



## 15-744: Computer Networking

L-22: P2P

## Peer-to-Peer Networks



- Typically each member stores/provides access to content
- Has quickly grown in popularity
  - Bulk of traffic from/to CMU is Kazaa!
- Basically a replication system for files
  - Always a tradeoff between possible location of files and searching difficulty
  - Peer-to-peer allow files to be anywhere → searching is the challenge
  - Dynamic member list makes it more difficult
- What other systems have similar goals?
  - Routing, DNS

Lecture 22: 11-12-2002

2

## Overview

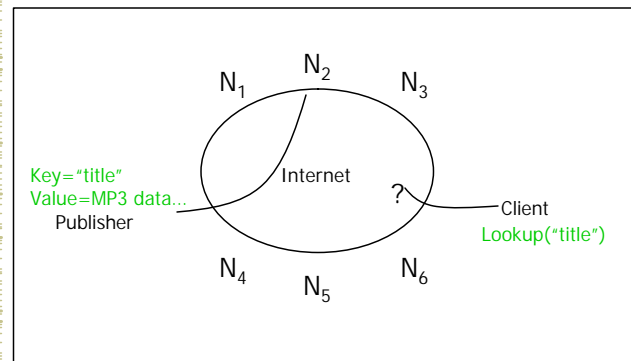


- P2P Lookup Overview
- Centralized/Flooded Lookups
- Routing-based Lookups
- CMU Research in P2P

Lecture 22: 11-12-2002

3

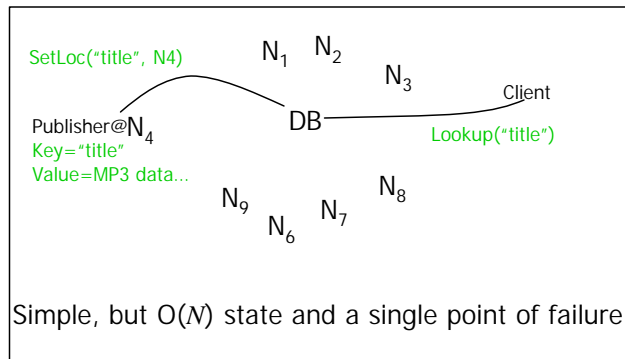
## The Lookup Problem



Lecture 22: 11-12-2002

4

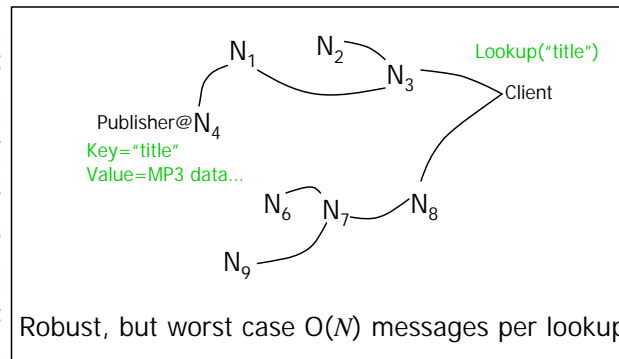
## Centralized Lookup (Napster)



Lecture 22: 11-12-2002

5


## Flooded Queries (Gnutella)



Lecture 22: 11-12-2002

6

## Routed Queries (Freenet, Chord, etc.)



Lecture 22: 11-12-2002

7

## Overview



- P2P Lookup Overview
- Centralized/Flooded Lookups
- Routing-based Lookups
- CMU Research in P2P

Lecture 22: 11-12-2002

8

## Centralized: Napster



- Simple centralized scheme → motivated by ability to sell/control
- How to find a file:
  - On startup, client contacts central server and reports list of files
  - Query the index system → return a machine that stores the required file
    - Ideally this is the closest/least-loaded machine
  - Fetch the file directly from peer

Lecture 22: 11-12-2002

9

## Centralized: Napster



- Advantages:
  - Simple
  - Easy to implement sophisticated search engines on top of the index system
- Disadvantages:
  - Robustness, scalability
  - Easy to sue!

