Peer-to-Peer Systems

Analysis of unstructured P2P systems

Michael Welzl michael.welzl@uibk.ac.at

DPS NSG Team http://dps.uibk.ac.at/nsg
Institute of Computer Science
University of Innsbruck, Austria

Some questions…

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

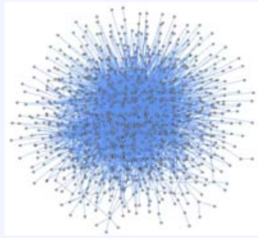

Gnutella snapshot, 2000

Scalability of Gnutella: quick answer

- Bandwidth Generated in Bytes (Message 83 bytes)
  - 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 open
- T = number of hops

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 $\{v, u\} = \{x, y\}$

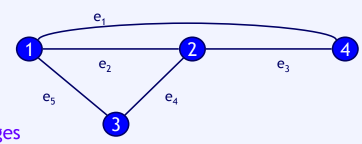

Example graph with
4 vertices and 5 edges

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), …, $(v_{k-1}, u) \in E$
- Graph G is connected if all v, $u \in V$ are connected
- Undirected, connected, acyclic graph is called a tree
  - Sidenote: Undirected, acyclic graph which is not connected is called a forest
- Directed, connected, acyclic graph is also called DAG
  - DAG = Directed Acyclic Graph (connected is assumed)
- An induced graph $G(V_C) = (V_C, E_C)$ is a graph $V_C \subseteq V$ and with edges $E_C = \{e = (i, j) \mid i, j \in V_C\}$
- An induced graph is a component if it is connected

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

Important Graph Metrics

- Distance: $d(v, u)$ between vertices v and u is the length of the shortest path between v and u
- Average path length: Sum of the distances over all pairs of nodes divided by the number of pairs
- Diameter: $d(G)$ of graph G is the maximum of $d(v, u)$ for all v, $u \in V$
- Order: the number of vertices in a graph
- Clustering coefficient: number of edges between neighbors divided by maximum number of edges between them

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

$E(N(i))$ = number of edges between neighbors of i
$d(i)$ = degree of i

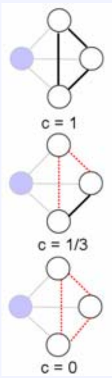

Source: Wikipedia

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 (we will use this)
- 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 will be one component of size $n*(f(c) + O(1))$ where $f(c) > 0$, with high probability. All other components have size $O(\log N)$

- In plain 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

- Before random graphs, regular graphs were popular
  - Regular: Every node has same degree
- Random graphs have two advantages over regular graphs
  1. Many interesting properties analytically solvable
  2. Much better for applications, e.g., social networks
- Note: Does not mean social networks are random graphs; just that the properties of social 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

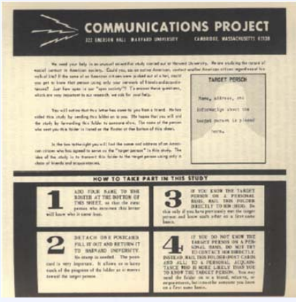
COMMUNICATIONS PROJECT
TARGET PERSON
HOW TO TAKE PART IN THIS STUDY

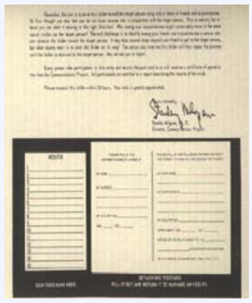

Small-World Networks

- 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 *C. elegans*
- Measured characteristics:
  - Clustering coefficient as a measure for 'regularity', or 'locality' of the network
    - If it is high, edges are rather build between neighbors than between far away nodes
  - The average path length between vertices
- Result
  - Most real-world networks have a **high clustering coefficient** (0.3-0.4), but nevertheless a **low average path length**
- Grid-like networks:
  - High clustering coefficient ⇒ high average path length (edges are not 'random', but rather 'local')

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

	$D_{\odot}$(real)	$D_{\odot}$(random)	C(real)	C(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

