

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 17: Logging

- We're working with atomicity today. What does it mean for an action to be atomic?
- Motivation: shadow copies (the topic of the last lecture) provide atomicity in some situations. What is the main problem with shadow copies?
- Basics of logging (*at this point we have no cell storage, no cache*)
  - What does each type of log record (e.g., UPDATE, COMMIT) contain?
    - *Given a sequence of transactions, you should be able to create the associated log*
  - How is a log used to read the value of a particular variable?
  - What is the performance of the log like?
- Adding cell storage
  - How do reads and writes work with the addition of cell storage?
  - Is cell storage on disk (and thus nonvolatile) or in memory (and thus volatile)? Why does this matter?
  - Why must logging be done before installing?
  - How do we repair cell storage after a crash?
    - *Given a log + cell-storage setup that crashes, you should be able to repair cell storage after recovery*
  - What is the performance of the log like with the addition of cell storage?
- Adding a cache
  - How do reads and writes work with the addition of cell storage *and* a cache?
  - Is the cache on disk (and thus nonvolatile) or in memory (and thus volatile)? Why does this matter?
  - How do we repair cell storage after a crash when there's a cache involved?
    - *Given a log + cell-storage + cache setup that crashes, you should be able to repair cell storage after recovery*
  - What is the performance of the log like with the addition of a cache?
  - What is the purpose of truncating the log + writing checkpoints?