

# CSE 444 Practice Problems

## Distributed DBMS

### 1. Two-phase Commit

- (a) In the two-phase commit protocol, what happens if the coordinator sends PREPARE messages and crashes before receiving any votes?
  - i. What is the sequence of operations at the coordinator after it recovers.

- ii. What is the sequence of operations at a subordinate that received the message and replied to it before the coordinator crashed.

iii. What is the sequence of operations at a subordinate that did not receive the message before the coordinator crashed.

(b) In the two-phase commit protocol, why do subordinates need to force-write a prepared log record before sending a YES VOTE? To answer this question, use an example failure scenario. Show what happens if a subordinate does NOT force-write the prepared log record, then show what happens if the subordinate does force-write the prepared log record.

## 2. Replication

- (a) In eager master replication, when the master fails, why does a group of replicas need to have the majority of nodes in order to elect a primary and continue processing requests?
  
- (b) What are the differences between eager and lazy replication? Please list differences in the approaches and differences in the properties that result. Discuss master vs group replication if appropriate.