

Peer-to-Peer Networks

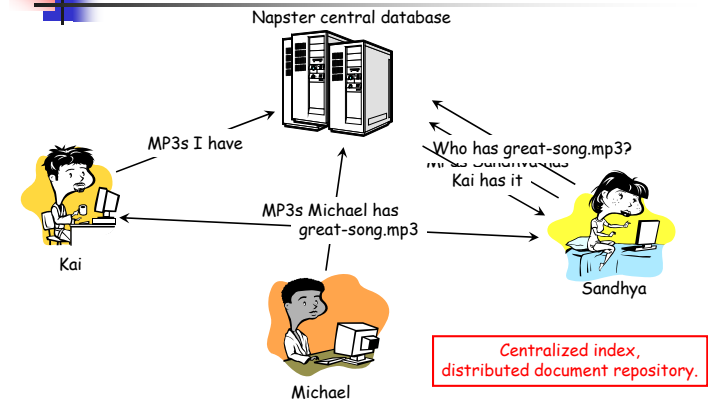
Kai Shen

11/2/2009

CSC 257/457 - Fall 2009

1

Distributed Search I: Napster (central index)

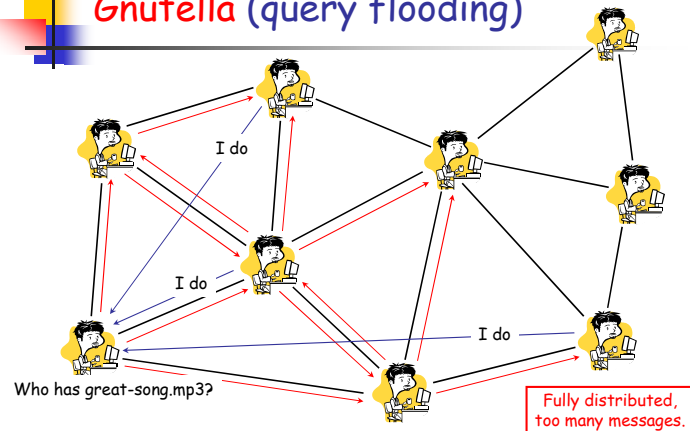


11/2/2009

CSC 257/457 - Fall 2009

2

Distributed Search II: Gnutella (query flooding)

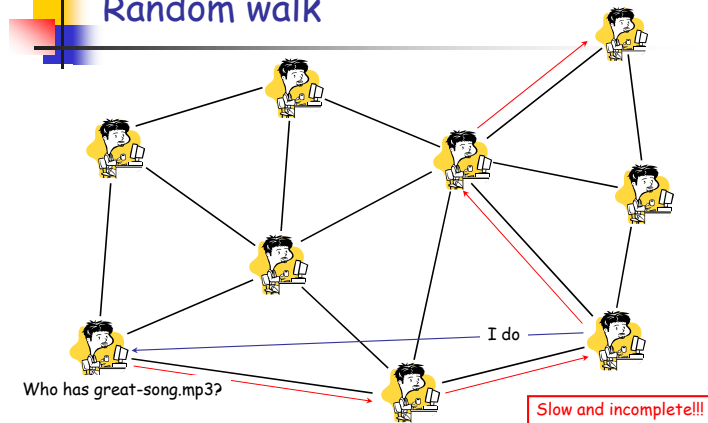


11/2/2009

CSC 257/457 - Fall 2009

3

Distributed Search III: Random walk



11/2/2009

CSC 257/457 - Fall 2009

4

Distributed Search IV: Object-addressable network

- "great-song.mp3" is deterministically hashed into location "L"
- every node forwards each query to a neighbor heuristically closer to "L"

Here is great-song.mp3, store it somewhere.

Can I have great-song.mp3?

How does routing work?

