Peer-to-Peer Networks

02: Napster & Gnutella

Christian Schindelhauer
Technical Faculty
Computer-Networks and Telematics
University of Freiburg

CoNe Freiburg

Napster

- Shawn (Napster) Fanning
  - published 1999 his beta version of the now legendary Napster P2P network
  - File-sharing-System
  - Used as mp3 distribution system
  - In autumn 1999 Napster has been called download of the year
- Copyright infringement lawsuit of the music industry in June 2000
- End of 2000: cooperation deal
  - between Fanning and Bertelsmann Ecommerce
- Since then Napster is a commercial file-sharing platform

UNI FREIBURG

How Did Napster Work?

- Client-Server
- Server stores
  - Index with meta-data
    - file name, date, etc
  - table of connections of participating clients
  - table of all files of participants
- Query
  - client queries file name
  - server looks up corresponding clients
  - server replies the owner of the file
  - querying client downloads the file from the file owning client

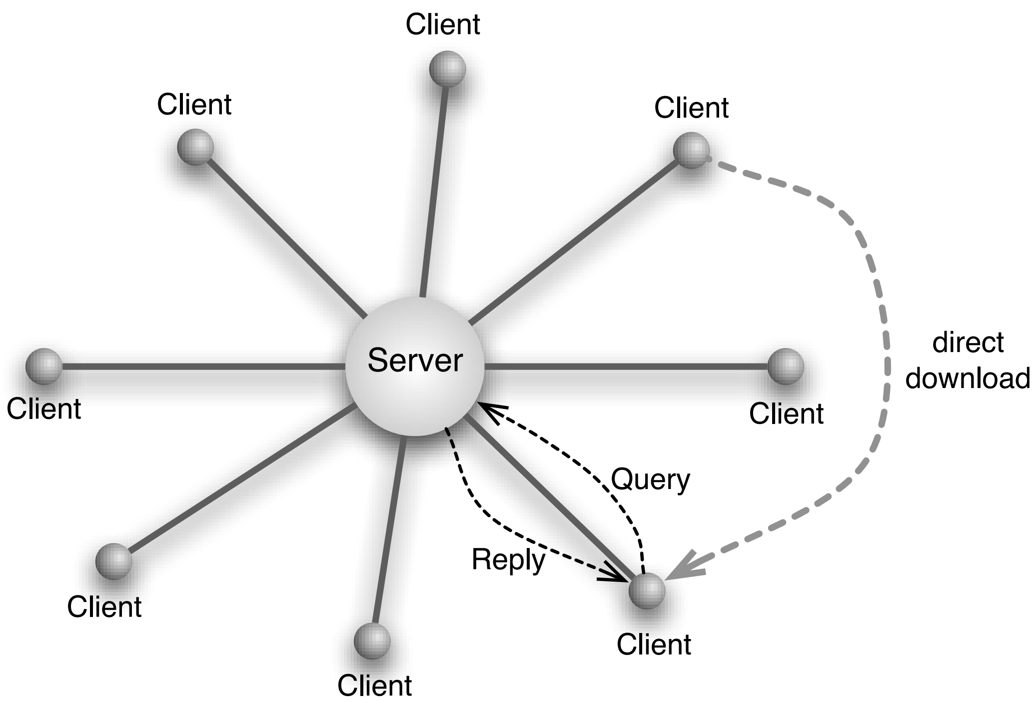

Discussion of Napster

- Advantages
  - Napster is simple
  - Files can be found fast and effective
- Disadvantages
  - Central structure eases censorship, hostile attacks and vulnerability against technical problems
    - e.g. denial of service (DOS) attack
  - Napster does not scale
    - i.e. increasing number of participants implies a decline in performance
    - bandwidth and memory of the server is limited
- Conclusion
  - Napster is not an acceptable P2P network solution
  - Except the download part Napster is not a real P2P network

History of Gnutella

- Gnutella
  - was released in March 2000 by Justin Frankel and Tom Pepper from Nullsoft
  - Since 1999 Nullsoft is owned by AOL
- File-Sharing system
  - Same goal as Napster
  - But without any central structures

Gnutella — Connecting

- Neighbor lists
  - Gnutella connects directly with other clients
  - the client software includes a list of usually online clients
  - the clients checks these clients until an active node has been found
- an active client publishes its neighbor list
  - the query (ping) is forwarded to other nodes
  - the answer (pong) is sent back
  - neighbor lists are extended and stored
  - the number of the forwarding is limited (typically: five)

Gnutella — Connecting

- Protokoll
  - Ping
    - participants query for neighbors
    - are forwarded according for TTL steps (time to live)
  - Pong
    - answers Ping
    - is forwarded backward on the query path
    - reports IP and port adress (socket pair)
    - number and size of available files

