Peer to Peer Computing

These slides are based on the slides made available by the authors of
Computer Networking: A Top Down Approach Featuring the Internet, 2nd edition.
Jim Kurose, Keith Ross
Addison-Wesley, July 2002.

1

Peer-peer computing and networking

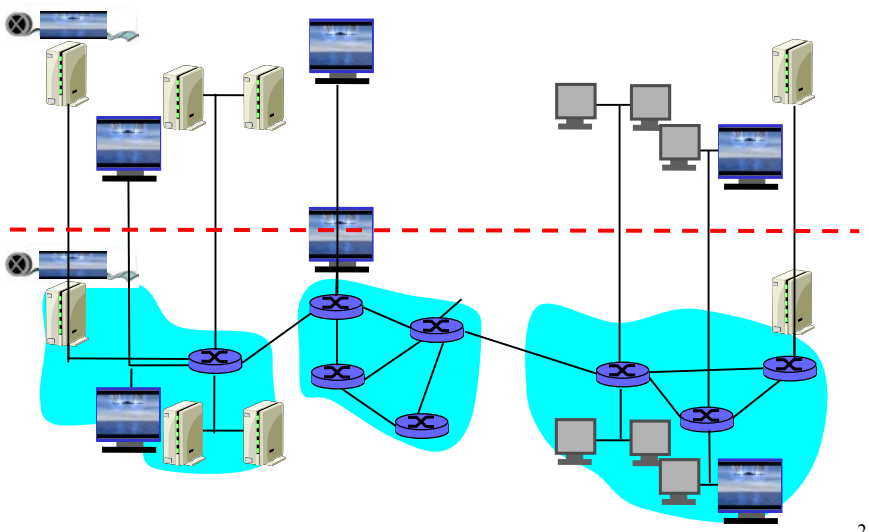

2

Peer-peer network

Focus at the application level

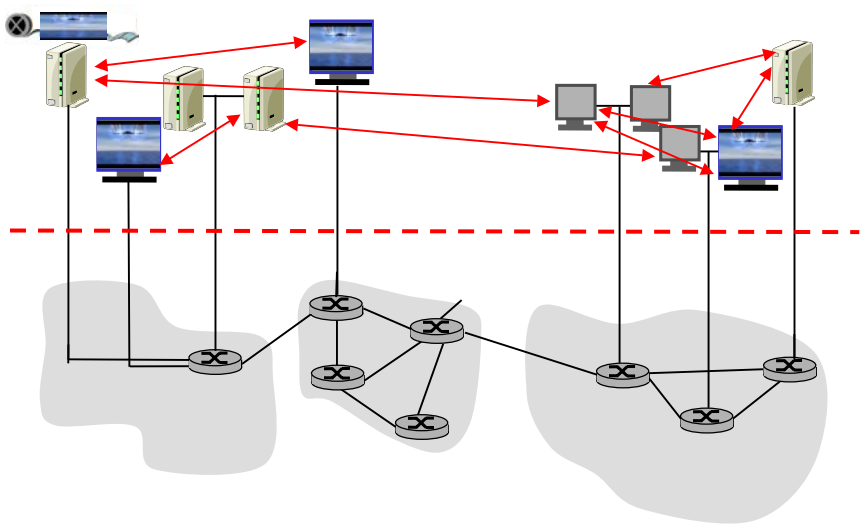

3

Peer-peer computing

Peer-peer applications

- File sharing
  - Napster, Gnutella, Freenet
  - Second generation projects
    - Oceanstore, PAST, Freehaven
- Distributed Computation
  - SETI@home, Entropia, Parabon, United Devices, Popular Power
- Other Apps?
  - Instant Messaging (Jabber), Anonymous Email
  - Groupware (Groove)
  - P2P Databases

4

P2P: Related Technologies

- Distributed computing.
  - How is P2P different from distributed computing?
- Grid computing.
  - How is the computational grid different from P2P networks?

