

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 11: TCP

- What is reliable transport? What does it mean to be a “reliable transport protocol”?
- In what ways can a network be unreliable? (i.e., what can go wrong)
- Basics of reliable transport:
 - What are sequence numbers for? How do they work?
 - What are acknowledgements (ACKs) for? How do they work?
 - How does a TCP sender use a timeout to infer that a packet has been lost?
 - How does a TCP receiver decide whether to deliver a packet to the receiving application?
- Congestion control: the basics
 - What is congestion control for?
 - What are the goals of congestion control?
 - What do we mean by the “window”?
 - TCP's congestion control is AIMD: additive increase multiplicative decrease
 - What does this mean? How do senders react to congestion?
 - What is the intuition behind AIMD? (example: why not multiplicative increase multiplicative decrease? Or additive increase additive decrease? etc.)
- Congestion control: additional mechanisms
 - How does slow start work, and why is it used?
 - How does fast retransmit/fast recovery work, and why is it used?
 - Why does TCP react to loss differently with fast retransmit/fast recovery than with a timeout?
- There is a lingering problem with TCP: what is it? (This is another way of asking: what are we going to talk about in the next lecture?)