

This homework is due at 11:59:59 PM on December 4, 2020 and is worth 3% of your grade.

Name: _____

CCIS Username: _____

Problem	Possible	Score
1	35	
2	10	
3	40	
4	15	
Total	100	

1a. State the minimal series of failures that are necessary to deadlock a system using two-phase commit (i.e., bring the system to a state where it cannot make progress or accept new updates). You may assume that replicas will elect a new leader if the old leader fails. (10 pts)

1b. Explain why the series of failures in the previous question deadlock the two-phase system. (10 pts)

1c. Briefly explain how a three-phase commit protocol addresses the deadlock issue raised in the previous two questions. (15 pts)

2. Recall Eric Brewer's CAP Theorem: No system can simultaneously achieve Consistency and Availability and Network Partition Tolerance. Explain why, in practice, you have to choose between Consistency and Availability? (10 pts)

3a. What is a Byzantine fault (or failure)? Why might a Byzantine failure be more difficult to deal with than fail-silent one? (10 pts)

3b. What is a possible outcome between a network partition fault? (5 pts)

3c. What is one way in which Raft differs from the Paxos implementation discussed in class? (5 pts)

3d. Why do Raft leaders send empty AppendEntry messages to each server? (5 pts)

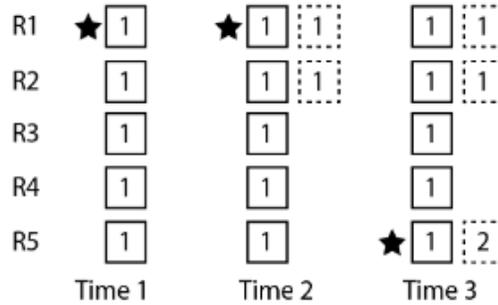
3e. What does a Raft follower do when an election timeout occurs?

(5 pts)

3f. How does a Raft election start? Who is generally the winner?

(10 pts)

4. The next questions refer to the following image, which shows a series of PUTs to a Raft system with five replicas (R1 through R5). Each box represents a PUT, and the number in the box represents the term under which the PUT happened. Solid boxes indicate that the PUT is fully committed (i.e. applied), while dashed boxes indicate that the PUT is not fully committed/applied yet. The stars indicate which replica was the leader during each time interval.



4a. Which replicas are eligible to become leader in the next time interval? (5 pts)

4b. Depending on who becomes leader in the next time interval, there are two possible outcomes for the global log. Draw the two possible final logs. You only need to show the log for a single replica in each case. Give the series of steps needed to produce each of the possible logs (e.g. “replica X is elected leader, append entry #Y, etc”). (10 pts)