

Peer-to-Peer Information Systems

Gerhard Weikum

weikum@mpi-sb.mpg.de

http://www.mpi-sb.mpg.de/units/ag5/teaching/ws04_05/p2p-seminar.html

Outline:

- ★ **History of P2P Systems**
- ★ **Future Applications and Research Topics**
- ★ **Seminar Organization**

Motivation for P2P

**exploit distributed computer resources
available through the Internet and mostly idle
→ tackle otherwise intractable problems
(e.g. SETI@home)**

make systems ultra-scalable & ultra-available

**break information monopolies,
exploit small-world phenomenon**

**replace admin-intensive server-centric systems
by self-organizing dynamically federated system
without any form of central control**

→ make complex systems manageable

„Autonomic Computing Laws“

Vision:

all computer systems must be self-managed, self-organizing, and self-healing (like biological systems?)

Eight laws:

- know thy self
- configure thy self
- optimize thy self
- heal thy self
- protect thy self
- grow thy self
- know thy neighbor
- help thy users

My interpretation:

need design for predictability:
self-inspection, self-analysis, self-tuning

1st-Generation P2P

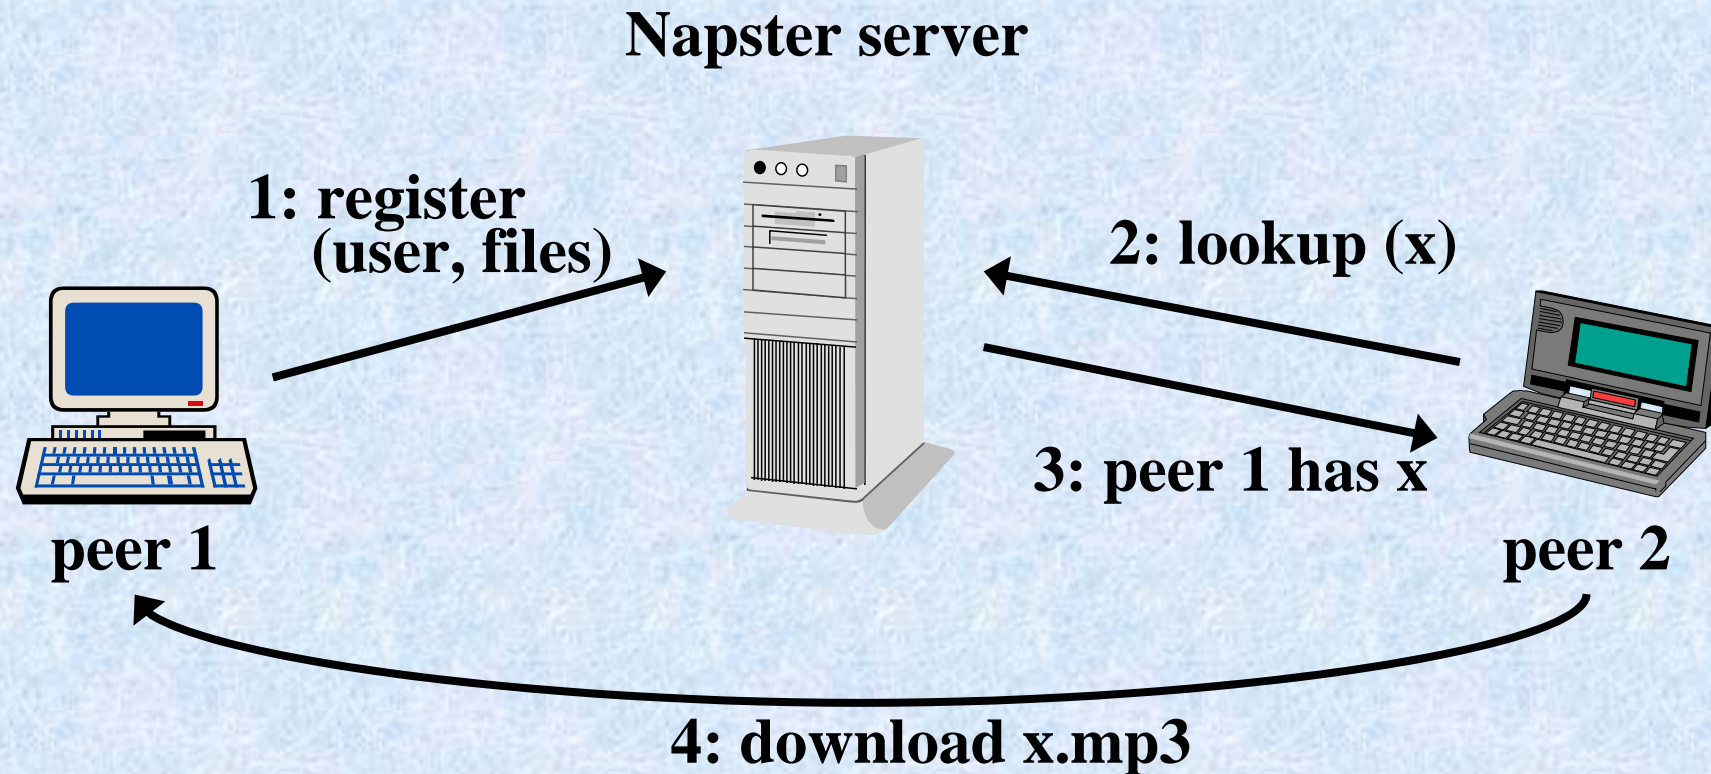
**Napster (1998-2001) and Gnutella (1999-now):
driven by file-sharing for MP3, etc.
very simple, extremely popular**