Gnutella — Connecting

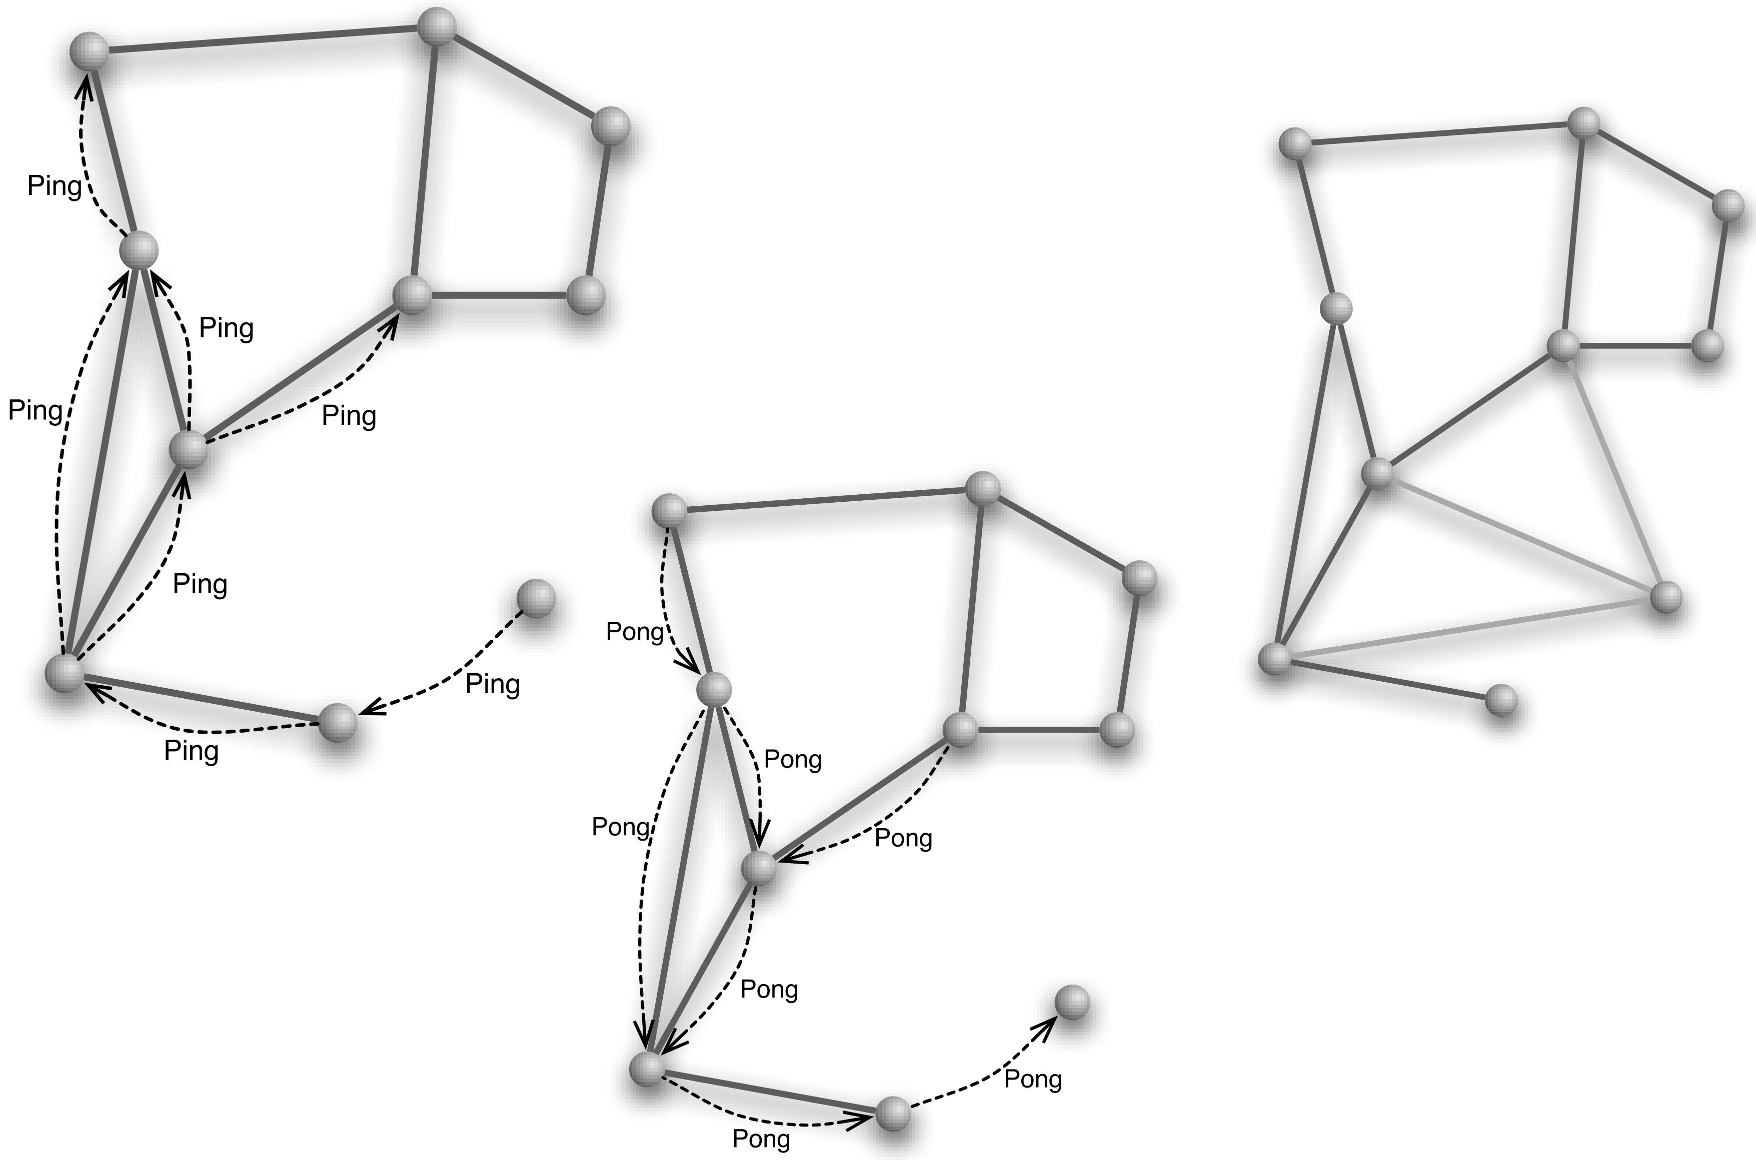

Gnutella — Graph Structure

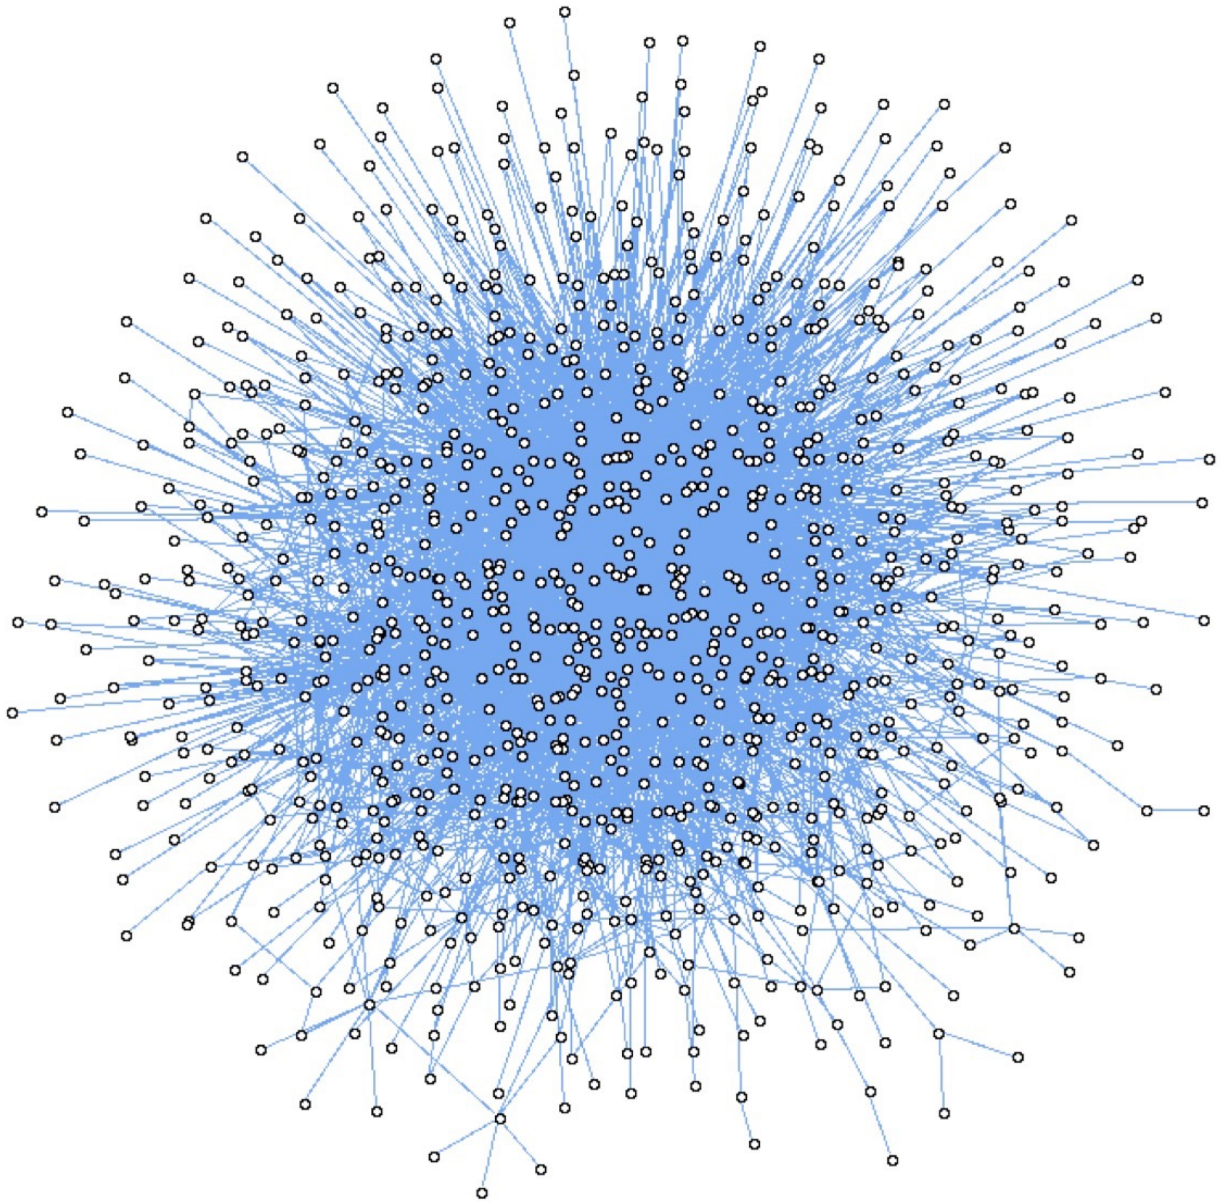

shot in 2000

Gnutella — Graph Structure

- Graph structure
  - constructed by random process
  - underlies power law
  - without control

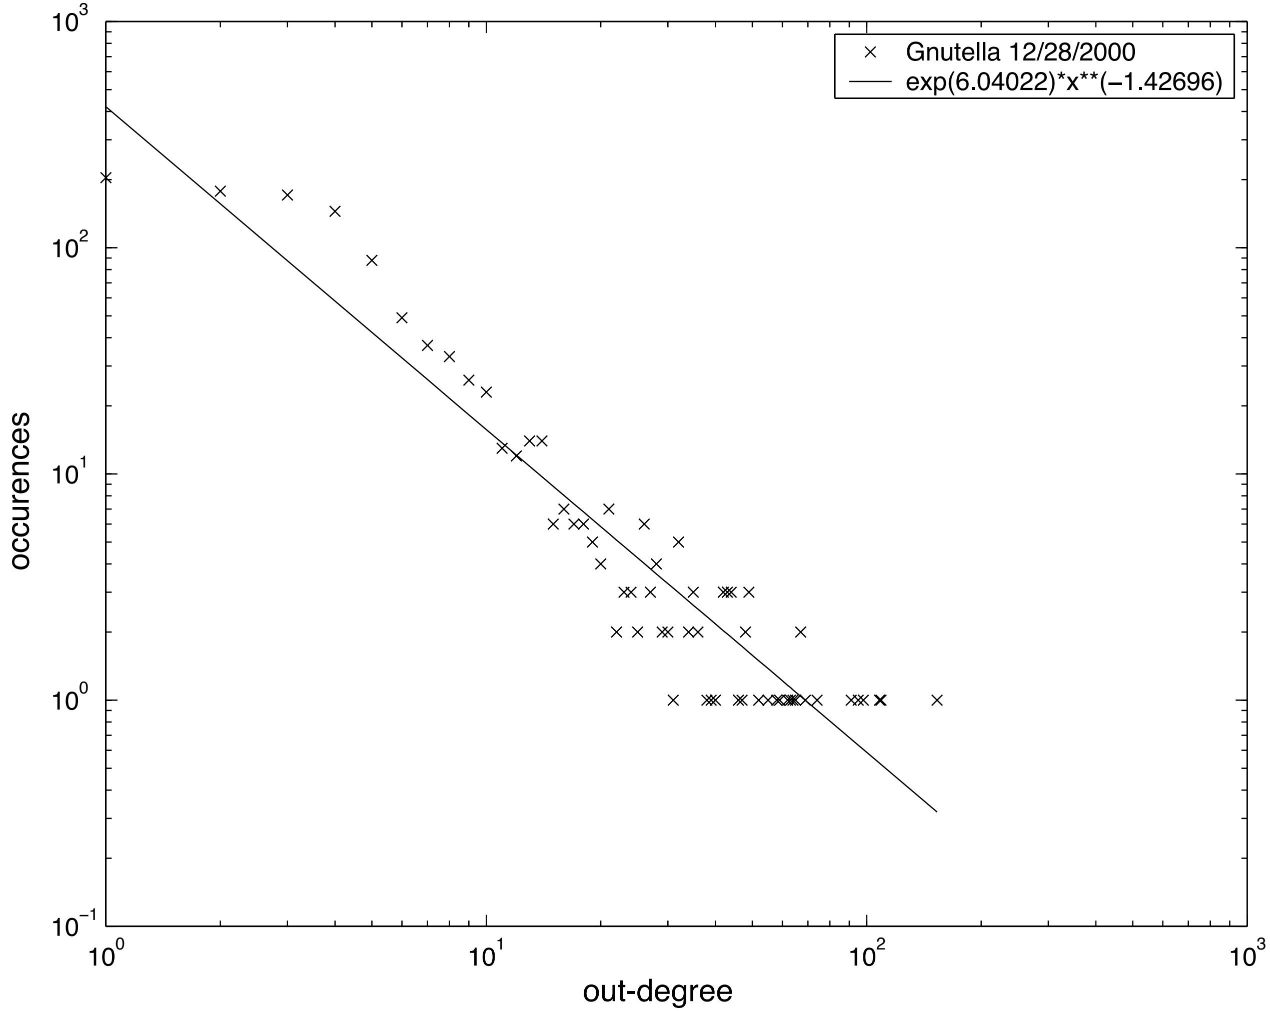

Gnutella snapshot in 2000

Gnutella — Query

- File Query
  - are sent to all neighbors
  - Neighbors forward to all neighbors
  - until the maximum hop distance has been reached
    - TTL-entry (time to live)
- Protocol
  - Query
    - for file for at most TTL hops
  - Query-hits
    - answers on the path backwards
- If file has been found, then initiate direct download

