

Distributed Systems

ID2201



P2P and DHT:s
Johan Montelius

Idéa

- use resources in edge of network
 - computing
 - storage
 - communication



Computing



seti@home

- central server
 - millions of clients
 - hundred of thousands active
- super computer
 - hundreds of TeraFLOPS
 - one of the largest computations performed
- continued in the BOINC project



File sharing

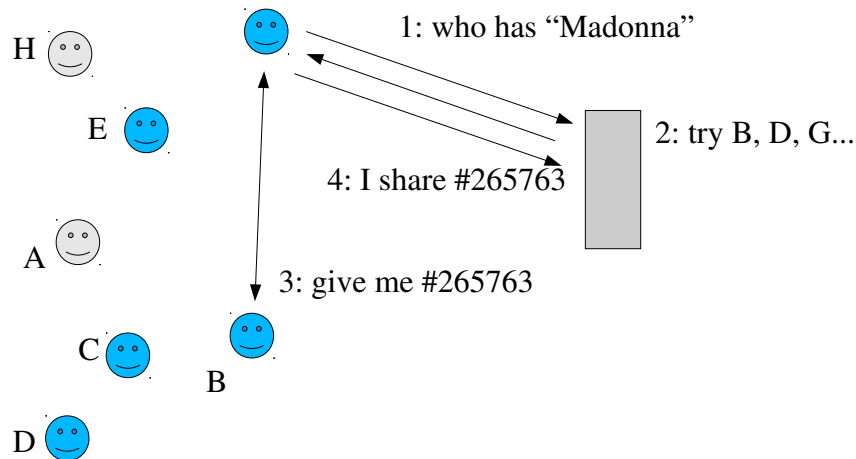


Napster



- First large scale peer-to-peer file sharing system - 1999
- Used a central server to store index of all files.
- Clients copied files peer-to-peer.

Napster



Napster




- Central server
 - knows everything
 - needs to be alive
 - can easily be replicated
- File transfer
 - limited by client upload capacity
- Problems
 - copyright issues
 - why share
 - is it the correct file

Next step



9

Kazaa



10


Kazaa

- FastTrack (closed protocol)
 - super nodes, unstructured, responsible for indexing
 - central server, blacklist of super nodes
 - regular nodes, connects to local super node
- Integrity is checked by hash function.
 - not very strong
- Money made on
 - advertising
 - ...



11

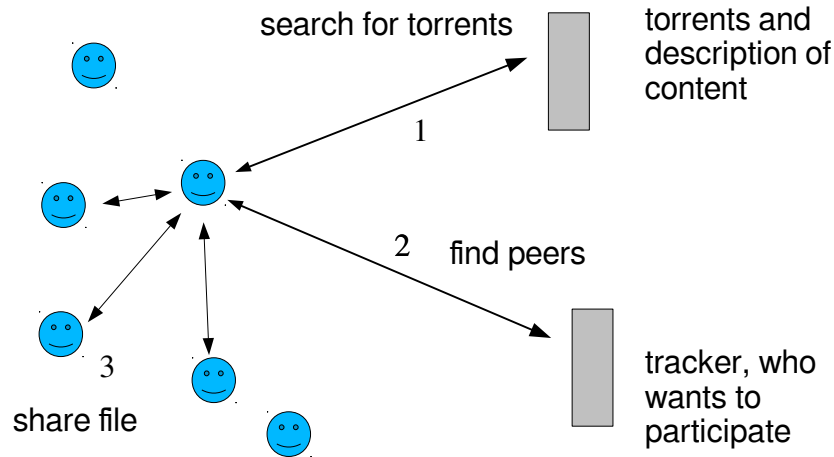
Time of the pirates



12



BitTorrent



13



BitTorrent

- torrent
 - trackers to use
 - name of content
 - size and number of segments
 - hash codes of segments
- tracker
 - provides list of peers
 - could be helpful in suggesting network close peers

Distributed Systems ID2201

14



BitTorrent

- Query peers to find who has what.
- Tit-for-tat
 - per file, not on total
 - if you don't get something, why share
- Rarest first
 - rare segments are valuable
- Multiple peers
 - change if connection is slow
 - choke if you don't get anything back

15



Magnet links and DHT

Magnet : ?

```
xt=urn:btih:d2438d70a205554c1270237b1dbcf&
dn=Rick+Astley+Never+Gonna+Give+You+Up.mp3&
```

How do we find the torrent?

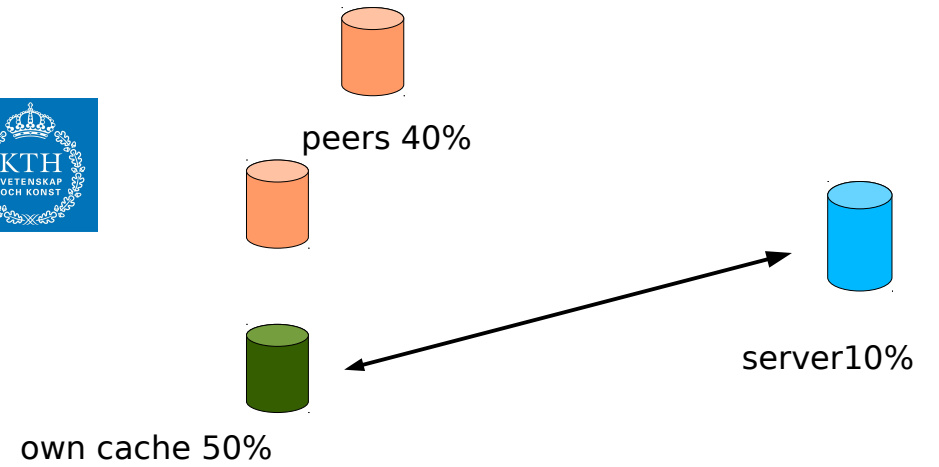
Distributed Systems ID2201

16


All the music, all the time.




Spotify



Some privacy



Tor - anonymous routing



P2P middleware



- Function
 - add, remove, locate and communicate with resources in a network
- Requirements
 - global scale, millions of nodes
 - dynamic availability
 - integrity, privacy, anonymity, deniability

21

Overlay routing




- name space: global unique identifier (GUID)
- structured or unstructured
 - pay when you add nodes and objects
 - pay when you search for objects
- fault tolerance and consistency
 - replication

22

Description, Identifier and Objects




Description



Is the description unique?

Identifier


How do we find a unique identifier?



Object

23

unstructured overlay



24

Unstructured




- No network structure.
- You know some other nodes.
- No fixed location of objects
- Easy to join.
- Hard to search.
- No guarantees

Searching



- Flood the network
 - there should be a limit!
- Expanding ring
 - iterative flooding
- Random walk
 - several independent searchers
- Gossip
 - hopefully they will know

structured overlay



Structured overlay



- Pros
 - makes searching easier
 - more predictable
- Cons
 - expensive to add nodes
 - expensive to add content

Routing overlay



- The problem of finding a node, object or resource in a network:
 - nodes can leave and join
 - nodes might fail
- Each object is described by a globally unique identifier (GUID).

29

Distributed Hash Tables



- Use a hash value as key/identifier
 - uniform distribution
- Each node chooses a random hash value as its identifier.
- Nodes form a ring and can forward request in the ring.

30

Summary



- P2P
 - computing
 - storage
 - network
- Overlay networks
 - unstructured
 - structured

31