

## Selected Questions

Exam 2  
Spring 2008

Page 1

### Question 1

Explain the difference between logical and physical clocks.

Physical clocks measure the time of day. Logical clocks are used to mark relationships among events in a distributed system.

Page 2

### Question 2

What problem with Lamport clocks to vector clocks solve?

With Lamport clocks, you cannot tell whether two events are causally related or concurrent by looking at the timestamps. Just because  $L(a) < L(b)$  does not mean that  $a \rightarrow b$ .

Vector clocks allow you to compare two vector timestamps to determine whether the events are concurrent or not.

Page 3

### Question 3

A client's clock reads 3:20:00. The server's clock reads 3:10:00 when they synchronize using the Berkeley algorithm. Assume message delays are negligible. What is the time at the client after synchronization?

The Berkeley algorithm averages clocks among the entire group. In this case, the group has two members: the client and the server. The average of the two clocks is  $(3:20:00 + 3:10:00)/2 = 3:15$ . Both the client and server will be set to 3:15.

Page 4

### Question 4

A client's clock reads 3:20:00. The server's clock reads 3:10:00 when they synchronize using Cristian's algorithm. Assume message delays are negligible. What is the time at the client after synchronization?

Cristian's algorithm assumes that the server has an accurate clock. The client requests the time and sets its clock to the server's time +  $\frac{1}{2}$ (round trip delay). In this case, the round-trip delay is negligible (i.e., 0), so the client just sets its clock to the server's time: 3:10:00.

Page 5

### Question 5a

Some networks, such as cable TV Internet service, provide an asymmetric bandwidth. For example, a cable modem may provide 12 Mbps downstream service but only 1 Mbps upstream. Cristian's algorithm assumes symmetric delays to and from the server.

(a) Reformulate the algorithm to accommodate asymmetric delays, where  $T_S$  is the server's time,  $T_a$  is the time the request was sent,  $T_r$  is the time the response was received,  $U$  is the upstream bandwidth, and  $D$  is the downstream bandwidth.

Assume the message size is  $x$ . The best-case time to send a message is  $x/U$ . The best-case time to receive a response is  $x/D$ .

The total time is  $(x/U + x/D) = x(D+U)/DU$ .

The ratio of the time to get a response to the total time is:

$[x/D] / [x(D+U)/DU] = xDU / x(D+U) = U/D+U$ . We assume that the response was generated this long ago compared to the total time.

The special-case Cristian's algorithm sets the clock to

$$T_{\text{client}} = T_S + (U/(D+U)) (T_r - T_a)$$

Page 6

### Question 5b

(b) You are on a network with a 9 Mbps downstream bandwidth and a 1 Mbps upstream bandwidth. Your client makes a request at 9:10:00.0 and gets a response 200 msec later. If the time on the server is 8:42:00.0, what would the time be set to on your client?

From (a):

$$U = 1 \text{ Mbps}, D = 9 \text{ Mbps}$$

$$T_{\text{client}} = T_s + (U/(D+U)) (T_b - T_a)$$

$$T_{\text{client}} = T_s + (1/10) (T_b - T_a)$$

$$= 8:42:00.0 + 0.1 (0:00:00.200)$$

$$= 8:42:00.02$$

Page 7

### Question 7

The Coda distributed file system does not support:

→ (a) Tokens for file locking.

(b) Disconnected clients.

(c) Replicated servers.

(d) Session semantics.

Page 8

### Question 8

The purpose of DFS tokens is to:

→ (a) Give the server control over which actions clients can perform and how they may cache data.

(b) Ensure mutual exclusion is granted when servers are accessed.

(c) Support read/write replication on servers through token exchange.

(d) To avoid the overhead of having clients cache file data.

Page 9

### Question 9

SMB is best characterized as a:

→ (a) Remote call, connection-oriented file system.

(b) Remote call, connectionless file system.

(c) File upload/download, connectionless file system.