Gnutella — Query

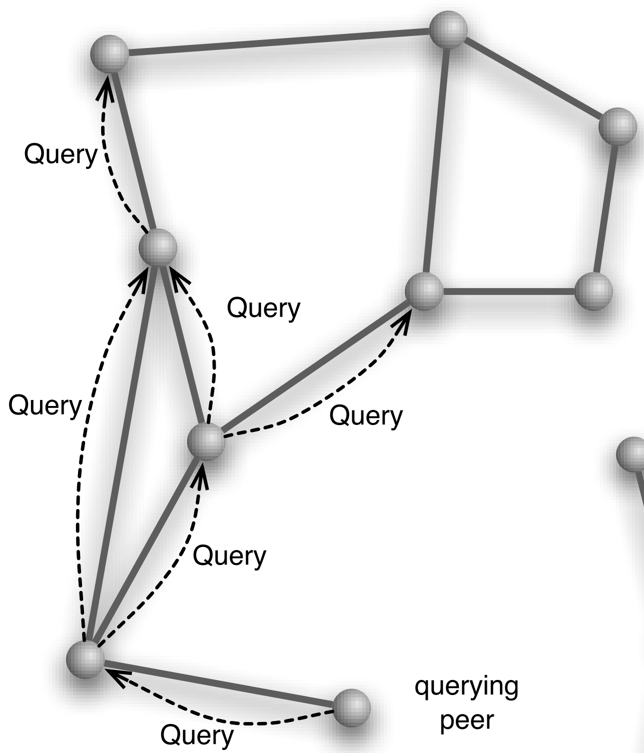

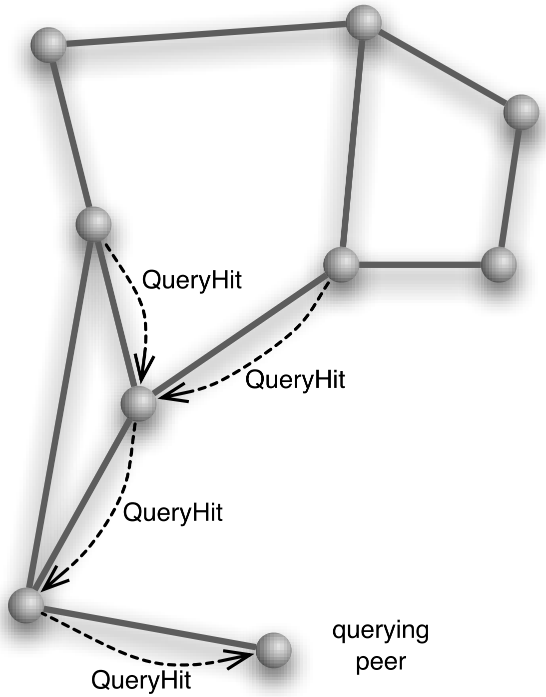

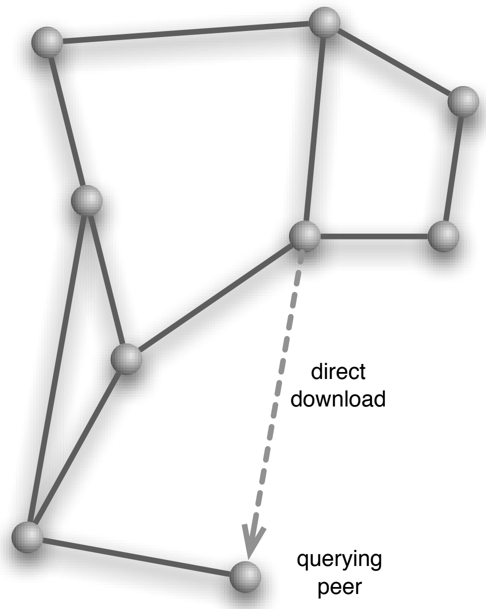

Gnutella - Discussion

- Advantages
  - distributed network structure
  - scalable network
- Disadvantages
  - bounded breadth depth search leads to implizit network partition
  - this reduces success probability
  - long paths, slow latency
- Suggested improvements
  - random walks instead broadcasting
  - passive replication of index information

FastTrack & Gnutella2

- Hybrid Structure
  - high bandwidth node are elected as P2P-servers, aka. super-nodes
  - super-nodes are connected using the original Gnutella protocol
  - client nodes are connected only to super-nodes
- Used in
  - FastTrack
  - Gnutella 2
- Advantages
  - improved scalabilty
  - smaller latency
- Disadvantages
  - still unreliable and slow
  - peers decline to serve as super-nodes

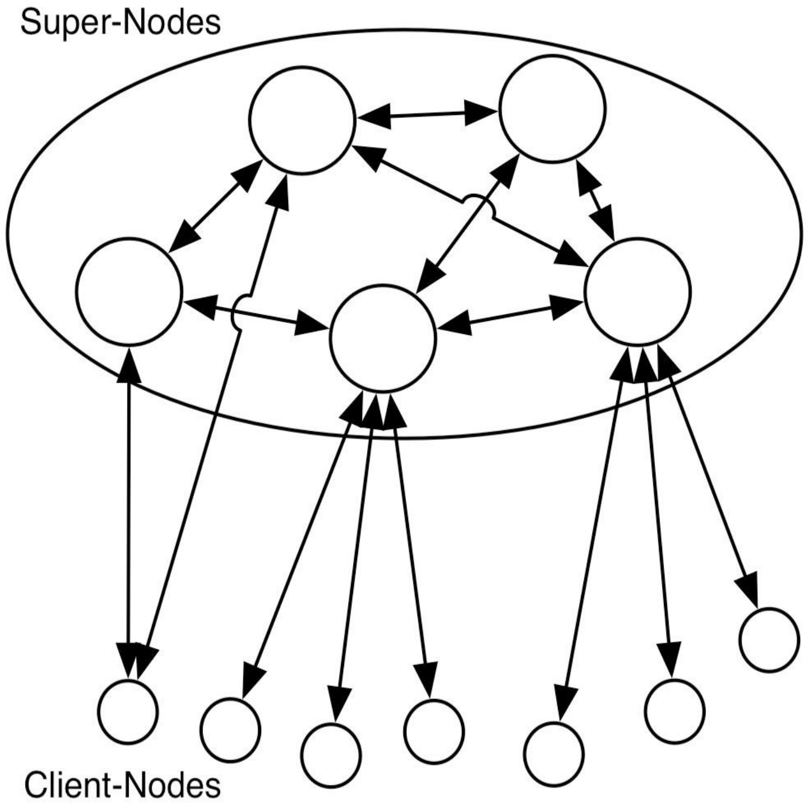

Degree Distribution in Gnutella

- Modeling Large-scale Peer-to-Peer Networks and a Case Study of Gnutella
  - Mihajlo A. Jovanovic, Master Thesis, 2001
- The number of neighbors is distributed according a power law (Pareto) distribution
  - log(#peers with degree d) = c - k log d
  - #peers with degree d = C/d^k

