


Internet Overlay Networks

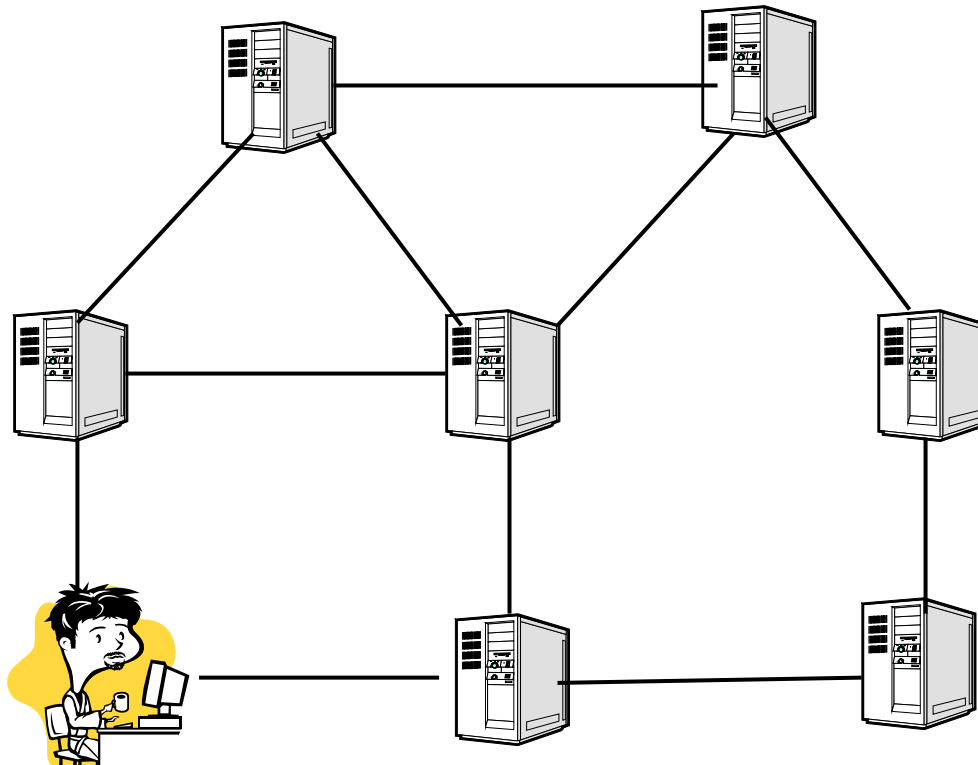


Dept. of Computer Science, University of Rochester

Overlay Networks


Overlay networks: A logical network laid on top of the Internet

- nodes are Internet end hosts
- links are virtual Internet paths



Multicast

- **Multicast:** act of sending data to multiple receivers.

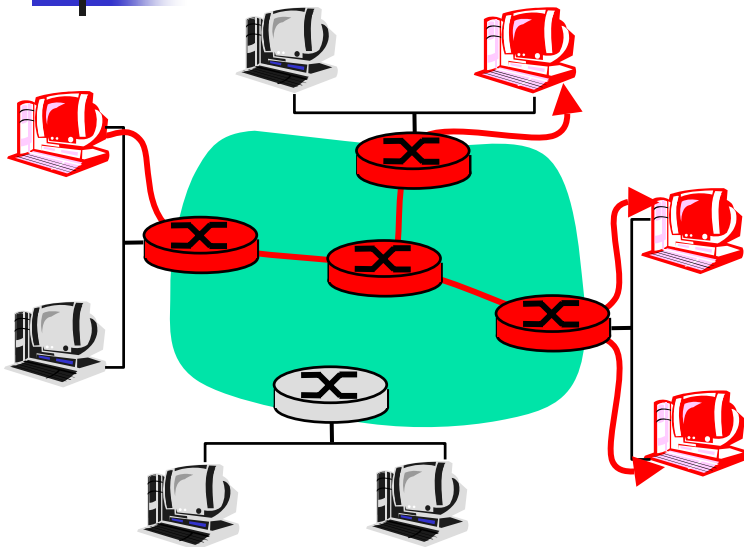


Remember the protocols?


