



P2P File Sharing



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 (Napster's Approach)

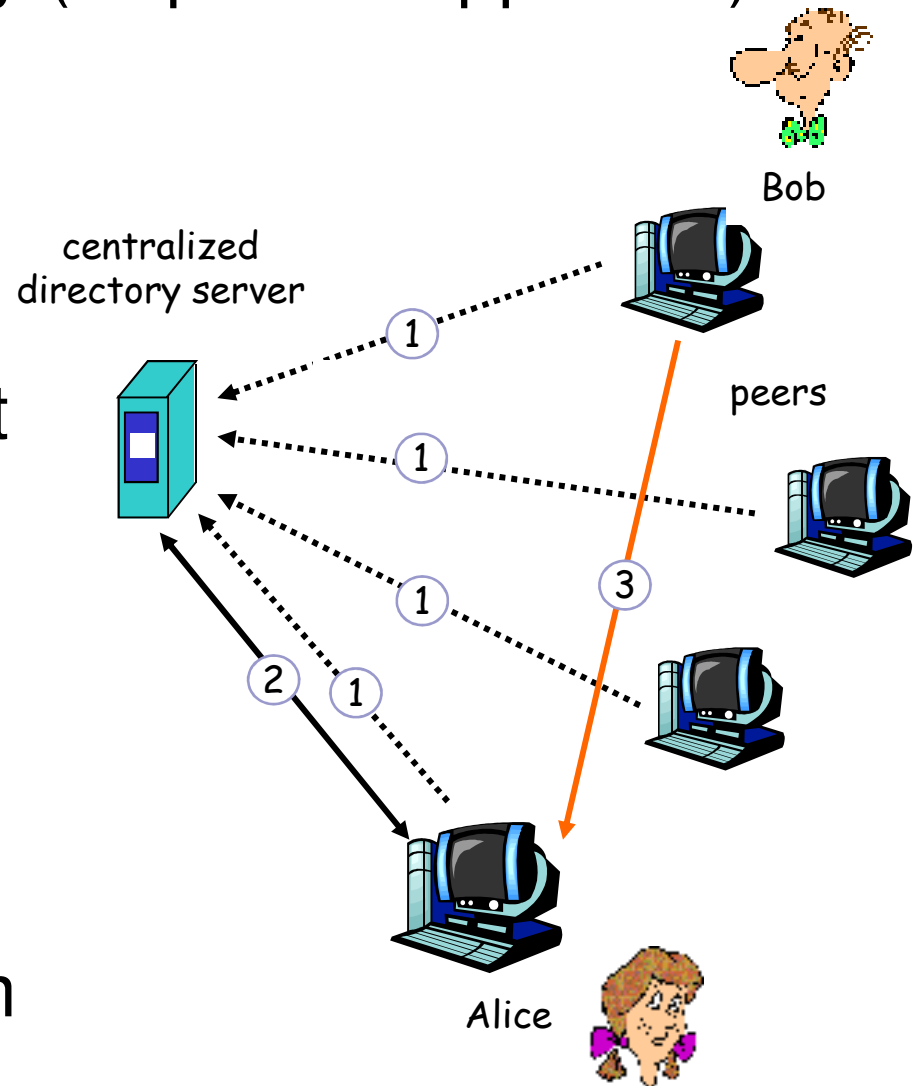
original “Napster” design

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



Query flooding: Gnutella

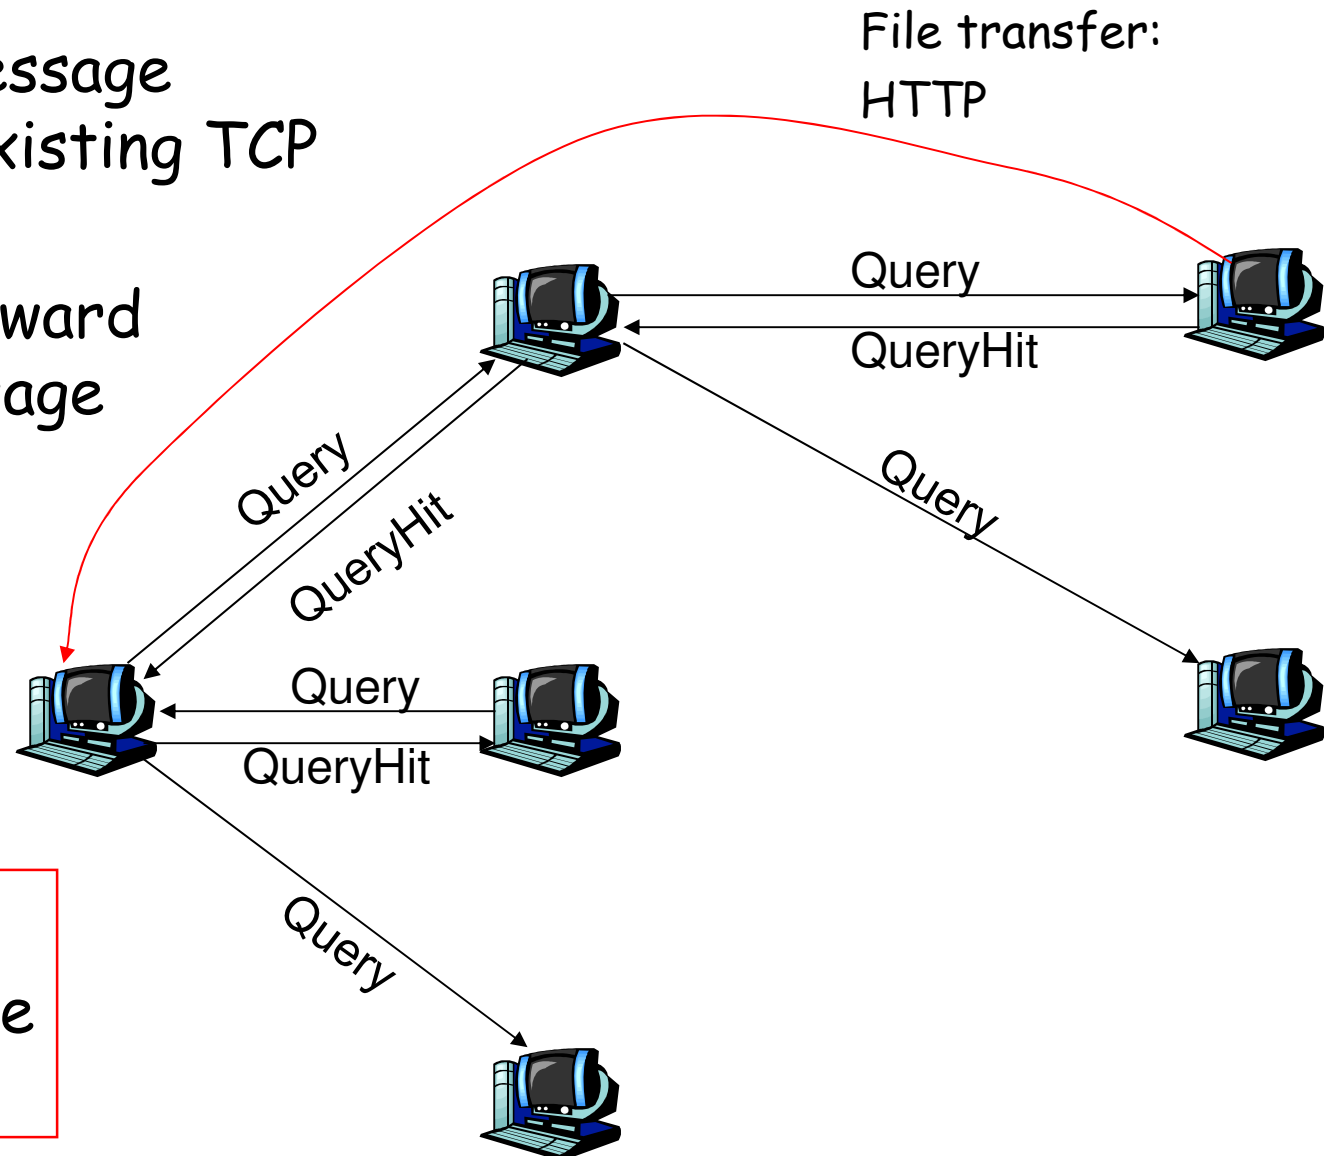
- fully distributed
 - no central server
- public domain protocol
- many Gnutella clients implementing protocol

overlay network: graph

- edge between peer X and Y if there's a TCP connection
- all active peers and edges is overlay net
- Edge is not a physical link
- Given peer will typically be connected with < 10 overlay neighbors

Gnutella: protocol

- ❑ Query message sent over existing TCP connections
- ❑ peers forward Query message
- ❑ QueryHit sent over reverse path



Scalability:
limited scope
flooding




Gnutella: Peer joining

1. Joining peer X must find some other peer in Gnutella network: use list of candidate peers
2. X sequentially attempts to make TCP with peers on list until connection setup with Y
3. X sends Ping message to Y; Y forwards Ping message.
4. All peers receiving Ping message respond with Pong message
5. X receives many Pong messages. It can then setup additional TCP connections

Exploiting heterogeneity: KaZaA

- Each peer is either a group leader or assigned to a group leader.
 - TCP connection between peer and its group leader.
 - TCP connections between some pairs of group leaders.
- Group leader tracks the content in all its children.





