Midterm Review

CS144 Review Session 5
May 2, 2008
Derrick Isaacson
Maria Kazandjieva
Ben Nham

Announcements

- Upcoming dates
  - Sat, 5/3: Lab 3 due with extension
  - Mon, 5/5: In-class midterm
  - Wed, 5/14: Lab 4 due
- Lab 4 is out and is more complex than labs 1, 2, and 3

Midterm Review

1. Applications & app programming
2. IP and forwarding
3. Routing
4. Transport & reliability
5. TCP & congestion control
6. DNS
7. DCCP & NAT
8. Queuing, caching, & content distribution
9. Questions

Applications & app programming

- Telnet
- Web/HTTP
- BitTorrent
- Skype

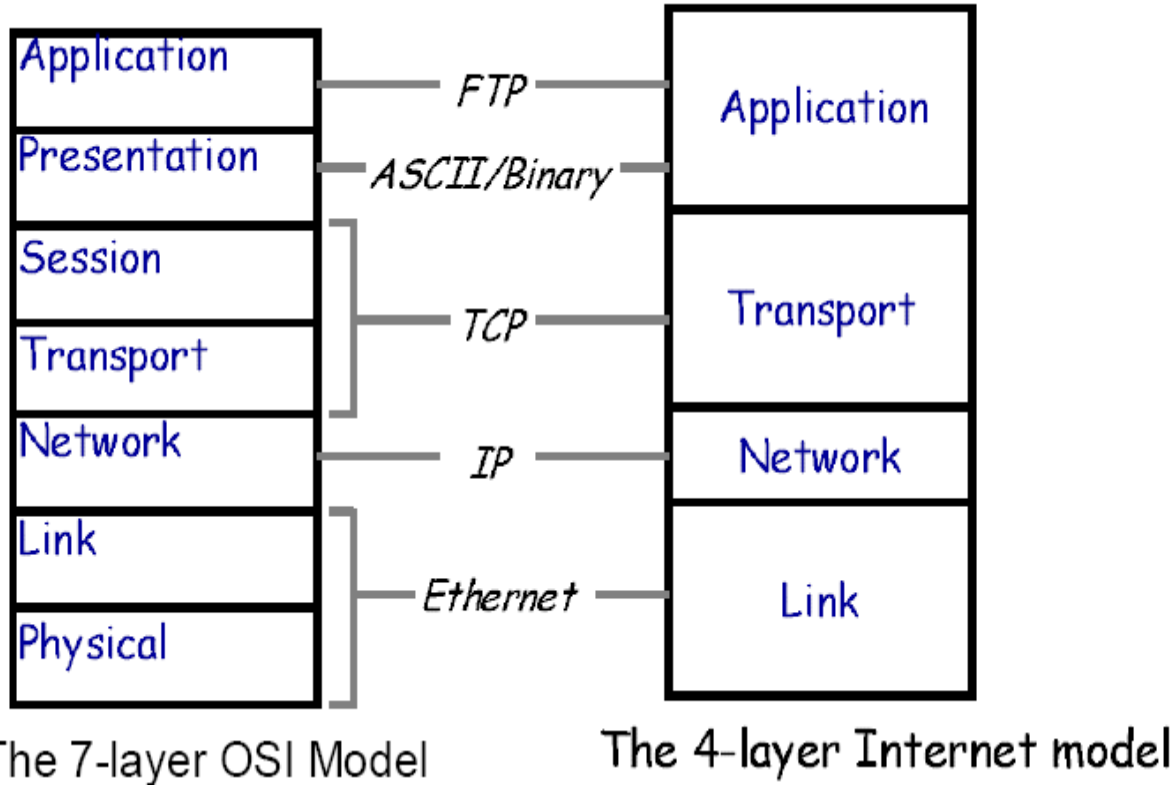

Definitions

- End-to-End principle - dumb network, smart terminals
- Throughput – bits/time
- Latency – time for message to cross network
- Goodput – application level throughput
- Jitter – variation in latency

P2P

- BitTorrent
  - Tracker
  - Tit-for-Tat - “choked,” “top,” and “probing” peers
  - BitTyrant
- Skype
  - Distributed index located at super-peers
  - Ch. 2 R.25 “Skype uses P2P techniques for two important functions. What are they?”
    1. ?
    2. ?

P2P

- BitTorrent
  - Tracker
  - Tit-for-Tat - “choked,” “top,” and “probing” peers
  - BitTyrant
- Skype
  - Distributed index located at super-peers
  - Ch. 2 R.25 “Skype uses P2P techniques for two important functions. What are they?”
    1. Query index (find someone)
    2. Relays – get around NATs

