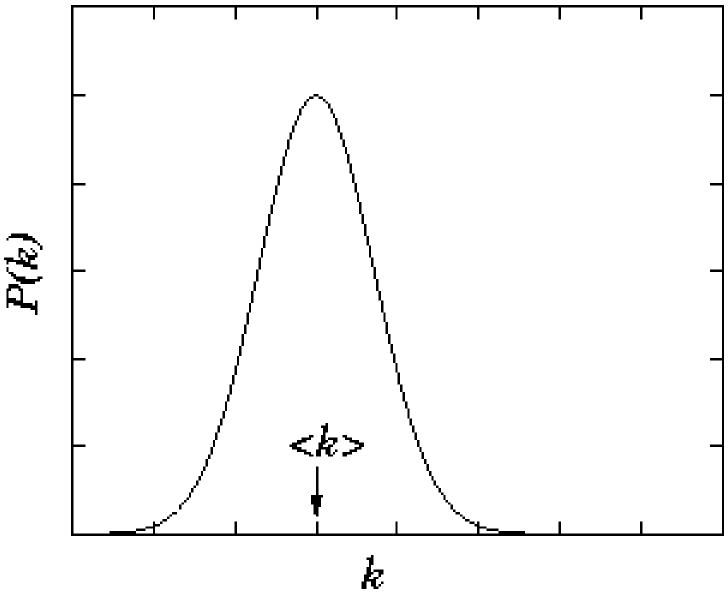

- *Probability "page has 500 links":* *$P(k=500) \sim 10^{-99}$*
- *Number of pages to which 500 links exist:* *$N(k=500) \sim 10^{-90}$*

WWW: result of investigation

P(page has k links)

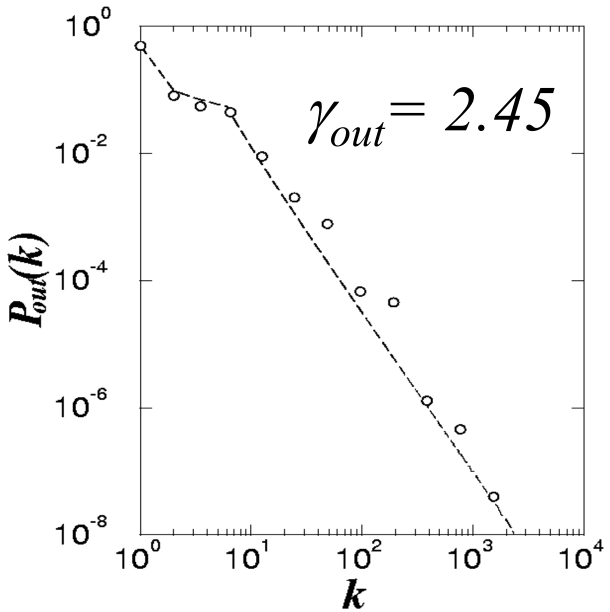

$$P_{out}(k) \sim k^{-\gamma_{out}}$$

P(k pages link to this page)

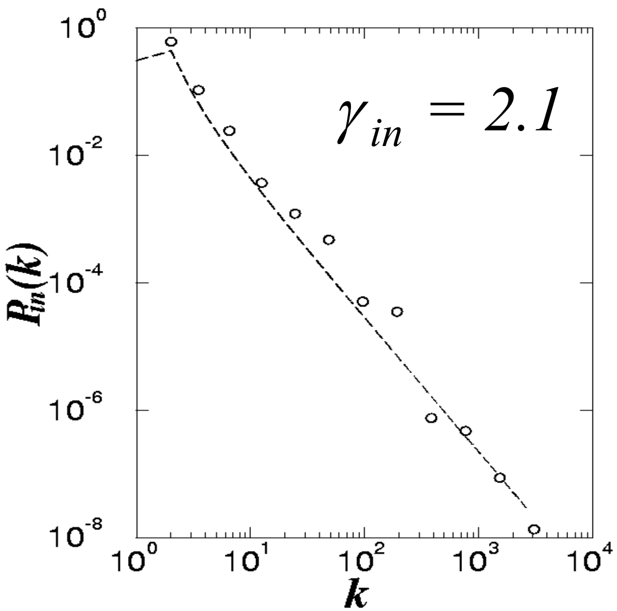

$$P_{in}(k) \sim k^{-\gamma_{in}}$$

$P(k=500) \sim 10^{-6}$

→ $N(k=500) \sim 10^{-3}$

Power Law Networks

- Also known as scale-free networks

- “Power Law” relationship for Web pages
  - The probability P(k) that a page has k links (or k other pages link to this page) is proportional to the number of links k to the power of y

- General ”Power Law” Relationships
  - A certain characteristic k is – independent of the growth of the system – always proportional to k^a, whereby a is a constant (often $-2 < a < -4$)

- Power laws very common (“natural”)
  - and power law networks exhibit small-world-effect
  - E.g. WWW: 19 degrees of separation
    (R. Albert et al, Nature (99); S. Lawrence et al, Nature (99))

Examples for Power Law Networks

- Economics
  - Pareto: income distribution
    (common simplification: 20% of population own 80% of the wealth)
  - Standardized price returns on individual stocks or stock indices
  - Sizes of companies and cities (Zipf's law)

- Human networks
  - professional (e.g. collaborations between actors, scientists)
  - social (friendship, acquaintances)
  - Sexual-contact networks

- Many other natural occurrences
  - Distribution of English words (Zipf's law again)
  - Areas burnt in forest fires
  - Meteor impacts on the moon

- Internet also follows some power laws
  - Popularity of Web pages (possibly related to Zipf's law for English words?)
  - Connectivity of routers and Autonomous Systems
  - Gnutella's topology!