Lecture 22: 11-12-2002

10

## Flooding: Gnutella



- On startup, client contacts any servent (**server + client**) in network
  - Servent interconnection used to forward control (queries, hits, etc)
- Idea: broadcast the request
- How to find a file:
  - Send request to all neighbors
  - Neighbors recursively forward the request
  - Eventually a machine that has the file receives the request, and it sends back the answer
  - Transfers are done with HTTP between peers

Lecture 22: 11-12-2002

11

## Flooding: Gnutella



- Advantages:
  - Totally decentralized, highly robust
- Disadvantages:
  - Not scalable; the entire network can be swamped with request (to alleviate this problem, each request has a TTL)
  - Especially hard on slow clients
    - At some point broadcast traffic on Gnutella exceeded 56kbps – what happened?
    - Modem users were effectively cut off!

Lecture 22: 11-12-2002

12

## Flooding: Gnutella Details



- Basic message header
  - Unique ID, TTL, Hops
- Message types
  - Ping – probes network for other servents
  - Pong – response to ping, contains IP addr, # of files, # of Kbytes shared
  - Query – search criteria + speed requirement of servent
  - QueryHit – successful response to Query, contains addr + port to transfer from, speed of servent, number of hits, hit results, servent ID
  - Push – request to servent ID to initiate connection, used to traverse firewalls
- Ping, Queries are flooded
- QueryHit, Pong, Push reverse path of previous message


Lecture 22: 11-12-2002

13

## Flooding: Gnutella Example



Assume: m1's neighbors are m2 and m3; m3's neighbors are m4 and m5;...



Lecture 22: 11-12-2002

14

## Flooding: FastTrack (aka Kazaa)



- Modifies the Gnutella protocol into two-level hierarchy
- Supernodes
  - Nodes that have better connection to Internet
  - Act as temporary indexing servers for other nodes
  - Help improve the stability of the network
- Standard nodes
  - Connect to supernodes and report list of files
  - Allows slower nodes to participate
- Search
  - Broadcast (Gnutella-style) search across supernodes
- Disadvantages
  - Kept a centralized registration → allowed for law suits ☹

Lecture 22: 11-12-2002

15

## Overview



- P2P Lookup Overview
- Centralized/Flooded Lookups
- Routing-based Lookups
- CMU Research in P2P

Lecture 22: 11-12-2002

16

## Routing: Freenet



- Addition goals to file location:
  - Provide publisher anonymity, security
- Files are stored according to associated key
  - Core idea: try to cluster information about similar keys
- Messages
  - Random 64bit ID used for loop detection
    - Each node maintains the list of query IDs that have traversed it → help to avoid looping messages
  - TTL
    - TTL is decremented each time the query message is forwarded

Lecture 22: 11-12-2002

17

## Routing: Freenet Routing Tables



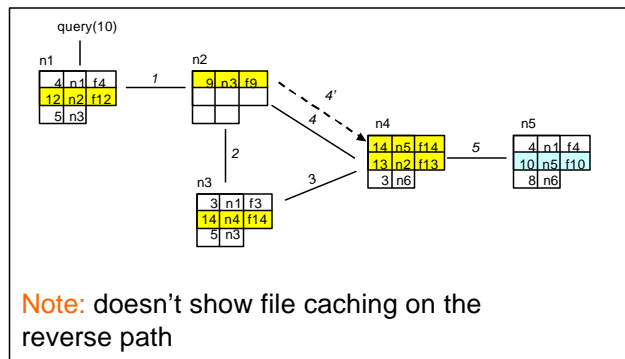
- *id* – file identifier
- *next\_hop* – another node that stores the file id
- *file* – file identified by *id* being stored on the local node
- Forwarding of query for file *id*
  - If file *id* stored locally, then stop
    - Forward data back to upstream requestor
    - Requestor adds file to cache, adds entry in routing table
  - If not, search for the “closest” *id* in the stack, and forward the message to the corresponding *next\_hop*
  - If data is not found, failure is reported back
    - Requestor then tries next closest match in routing table

<i>id</i>	<i>next_hop</i>	<i>file</i>
	:	
	:	
	:	
	:	
	:	

Lecture 22: 11-12-2002

18

## Routing: Freenet Example



Lecture 22: 11-12-2002

19

## Routing: Structured Approaches



- Goal: make sure that an item (file) identified is always found in a reasonable # of steps
- Abstraction: a distributed hash-table (DHT) data structure
  - insert(*id*, item);
  - item = query(*id*);
  - Note: item can be anything: a data object, document, file, pointer to a file...
- Proposals
  - CAN (ICIR/Berkeley)
  - Chord (MIT/Berkeley)
  - Pastry (Rice)
  - Tapestry (Berkeley)

