Peer-to-Peer Internet Applications: A Review

Davide Quaglia

Outline

- Introduction
- Key points
- Lookup task
  - Centralized (Napster)
  - Query flooding (Gnutella)
  - Distributed Hash Table (Chord)
- Simulation Tools
- Streaming P2P
- Potential research interests

Introduction

- All nodes in a network join together dynamically to participate in both traffic routing and processing tasks.
- Flat architecture
  - No central server
  - All nodes are equal: each node acts both as client and server
- Applications:
  - distributed DB
  - distributed services
  - network games

Some names

- Real-world applications:
  - Napster,
  - Freenet,
  - WinMX,
  - Gnutella,
  - KaZaA,
  - BitTorrent,
  - Skype

Some names (2)

- Research:
  - Pastry,
  - CAN,
  - Chord,
  - Oceanstore,
  - Yallcast,
  - Kademlia,
  - Tapestry,
  - Koorde,
  - JXTA

Overlay network

- A new application-level network build over the hosts of the traditional IP network (hosts+routers)
- Nodes have IDs (application-level addresses)
  - IDs are then mapped on IP addresses
- A neighborhood function is defined
  - 2 nodes that are neighbors in the ID space may be very far in the IP space !

Overlay network (2)

- Transport-level links between hosts
- Advantages
  - No need for public IP addresses (IDs are used)
  - It works with firewalls and NAT
  - No changes in the IP network (no need for agreements with administrators)
- Disadvantages
  - Waste of bandwidth due to suboptimal data transfers

Overlay network (3)

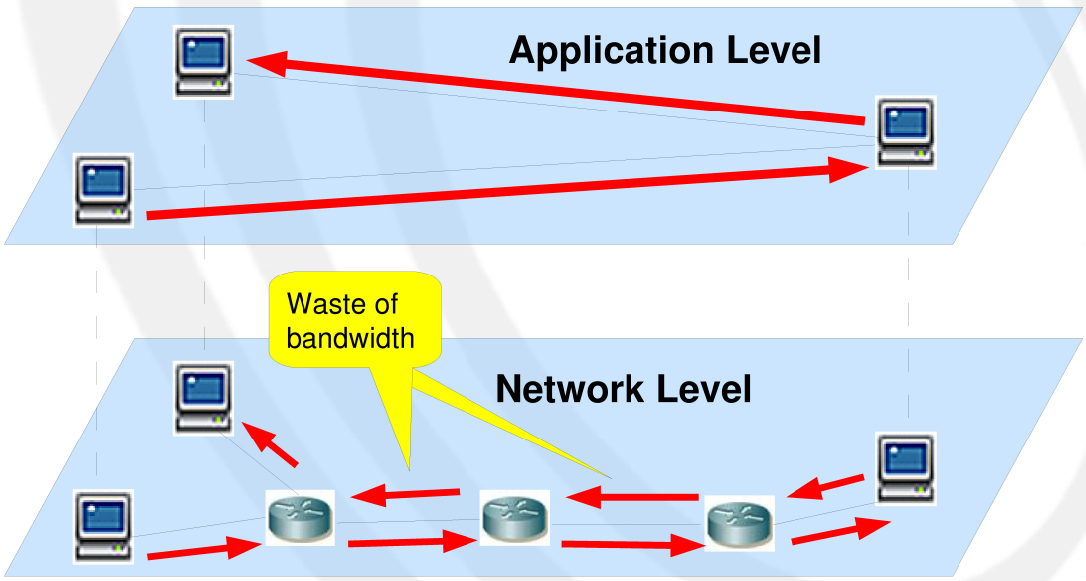

Network dynamics

- Number of nodes
- Mean degree
- Turnover
- Half-life
- Persistence
- Size-change

Milestones

- 1997: paper by Plaxton et al.
- 1999: Napster
- 2000: Gnutella
- 2001: Chord, CAN (Content Addressable Network), Pastry, Tapestry
- 2002: Kademlia, Viceroy
- 2003: Skype
- SuperNode (KaZaA) ? Joltid ?

Main functions

- Boot/Topology
  - how to choose the node ID and to create connections in order to keep the graph balanced and to minimize communication costs ?
    - Unstructured vs Structured P2P networks
- Lookup (and routing)
  - given a search topic, the nodes responsible for the topic must be identified.
- Application-level multicast

Main functions (2)

- Security
- Anonimity
- Reliability
- Robustness to
  - Attacks (e.g., Denial of Service, Poisoning)
  - Freeriders
- Reputation management (virt. Currency)
- Node heterogeneity and networks changes

Lookup

- Centralized index: Napster
- Query flooding: Gnutella
- Super-peer: KaZaA (FastTrack)
- Distributed Hash Table (DHT): Chord
- Efficiency metrics:
  - Size of the routing tables k
  - Number of hops h