  KEY DIFFERENCES: Peers are on the edges of the Internet, are autonomous, have variable connectivity, and temporary network addresses
- Application-level networking.
  - Resilient overlay networks for multicast, video distribution, etc.

5

P2P: Related Technologies

- Wireless ad-hoc networks.
- Sensor networks.
- P2P devices/ubiquitous computing.
  - JINI.
- Web services.
  - .NET framework, SOAP, UDDI.

6

P2P file sharing

Example

- Alice runs P2P client application on her notebook computer
- Intermittently connects to Internet; gets new IP address for each connection
- Asks for "Hey Jude"
- Application displays other peers that have copy of Hey Jude.

- Alice chooses one of the peers, Bob.
- File is copied from Bob's PC to Alice's notebook: HTTP
- While Alice downloads, other users uploading from Alice.
- Alice's peer is both a Web client and a transient Web server.

All peers are servers = highly scalable!

P2P: centralized directory

original "Napster" design

1) when peer connects, it informs central server:
   - IP address
   - content
2) Alice queries for "Hey Jude"
3) Alice requests file from Bob

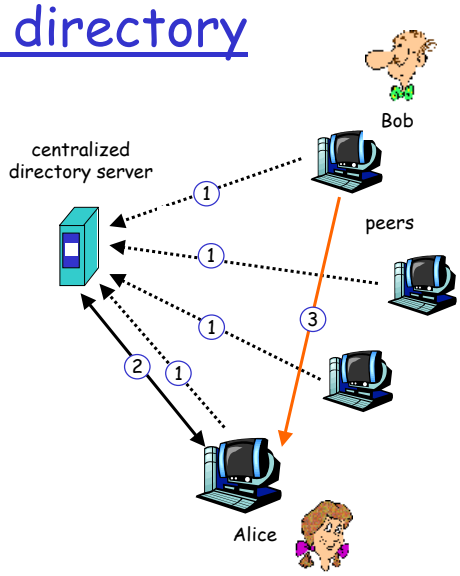

P2P: problems with centralized directory

- Single point of failure
- Performance bottleneck
- Copyright infringement

file transfer is decentralized, but locating content is highly decentralized

9

Napster

- program for sharing files over the Internet
- a killer application?
- history:
  - 5/99: Shawn Fanning (freshman, Northeasten U.) founds Napster Online music service
  - 12/99: first lawsuit
  - 3/00: 25% UWisc traffic Napster
  - 2000: est. 60M users
  - 2/01: US Circuit Court of Appeals: Napster knew users violating copyright laws
  - 7/01: # simultaneous online users:
    Napster 160K, Gnutella: 40K, Morpheus: 300K
  - 2001: Napster shut down; Bertelsmann acquire assets, etc.
- Today?: OpenNap (open source napster server)

10

Napster: how did it work

Application-level, client-server protocol over point-to-point TCP

Four steps:

- Connect to Napster server
- Upload your list of files (push) to server.
- Give server keywords to search the full list with.
- Select "best" of correct answers. (pings)

Napster

1. File list is uploaded

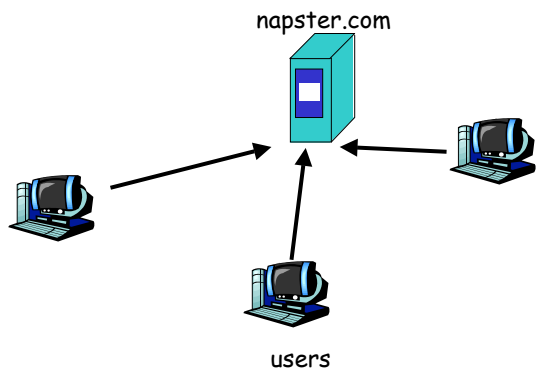

Napster

2. User requests search at server.

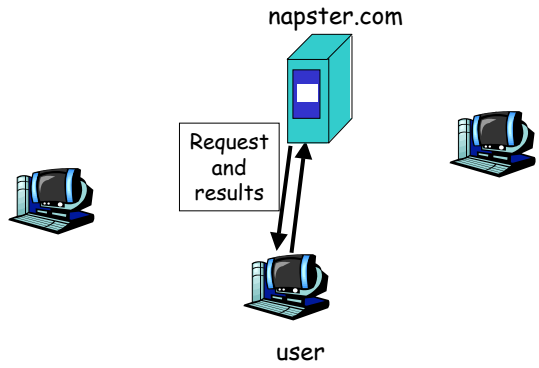

13

Napster

3. User pings hosts that apparently have data.

   Looks for best transfer rate.

