

Each 6.1800 lecture will come with an outline. You can fill this in during lecture, after lecture, or not at all — it's entirely up to you how you use it. The goal of these outlines is to help you understand the main points that you should be taking away from each lecture. In some cases we will also include examples of things you should be able to do after each lecture.

In the past, these outlines have proved to be an effective tool for studying for the exams. Note that the outlines are **not exhaustive**; there will be topics and nuances in lecture that aren't captured by the outline.

Lecture 12: In-network resource management

- Queue management
 - How does DropTail queueing work? What are its pros and cons?
 - What do we mean by flow synchronization? Why does this happen with DropTail queueing?
 - How does RED work? What are its pros and cons?
 - What limitation of DropTail does RED address?
 - How does ECN work? What is the point of it?
- Delay-based Scheduling
 - How can priority queueing help serve applications that require low latency?
- Bandwidth-based scheduling

For all of these: you should be able to work through the sorts of examples that Katrina will do in class.

 - How does round-robin scheduling work?
 - What is the main problem with round-robin scheduling?
 - How does deficit round-robin scheduling work?
 - How does one set the values of the quantum?
 - Why doesn't credit accumulate for a queue if the queue is empty?
- Is all of this a good idea on the Internet? Who should make the relevant decisions?