Network layer:
IP, forwarding, routing

Big Picture

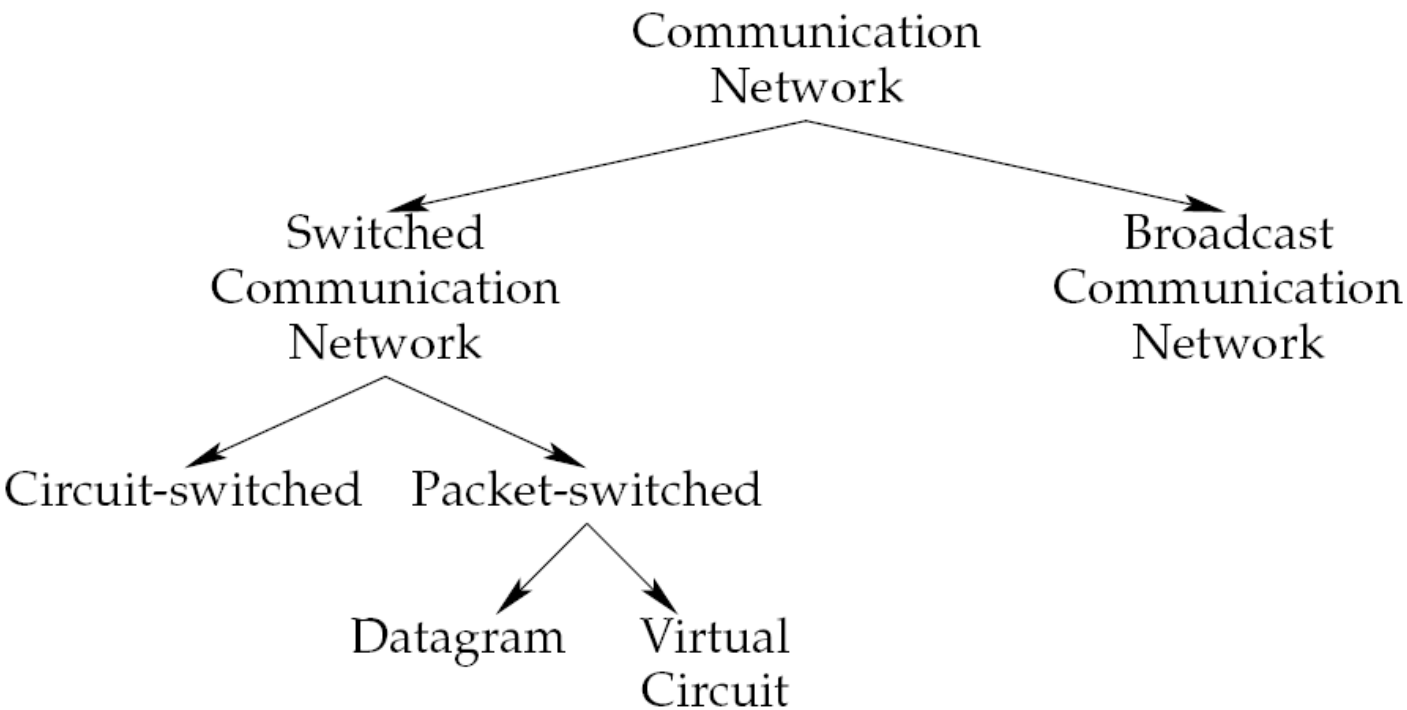

IP

- Skinny waist of internet
- Connectionless
- Best-effort delivery (unreliable)
- TTL
- Framentation
- Classful networks – class A, B, C
- Classless Inter-Domain Routing (CIDR)

Network layer supporting protocols

- ARP – map IP to link layer ethernet address
- ICMP – ping, TTL exceeded, etc.
- DHCP – dynamic IP address assignment
- DNS (later)

Switching

- Optical switch
- Bridge – connect multiple LANs
- VLAN – broadcast to portion of LAN
- Banyan networks – one path from input to output
- Bisection bandwidth – worst-case aggregate bandwidth between two equal-sized sets

Routing

- Forwarding – moving packets from input port to output port
- Routing – populating forwarding table
- Goal – find lowest cost path between two nodes
- Distance Vector (Bellman-Ford) algorithm – local, issues dealing with failures, RIP
- Link State (Dijkstra's) algorithm – global, broadcasts LSP, OSPF
- Path Vector – next slide

Autonomous Systems

- Each AS corresponds to an administrative domain
- Want own intra-domain routing protocol
- Want to set inter-domain routing based on policies (financial, legal, …)
- Issues – transit/peering relationships, local/transit traffic, multihomed,
- Path Vector – BGP, ASPATH announces, allows policy choices based on ASes in path, Multi-Exit Discriminators

Datagram Congestion Control Protocol

- Make UDP play well with TCP
- Connection oriented – Acks for congestion control NOT for reliable transport
- Acks are for last packet received rather than cumulative
- Uses sequence number windows to protect against attacks – need to resynchronize when a large burst of losses cause packets to fall past window
- Congestion Conrol IDs (CCIDs) – what type of congestion control you use
  - CCID 2: TCP CC (AIMD)
  - CCID 3: TCP-friendly CC – uses sending rate rather than congestion window – receiver reports loss rate once per RTT

NAT

- Expand 32 bit address space
- Translates local IP addresses to globally routable IP addresses, and vice versa