(d) File upload/download, connection-oriented file system.

Page 10

### Question 10

Opportunistic locks, or oplocks:

→ (a) Tell the client how it may cache file data and attributes.

(b) Are a first-come, first-served mechanism for locking a file.

(c) Allow clients to grab locks on a remote file under CIFS.

(d) Guarantee exclusive access to a file once it is open.

Page 11

### Question 11

Alice has a remote file open on her system with a Level 1 oplock. Bob tries to open the file from another system. The most likely scenario is:

→ (a) Alice's oplock will be revoked and Bob will be able to share the file.

(b) Bob's request will fail.

(c) Bob will be blocked until Alice releases her lock.

(d) Alice's oplock will be revoked and she will no longer have access to the file.

Page 12

### Question 15

A synchronization subnet is:

- (a) The collection of NTP servers on the Internet.
- (b) An IP multicast group used for clock synchronization
- (c) The set of machines addressed by an NTP server operating in multicast mode
- (d) The set of NTP servers with which you are currently synchronizing.

Page 13

### Question 16

IP multicast uses:

- (a) Unreliable multicast.
- (b) Reliable multicast.
- (c) Atomic multicast.
- (d) User-selectable reliability.

Page 14

### Question 17

False sharing in a distributed shared memory system arises because:

- (a) Unrelated data is shared between multiple systems but resides on the same memory page.
- (b) One distributed memory page is mapped onto two or more systems at the same time.
- (c) Two or more distributed memory pages are mapped onto the same system.
- (d) Bad message ordering makes it appear that data is shared when, in reality, it is not.

Page 15

### Question 18

A distributed directory in a distributed shared memory system:

- (a) Ensures cache coherence in a sequentially coherent memory system.
- (b) Allows concurrent queries and updates to a memory page.
- (c) Is necessary for systems that employ caching.
- (d) Distributes the load of looking up page owners.

Page 16

### Question 21

The ring election algorithm works by:

- (a) Having all nodes in a ring of processors send a message to a coordinator who will elect the leader.
- (b) Sending a token around a set of nodes. Whoever has the token is the coordinator.
- (c) Sending a message around all available nodes and choosing the first one on the resultant list.
- (d) Building a list of all live nodes and choosing the largest numbered node in the list.

Page 17

### Question 22

Which event is concurrent with the vector clock (2, 8, 4)?

- (a) (3, 9, 5)
- (b) (3, 8, 4)
- (c) (1, 7, 3)
- (d) (4, 8, 2)

Page 18

### Question 23

If a DES key is 56 bits, a triple-DES key is:

- (a) 56 bits
- (b) 112 bits
- (c) 144 bits
- (d) 224 bits

Page 19

### Question 25

Which statement about the Vigenère ciphers is not true?

- (a) It is a block cipher that uses a secret key.
- (b) It is a polyalphabetic substitution cipher.
- (c) It is a collection of Caesar ciphers.
- (d) It attempts to be resilient to frequency analysis.

Page 20

### Question 27

The Ricart & Agrawala distributed mutual exclusion algorithm is:

- (a) More efficient and more fault tolerant than a centralized algorithm.
- (b) More efficient but less fault tolerant than a centralized algorithm.
- (c) Less efficient but more fault tolerant than a centralized algorithm.
- (d) Less efficient and less fault tolerant than a centralized algorithm.

Page 21

## Selected Questions

Exam 2  
Spring 2008

Page 22

### Question 5

The clock in the clock tower in the town of Chronos broke. It was repaired but now the clock needs to be set. A train leaves for the nearest town, Temporis, 100 miles away. It returns  $4\frac{1}{2}$  hours later with a report that the time according to the clock tower in Temporis is 4:05.

(a) To what value should the time be set on the clock tower?

