Exam 2 Review
Fall 2011

Question 1

What is a drawback of the token ring election algorithm?

Bad question! Token ring mutex vs. Ring election!

Ring election:
- multiple concurrent elections
- message size grows to contain a list of all processes: particularly bad if a process comes up midway through an election – two elections can elect multiple leaders.
- split brain can occur and result in multiple leaders

Token ring mutex:
- Token can get lost. Need reliable mechanism to detect and regenerate.
- Access is well-ordered but not first-come, first-served.

Question 2

How does a process in a ring election algorithm know that the election message that it receives is the one it initiated?

The head of the list in the received message contains that process' ID

Question 3

Explain the purpose of the two phases in a two-phase commit protocol.

1. Consensus: Get a "ready to commit" or "abort" message from all cohorts.
2. Commit/abort: Send a "commit"/"abort" message to all.

NOT:
If timeouts occur then abort. 2PC never times out; 3PC does.

NOT:
The 2PC protocol makes the transaction atomic. Locking does that.

Question 4a

a) What is the advantage of a lease over a lock?

A lease is a lock that expires after a specific time.

Bad answer: prevents deadlocks.
Yes, it can cause a deadlock to break but at the loss of ACID guarantees

Question 4b

a) What problems can arise? Use centralized mutual exclusion as an example.

1. Slow message transmission of renew can cause the lock to expire and be granted to someone else.
2. Unnecessary delays if a release message is lost and others have to wait for a timeout.

Poor but ok: overhead in renewing a lease.

NOT: process hogs resource for a long time. That only happens if it forgets to release the resource.

Question 5

What do you have to give up if you want a highly available system that withstands network partitioning?

Brewer's CAP Theorem:

You can have only two of Consistency, Availability, and Partition tolerance

To get high availability and partition tolerance, you have to give up consistency and go for eventual consistency.

Question 6

Explain the problem with two-phase locking that is fixed by strict two-phase locking.

Avoids having cascading aborts.

Two-phase locking releases locks when a resource is no longer needed and no new locks will be acquired. This means that other transactions can access uncommitted data. If the transaction that released the locks aborts, then all transactions that accessed the uncommitted data will have to abort.

Question 7

Lamport's mutual exclusion algorithm is:

(a) Token-based.

(b) Contention-based.

(c) Centralized.

(d) Ring-based.

Question 8

A process gives you access to a resource in Lamport's algorithm based on:

(a) Receiving a grant message from a majority of processes.

(b) Ownership of a token.

(c) Receiving permission from all other processes.

(d) Being first in a time-ordered queue.

Question 9

hich mutual exclusion algorithm is the most efficient (i.e., requires the fewest messages)?

(a) Centralized.

(b) Ricart & Agrawala.

(c) Lamport's.

(d) Token Ring.

Question 10

The Chang & Roberts ring algorithm does not:

(a) Ensure that an election message contains at most one process ID.

(b) Circulate an election message, skipping non-responding processes.

(c) Resolve competing elections by comparing timestamps.

(d) Try to kill off multiple concurrent elections.

The algorithm compares process IDs, not time stamps.

Question 11

Split brain is a condition that arises when

(a) A network is partitioned.

(b) An election algorithm fails to elect any leader.

(c) Processes send faulty messages.

(d) A process shares multiple process ID numbers.

Question 12

A process that sends deceptive messages to another process exhibits this fault:

(a) Fail-stop.

(b) Fail-silent.

(c) Transient.

(d) Byzantine.

Question 13

The two-army problem states that, over a faulty communication channel:

(a) You need to send message acknowledgements.

(b) You need to use two-way message acknowledgements.

(c) You need at least 3n + 1 message exchanges for n processes to communicate reliably.

(d) You can never guarantee reliable communication.

Question 14

For N acceptors, Paxos will reach consensus if a proposer can communicate with at least:

(a) N of them.

(b) A majority of them.

(c) $(2/3)N$

(d) $(2N+3)/2$

Question 15

Hierarchical leases are useful for:

(a) Getting the advantages of fast lease management while using a fault-tolerant algorithm to elect coordinators.

(b) Allowing sub-transactions to get leases from the top-level transaction.

(c) Processes that need to get leases on more than one resource; they can do so with a single request.

(d) Processes that may need to use a combination of locks and leases.

Question 16

In transactions, a write-ahead log is used for:

(a) Ensuring that a process can undo any modifications it made to stored data.

