

*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 07: Performance, Storage**

- What do we mean by a performance bottleneck?
- Why is it helpful to have a model of a system when we're thinking about its performance?
- What are common performance metrics? What do each of them mean? How do they relate to one another?
- What are examples of systems-level techniques that often improve performance?
- How do reads/writes to a disk work...
  - ...for hard-disk drives (HDDs)?
  - ...for solid-state drives (SSDs)?
- Why does batching reads on HDDs improve performance?
- Suppose we were going to implement a huge database as a series of files (so in some ways implementing a database "on top" of a filesystem). What sorts of things would we have to consider?
- What is a database management system (DBMS) good at?

*After this lecture, you will **not** know everything about how a database works! This was a quick introduction in the context of storage and performance.*

*You should also consider this lecture alongside the following recitation. In lecture, we (mostly) focused on performance improvements for a single machine. In recitation, you'll focus on performance improvements for distributed systems.*