Cristian's algorithm:  $T_{\text{server}} + \frac{1}{2}(\text{round-trip delay})$

Round-trip delay = 4:30

$4:05 + (4:30/2) = 4:05 + 2:15 = 6:20$

Page 23

### Question 5

The clock in the clock tower in the town of Chronos broke. It was repaired but now the clock needs to be set. A train leaves for the nearest town, Temporis, 100 miles away. It returns  $4\frac{1}{2}$  hours later with a report that the time according to the clock tower in Temporis is 4:05.

(b) The maximum speed that the train can travel is 50 mph. What is the error of the clock?

Best case: 100 miles @ 50 mph = 2 hours

Minimum time traveling: 4 hours

Uncertainty window: 4:30 - 4:00 = 30 min

Error =  $\pm 15$  min

Page 24

### Question 10

Event *a* has a Lamport timestamp of 4.  
Event *b* has a Lamport timestamp of 8.  
What can we tell about events *a* and *b*?

- a) Events *a* and *b* are causally related.
- b) Events *a* and *b* are concurrent.
- c) Event *a* happened before event *b*.
- d) If events *a* and *b* are causally related, then event *a* happened before event *b*.

Page 25

### Question 11

Which event is concurrent with the vector timestamp (2, 4, 6)?

- a) (3, 5, 7)
- b) (1, 3, 5)
- c) (1, 4, 6)
- d) (1, 4, 7)

Page 26

### Question 12

A client has a time of 5:05 and a server has a time of 5:25. Using the Berkeley algorithm, the client's clock will be set to:

- a) 5:15  $(5:05 + 5:25)/2 = 5:30/2 = 5:15$
- b) 5:20
- c) 5:25
- d) 5:30

Page 27

### Question 13

Which offers the most fault-tolerant message delivery?

- a) Atomic multicast.
- b) Totally ordered reliable multicast.
- c) Causally ordered reliable multicast.
- d) Hardware multicast.

Page 28

### Question 14

When we had to send an IP packet over an ethernet network, we found the corresponding ethernet address via the Address Resolution Protocol (ARP). How do we find the ethernet address that corresponds to a particular IP multicast address?

- a) All IP multicast packets use the same reserved ethernet address.
- b) ARP is also responsible for handling multicast addresses.
- c) By mapping some of the bits of an IP multicast address onto an ethernet address.
- d) By mapping all of the bits of an IP multicast address onto an ethernet address.

Page 29

### Question 14



Page 30

### Question 15

How does IP multicast achieve reliable packet delivery?

- a) It doesn't
- b) Via best-effort multicast.
- c) Via atomic multicast.
- d) Via IGMP (Internet Group Management Protocol)

Page 31

### Question 16

A copyset is used:

- a) In a DSM system that employs page caching.
- b) In a DSM system as an alternative to page caching.
- c) To implement lazy release consistency.
- d) To implement entry consistency.

Page 32

### Question 17

An ethernet card may support all of the following mechanisms for multicast except:

- a) Filtering packets based on a hash of the multicast address.
- b) Searching through a small table for a matching ethernet address.
- c) Searching through a larger table of ethernet addresses in main memory.
- d) Accepting all multicast packets via multicast promiscuous mode.

Page 33

### Question 18

If a frequency analysis of ciphertext reveals a statistical equivalence to the corresponding frequencies of plain English text, we can guess that the cipher is:

- a) a monoalphabetic substitution cipher.
- b) a polyalphabetic substitution cipher.
- c) a transposition cipher.
- d) a rotor machine.

Page 34

### Question 19

A rotor machine implements:

- a) a monoalphabetic substitution cipher.
- b) a polyalphabetic substitution cipher.
- c) a transposition cipher.
- d) a one-time pad.

Page 35

### Question 20

False sharing in a DSM system refers to:

- a) an unwanted page being cached on another system
- b) unrelated data being resident on the same page.
- c) invalid data in the directory, leading the system to falsely believe the page is shared.
- d) two versions of the same page being present on one system.