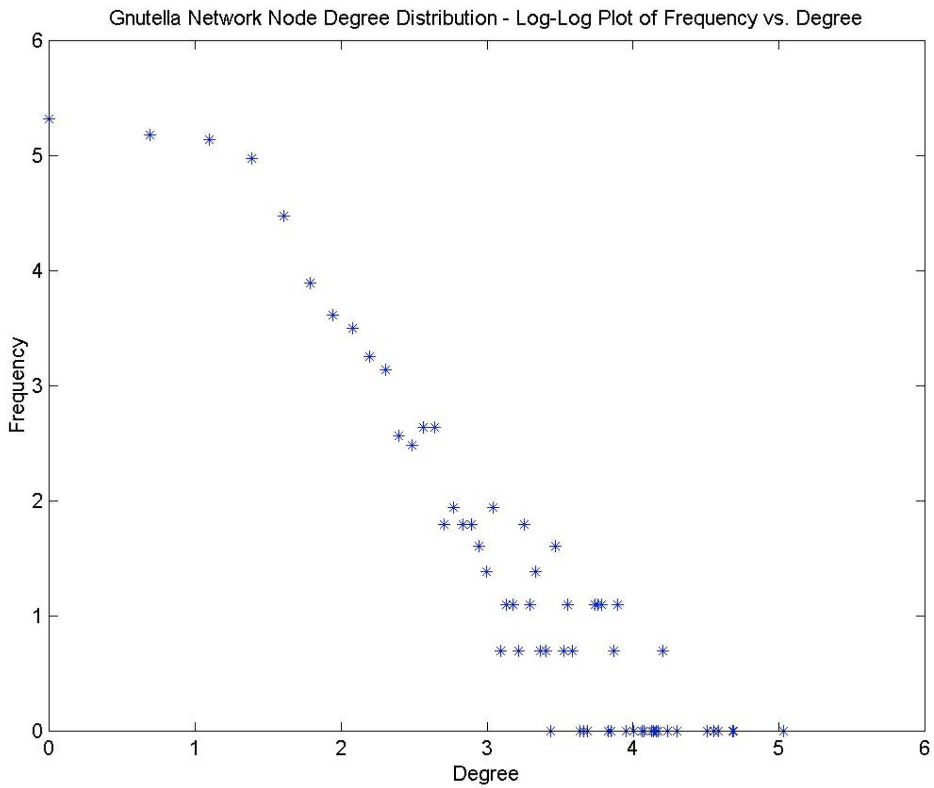

Pareto-Distribution Examples

- Pareto 1897: Distribution of wealth in the population
- Yule 1944: frequency of words in texts
- Zipf 1949: size of towns
- length of molecule chains
- file length of Unix-system files
- ….

Pareto Verteilung

CoNe Freiburg

- Discreet Pareto-Distribution for $x \in \{1,2,3,\ldots\}$

$$\mathbf{P}[X = x] = \frac{1}{\zeta(\alpha) \cdot x^\alpha}$$

  - with constant factor

$$\zeta(\alpha) = \sum_{i=1}^{\infty} \frac{1}{i^\alpha}$$

  - (also known as Riemann´s Zeta-function)
- Heavy tail property
  - not all moments $E[X^k]$ exist
  - the expectation exists if and only if (iff) $\alpha>2$
  - variance and $E[X^2]$ exist iff $\alpha>3$
  - $E[X^k]$ exists iff $\alpha>k+1$
- Density function of the continuous function for $x>x_0$

