

# CSCI-1380: Distributed Computer Systems

## Homework #2

Assigned: 03/05/2020

Due: 03/16/2020

### 1 Tapestry

1. In the table, we specify the size of the ID space (i.e.,  $N$ ) and the Base used within a specific Tapestry network. For each such network, what is the size of the table (assume each cell only maintain one entry, i.e.,  $c=1$ )? What is the maximum number of hops (RPC calls)?

N	Base	Size of table	max # of hops
128	2		
256	4		
1024	4		
4096	16		

2. What is the soft-state? How does Tapestry use soft-state to maintain fault tolerance during network failures and membership changes?
3. Given a tapestry network with Base 8 and ID space of 512. The current node IDs are: 027, 712, 621, 521, 000, 121, 143, 164, 167.

(a) Write down a possible routing table for node 000 (if there are multiple possibilities for an entry, you can choose arbitrarily).

(b) Find the root for each key from certain starting node. List the path (including the starting node):

Key ID (for look)	Starting node	Root	Lookup path
777	000		
124	712		
452	121		
672	621		

(c) Which nodes are in the backpointers for node 712?

(d) If node with ID 531 were added into the network, in which nodes will 531 fill an empty hole?

## 2 Replication and Consistency

1. Given a replication system with  $N$  nodes, how many messages are exchanged? (assuming no packet lost or re-transmitted)

N	Active Replication	Passive Replication
7		
19		
3		
101		

2. Your replication scheme uses active replication and your network re-orders over 50% of the packets. How will active replication deal with packet reordering? Will this impact consistency or correctness?
3. All the servers in your cluster can handle 1000 requests per second. You have setup the DHT in your cluster to use passive replication with 5 replicas. There are 100 servers in your cluster. Assuming equal distributed of request across all keys, what do you anticipate the cluster throughput to be. Why?

### 3 Raft

1. 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](#))
2. 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)
3. You have a raft cluster with 5 nodes,  $N_1, N_2, N_3, N_4, N_5$ . The leader is  $N_5$ . A network partition occurs dividing the cluster into two groups. In one group  $N_1, N_5$ , in the second group  $N_2, N_3, N_4$ .
  - (a) After this partition occurs, what happens to  $N_5$ ? Is it still a leader and will it be able to commit?
  - (b)  $N_4$  is first to have its timeout fired. How many votes does  $N_4$  need to become the leader (including the vote from itself)?
  - (c)  $N_4$  became the leader. After the network partition is healed and all nodes are able to communicate, what does  $N_4$  do to  $N_1$  and  $N_5$ ?
4. Your Raft cluster has no bugs in it, yet you notice that there are many terms in which no leader is elected. What could be a problem (other than network partitions)? How can you fix this problem?
5. How does Raft ensure that once data is committed to log of a leader, the data is always in the log of future leaders?
6. The figure below shows the state of the logs in a cluster of 5 servers.

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		

- (a) Who can become the leader? Who will vote for them?
- (b) Which log entries may have been committed? Explain your answer.
- (c) S1 is currently the leader. Several entries are in the follower's log but not in the leader's log. Can they get into the leader's log? Why or why not?

### 4 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 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: <http://cs.brown.edu/courses/cs138/s20/feedback.html>