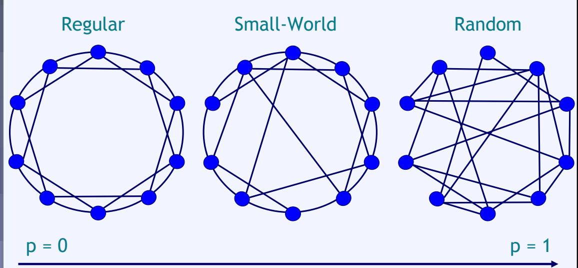

Problems with Small-World Graphs

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%

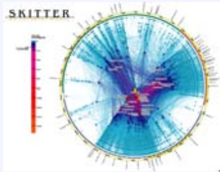

- Motivation:
  - What does the Internet look like?
  - Are there any topological properties that don't change over time?
  - How to generate Internet-like graphs for simulations?
- 4 key properties found, each follows a power-law;
  Sort nodes according to their (out)degree
  1. *Outdegree of a node 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: $N_{WWW} \sim 10^9$

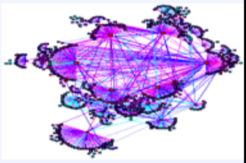

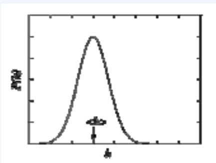

- Probability "page has 500 links": $P(k{=}500) \sim 10^{-99}$
- Number of pages with 500 links: $N(k{=}500) \sim 10^{-90}$

WWW: result of investigation

P(page has k links)

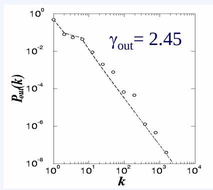

$P_{out}(k) \sim k^{-\gamma\ out}$

P(k pages link to this page)

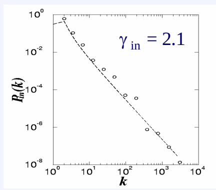

$P_{in}(k) \sim k^{-\gamma\ in}$

$P(k{=}500) \sim 10^{-6}$ $N_{WWW} \sim 10^9$ $\rightarrow N(k{=}500) \sim 10^3$

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)

Pareto, Zipf distributions are the same (converted)

- 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!

Barabasi-Albert-Model

- How do power law networks emerge?
  - In a network where new vertices (nodes) are added and new nodes tend to connect to well-connected nodes, the vertex connectivities follow a power-law
- Barabasi-Albert-Model: power-law network is constructed with two rules
  1. Network grows in time
  2. New node has preferences to whom it wants to connect
- Preferential connectivity modeled as
  - Each new node wants to connect to *m* other nodes
  - Probability that an existing node *j* gets one of the *m* connections is proportional to its degree *d(j)*
- New nodes tend to connect to well-connected nodes
- Another way of saying this: "the rich get richer"

Copying model

- Alternative generative model *(R. Kumar, P. Raghavan, et al. 2000)*
  - In each time step randomly copy one of the existing nodes keeping all its links
  - Connect the original node and the copy
  - Then randomly remove edges from both nodes with a very small probability, and for each removed edge randomly draw new target nodes
- In this model the probability of node *v* getting a new edge in some time step is proportional to its degree at that time
  - The more edges it has, the more likely it is that one of its neighbors is chosen for copying in the next time step
- In contrast to random networks, scale-free networks show a small number of well-connected hubs and many nodes with few connections

(a) Random network (b) Scale-free network

Robustness of Scale Free Networks

- Experiment: take network of 10000 nodes (random and power-law) and remove nodes randomly
- Random graph:
  - Take out 5% of nodes: Biggest component 9000 nodes
  - Take out 18% of nodes: No biggest component, all components between 1 and 100 nodes
  - Take out 45% of nodes: Only groups of 1 or 2 survive
- Power-law graph:
  - Take out 5% of nodes: Only isolated nodes break off
  - Take out 18% of nodes: Biggest component 8000 nodes
  - Take out 45% of nodes: Large cluster persists, fragments small
- Networks with power law exponent < 3 are very robust against random node failures
  - ONLY true for random failures!