$$f(x) = \frac{\alpha - 1}{x_0} \left(\frac{x_0}{x}\right)^\alpha$$

UNI FREIBURG

Indegree and Outdegree of Web-Pages

- are described by a power law (Pareto) distribution

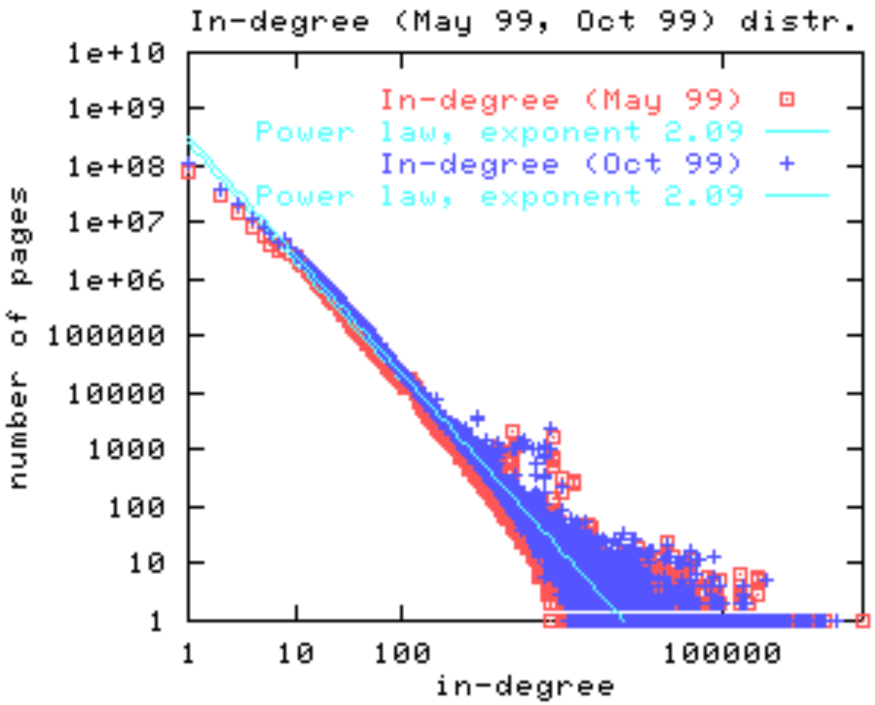

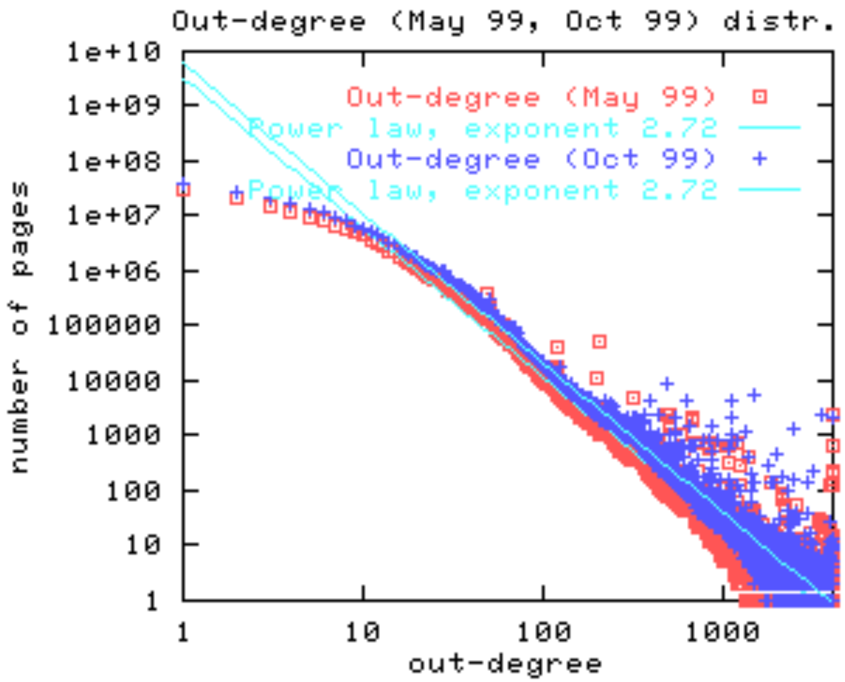

- Experiments of
  - Kumar et al 97: 40 millions Webpages
  - Barabasi et al 99: Domain *.nd.edu + Web-pages in distance 3
  - Broder et al 00: 204 millions web pages (Scan Mai und Okt. 1999)

Connectivity of Pareto Graphs

- William Aiello, Fan Chung, Linyuan Lu, A Random Graph Model for Massive Graphs, STOC 2000
- Undirected graph with n nodes where
  - the probability of k neighbors for a node is p_k
  - where $p_k = c\, k^{-\tau}$ for some normalizing factor c
- Theorem
  - For sufficient large n such Pareto-Graphs with exponent τ we observe
    - for $\tau < 1$ the graph is connected with probability 1-o(1)
    - for $\tau > 1$ the graph is nont connected with probability 1-o(1)
    - for $1 < \tau < 2$ there is a connected component of size $\Theta(n)$
    - for $2 < \tau < 3.4785$ there is only one connected component of size $\Theta(n)$ and all others have size $O(\log n)$
    - for $\tau > 3.4785$: there is no large connected component of size $\Theta(n)$ with probability 1-o(1)
    - For $\tau > 4$: no large connected components which size can be described by a power law (Pareto) distribution

Zipf Distribution as a Variant of Power Laws

- George Kinsley Zipf claimed
  - that the frequency of the n most frequent word f(n)
  - satisfies the equation n f(n) = c.
- Zipf probability distribution for x ∈ {1,2,3,…}

$$\mathbf{P}[X = x] = \frac{c}{x}$$

  - with constant factor c only defined for connstant sized sets, since

$$\ln n \leq \sum_{i=1}^{n} \frac{1}{i} \leq 1 + \ln n$$

  - is unbounded
- Zipf distribution relate to the rank
  - The Zipf exponent α may be larger than 1, i.e. $f(n) = c/n^{\alpha}$
- Pareto distribution realte the absolute size, e.g. the number of inhabitants

Size of towns
Scaling Laws and Urban Distributions, Denise Pumain, 2003

Figure 1 The hierarchical differentiation in urban systems:
Rank-size distribution of French agglomerations (1831-1999)