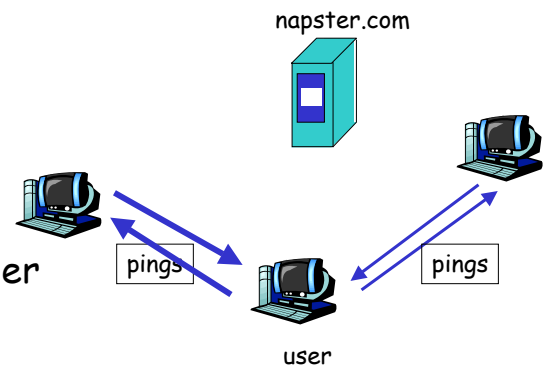

14

Napster

4. User retrieves file

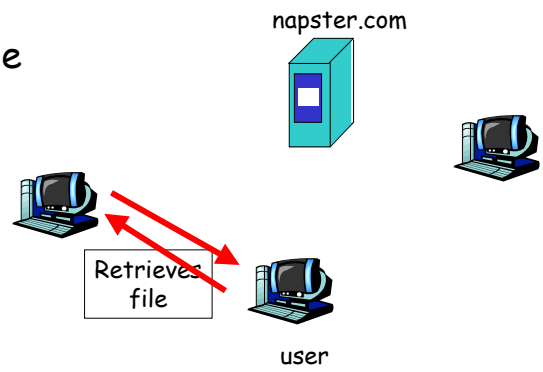

Napster: architecture notes

- centralized server:
  - single logical point of failure
  - can load balance among servers using DNS rotation
  - potential for congestion
- no security:
  - passwords in plain text
  - no authentication
  - no anonymity

P2P: decentralized directory

- Each peer is either a group leader or assigned to a group leader.
- Group leader tracks the content in all its children.
- Peer queries group leader; group leader may query other group leaders.

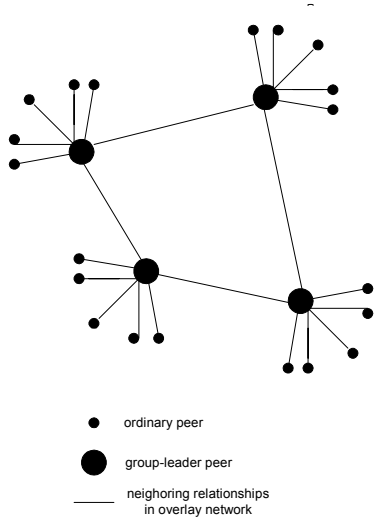

17

More about decentralized directory

overlay network

- peers are nodes
- edges between peers and their group leaders
- edges between some pairs of group leaders
- virtual neighbors

bootstrap node

- connecting peer is either assigned to a group leader or designated as leader

advantages of approach

- no centralized directory server
  - location service distributed over peers
  - more difficult to shut down

disadvantages of approach

- bootstrap node needed
- group leaders can get overloaded

18

P2P: Query flooding

- Gnutella
- no hierarchy
- use bootstrap node to learn about others
- join message

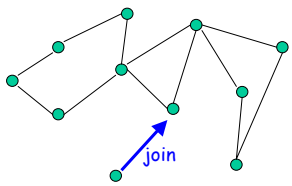

- Send query to neighbors
- Neighbors forward query
- If queried peer has object, it sends message back to querying peer