Robustness of Scale-Free Networks /2

- Robustness against random failures = important property of networks with scale-free degree distribution
  - Remove a randomly chosen vertex v from a scale-free network: with high probability, it will be a low-degree vertex and thus the damage to the network will not be high
- But scale-free networks are very sensitive against attacks
  - If a malicious attacker removes the highest degree vertices first, the network will quickly decompose in very small components
- Note: random graphs are not robust against random failures, but not sensitive against attacks either (because all vertices more or less have the same degree)

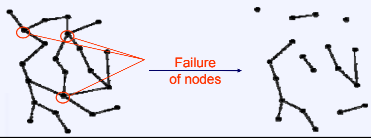

Robustness of Scale-Free Networks /3

- Random failures vs. directed attacks

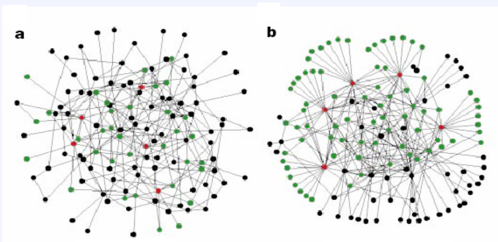

Random Graph „Power Law" Graph

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 $n \times n$ grid
  - Distance is the number of "steps" separating points
    - $d(i, j) = |x_i - x_j| + |y_i - y_j|$
- Construct graph as follows:
  - Every node *i* is connected to node *j* within distance *q*
  - For every node *i*, additional *q* edges are added. Probability that node *j* is selected is proportional to $d(i, j)^{-r}$, for some constant *r*

Navigation in Kleinberg's Model

- Simple greedy routing
  - If r=2, expected lookup time is $O(\log^2 n)$
  - If $r \neq 2$, expected lookup time is $\Omega(n^{\varepsilon})$, ε depends on r
- Decentralized algorithm: sending node only knows its local neighbors, position of the target node on the grid, locations and long-range contacts of all nodes who come in contact of the message
- Can be shown: Number of messages needed is proportional to *O(log 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
- 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 DHT (later)

Unstructured P2P Networks

- What do real (unstructured) Peer-to-Peer Networks look like?
- Depends on the protocols used
  - It has been found that some peer-to-peer networks, e.g., Freenet, evolve voluntarily in a small-world with a high clustering coefficient and a small diameter
  - Analogously, some protocols, e.g., Gnutella, will implicitly generate a scale-free degree distribution
- Case study: Gnutella network
- How does the Gnutella network evolve?
  - Nodes with high degree answer more likely to Ping messages
  - Thus, they are more likely chosen as neighbor
  - Host caches always/often provide addresses of well connected nodes

Gnutella

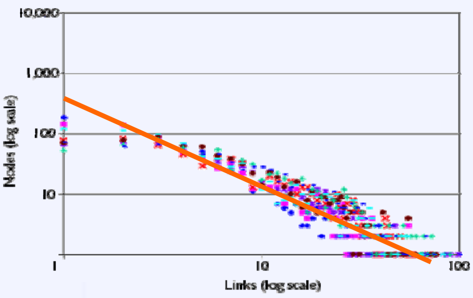

Node degrees in Gnutella follow Power-Law rule

Gnutella /2

- Network diameter stayed nearly constant, though the network grew by one order of magnitude
- Robustness
  - Remember: we said that networks with power law exponent < 3 are very robust against random node failures
  - In Gnutella's case, the exponent is 2.3
- Theoretical experiment
  - Subset of Gnutella with 1771 nodes
  - Take out random 30% of nodes, network still survives
  - Take out 4% of best connected nodes, network splinters
- For more information on Gnutella, see:
  - Matei Ripeanu, Adriana Iamnitchi, Ian Foster: Mapping the Gnutella Network, IEEE Internet Computing, Jan/Feb 2002
  - Zeinalipour-Yazti, Folias, Faloutsos, "A Quantitative Analysis of the Gnutella Network Traffic", Tech. Rep. May 2002

Summary

- The network structure of a peer-to-peer system influences:
  - average necessary number of hops (path length)
  - possibility of greedy, decentralized routing algorithms
  - stability against random failures
  - sensitivity against attacks
  - redundancy of routing table entries (edges)
  - many other properties of the system build onto this network
- Important measures of a network structure are:
  - average path length
  - clustering coefficient
  - the degree distribution
- Next: how to build systems based on edge generation rules such that a network structure arises supporting the desired properties of the system

References / acknowledgments

- Slides from:
  - Jussi Kangasharju
  - Christian Schindelhauer
  - Klaus Wehrle