

Distributed Software Development
Introduction to P2P

Chris Brooks

Department of Computer Science
University of San Francisco

Department of Computer Science — University of San Francisco — p. 177

14-2: Peer-to-Peer

- Peer-to-peer is often touted as a hot new technology
 - Really, it's as old as the Internet (or earlier)
- Fundamental idea: all nodes in a network should be capable of the same functionality.
- Contrast this with client-server.
- In practice, many systems are a hybrid of p2p and client-server.

Department of Computer Science — University of San Francisco — p. 277

14-3: Examples

- What are some examples of p2p systems?

Department of Computer Science — University of San Francisco — p. 277

14-4: Examples

- File sharing apps (Morpheus, Kazaa, Limewire, etc)
- IRC/ICQ
- SETI@Home/distributed.net
- Web caching
- VoIP
- SMTP
- etc

Department of Computer Science — University of San Francisco — p. 477

14-5: Historical examples

- USENET
 - Peers within a domain would collect usenet posts and forward them via UUCP, then NNTP.
 - No central server or repository.
- DNS
 - A hierarchical p2p system

Department of Computer Science — University of San Francisco — p. 577

14-6: Why did client/server become so popular?

- Asynchronous network structure
- Ease of discovery
- Centralized content development
- Firewalls, NAT make it harder for users to host content.
- Model changes: fewer developers, more browsers.

Department of Computer Science — University of San Francisco — p. 577

14-7: Potential advantages of peer-to-peer

- Reduce server bottlenecks
- Move content closer to users
- Allow users to publish as well as consume information
 - Note that publishing and authoring can be different things.
- Scalability
- More able to deal with adaptive, dynamic networks.

14-8: Challenges of peer-to-peer

- Finding peers (addressing)
 - How does a peer join the network?
 - How do peers identify each other?
 - Is the network constructed statically, or dynamically?
 - Does a peer always connect to the same set of peers?

14-9: Challenges of Peer-to-Peer

- Searching for data
 - Naming conventions
 - How are queries distributed?
 - Are all peers searched, or just a subset?
 - Do peers keep track of the contents of other peers?

14-10: Challenges of Peer-to-Peer

- Other challenges:
 - Interoperability and standards
 - Currently, every P2P network has its own protocol - you can't use a BitTorrent client to search on Gnutella.
 - Dealing with dynamic networks
 - Security
 - This includes the distribution of malicious files, protection from snooping, and user authentication.

14-11: First-generation applications

- So-called first-generation applications use a centralized name server.
- All user addresses and file indexes are kept there.
 - This can be used to build an application-level address-resolution protocol.
- Napster is the canonical example of this.
- Most IM programs also work this way.
- Central server is used to search.
- Files are transferred between peers.

14-12: Weaknesses of the first-generation model

- Central server provides a bottleneck.
- Difficult to implement in a closed or dynamic environment.
- Scalability can be a problem.
- (also, legal issues in the case of Napster)

14-13: Second-generation models

- Second-generation models remove the central server.
- All file information is distributed.
- Advantage:
 - Potentially more scalable
 - Legal issues avoided.
- Disadvantages:
 - Search is more complex
 - Discovery of peers more complicated.

14-14: Example: Gnutella

- Gnutella is an example of an open P2P system.
- Many clients use the same protocol
 - LimeWire, BearShare, GnuCircus, etc
- Gnutella uses no centralized server - all information is stored at the peers.

14-15: Example: Gnutella

- The Gnutella network is completely decentralized.
- When a peer comes onto the network, it sends a ping to all known peers.
- Those peers respond with a pong, which identifies all their known peers.
- Your peer can also remember previous connections.

14-16: Example: Gnutella

- When a peer wants to perform a search, it sends a query to all known nodes.
 - This request is then forwarded on to nodes one level away, who forward it to all nodes they know about, and so on.
 - This is known as query flooding.
 - If a hit is found, the node containing the file contacts the searcher and download begins.

14-17: Example: Gnutella

- Problems with Gnutella:
 - Searching is unreliable - the network is often partitioned.
 - Search based only on keywords - this produces naming issues.
 - If I want "Lost S1, Ep 3", what do I search for?
 - Wastes bandwidth forwarding searches to all nodes.
 - Does not take advantage of network structure
 - Some peers may have better bandwidth, more information, or a more central network structure.

14-18: Hierarchical p2p

- Kazaa is an example of a hierarchical p2p system
- When a user enters the network, it connects to a meganode.
- This meganode forwards queries to all other nodes connected to it.
- It also collects file information for each client.
- Periodically, it also exchanges information with other meganodes.
- Again, searches typically do not reach the entire network.
- Higher-bandwidth nodes act as meganodes, process more traffic.

14-19: Hierarchical P2P

- This allows peers with greater resources to do more.
- Searches don't need to be propagated more than is necessary.
- Kazaa is a closed network, so most of its inner workings have been figured out through reverse engineering.

14-20: Issues with file distribution

- File integrity - how do you guarantee you're getting the file you think you're getting?
 - md5 or SHA hash of file contents
- How can files be shared anonymously?
 - Virtual addresses and randomized routing can hide identity.
- How to distribute large files?
 - Download chunks simultaneously from many peers at once.

14-21: Example: BitTorrent

- BitTorrent is a p2p network specifically designed for sharing large files.
 - Example: videos from the Jan 2005 tsunami
- Key ideas:
 - File is broken into lots of small pieces that can be independently downloaded
 - Users can upload pieces to other users while downloading.
 - 'Free riders' are punished, sharers are rewarded.

14-22: Example: BitTorrent

- Someone who wants to share a file creates a .torrent file.
 - File name, size, hash of each block.
 - Address of a tracker.
- The torrent is downloaded and the peer registers with the tracker, which provides a list of available peers and seeds.
- The peer begins requesting blocks, starting with the rarest available block.
- As it finishes receiving a block, it begins uploading those blocks to other peers.
- BitTorrent uses a pretty simple tit-for-tat approach to sharing
 - Those four peers who have shared the most data in the last 10 seconds are unchoked.

14-23: Example: BitTorrent

- Difficulties with BitTorrent
 - Works best with files that are widely copied on the network.
 - In practice, files appear and disappear.
 - There is no permanent archive, or incentive for users to keep old files.
 - How to find torrents?
 - No central torrent repository
 - No metadata standard for searching and indexing torrents.
 - Since blocks are not downloaded sequentially, a partial file is not useful.

14-24: Current challenges in p2p

- (Many of these would be interesting projects)
- Efficient search
- Interoperability
- Better metadata and naming schemes
- Punishing free riders
- Anonymity
- User clustering and grouping
- Other interfaces beyond keyword search

14-25: Coming Attractions

- Distributed Hash Tables
- Semantic Overlay Networks
- P2P applied to storage
- Trust in P2P and distributed networks.