

# CSCI-1380: Distributed Computer Systems

## Homework #4

Assigned: 03/20/2021

Due: 03/27/2021

### 1 Replication (Ordering)

As the distributed systems expert and leader of the platforms team in your startup, you are tasked with developing a new replication scheme (Passive and Active) for your company's storage layer. In your startup, external clients will interact with the web service (i.e., FEs) which interacts with the storage layer (i.e., RMs). In your answers please use FE and RM terminology.

- Your application developers (i.e., FE developers) are quite concerned about general write latency between the FE and the RMs: they want minimal latency during failure and none failure scenarios. Which technique will you use? Please explain your choice.
- The application team is concerned about ordering. They want an ordering of events across databases that is total ordered, FIFO-ed, and **also real time**. Here real-time is defined: as the time that the FE generates the request. Essentially, the database should process request in a manner that is total ordered (i.e., all RMs process requests in same order), FIFO (i.e., different requests from a specific FE are processed in the order the FE generates them), and real time (i.e., requests from different FE are ordered based on the real-time). Does your replication scheme provide all of these (Total, FIFO, real time)? If not, which subset does it provide? Explain.

## 2 Raft

- Given two log entries, how does Raft determine which one is more up-to-date? Your answer can directly copy from the paper. (Refer to Section 5.4.1 of the [Raft Paper](#))
- Describe the process in which the leader brings all of its follower's logs into consistency with its own when it first comes into power. Use a list that shows each step in the process. (Refer to Section 5.3 of the Raft Paper)
- Raft has some clear rules about what can be considered committed and also which candidate can be elected as a leader.

- The figure below shows the state of the logs in a cluster of 5 servers: S1 was the leader for term 5, now consider that S1 crashes.

log index	1	2	3	4	5	6	7	8
s1	1	1	2	2	2	2	5	
s2	1	1	1	3	4			
s3	1	1	2	2	2			
s4	1							
s5	1	1	2	2	2	2		

- Which nodes can become leader?
  - If node S5 tries to become leader, which nodes will vote for it? Explain your answer.
  - What is the definition of a majority in this scenario? Explain your answer.
  - In the current scenario, which entries are considered committed? Explain your answer.
- The figure below shows the state of the logs in a cluster of 7 servers: S5 was the leader for term 7, now consider that S5 crashes.

	1	2	3	4	5	6	7	8	9	10	11	12
S1	1	1	1	4	4	5	5	6	6	6		
S2	1	1	1	4	4	5	5	6	6			
S3	1	1	1	4								
S4	1	1	1	4	4	5	5	6	6	6		
S5	1	1	1	4	4	5	5	6	6	7	7	
S6	1	1	1	4	4	4	4					
S7	1	1	1									

- Which nodes can become leader?
- If node S4 tries to become leader, which nodes will vote for it? Explain your answer.
- What is the definition of a majority in this scenario? Explain your answer.
- In the current scenario, when S4 tries to repair the log for S5, which indices will be replaced..

### 3 Handing In

Once finished, you should hand in a PDF with your answers on Gradescope. Gradescope will allow you to select which pages contain your answers for each part of each question.

**Please do not put your name on any page of your handin!** This will allow us to do fully anonymized grading through Gradescope.

Please let us know if you find any mistakes, inconsistencies, or confusing language in this or any other CS138 document by filling out the [anonymous feedback form](#).