11/2/2009 CSC 257/457 - Fall 2009 5

Scalable Object-addressable Net

Content-addressable net:

- each node owns a region in a virtual 2-D space
- "great-song.mp3" is hashed into a virtual location "L" (1.3, 2.4)
- every node forwards each query to a neighbor heuristically closer to "L"

Here is great-song.mp3, store it somewhere.

Can I have great-song.mp3?

11/2/2009 CSC 257/457 - Fall 2009 6

Analysis

Content-addressable network [Ratnasamy et al. 2001]:

- each node owns a region in a virtual 2-D space
- each object is hashed into a virtual location "L"
- every node forwards each query to a neighbor heuristically closer to "L"

Questions:

- Space consumption? Lookup cost for a network of N nodes?
 - 4 links per node; $O(N^{1/2})$
- Can the lookup cost be improved?
- How to take into account the actual link latency?

11/2/2009 CSC 257/457 - Fall 2009 7

Chord [Stoica et al. 2001]

- Each object is mapped to the first node in the clockwise direction on the ring.
- Every node maintains links to
 - the node half-way in the circular ID space
 - the node a quarter-way across the circular ID space
 - ...
- Space consumption? Lookup cost for a network of N nodes?

11/2/2009 CSC 257/457 - Fall 2009 8

Distributed Hashtables

- Content-addressable network (CAN) [Ratnasamy et al. 2001] and Chord [Stoica et al. 2001] are also called scalable distributed hash tables (DHTs)
- There are other scalable DHT protocols.

11/2/2009

CSC 257/457 - Fall 2009

9

Peer-to-peer Networks

- Peer-to-peer networks: distributed systems of **no hierarchy** - each node is a peer of any other in terms of equal functionality
 - DNS? Napster? Gnutella? Random walk? DHTs (CAN and Chord)?
- Fundamental advantage of p2p networks
 - better scalability?
 - no performance bottleneck
 - better robustness?
 - more tolerant to intentional attacks; better worst-case scenario under random failures

11/2/2009

CSC 257/457 - Fall 2009

10

More on Scalability & Robustness

- Scalability:** able to support large number of nodes
 - cost of each operation is below linear-scaling - goes up slowly when the system size increases: $O(n)$ is terrible, $O(n^{1/2})$ is OK, $O(\log n)$ is good.
 - space requirement at each node is below linear-scaling.
 - Gnutella? Random walk? CAN? Chord?
- Robustness:** complexity of repair mechanism is a very important issue
 - Gnutella? Random walk? CAN? Chord?

11/2/2009

CSC 257/457 - Fall 2009

11

Where are we now?

- Four approaches for searching large, distributed systems:
 - central index
 - query flooding
 - random walk
 - hash + heuristics-based routing (DHT)
- Concept of peer-to-peer networks.
- DHT can serve as a fundamental building block for other network services
 - p2p distributed file systems
 - DNS
 - Keyword search in peer-to-peer fashion.

11/2/2009

CSC 257/457 - Fall 2009

12

Keyword Search

- user inputs a few keywords, the system returns a list of documents matching the keywords - **Google**
- Google maintains a central search index:
 - a search index contains a list of all searchable words, each of which contains a list of documents relevant to the word
 - intersection of document lists for multiple-word queries

Java:

Page #123	Page #157
-----------	-----------	-----	-----

Sun:

Page #157	Page #468
-----------	-----------	-----	-----

...

11/2/2009 CSC 257/457 - Fall 2009 13

Peer-to-peer Keyword Search

Solution 1: Split based on keywords

Split the index database to many pieces based on **keywords** and distribute them to many nodes in the network.

Central index - Google

Distributed index

Weakness: transferring large index for multiple-word queries.

11/2/2009 CSC 257/457 - Fall 2009 14

Peer-to-peer Keyword Search

Solution 2: Split based on documents

Split the index database to many pieces based on **documents** and distribute them to many nodes.

Central index - Google

UR

Joe's Web site

Weakness: too many sites to visit for each query.

11/2/2009 CSC 257/457 - Fall 2009 15