KaZaA: Querying

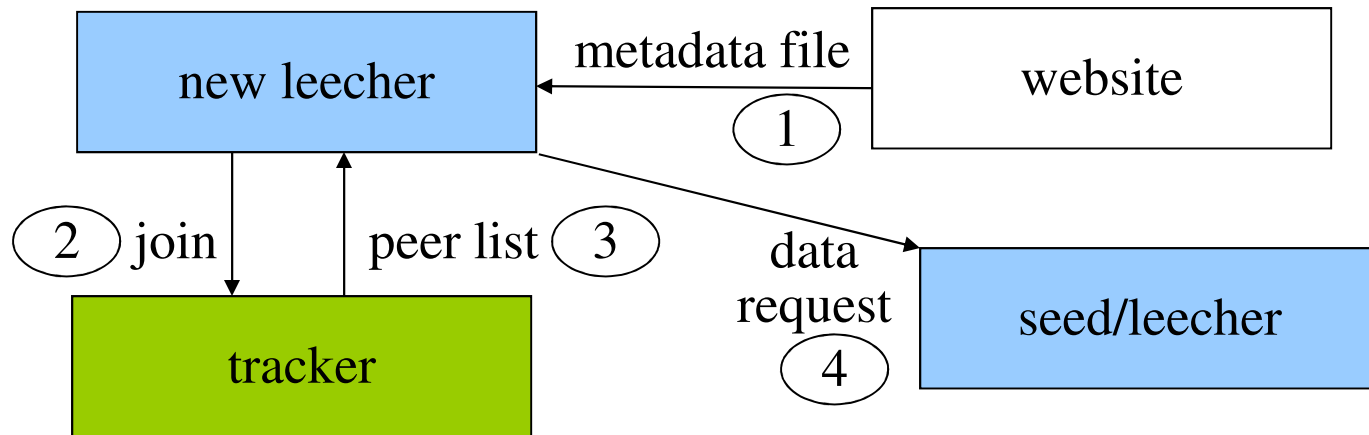
- Each file has a hash and a descriptor
- Client sends keyword query to its group leader
- Group leader responds with matches:
 - For each match: metadata, hash, IP address
- If group leader forwards query to other group leaders, they respond with matches
- Client then selects files for downloading
 - HTTP requests using hash as identifier sent to peers holding desired file



Kazaa tricks

- Limitations on simultaneous uploads
- Request queuing
- Incentive priorities
- Parallel downloading

