

# Indian Institute of Technology, Delhi

---

SPRING, 2013

---

## CSL 860: ADVANCED DISTRIBUTED SYSTEMS

### Midterm

### Two Hours

- NOTE:** – All answers need to be brief and to the point.  
– Please make any assumptions that you deem to be reasonable.

**Total Marks: 100**

**Total Number of Pages : 2**

1. Explain how we can achieve totally ordered multicasting with Lamport clocks? Assume FIFO channels. (HINT: try to modify the algorithm for totally ordered mutual exclusion) (10 marks)
2. Is the following sequence of events allowed with a sequentially-consistent data store? What about a causally-consistent data store? Explain your answer. (10 marks)

|     |       |       |       |       |
|-----|-------|-------|-------|-------|
| P1: | W(x)a |       | W(x)c |       |
| P2: | R(x)a | W(x)b |       |       |
| P3: | R(x)a |       | R(x)c | R(x)b |
| P4: | R(x)a |       | R(x)b | R(x)c |

[W(x)a means that we write  $a$  to location  $x$ , and R(x)b means that we read the value  $b$  from location  $x$ ]

3. We wish to prevent read-write conflicts and write-write conflicts in a quorum based system. If there are two concurrent writes, then the system needs to order one write before the other. Assume that an ordered multicast primitive is available across the servers. Assume there are  $N$  servers. While reading, we wish to read from  $N_R$  servers, and while writing, we wish to write to  $N_W$  servers. What is the relationship between  $N_R$ ,  $N_W$  and  $N$ ? (Note that  $N_R = N_W = N$  is a trivial solution, and is not allowed) (10 marks)
4. What is the algorithm for a node to join a Chord ring? (10 marks)
5. Give an epidemic algorithm (anti-entropy or rumor mongering) to estimate the size of the network for these two scenarios. (10 marks)
  - (a) Assume that any node can communicate with any other node.
  - (b) Assume that a node can communicate with only its neighbors.Estimate its size, complexity, and analyze the tradeoffs.

6. What is the difference between data centric and client centric consistency models. (10 marks)

7. (a) In Chord, why does the lookup algorithm start from the last finger and proceed towards the first finger. Why not proceed in the reverse order? (5 marks)
- (b) Why is the average lookup time in Pastry  $O(\log_{2^b}(N))$ ? If we are using a  $k$  bit hash, shouldn't it be equal to  $O(\log_{2^b}(2^k))$ ? Note that while constructing the routing table, we do not have an estimate of the number of nodes in the network. (5 marks)
8. Assume a network with a star like topology. There is a central node, that sends messages to child nodes. There are no interconnections between child nodes. Provide an algorithm to take a consistent snapshot of the entire system with the central node as the initiator. Assume that the channels do not obey the FIFO (first in, first out) property. (10 marks)
9. (a) In the GHS (minimum spanning tree) algorithm, we rely on the concept of a *core edge*. Let us say that instead of relying on the concept of a core edge, we want to rely on the concept of a *core vertex*, what changes do we need to make in the basic GHS algorithm? (5 marks)
- (b) Why is it important to have the notion of *level* in the GHS algorithm? (5 marks)
10. Prove that the Ricart-Agarwala algorithm achieves mutual exclusion (Brief sketch of the proof is required). (10 marks)
-