22C: 166 Distributed Systems and Algorithms

Homework 5, Total points = 60

Assigned 4/26/11, due 5/3/11

Please submit typewritten solutions through ICON, preferably in the pdf format. Late assignments will not be accepted without prior approval. Feel free to use pseudo-codes as necessary, but the clarity of your answer reflecting the main idea is very important.

Question 1. (15 points) Failure Detectors: In class, we discussed about converting a failure detector that was *weakly complete* to one that was *strongly complete*.

(a) In such a conversion, how is accuracy affected if the detector was weakly accurate? What if the detector was strongly accurate?

(b) Can we design a similar conversion to go from *weak* to *strong accuracy*? If so, what impact does this have on completeness? If not, what is different about accuracy from completeness that makes the conversion impossible?

Question 2. (10 points) Group Communication: Chapter 15, Exercise 1. Assume that the game is being played using the Transis group communication system. Assume that each player has a copy of the game board in a local cache. As a player makes a move, the move must be correctly reflected in every player's local copy. Your answer should refer to the least expensive form of multicast needed in such a case. Briefly explain your answer. (For extra help, refer to http://www.cs.huji.ac.il/labs/transis/lab-projects/guide/chap2.)

Question 3. (3 x 5 = 15 points) Consistency: Answer the following questions concerning the diagram shown below. The timelines are drawn in the same time scale (that is, if an event B is drawn to the right of some event A, then A occurred at an earlier global time than B). Initially, *x* and *y* are equal to 0.

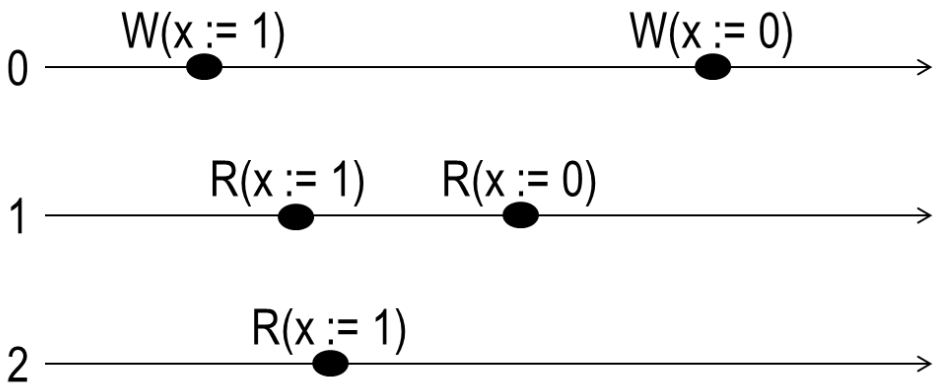

(a) Do the reads and writes shown in the diagram satisfy *linearizability*? Why or why not?

(b) Do the reads and writes shown in the diagram satisfy *sequential consistency*? Why or why not?

(c) Provide a modification to the example such that *causal consistency* is **violated**.

Question 4. (10 points) Distributed Storage and Replication: As mentioned in class, there are many different types of consistency, depending upon the goals of the system and expectations of the users. Consider Amazon's Dynamo system (information on this system can be found here: http://www.allthingsdistributed.com/files/amazon-dynamo-sosp2007.pdf), which implements a type of "eventual consistency", which is described in Section 4.4 of the above paper. Answer the following questions concerning Dynamo.

(a) What purpose does replication serve in Dynamo? (see Section 4.3)

(b) In one paragraph, describe Dynamo's consistency model – when are values the same, when are they different, why was it designed that way, etc.

Question 5. (10 points) CAP Theorem: At the Principles of Distributed Computing (PODC) 2000 conference, Eric Brewer gave a keynote talk where he presented the "CAP Theorem" concerning distributed data stores (slides of the talk can be found at http://www.cs.berkeley.edu/~brewer/cs262b-2004/PODC-keynote.pdf. Another nice reference is available at http://www.julianbrowne.com/article/viewer/brewers-cap-theorem). Nancy Lynch (of FLP fame) and Seth Gilbert later proved the theorem. Describe in a single paragraph the main point of the theorem. Then give three examples of distributed systems – one where C is sacrificed, one where A is sacrificed, and one where P is sacrificed (this will make more sense after reading about the CAP theorem).