Lookup: centralized index

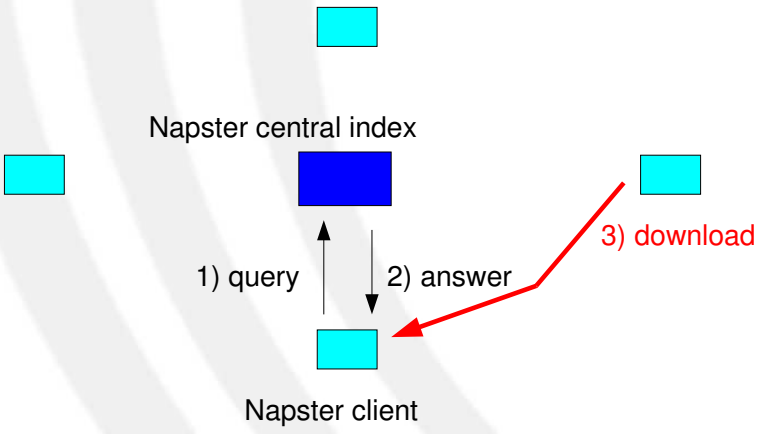

k = O(N) h = O(1)

Lookup: query flooding

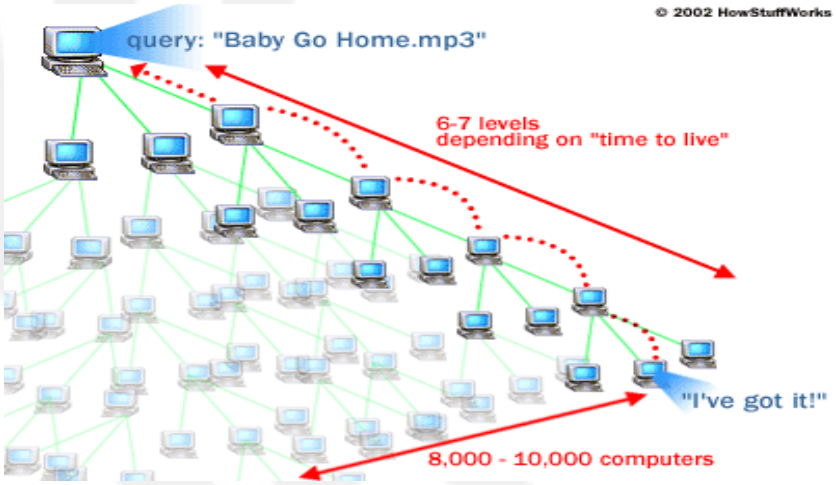

k = O(1) h = O(N)

Lookup: super-peer

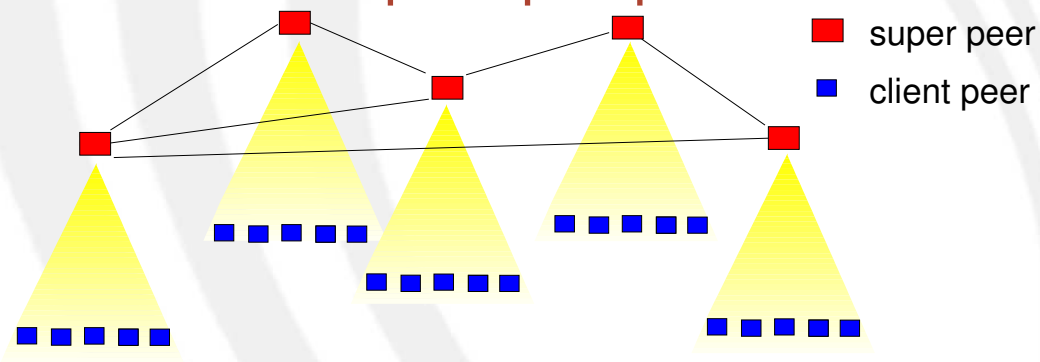

- Query flooding among super-peers
- Centralized index with respect to client peers
- k = O(1) h = O(N) but with smaller constants

Lookup: super-peer

- The most used approach in widespread peer-to-peer applications
- Various names for the super-peer:
  - Hub (Direct Connect)
  - SuperPeer , UltraPeer (Gnutella2)
  - SuperNode (KaZaA)
  - RendezVousNode (JXTA)
  - MainPeer(EDonkey)
  - Server (WinMX)

Lookup: distributed hash table

- Each resource is associated to a key by a hashing function.
- Keys belong to the same space of node IDs.
- A node contains resources whose keys are close to its ID.
- h = O(logN) k = O(logN)
- Problem: uneven distribution of resources.

Application-level multicast

- Multicast is reproduced at application level
- No need to enable multicast in layer 3 routers
  - No need to reach agreement with providers
- Waste of bandwidth due to data replication

Simulation tools

- Microsoft Research Pastry Simulator 3.0
- GnutellaSim (tool for NS)
  - www.cc.gatech.edu/computing/compass/gnutella/

P2P streaming

- P2P-radio, Peercast: basati su Icecast
- Skype: voice over UDP (no RTP), session open/close with TCP packets, ~20kb/s per talker.
- BitTorrent: file sharing (not streaming), the tracker is responsible for tracking nodes which have a piece of the resource, MIT license.
- GnuStream, PALS, PROMISE: multi-sender aggregation

Potential research interests

- Multimedia-specific issues
  - Reliability through retransmission and correcting codes
- Multicast over P2P
- Multisender aggregation
  - Multiple description coding
- Common ideas as with sensor and ad-hoc networks
- Simulation tools