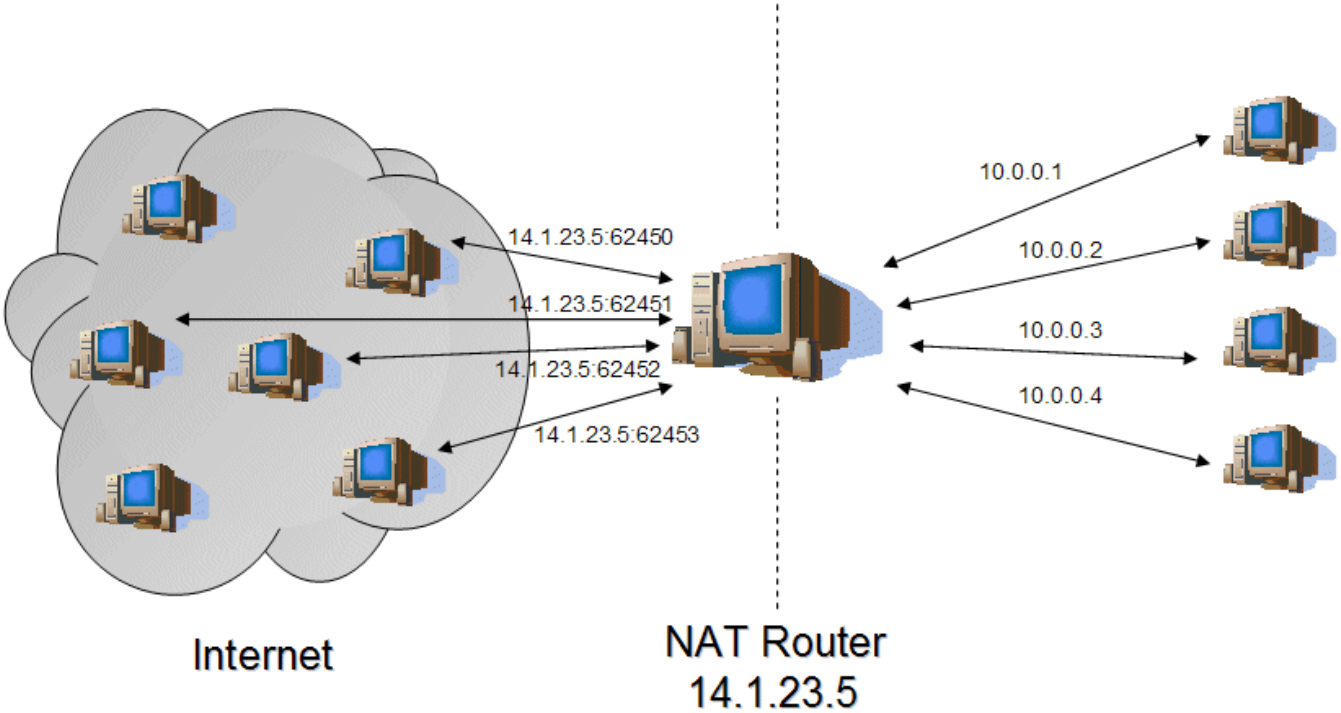

NAT types & problem

- Full cone, restricted cone, port restricted, symmetric
- Problems
  - Incoming connections
  - Port mapping constraints
- Solutions
  - Rendezvous servers (Skype)
  - STUN
  - NAT Hole-Punchingb

Queuing

- Routers handle burstiness of traffic by queuing incoming packets
- Routers try to provide fairness, high throughput, and low delay
- Fair Queuing – provide each flow an equal portion
  - Packets placed in per flow FIFO
  - Calculate F_i for each packet, or time it would complete being transmitted if we served just that flow
  - Next packet to transmit is one with lowest F_i timestamp
  - Packets will be sent out within one maximum packet transmission time of perfect “fairness”

Congestion avoidance

- Random Early Detection (RED)
  - Dropped packets imply congestion in the internet
  - Drop random packets early to fairly signal to all flows that congestion is increasing
  - Use *average* queue length
  - Make drop probability a function of time since last drop to avoid over penalizing one flow.

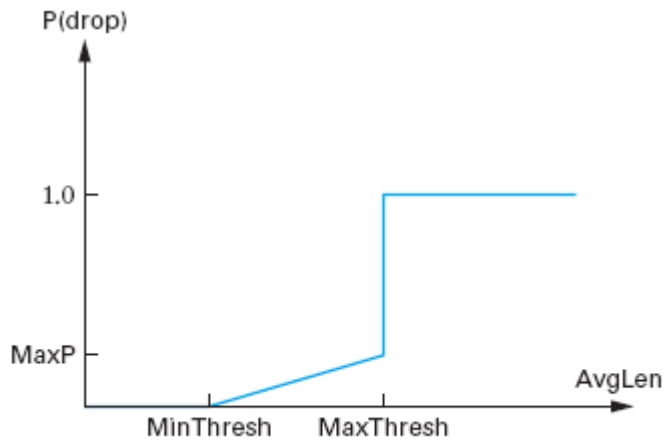

Caching & content distribution

- Web proxies, DNS servers, network file systems use to decrease latency and save bandwidth
- TTL, polling, callbacks, leases to deal with consistency
- Internet Cache Protocol – allow proxies to query each other
- Reverse proxies – content distribution networks like Akamai
- Use some kind of hashing to locate cached data