


## MapReduce

- Do you like it? Dislike it? Why?
- Map()/Reduce() interface
  - Is it the right interface?
  - What are the important criteria?
- MapReduce Power
  - Correlated high load is likely at beginning of map, which means high overall-system peak power

4/18/2011 CSC 258/458 - Spring 2011 1



## Data Center Computing

- New issues
  - Unprecedented scale, simple commodity hardware
- Revisiting parallelism
  - MapReduce
- Revisiting structured data management
  - BigTable
- Revisiting unstructured data management
  - Google File System
- Revisiting distributed fault tolerance
  - ChubbyLock - distributed lock, an implementation of the Paxos algorithm

4/18/2011 CSC 258/458 - Spring 2011 2