

MIT OpenCourseWare  
<http://ocw.mit.edu>

6.033 Computer System Engineering  
Spring 2009

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.

## Preparation for Recitation 8

- Read MapReduce.
- Skip sections 4 and 7.

This paper was published at the biennial Usenix Symposium on Operating Systems Design and Implementation (OSDI) in 2004, one of the premier conferences in computer systems. (OSDI alternates with the equally prestigious ACM Symposium on Operating Systems Principles (SOSP), at which appeared Eraser, the paper you read for the last recitation.)

As you read the paper, keep the following questions in mind:

- At first glance, the map/reduce model of computation seems limited. Did the paper persuade you that their model of computation has practical use?
- Are the authors trying to solve a technological problem (one that will be solved with faster computation), or an intrinsic problem?
- What assumptions do the authors make about how machines fail, what machines fail, and what they do when they fail? What happens to the system when a given machine fails?
- What exactly would happen if one block of one hard drive got erased during a map/reduce computation? What parts of the system would fix the error (if any), and what parts of the system would be oblivious (if any)?
- How do the authors evaluate the performance of their system? What are "Input," "Output," and "Shuffle?"
- How do "stragglers" impact performance?

Here are some points to keep in mind as you read:

- In the functional programming notation used in Section 2.2, the function takes the arguments shown to the left of the arrow and returns the type shown to the right of the arrow.
- After you read Section 3.1, you should be able to instantly recall the following terms: "split," "map worker," "reduce worker," "master."