

Final Exam Study Guide

The following is meant to give you an idea of which topics from the second half of the course might potentially appear on the final. The first half of the course is covered in the “Midterm Study Guide” document and those might appear on the final as well.

Switching

- How the Take-a-Ticket approach to switching works
- Why Head-of-Line Blocking is bad and how to avoid it
- Parallel Iterative Matching (PIM) and iSLIP algorithms for switching

TCP

- High-level goals of TCP
- Purpose of port numbers and how they are used in conjunction with IP addresses to identify a TCP connection (including within the context of Network Address Translation, or NAT (where we have private and public networks))
- How sequence and ack numbers are used to tell when a packet is lost or reordered
- Purpose of a few TCP header flags: SYN, ACK, FIN
- Process to start a TCP connection (3-way handshake) and stop a TCP connection (4-way handshake)
- Congestion control: how window size changes with each successful (acknowledged) packet transmission (be able to understand graph, how window size changes depending on whether congestion was detected via duplicate acks or timeout)

Output Scheduling

- Random Early Detect (RED): how algorithm decides with what probability to drop and why it's beneficial to drop packets even if buffer is not full
- Token Bucket: shaping vs. policing – how token bucket algorithm works at a high level, and when you would use shaping vs. policing (is there a single output queue or one for each flow?)
- Deficit Round Robin: fair management of multiple flows where each gets a “salary”

Sockets, Application Layer, DNS, and HTTP

- Goals of sockets and application layer
- Each of the servers that will be searched (and in what order) when a DNS request is issued, and what will be cached
- HTTP: format of request, response messages
- Structure of a URL