

CSE 552: Distributed and Parallel Systems

Problem Set 1

Due: October 4, 5pm

1. A standard implementation of a remote procedure call is to provide “at most once” semantics. Explain why an RPC designer might not implement the more natural sounding “exactly once” semantics instead.
2. Provide high-level pseudo-code for the coordinator of a two-phase commit protocol. (This should take less than a dozen lines.) Annotate each line with the recovery procedure that is needed if the coordinator fails immediately before the line executes. Mark the earliest point at which the transaction commits.
3. Explain how MapReduce can be used to compute PageRank.
4. Are updates to DNS sequentially consistent? Why or why not?