**can be seen as a mega-scale but very simple
publish-subscribe system:**

- **owner of a file makes it available under name x**
- **others can search for x, find copy, download it**

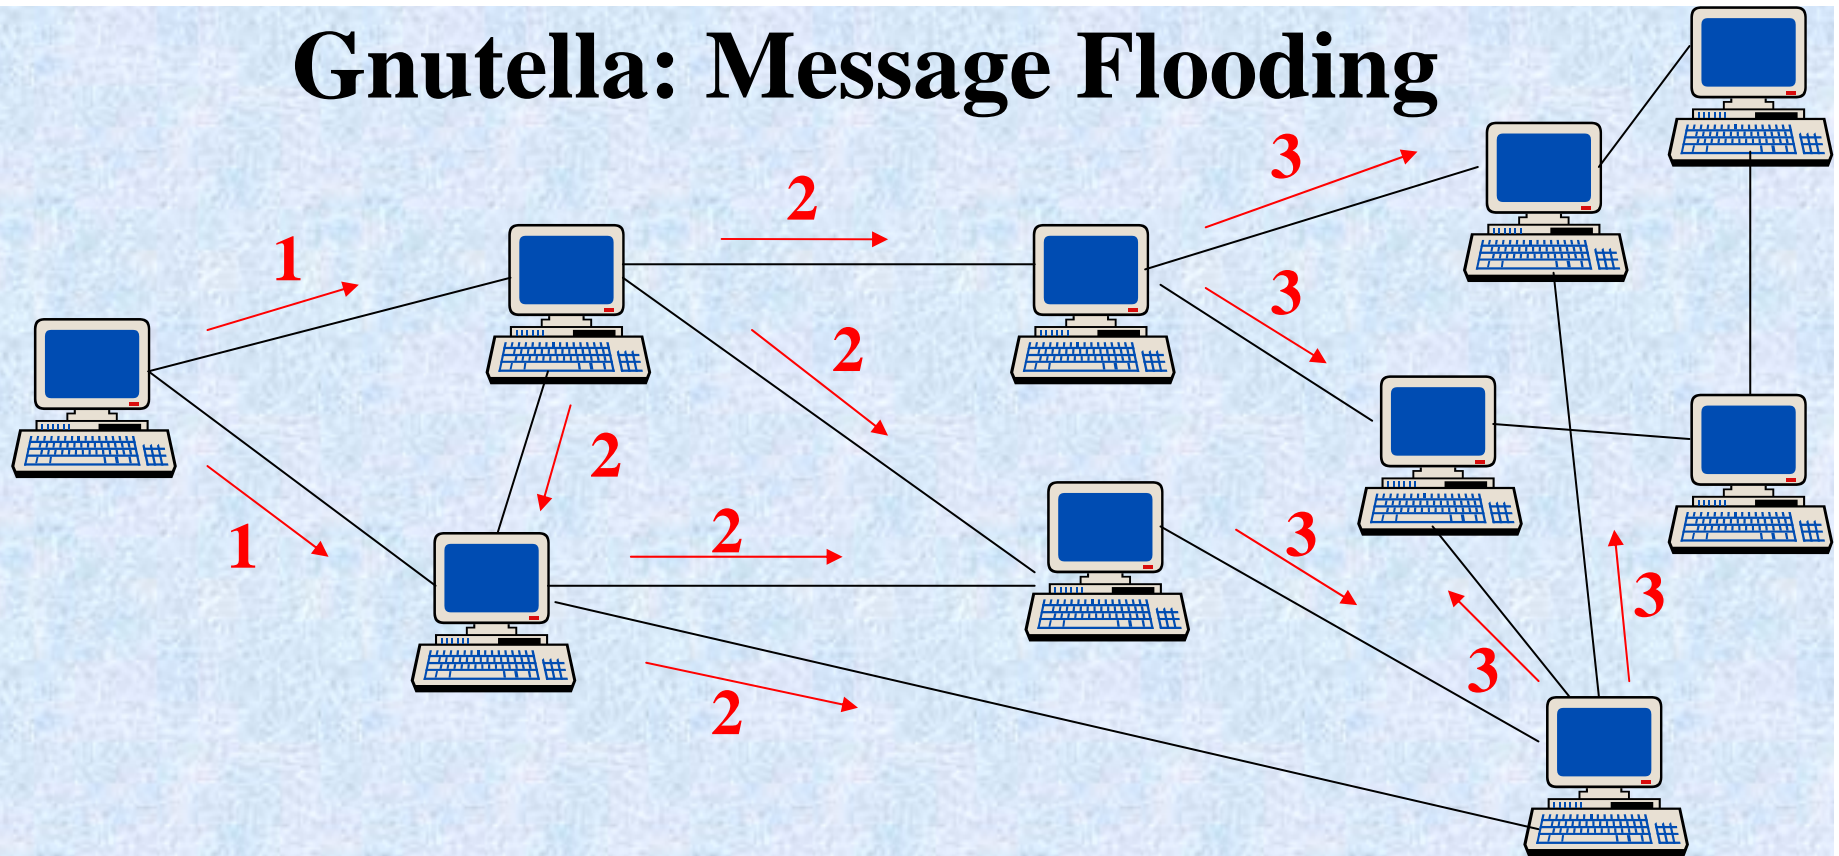
invitation to break the law (piracy, etc.) ?