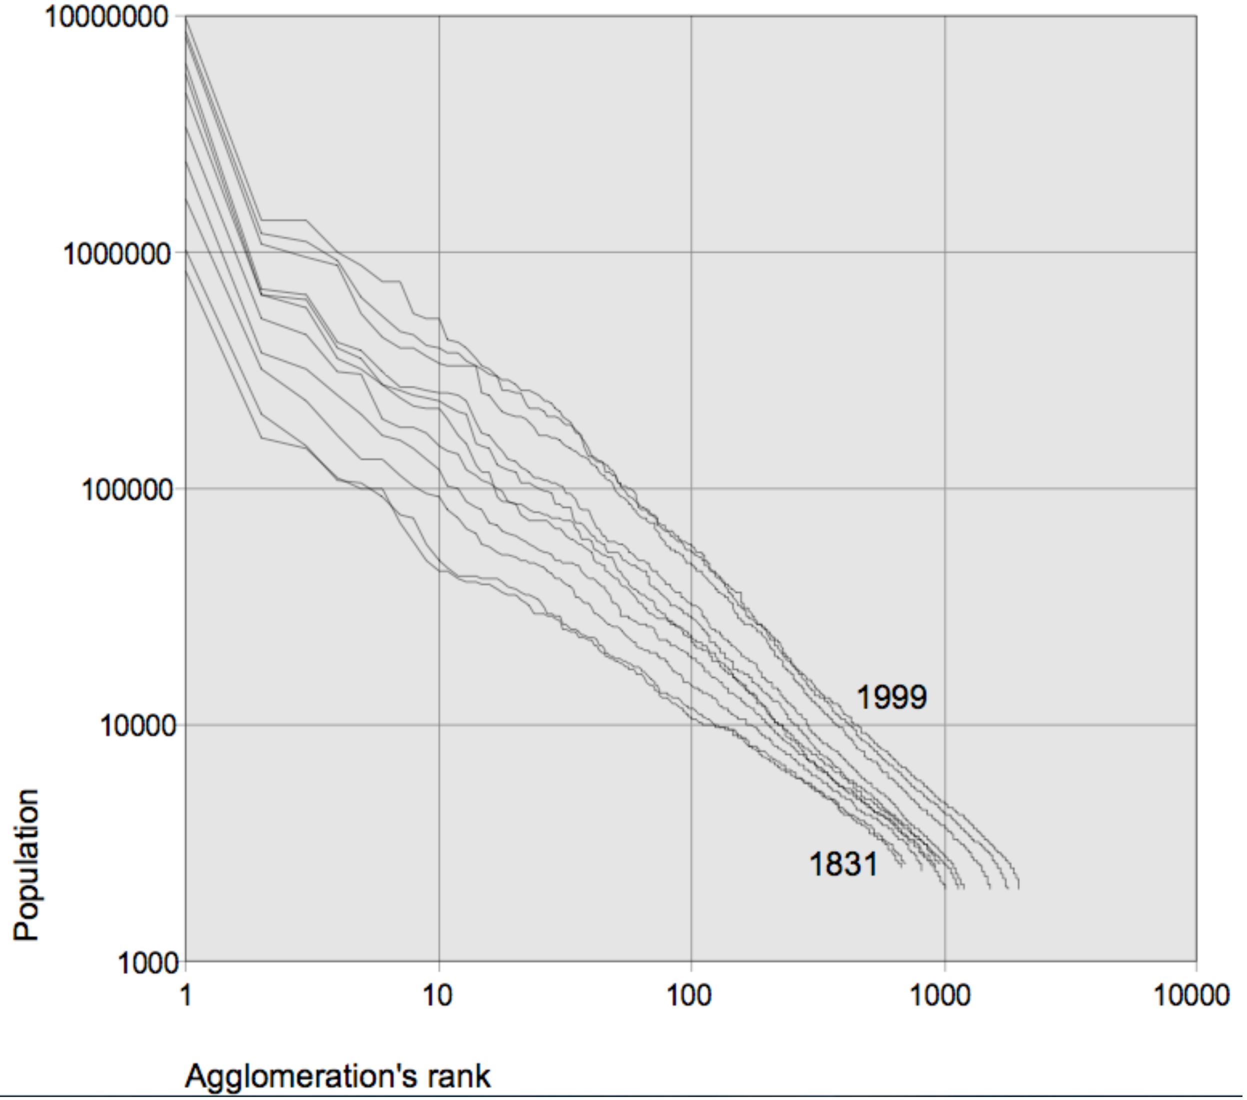

Zipf distribution

Zipf's Law and the Internet

Lada A. Adamic, Bernardo A. Huberman, 2002

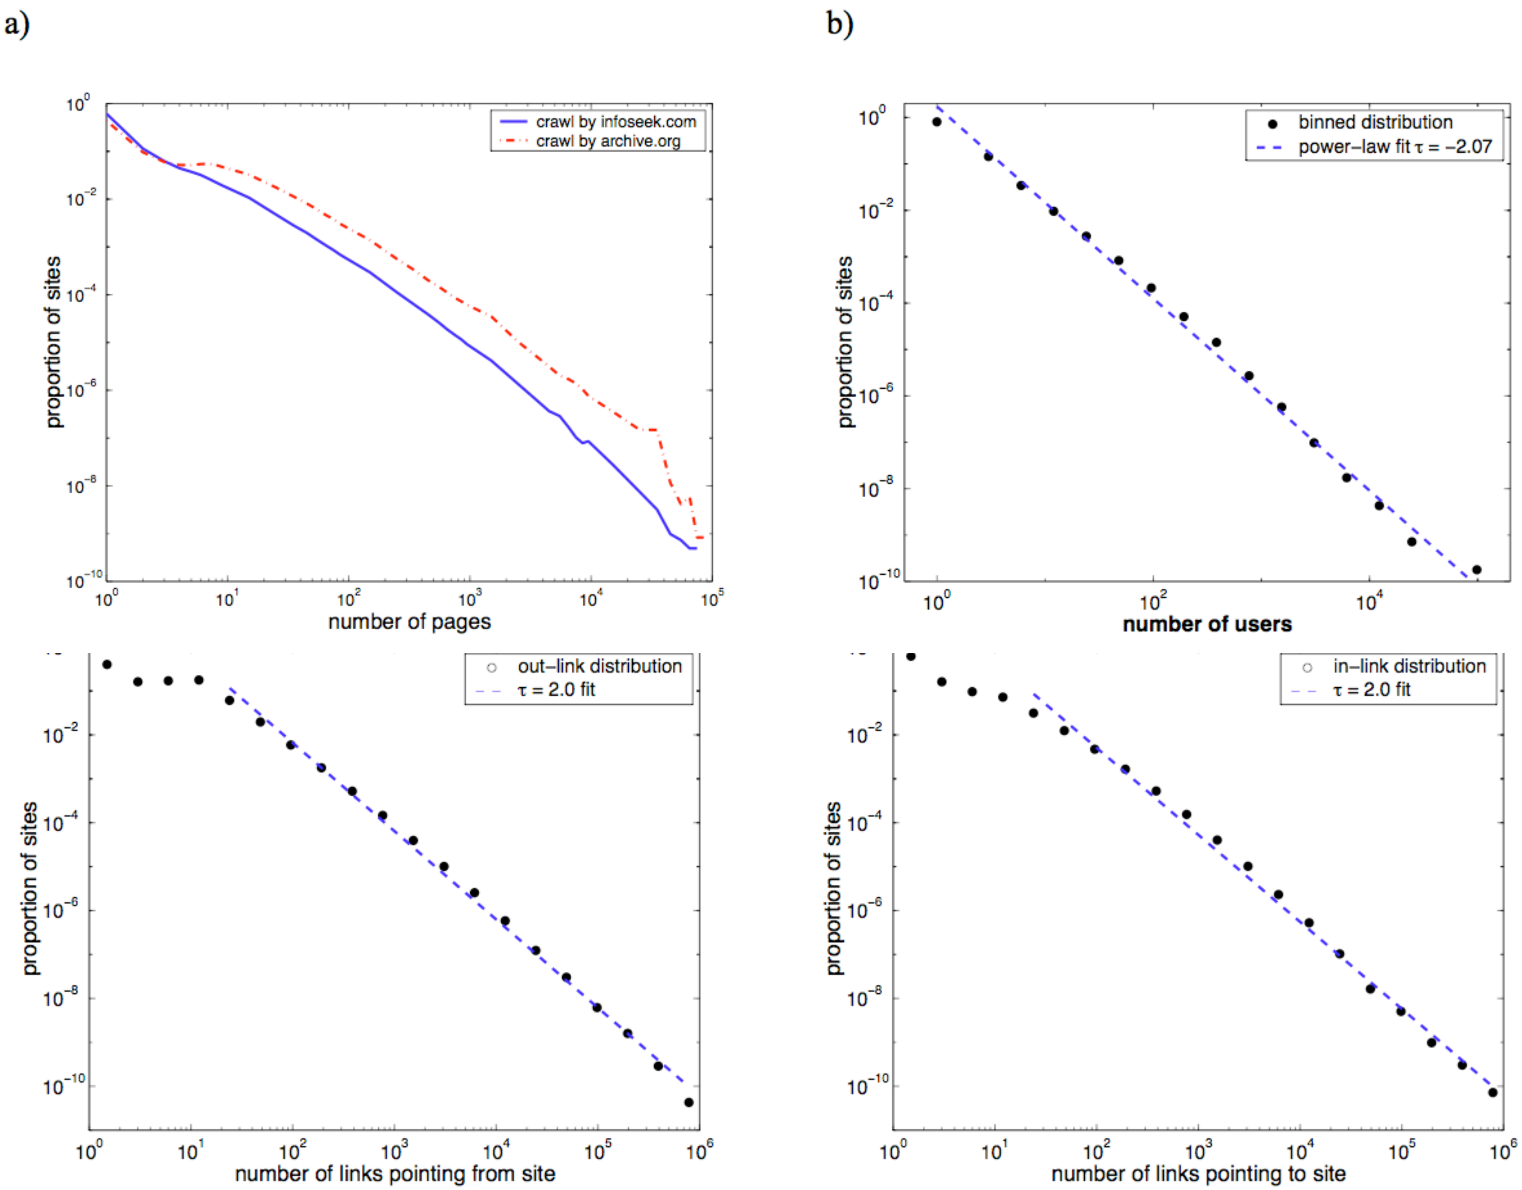

Figure 1. Fitted power law distributions of the number of site a) pages, b) visitors, c) out links, and d) in links, measured in 1997.