Practical problems of network-layer multicast:

- requires routers to maintain per-group state; complicates the router design and implementation;
- slow to be widely adopted.

Network-layer vs. Overlay Multicast



Network-layer multicast




Overlay multicast

- performance issues for overlay multicast
 - Limit the bandwidth consumption of individual physical links
 - Minimize the source-to-end latency

Overlay Multicast: One Approach

Two steps:

- Forming an overlay mesh network with good links, don't worry about loops at this step
 - too dense a mesh results in too much overhead in the second step; more probable to have high link stress
 - too sparse a mesh restricts what the second step can do
- For each multicast source, generating a shortest path tree or "reverse-path forwarding" on top of the overlay mesh





Revisit the Concept of Overlay Networks

Overlay networks: A logical network laid on top of the Internet.

Key advantages over networks with physical links:

- can be formed spontaneously
- can be easily maintained and changed

Key disadvantage compared with networks with physical links:

- performance

Overlay networks = peer-to-peer networks?



Applications to Overlay Networks


- multicast: network conferencing
- peer-to-peer file sharing
- multi-player interactive games
- software/news/files distribution
-



Multiplayer Interactive Games over Wide-area Networks

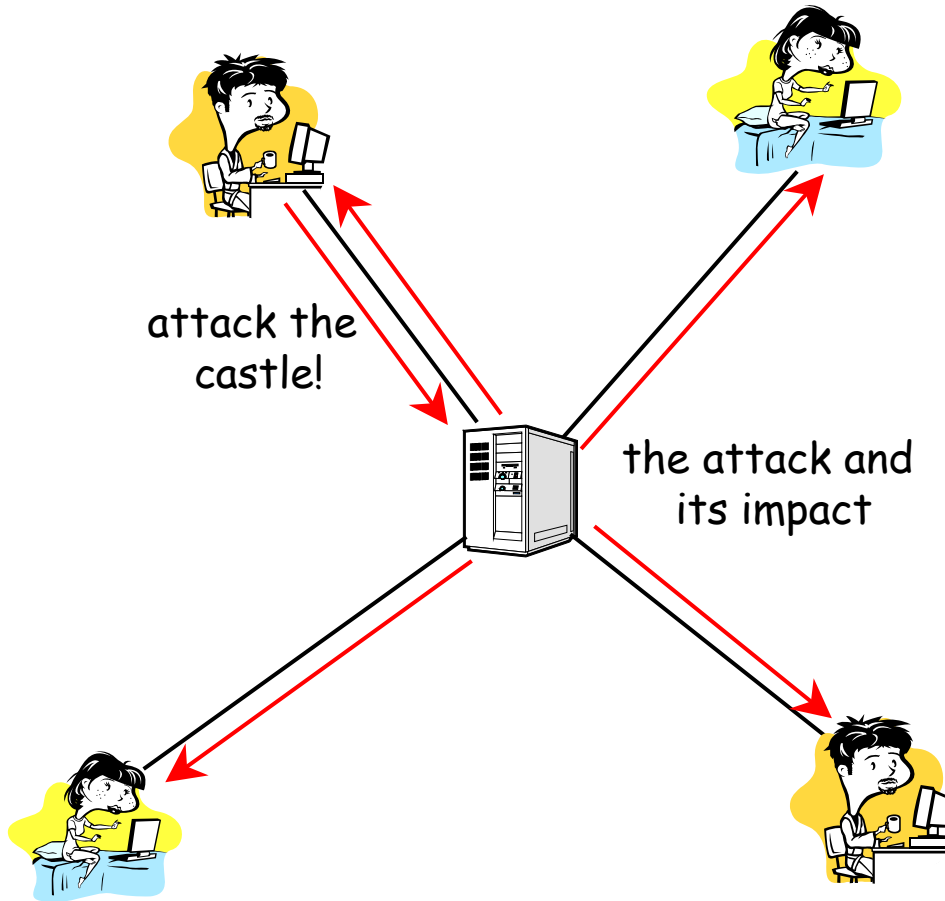
- Some basics:
 - A virtual 2/3-dimensional gaming space with many objects
 - Player-initiated actions: moving your objects, attacking some other objects, etc.
- Main challenges:
 - **Scalability**: supporting many players
 - **Interactive**: real-time constraint in synchronizing the actions and their impact on the gaming space
 - **Wide-area network**: limited network resources

A University of Rochester example



- Alto Trek
- circa 1979
- Authors: Eugene Ball and Rick Rashid at University of Rochester
- Status updates: broadcast via 3-Mbit/sec experimental Ethernet

A straightforward way to do it



- a server and multiple client (each for a player)
- all player-initiated actions get to the server and its impact is propagated to all players (with bounded delay)
- actions can be aggregated when too many come too quickly


any other ideas about how it works?



