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 15: Reliability

- What is our general, high-level approach for handling faults?
- How do we measure availability using mean-time-to-failure and mean-time-to-repair?
- We're focusing on improving the reliability of disks in this lecture. Why disks?
- RAID
    - Given two bitstrings, A and B, how do we compute A XOR B? Given A XOR B and A, how do we compute B?
    - What about given three bitstrings, A, B, and C; how do we compute A XOR B XOR C?
    - How do these ideas extend to RAID?
    - How does a RAID system recover from a single failed disk?
    - How do reads and writes work in…
        - … RAID 1?
        - … RAID 4?
        - … RAID 5?
    - What are the performance trade-offs between the different levels of RAID?
    - Can the RAID systems described in lecture recover if two disks fail at the exact same moment? Why or why not?
    - *Bonus question: Can they recover if one disk fails, and then another fails one second later? Ten seconds later? In general, how would you think about the required time-between-failures for these systems to recover?*