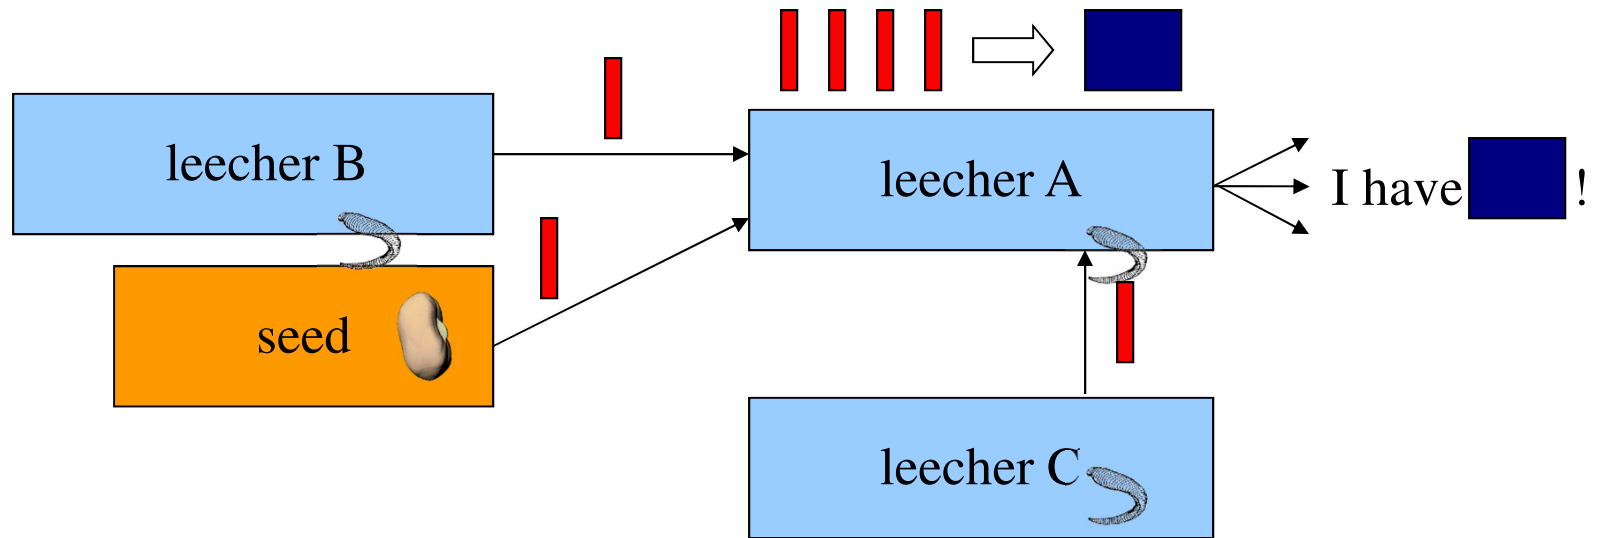
BitTorrent – joining a torrent



Peers divided into:

- 🌱 *seeds*: have the entire file
 - 🐛 *leechers*: still downloading
1. obtain the *metadata file*
 2. contact the *tracker*
 3. obtain a *peer list* (contains seeds & leechers)
 4. contact peers from that list for data

BitTorrent – exchanging data



- Verify *pieces* [blue square] using hashes
- Download sub-pieces [red bar] *in parallel*
- *Advertise* received pieces to the entire peer list
- Look for the *rarest* pieces



Distributed Hash Table (DHT)

- DHT: a *distributed P2P database*
- database has (key, value) pairs; examples:
 - key: ss number; value: human name
 - key: movie title; value: IP address
- Distribute the (key, value) pairs over the (millions of peers)
- a peer **queries** DHT with key
 - DHT returns values that match the key
- peers can also **insert** (key, value) pairs



Q: how to assign keys to peers?

- central issue:

- ☐ assigning (key, value) pairs to peers.

- basic idea:

- ☐ convert each key to an integer
- ☐ Assign an integer to each peer
- ☐ put (key,value) pair in the peer that is **closest** to the key



DHT identifiers


- assign integer identifier to each peer in range $[0, 2^n - 1]$ for some n .
 - each identifier represented by n bits.
- require each key to be an integer in same range
- to get integer key, hash original key
 - e.g., key = **hash**("Led Zeppelin IV")
 - this is why its is referred to as a ***distributed "hash" table***



Assign keys to peers

- rule: assign key to the peer that has the *closest* ID.
- convention in lecture: closest is the *immediate successor* of the key.
- e.g., $n=4$; peers: 1,3,4,5,8,10,12,14;
 - key = 13, then successor peer = 14
 - key = 15, then successor peer = 1


Circular DHT (I)




- each peer *only* aware of immediate successor and predecessor.
- “overlay network”

Circular DHT (I)

$O(N)$ messages
on average to resolve
query, when there
are N peers




Circular DHT with shortcuts



- each peer keeps track of IP addresses of predecessor, successor, short cuts.
- reduced from 6 to 2 messages.
- possible to design shortcuts so $O(\log N)$ neighbors, $O(\log N)$ messages in query

Peer churn



handling peer churn:


- ❖ peers may come and go (churn)
- ❖ each peer knows address of its two successors
- ❖ each peer periodically pings its two successors to check aliveness
- ❖ if immediate successor leaves, choose next successor as new immediate successor

example: peer 5 abruptly leaves

- peer 4 detects peer 5 departure; makes 8 its immediate successor; asks 8 who its immediate successor is; makes 8's immediate successor its second successor.
- what if peer 13 wants to join?

P2P Case study: Skype

- inherently P2P: pairs of users communicate.
- proprietary application-layer protocol
 - inferred via reverse engineering
- Index maps usernames to IP addresses; distributed over SNs
- hierarchical overlay with SNs



Peers as relays

- problem when both Alice and Bob are behind “NATs”.
 - NAT prevents an outside peer from initiating a call to insider peer
- solution:
 - using Alice’s and Bob’s SNs, *relay* is chosen
 - each peer initiates session with relay.
 - peers can now communicate through NATs via relay