(b) Optimizing remote file system writes over a network link.

(c) Creating a record of every transaction performed by a process.

(d) Ensuring that two transactions do not modify the same data concurrently.

Question 17

The three-phase commit protocol was created to:

(a) Ensure that the protocol completes within a certain time limit.

(b) Be more reliable than the two-phase commit protocol by sending a pre-commit message.

(c) Work reliably over unreliable communication lines.

(d) Not rely on a central coordinator for consensus.

Question 18

Paxos commit is most closely described as a protocol that:

(a) Ensures that transactions are run in a serial order.

(b) Removes the need for commits and aborts by providing an isolated execution environment.

(c) Allows a commit to take place as long as a majority of sub-transactions are alive.

(d) Replaces the coordinator of the two-phase commit protocol with a fault-tolerant one.

Question 19

False deadlock is due to:

(a) Bad message ordering.

(b) A failure to obtain locks.

(c) The use of optimistic algorithms.

(d) A failure to release locks.

Question 20

The following is not a condition for deadlock:

(a) Preemptive resource access.

(b) Hold and wait.

(c) Circular wait.

(d) Mutual exclusion.

Question 21

The Chandy, Misra, Haas algorithm is used for:

(a) Deadlock detection.

(b) Deadlock prevention.

(c) Deadlock avoidance.

(d) All of the above.

Probe message detects a deadlock condition.

Deadlock prevention = restriction on resource access so that deadlock cannot occur.

Deadlock avoidance = schedule transactions so deadlocks cannot happen – not feasible for arbitrary transactions

Question 22

The wait-die algorithm:

(a) Ensures that a circular wait can never develop.

(b) Has a process give up after waiting for a resource beyond a specific time.

(c) Requires a process to kill any process that is using resources it needs.

(d) Uses a coordinator to detect and kill any process that is waiting for a resource.

Wait-die: if a younger process is using the resource, then the older process waits. If an older process is using the resource, the younger kills itself (and restarts later).

Wait order is always *OLD WAITS ON YOUNG*, so no circular wait is possible

Question 23

Which file system supports a whole file download model?

(a) NFS

(b) Coda

(c) SMB

(d) GFS (Google File System)

Question 24

An SMB oplock gives clients:

(a) The possibility of caching file attributes.

(b) The ability to lock a remote file.

(c) The ability to check out a file for disconnected use.

(d) The ability to synchronize local copies of remote files.

Question 25

The main difference between the Google File System (GFS) and the Hadoop Distributed File System (HDFS) is that HDFS:

(a) Uses a server dedicated for keeping track of file names and locations.

(b) Stores files in replicated 64 MB sections.

(c) Does not allow reading from an existing file.

(d) Does not allow writing to an existing file.

Question 26

Optimistic concurrency control:

(a) Ensures that a transaction can never access data of an uncommitted transaction.

(b) Optimize the schedule of transactions to minimize resource contention.

(c) Divides a transaction into a growing phase and a shrinking phase.

(d) Tend to assume that transactions are unlikely to abort.

Question 27

Which file system is not implemented as a driver within the operating system?

(a) NFS

(b) SMB

(c) Coda

(d) Hadoop Distributed File System (HDFS)

Question 28

A client modification log in Coda is used to:

(a) Allow a client to send changes to several replicas of remote volume servers.

(b) Allow clients to roll back their file modifications to a previous state.

(c) Notify clients of file changes on a server.

(d) Allow clients to work without network connectivity to file servers.

Question 29

A bully election algorithm picks the process that:

(a) First started the election.

(b) Is nominated by the most processes.

(c) Has the highest numbered Lamport timestamp of all election messages.

(d) Has the highest numbered process ID.

Question 30

The following is not a property of transactions:

(a) Permanent: once committed, the results of a transaction are made permanent.

(b) Serializable: the result of concurrent transactions must be the same as if they were run in serial order.

(c) Indivisible: the transaction appears as an indivisible action.

(d) Bounded: a transaction must have an upper bound on the time it takes to complete.

Question 31

Akamai routes a content request for a particular machine name to an appropriate server by:

(a) Configuring IP routers to route requests to the nearest server.

(b) Having multiple machines respond to the same IP address.

(c) Having a domain name server return an IP address based on the location of the requestor.

(d) Mapping a machine name to one unique IP address via DNS.

The End