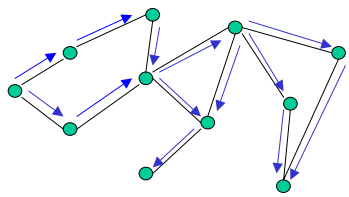

19

P2P: more on query flooding

Pros

- peers have similar responsibilities: no group leaders
- highly decentralized
- no peer maintains directory info

Cons

- excessive query traffic
- query radius: may not have content when present
- bootstrap node
- maintenance of overlay network

20

Gnutella

- peer-to-peer networking: applications connect to peer applications
- focus: decentralized method of searching for files
- each application instance serves to:
  - store selected files
  - route queries (file searches) from and to its neighboring peers
  - respond to queries (serve file) if file stored locally
- Gnutella history:
  - 3/14/00: release by AOL, almost immediately withdrawn
  - too late
  - many iterations to fix poor initial design (poor design turned many people off)

21

Gnutella: how it works

Searching by flooding:

- If you don't have the file you want, query 7 of your partners.
- If they don't have it, they contact 7 of their partners, for a maximum hop count of 10.
- Requests are flooded, but there is no tree structure.
- No looping but packets may be received twice.
- Reverse path forwarding

Note: Play gnutella animation at:
http://www.limewire.com/index.jsp/p2p

22

Flooding in Gnutella: loop prevention

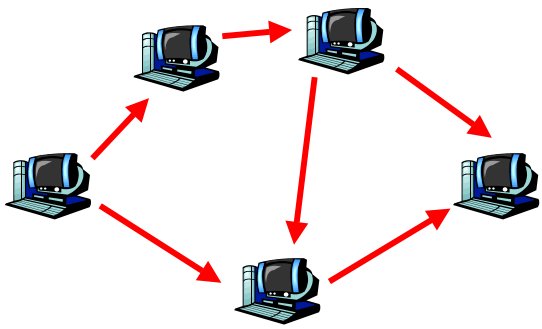

Seen already list: "A"

Distributed Computing

- Current supercomputers are too expensive
  - ASCI White (#1 in TOP500) costs more than $110 million and needed a new building
  - Few institutions or research groups can afford this level of investment
- There are more than 500 million PCs around the world
  - some as powerful as early 90s supercomputers
  - they are idle most of the time (60% to 90%), even when being used (spreadsheet, typing, printing,...)
  - corporations and institutions have hundreds or thousands of PCs on their networks

⇩

Try to harness idle PCs on a network and use them on computationally intensive problems

How it works

- Embarrassingly parallel applications
  - Large computation to communication ratio
  - Master/worker model
  - Applications can use local disk for checkpointing
- Provider farms out work to idle PCs across the internet
  - PC owners volunteer idle cycles (for money or altruistic purposes)

25

Entropia network

- Born in 1997 to apply idle computers worldwide to problems of scientific interest
- In 2 years grew to more than 30,000 computers with aggregate speed of over 1 Tflop/second
- Several scientific achievements, e.g. Identification of largest known prime number
- Gone commercial: www.entropia.com and used for applications from:
  - Life sciences
  - Financial services
  - Product design, etc.

26

SETI @ home project

setiathome.ssl.berkeley.edu

- SETI = Search for Extraterrestrial Intelligence
- Started in 1996 to enlist PCs to work on analyzing data from the Arecibo radio telescope
- Good mix of popular appeal and good technology

⇨
- Now running on more than ½ million PCs
- delivering ~ 1,200 CPU years per day
- ~ 35 Tflops/sec
- fastest (but special-purpose) computer in the world

27

Folding @ home project

www.stanford.edu/group/pandegroup/Cosm

- Enlists PCs to work on the protein folding problem
- most important problem in modern molecular biology
- From genome to structure:
  - Genome sequence of DNA specifies amino acids that make up proteins, but says little about their functions: what is needed is how a protein fold (3D structure)
  - Protein folding is very fast (microseconds) and complex
  - Simulation timescale is of the order of nanoseconds
    → 10^3 gap → distributed computing
- Currently around 20,000 users

28