Pareto Distribution!!

Heavy-Tailed Probability Distributions in the World Wide Web
Mark Crovella, Murad, Taqqu, Azer Bestavros, 1996

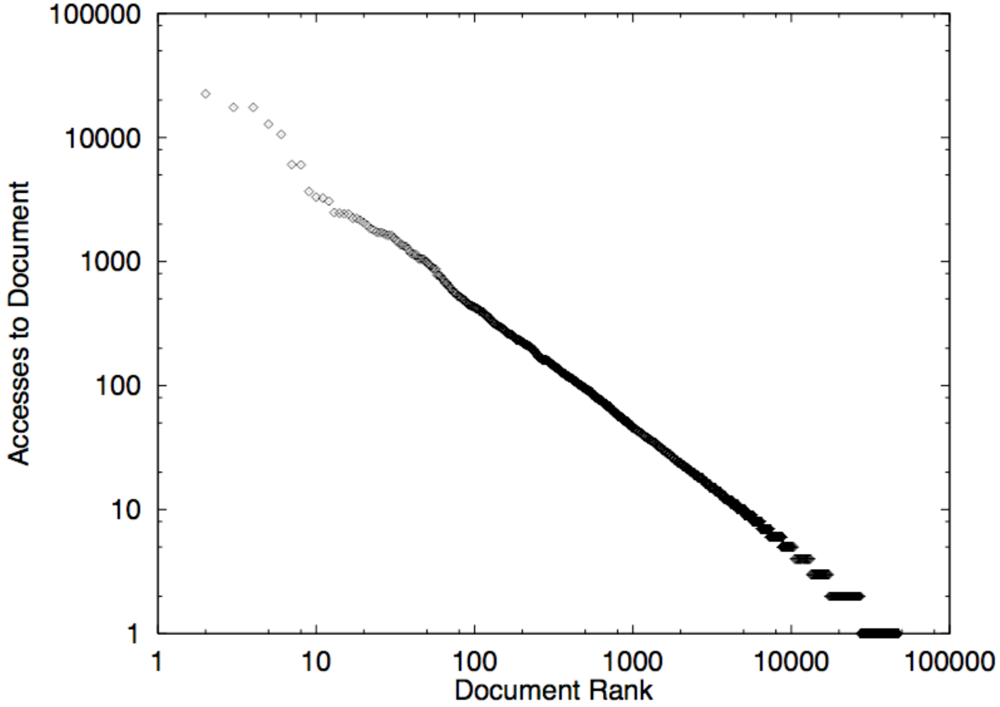

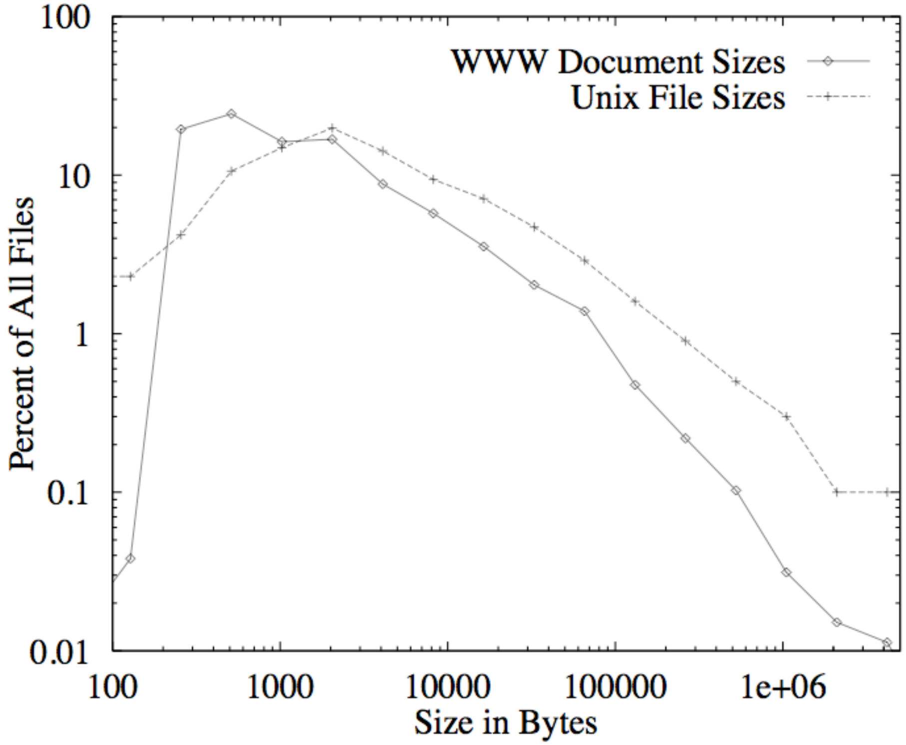

Small World Phenomenon

- Milgram's experiment 1967
  - 60 random chosen participants in Wichita, Kansas had to send a packet to an unknown address
  - They were only allowed to send the packet to friends
    - likewise the friends
- The majority of packets arrived within six hops
- Small-World-Networks
  - are networks with Pareto distributed node degree
  - with small diameter (i.e. $O(\log^c n)$)
  - and relatively many cliques
- Small-World-Networks
  - Internet, World-Wide-Web, nervous systems, social networks

How do Small World Networks Come into Existence?

- Emergence of scaling in random networks, Albert-Laszlo Barabasi, Reka Albert, 1999
- Preferential Attachment-Model (Barabasi-Albert):
  - Starting from a small starting graph successively nodes are inserted with m edges each (m is a parameter)
  - The probability to choose an existing node as a neighbor is proportional to the current degree of a node
- This leads to a Pareto network with exponent 2,9 ± 0,1
  - however cliques are very seldom
- Watts-Strogatz (1998)
  - Start with a ring and connections to the m nearest neighbors
  - With probability p every edge is replaced with a random edge
  - Allows continuous transition from an ordered graph to chaos
- Extended by Kleinberg (1999) for the theoretical verification of Milgram‘s experiment

CoNe Freiburg

IIF INSTITUT FÜR INFORMATIK FREIBURG

Peer-to-Peer Networks

02: Napster & Gnutella

Christian Schindelhauer
Technical Faculty
Computer-Networks and Telematics
University of Freiburg

UNI FREIBURG