Lecture 22: 11-12-2002

20

## Routing: Chord




- Associate to each node and item a unique *id* in an *uni*-dimensional space
- Properties
  - Routing table size  $O(\log(M))$ , where  $N$  is the total number of nodes
  - Guarantees that a file is found in  $O(\log(M))$  steps

Lecture 22: 11-12-2002

21

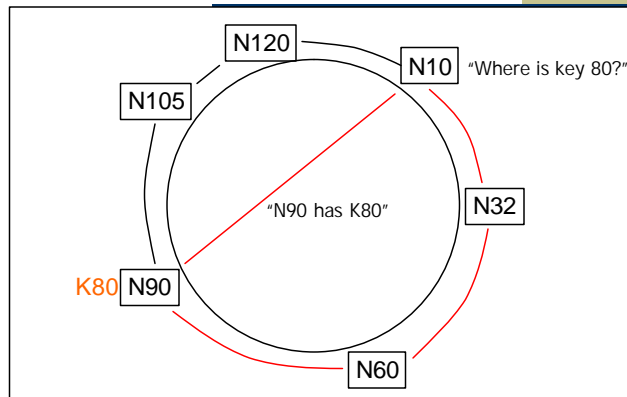
## Aside: Consistent Hashing [Karger 97]



Lecture 22: 11-12-2002

22


## Routing: Chord Basic Lookup



Lecture 22: 11-12-2002

23

## Routing: "Finger table" - Faster Lookups



Lecture 22: 11-12-2002

24

## Routing: Chord Summary



- Assume identifier space is  $0 \dots 2^m$
- Each node maintains
  - Finger table
    - Entry  $i$  in the finger table of  $n$  is the first node that succeeds or equals  $n + 2^i$
  - Predecessor node
- An item identified by  $id$  is stored on the successor node of  $id$


Lecture 22: 11-12-2002

25

## Routing: Chord Example



- Assume an identifier space  $0..8$
- Node  $n_1:(1)$  joins  $\rightarrow$  all entries in its finger table are initialized to itself




Lecture 22: 11-12-2002

26

## Routing: Chord Example



- Node  $n_2:(3)$  joins



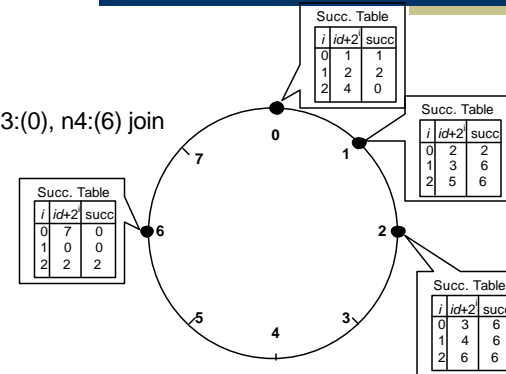
Lecture 22: 11-12-2002

27

## Routing: Chord Example



- Nodes  $n_3:(0)$ ,  $n_4:(6)$  join



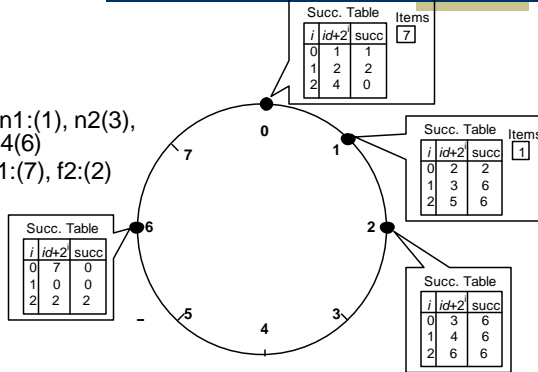
Lecture 22: 11-12-2002

28

## Routing: Chord Examples



- Nodes: n1:(1), n2(3), n3(0), n4(6)
- Items: f1:(7), f2:(2)