Napster: Centralized Index



+ chat room, instant messaging, firewall handling, etc.

Gnutella: Message Flooding



all forward messages carry a TTL tag (time-to-live)

- 1) contact neighborhood and establish virtual topology (on-demand + periodically): *Ping, Pong***
- 2) search file: *Query, QueryHit***
- 3) download file: *Get or Push* (behind firewall)**

2nd-Generation P2P

Freenet

emphasizes anonymity

eDonkey, KaZaA (based on FastTrack), Morpheus, MojoNation, AudioGalaxy, etc. etc.

**commercial, typically no longer open source;
often based on super-peers**

JXTA

(Sun-sponsored) open API

Research prototypes (with much more refined architecture and advanced algorithms):

Chord (MIT), CAN (Berkeley), OceanStore/Tapestry (Berkeley), Farsite (MSR), Spinglass/Pepper (Cornell), Pastry/PAST (Rice, MSR), Viceroy (Hebrew U), P-Grid (EPFL), P2P-Net (Magdeburg), Pier (Berkeley), Peers (Stanford), Kademia (NYU), Bestpeer (Singapore), YouServ (IBM Almaden), Hyperion (Toronto), Piazza (UW Seattle), PlanetP (Rutgers), SkipNet (MSR), Galanx (U Wisconsin), Minerva (MPII), etc. etc.

The Future of P2P: New Applications

Beyond file-sharing & name lookups:

- **partial-match search, keyword search**
(tradeoff efficiency vs. completeness)
- **Web search engines**
- **publish-subscribe with eventing (e.g., marketplaces)**
- **collaborative work (incl. games)**
- **collaborative data mining**
- **dynamic fusion of (scientific) databases with SQL**
- **smart tags (e.g., RFID) on consumer products**

The Future of P2P: More Challenging Requirements

**Unlimited scalability with millions of nodes
($O(\log n)$ hops to target, $O(\log n)$ state per node)**

**Failure resilience, high availability, self-stabilization
(many failures & high dynamics)**

**Data placement, routing, load management, etc.
in overlay networks**

Robustness to DoS attacks & other traffic anomalies

Trustworthy computing and data sharing

**Incentive mechanisms to reconcile selfish behavior
of individual nodes with strategic global goals**

Related Technologies

Web Services (SOAP, WSDL, etc.)

for e-business interoperability (supply chains, etc.)

Grid Computing

for scientific data interoperability

Autonomic / Organic / Introspective Computing

for self-organizing, zero-admin operation

Multi-Agent Technology

for interaction of autonomous, mobile agents

Sensor Networks

for data streams from measurement devices etc.

Content-Delivery Networks (e.g., Akamai)

for large content of popular Web sites

Seminar Organization

Each participant

- reads one paper (plus background literature)
- gives a 30-minute presentation, followed by up to 15 minutes discussion
- produces a 10-to-20-pages write-up, due one week after the presentation

Participants should work in 3 phases:

- now until -3 weeks:
understand literature, interact with tutor
- until -2 weeks:
work out content and organization of your talk
- until -1 week:
work out presentation (ready for rehearsal)

Seminar Topics

Nov 23: Scalable Routing and Object Localization

Nov 30: Performance

Dec 7: Semantic Overlay Networks

Dec 14: P2P Algorithms

Dec 21: Replication

Jan 11: Information Search on Web Data

Jan 18: Incentives and Fairness

Jan 25: Privacy, Security, and Trust