CS244A Midterm Review

Ben Nham

Some slides derived from:
David Erickson (2007)
Paul Tarjan (2007)

Midterm Information

- In class on Tuesday
  - E-mail staff if you need information on taking it offsite (SCPD)
- One side of notes on letter-sized paper
- Calculator
- Material: Lectures and reading up to Tuesday’s class on congestion avoidance
- Sample midterm on eeclass

Foundations and Basic Concepts

- Layers: Application, transport, network, link
- Circuit switching vs. packet switching
  - Circuit switching: plain old telephone service
  - Two forms of packet switched networks: virtual circuits and datagram networks
- Delays: Propagation , transmission, queuing, processing
- Basic queuing theory
  - Little's Result: $L = \lambda d$

Sample Problem: Latency

- Suppose A is directly connected to B and sends a 10 MB file to B over a 10 km, 100 Mbps fiber link. What is the latency of the transmission?
- Latency = Time from when first bit is transmitted until last bit is received = PROP + TRANSP
- PROP = (10000 m) / (2 * 10^8 m/s) = 0.05 ms
- TRANSP = (10* 2^20 * 8 bits) / (100 * 10^6 bits/s) = 83.89 ms
- Latency = PROP + TRANSP = 83.94 ms

Sample Problem: Little's Result

- You are hungry and get in your car and head to Burger Land. They are a high tech company and have an sign board showing the average arrival rate of cars to their drive through line, and the average number of cars in line. The arrival rate is 2 cars per minute, and there are on average 8 cars in line, and no one leaves the line before they are finished being serviced. How long can you expect to wait in line?
- Little's Result: $L = \lambda d$ → Average queue length = Arrival rate * average wait time
- We will wait for 8 cars / (2 cars/min) = 4 min

Network Layer

- Virtual circuit networks
  - Connection-oriented service
  - Per-flow startup and teardown
- IP
  - Unreliable, connectionless service
  - Address allocation
    - Classful vs. classless (CIDR)
    - Subnetting for internal networks
  - Forwarding
    - Longest prefix match for classless addressing

Routing Algorithms

- Unicast
  - Distance vector (Bellman-Ford): RIP (intra-AS)
  - Link state (Djikstra's): OSPF (intra-AS)
  - Path vector: BGP (for inter-AS)
  - Compare: communication that is needed, convergence issues, centralized vs. distributed
  - Be sure you know how to use the algorithms on paper (examples in routing review session, PS2)
- Broadcast and multicast
  - Sequence number based (e.g., link state advertisements)
  - Reverse path broadcast/multicast
  - Spanning tree-based

Transport Layer

- TCP
  - Connection-oriented: 3 way handshake for setup, 2x2 way handshake at teardown
  - Sequence number: acknowledge up to one past last contiguous byte received
    - TCP is Go-Back-N, not selective repeat
  - Window = min(cwnd, rwnd)
    - Keep sending unacked data up to the window size
  - Transmission rate (ideal): W/RTT

Congestion Control

- Detect drop as not getting ACK before RTO
  - RTO = RTT + guardband (usually some exponentially decaying moving average, perhaps with variance estimate)
- TCP uses Additive Increase, Multiplicative Decrease
  - Increase window size by one maximum segment size for every RTT
  - If we detect a drop, cut window size in half
- "Slow start" optimization
  - Start cwnd at 1 MSS
  - Increase exponentially until a drop, record this window size as *threshold*
  - Drop cwnd to 1 MSS again
  - Increase window exponentially to *threshold/2*, then start AIMD

Sample Problem: TCP Congestion Control

- Assume this TCP implementation:
  - MSS = 125 bytes
  - RTT is fixed at 100 ms (even when buffers start filling)
  - Uses slow start with AIMD
  - Analyze one flow between A and B, where bottleneck link is 10 Mbps
  - Ignore receiver window
- What is the maximum congestion window size?
  - For one flow (ideally), W/RTT = rate
  - W = (100 ms * 10 Mbps) / (8 bits/byte) = 125000 bytes
- How long does it take to reach this size?
  - Slow start grows cwnd exponentially, starting from one MSS
  - Find n s.t. 125 * 2^n = 125000 => n = 10
  - Then it takes n * RTT = 1s to reach the max cwnd size

Congestion Avoidance

- DECbit
  - If average queue length over last busy + idle + current period is > 1, set a bit in reply
  - If congestion bit is on >50% time: multiplicative decrease
  - Else: Additive increase
- Random Early Detection (RED)
  - Set min and max thresholds
  - If average queue size > max, drop
  - Else if average queue size < min, service
  - Else if average queue size between min and max:
    - Randomly drop packet with increasing probability as queue size grows to max