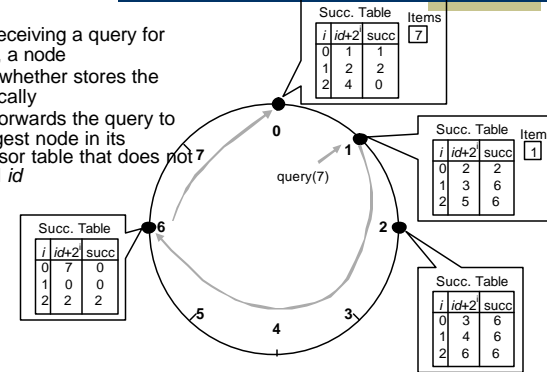
Lecture 22: 11-12-2002

29

## Routing: Query



- Upon receiving a query for item  $id$ , a node
- Check whether stores the item locally
- If not, forwards the query to the largest node in its successor table that does not exceed  $id$



Lecture 22: 11-12-2002

30

## Performance Concerns



- Each hop in a routing-based P2P network can be expensive
  - No correlation between neighbors and their location
  - A query can repeatedly jump from Europe to North America, though both the initiator and the node that store the item are in Europe!

Lecture 22: 11-12-2002

31

## Summary



- The key challenge of building wide area P2P systems is a scalable and robust location service
- Solutions covered in this lecture
  - Naptser: centralized location service
  - Gnutella: broadcast-based decentralized location service
  - Freenet: intelligent-routing decentralized solution (but correctness not guaranteed; queries for existing items may fail)
  - CAN, Chord, Tapestry, Pastry: intelligent-routing decentralized solution
    - Guarantee correctness
    - Tapestry (Pastry ?) provide efficient routing, but more complex

Lecture 22: 11-12-2002

32



## Overview



- P2P Lookup Overview
- Centralized/Flooded Lookups
- Routing-based Lookups
- CMU Research in P2P


Lecture 22: 11-12-2002

33

## What Do Games Look Like?



- Large shared world
  - Composed of map information, textures, etc
  - Populated by active entities: user avatars, computer AI's, etc
- Only parts of world relevant to particular user/player



Lecture 22: 11-12-2002

34

## Individual Player's View



- Interactive environment (e.g. door, rooms)
- Live ammo
- Monsters
- Players
- Game state



Lecture 22: 11-12-2002

35

## Current Game Architectures



- Centralized client-server (e.g., Quake)
  - Every update sent to server who maintains "true" state
- Advantages/disadvantages
  - + Reduces client bandwidth requirements
  - + State management, cheat proofing much easier
  - Server bottleneck for computation and bandwidth → current games limited to about 6000 players
  - Single point of failure
  - Response time limited by client-server latency

Doesn't scale well

Lecture 22: 11-12-2002

36

## Goal: A P2P Multiplayer Game




- Allow 1000's of people to interact in a single virtual world
- Key requirements
  - Robustness: node failures
  - Scalability: number of participants & size of virtual world
  - Performance: interaction quality should only be limited by capabilities of nodes and connectivity between them
  - Extensible: should be possible to add to virtual world

Lecture 22: 11-12-2002

37

## What is Publish-Subscribe ?



- Publishers produce events or *publications*
- Subscribers register their interests via *subscriptions*
- Network performs routing such that
  - Publications “meet” subscriptions
  - Publications delivered to appropriate subscribers

Lecture 22: 11-12-2002

38

## Mercury

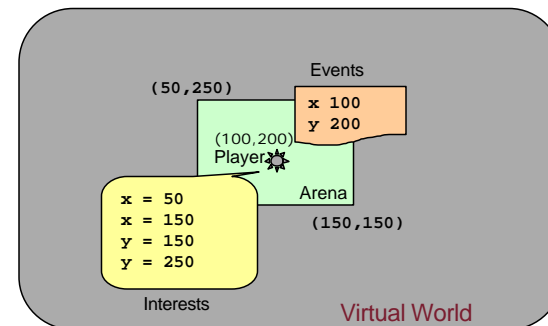


- A P2P publish-subscribe system
- Query language
  - Type, attribute name, operator, value
    - Example: `int x = 200`
  - Attribute-values are sortable
- Sufficient for modeling games
  - Game arenas
  - Player statistics, etc

Lecture 22: 11-12-2002

39

## Modeling a Game



Lecture 22: 11-12-2002

40