Scalability


- What are the potential problems when the number of players scales up?
- Network bandwidth connecting the server is the main problem
- How to solve this?
 - multicast
 - distribute the server

Distribute the Server




How do we keep track of which part belongs to which server?

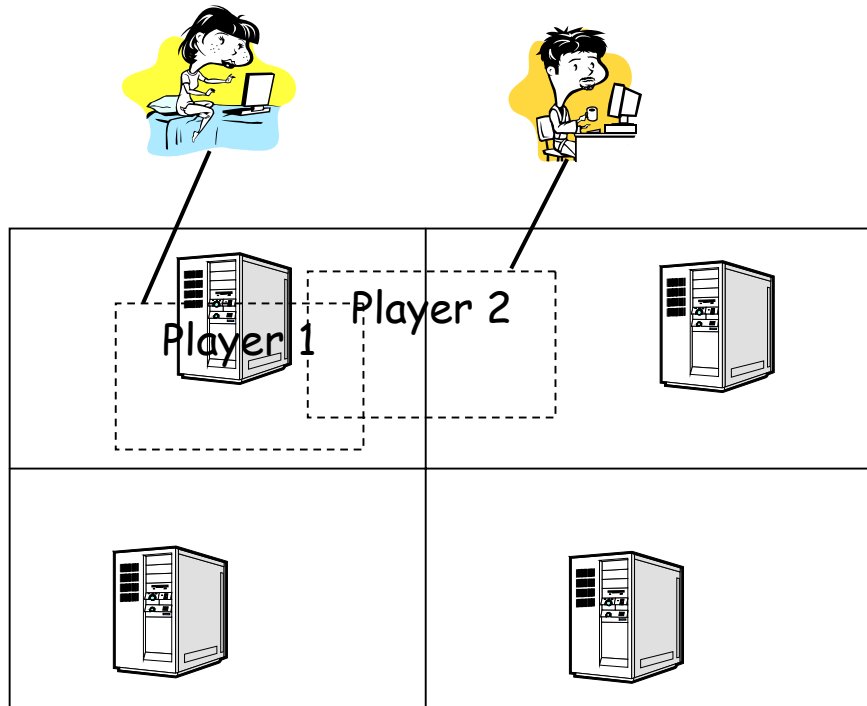
Chord



Content-addressable network



Distribute the Server (cont.)



Additional problems:

- load imbalance
- handoff between partitions is tricky



Files/News/Software Distribution on Demand

- Content distribution:
 - Popular web sites like Yahoo;
 - Breaking news from CNN;
 - Microsoft is distributing a new version of Windows online.
- All downloading from a central content site creates performance problems:
 - Limited bandwidth at the central site;
 - Many users are far away from the central site.



Means of Content Distribution

- By content providers, e.g., www.yahoo.com
 - replicated, distributed Internet sites
- By content consumers (or clients), e.g., UR
 - reduce bandwidth consumption of an organization through caching - Web caching
- By a third-party
 - content distribution network

Replicated Internet Sites

- Content providers build a set of distributed sites on the Internet.
- DNS answers the IP address of the site close to the client

Is DNS caching a problem?

limited scale!

a client at
South America



HTTP request

Yahoo at South
America

Yahoo at North America

Yahoo at
Europe

Yahoo
at Asia

Yahoo authoritative
DNS server

DNS query

IP address of S. America site



Determining Nearby Content Server

- We create a "map", indicating distances between ISPs and content servers
- When a DNS query arrives at authoritative DNS server:
 - server determines ISP from which query originates
 - uses "map" to determine a nearby content server
- How to build the map?
 - pings from content servers to all possible ISPs