CS378H Honors Concurrency
Midterm 2 Faux Quiz Compilation
Fall 2018

- Why/when might one prefer an FPGA over an ASIC, CPU, or GPU?
- Define CLB, BRAM, and LUT. What role do these things play in FPGA programming?
- Describe the FPGA build process; which phases are relatively short vs. relatively long?
- Define CLB, BRAM, and LUT. What role do these things play in FPGA programming?
- What is the difference between blocking and non-blocking assignment in Verilog?
- What is the difference between structural and behavioral modeling?
- How is synthesizable Verilog different from un-synthesizable? Give an example of each?
- What is discrete event simulation?
- What is the difference between horizontal and vertical scaling? Describe a setting or application in which horizontal might be preferable and one in which vertical scaling makes more sense.
- What is PGAS?
- What is bisection bandwidth?
- What is a "shared nothing" architecture?
- What is an All-reduce operation? In MPI can it be implemented with other primitives (e.g. send/recv)? Why does MPI support it as a first class API?
- What is distributed shared memory? Suggest some implementation techniques.
- What is a collective operation? Give an example and explain why it is a collective.
- What is a 3-Tier architecture? What application (s) might be it good for and why?
- What are some advantages and disadvantage of distributed memory architectures?
- What is the difference between 1-sided and cooperative MPI operations? Advantages/disadvantages of each?
- What phenomena can slow down a map task?
- Do reducers wait for all their mappers before starting? Why/why not?
- What machine resources does the shuffle phase consume most?
- Is it safe to re-execute a failed map/reduce job sans cleanup? Why [not]?
- How does MR handle master failure? What are the alternatives?
- What is the difference between *transformations* and *actions* in Spark?
- Spark supports a persist API. When should a programmer want to use it? When should she [not] use the "*RELIABLE*" flag?
- Compare and contrast fault tolerance guarantees of Spark to those of MapReduce. How are[n't] the mechanisms different?

- Is Spark a good system for indexing the web? For computing page rank over a web index? Why [not]?
- List aspects of Spark's design that help/hinder multi-core parallelism relative to MapReduce. If the issue is orthogonal, explain why.
- What is the CAP theorem? What does "PACELC" stand for; how does it relate to CAP?
- What is the difference between ACID and BASE?
- What is causal consistency?
- What is chain replication?
- Why do NoSQL systems claim to be more horizontally scalable than RDMBSes? List some features NoSQL systems give up toward this goal.
- What is eventual consistency? Give a concrete example of how of why it causes a complex programming model (relative to a strongly consistent model).
- Compare and contrast Key-Value, Document, and Wide-column Stores
- Define and contrast the following consistency properties: strong consistency, eventual consistency, consistent prefix, monotonic reads, read-my-writes, bounded staleness
- What is obstruction freedom, wait freedom, lock freedom?
- How can one compose lock free data structures?
- Why should I want a lock free hash table instead of a fine-grain lock-based one?
- What is the difference between linearizability and strong consistency? Between linearizability and serializability?
- What is the ABA problem? Give an example.
- How do lock-free data structures deal with the "inconsistent view" problem?
- What is a data race? What kinds of conditions make them difficult to detect automatically?
- What is a consistent cut in a distributed causality interaction graph?
- List some tradeoffs between static and dynamic race detection
- What are some pros and cons of happens-before analysis for race detection? Same for lockset analysis?
- Why might one use a vector clock instead of a logical clock?
- What are some advantages and disadvantages of combined lock-set and happens-before analysis?