ICS 351: Today's plan

- Performance
- Peer-to-Peer Networking

Performance

- performance is measured as both latency and throughput
- latency matters when there are lots of round-trips
- throughput matters when lots of data must be sent
- at each level in the network, it is usually cheaper to provide less performance than more, or more profitable to serve more people with the same bandwidth
- HTTP headers consume multiple hundred bytes, as opposed to about 40-60 bytes for TCP/IP headers
- but nobody is concerned enough to try to compress them
- so people must have lots of bandwidth available!

Improving Throughput

- TCP adjusts to available bandwith,
- so, increase bottleneck bandwidth
- also, reduce congestion:
  - separate traffic as much as possible
  - provide more links and more servers

Improving Latency

- ultimately limited by the speed of light
- identify slow links with traceroute
- find more direct connections
- maybe reconfigure BGP

Improving Reliability and Availability

- redundant links
- automatic failover:
  - routing protocols
  - spanning tree
- redundant power supply, servers, etc (e.g., several IP address per domain name)

Peer-to-peer Networking

- in this class, it has been clear that a PC can perform as a router, and vice-versa
- why not use our PCs as routers?
- for many purposes this is adequate
- if we are not restricted to forwarding to IP addresses, there is less need for hierarchical control
- in fact, we can have content-addressable networks
- perhaps, to increase the incentive, every node that has content and makes it available gets priority for obtaining more content
- the decentralization and lack of control are one of the attractions for many of the users

Peer-to-peer Networking Management

- management of a network requires some authorities that cooperate
- in a peer-to-peer network, both the management and the authority are made as small as possible
- because there is no hierarchical assignment of addresses, each peer can decide what content to provide, i.e. what "address" to use
- this minimizes management
- the network is connected by having one peer exchange addresses about other peers
- then, if the original peer stops collaborating, the other peers can be used to connect to the network
- the expectation is that some useful data exchange will take place, not necessarily that there will be continuous end-to-end connectivity with all parties involved

Peer-to-peer vs. Client-Server

- a server is a program that provides a service
- typically, a server is found at a given address and port number, and is maintained by an individual or an organization
- clients are often anonymous or at least unidentified
- peer-to-peer often brings the same anonymity of a client to the infrastructure
- for a very different example, routing is a peer-to-peer process, but the routers are identified by IP address if nothing else
- there could also be authenticated peer-to-peer networks, set up for the reliability rather than the anonymity aspect
- the ultimate appeal of peer-to-peer for some networking people is a self-organizing, self-managing scalable network
- peer-to-peer networks seem to be widely accepted

Peer-to-Peer Examples

- Bittorrent, eDonkey – 50% of Internet traffic worldwide in 2008/09 (http://www.ipoque.com/sites/default/files/mediafiles/documents/internet-study-2008-2009.pdf)
  - share content
  - unregulated
- FreeNet
  - share content
  - some cryptographic protection