CSE 486/586 Distributed Systems
Sample Final Exam Questions

1. (2 points) Which of the following are true for Paxos?

   (a) It takes a minimum of two network round trips for any proposer to have its proposal accepted.

   (b) Once a proposer chooses a proposal number N and sends a prepare request (with proposal number N) to acceptors, a proposal (N, someValue) will be chosen eventually.

   (c) In Paxos phase 1, if there are two concurrent proposers, Paxos will forbid both proposers from getting a majority of replies.

   (d) If there are $2f+1$ acceptors, Paxos is able to work as long as at least $f+1$ acceptors are alive and reachable.

   Answer: a, d

2. (2 points) Which of the following are true for distributed mutual exclusion algorithms?

   (a) The centralized algorithm does not require a process to know anything about other processes in the group.

   (b) Both the centralized algorithm and the token ring algorithm are deadlock-free.

   (c) None of the above

   Answer: a, b

3. (2 points) Until v4, NFS used a stateless server design. Please choose all correct explanation of stateless and stateful design choices regarding NFS.

   (a) File locking is not straightforward to implement with stateless servers.

   (b) It is easier to deal with server failures with a stateful server than a stateless server.

   (c) Client-side implementation can be more complicated with a stateful server.

   (d) In stateless approaches, each request from a client contains complete information about the request (e.g., file name, offset, etc.).

   (e) None of the above

   Answer: a, d

4. (2 points) Consider the following statements about Byzantine faults.

   (a) No process deviates from the protocol.
   (b) The Paxos algorithm may be used to tolerate Byzantine faults.
   (c) With f faulty nodes, we need $2f+1$ nodes to tolerate the faulty nodes' Byzantine behavior.
   (d) Processes or channels behave arbitrarily; they do not just fail-stop or crash.

   Choose *all* correct statements. Or, if there is no correct statement, choose:

   (e) There is no correct statement.

   Answer: d

5. (2 points) Choose all **correct** statements about digital certificates:

   (a) Digital certificates have a standard format.

   (b) To be useful, a digital certificate relies on a chain of trust.

   (c) To be useful, a digital certificate needs a trusted authority.

   Or, if there is no correct statement, choose:

   (d) There is no correct statement.

   Answer: a, b, c

6. (3 points) Suppose we modified the chain replication strategy discussed in class such that query requests & replies occur on the head node and all other functionality remains the same. Given this modification, what are the consequences? Choose all **correct** statements.

   (a) The modified chain replication and the original chain replication provide the same consistency model.

   (b) The modified chain replication and the original chain replication do not provide the same consistency model.

   (c) A system implementing this is not consistent.

   (d) A system implementing this provides linearizability.

   Or, if there is no correct statement, choose:

   (e) There is no correct statement.

   Answer: a, d

7. (3 points) After taking CSE 486/586 at UB, Anna joined a web service company. Her first task as a distributed systems developer is to design a failure detection service for the company's web servers. The web servers are distributed over multiple data centers (one in the west coast and the other in the east coast) and the data centers are connected through the Internet. Anna decides to use a gossip-based protocol and comes up with the following algorithm.

   - Each member maintains a list of all members (including itself) with some additional information. For each member, this additional information includes the member's IP address, the heartbeat counter, and the timestamp for the heartbeat counter.
   - Every T_{gossip} seconds, each member increments its own heartbeat counter and its timestamp. It then selects one other member at random to send its list to.
   - Upon receiving a list from another member, each member updates its own list. It compares its own list to the received list, and updates each entry with the latest information using the timestamps.
   - For each entry, if the heartbeat counter has not increased for more than T_{fail} seconds, then the member is considered failed and removed from the list.

   Anna presented this algorithm to her manager, Elsa. Elsa found out several other issues to consider, so asked Anna to think about them before implementing it. What would be the other issues to consider in the algorithm? Please choose all.

   (a) The heartbeat counter can overflow.
   (b) A member which is considered failed should not be removed immediately after T_{fail} seconds.
   (c) There is a possibility that a member can be erroneously determined as failed due to slow/failed network links.

   Or, if there is no issue, choose:

   (d) There is no issue.

   Answer: a, b, c

8. (3 points) Consider the following two transactions and the order in which all operations are processed. $r[X]$ means reading from an object X, $w[X]$ means writing to an object X.

Transaction T_0	Transaction T_1
$r[X]$	
$w[X]$	
	$r[X]$
	$r[Y]$
	commit
$r[Y]$	
$w[Y]$	
commit	

Choose all correct statements:

(a) These transactions have conflicting operations.

(b) These transactions do not have conflicting operations.

(c) This order provides serial equivalence.

(d) This order does not provide serial equivalence.

Or, if there is no correct statement, choose:

(e) There is no correct statement.

Answer: a, d

9. (4 points) Assume that the following happened with object x over a linearizable storage:

- P1: x.write(A)
- P2: x.write(B), x.read() $\Rightarrow$ C
- P3: x.read() $\Rightarrow$ B, x.read() $\Rightarrow$ A, x.write(C)
- P4: x.read() $\Rightarrow$ B, x.read() $\Rightarrow$ A

What would be an actual-time ordering of the event? Choose all possible scenarios.

(a) P2: x.write(B)
P3: x.read() $\Rightarrow$ B
P4: x.read() $\Rightarrow$ B
P1: x.write(A)
P3: x.read() $\Rightarrow$ A
P4: x.read() $\Rightarrow$ A
P3: x.write(C)
P2: x.read() $\Rightarrow$ C

(b) P1: x.write(A)
P2: x.write(B)
P3: x.read() $\Rightarrow$ B
P4: x.read() $\Rightarrow$ B
P4: x.read() $\Rightarrow$ A
P3: x.read() $\Rightarrow$ A
P3: x.write(C)
P2: x.read() $\Rightarrow$ C

(c) P2: x.write(B)
P3: x.read() $\Rightarrow$ B
P4: x.read() $\Rightarrow$ B
P1: x.write(A)
P4: x.read() $\Rightarrow$ A
P3: x.read() $\Rightarrow$ A
P3: x.write(C)
P2: x.read() $\Rightarrow$ C

Or, if there is no correct answer, choose:

(d) There is no correct answer.

Answer: a and c

10. (4 points) *Read-your-own (RYO) consistency* is yet another consistency model where a process is always guaranteed to read the same value that its last operation produced. Specifically, it considers the following three cases for a process P issuing a read operation R for a data item X. Note that these are the only possible cases, i.e., RYO covers all scenarios. In all cases, the definition of "latest" uses actual time.

    - Case 1: There was no prior operation for X from P before R. In this case, R returns the result of the latest write for X, no matter which process did the latest write.
    - Case 2: There was at least one prior operation for X from P before R, and the last operation (again for X from P) was a write. In this case, R returns the result of the last write of its own for X.
    - Case 3: There was at least one prior operation for X from P before R, and the last operation (again for X from P) was a read. In this case, R returns the same value that the last read of its own returned.

    Choose all **correct** statements about the relationships among the consistency models:

    (a) Linearizability is stronger than RYO, i.e., any system that provides linearizability also provides RYO always.

    (b) RYO is stronger than linearizability, i.e., any system that provides RYO also provides linearizability always.

    (c) Sequential consistency is stronger than RYO, i.e., any system that provides sequential consistency also provides RYO always.

    (d) RYO is stronger than causal consistency, i.e., any system that provides RYO also provides causal consistency always.

    Or, if there is no correct statement, choose:

    (e) There is no correct statement.

    Answer: d