Page 36

### Question 21

A bully election algorithm:

- a) picks the first process to respond to an election request.
- b) relies on majority vote to pick the winning process.
- c) assigns the role of coordinator to the process holding the token at the time of election.
- d) picks the process with the largest ID.

Page 37

### Question 22

Which mutual exclusion algorithm works when the membership of the group is unknown?

- a) Centralized.
- b) Ricart-Agrawala.
- c) Lamport.
- d) Token Ring.

Page 38

### Question 23

Weak consistency models offer an advantage over sequential consistency because:

- a) memory synchronization does not have to occur with each memory operation.
- b) caching may be used.
- c) only the needed data needs to be shared, not the entire memory system.
- d) they ensure that each write operation invalidates or updates all cached copies before the next instruction is executed.

Page 39

### Selected Questions

Spring 2006  
Exam 2

Page 40

### Question 5

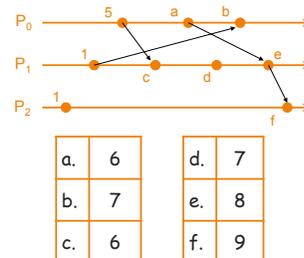
What is *false sharing*?

- When processes are accessing unrelated data that happens to reside on the same shared page. This leads to thrashing.
- *Not*: when different data resides on the same page.

Page 41

### Question 6

Assign Lamport timestamps to the following events:



Page 42

Question 7

You have several computers cooperating on a parallel program. CPU 1 tries to write to page 302, which is currently owned by CPU 0. Explain the sequence of operations that takes place in a DSM system.

- CPU 1 attempts to write and get a page fault.
- Page fault handler contacts directory to find current owner of page 302.
- CPU 1 fetches page from CPU 0 (CPU 0 sends page).
- CPU 0 creates a free page frame and loads page 302 into it.
- CPU 0 adjusts PTE in MMU and restarts instruction.
- CPU 1 writes to page 302 successfully.

Page 43

11. Which set of events is concurrent (all events are concurrent with each other)?

- a. (3, 1, 5, 7), (3, 2, 6, 7), (2, 1, 6, 8)
- b. (2, 1, 3, 4), (2, 2, 3, 3), (3, 3, 2, 5)
- c. (1, 2, 3, 4), (2, 3, 4, 5), (3, 4, 5, 6)
- d. (1, 5, 6, 7), (1, 4, 5, 7), (1, 3, 2, 2)

If we do an element-by-element comparison:  
(2, 1, 3, 4) is neither  $\geq$  nor  $\leq$  to (2, 2, 3, 3)  
 $2=2$ ;  $1<2$ ;  $3=3$ ;  $4>3$ .  
(2, 2, 3, 3) is neither  $\geq$  nor  $\leq$  to (3, 3, 2, 5)  
 $2<3$ ;  $2<3$ ;  $3>2$ ;  $3<5$

Page 44

13. A linear compensating function adjusts a clock to:

- a. Tick at constantly decreasing or increasing intervals until synchronization is achieved.
- b. Compensate for variations in a clock's frequency to make it appear to tick at a constant rate.
- c. Tick at a constant faster or slower interval.
- d. Compensate for sudden changes such as leap seconds or daylight savings time.

Page 45

14. A client gets a timestamp of 4:12:30.500 from a time server. The elapsed time between the request and response was 20 msec (0.020 sec). The current time on the client is 4:12:30.510. Using Cristian's algorithm, what is the time set to on the client?

- a. 4:12:30.480
- b. 4:12:30.490
- c. 4:12:30.510
- d. 4:12:30.520

Cristian's algorithm sets the time to:  
 $server\ time + \frac{1}{2}(elapsed\ time)$   
 $= 3:12:30.500 + 0.020/2$   
 $= 3:12:30.500 + 0.010$   
 $= 3:12:30.510$

Page 46

15. Which mutual exclusion algorithm is useful when the membership of the group is unknown?

- a. Centralized.
- b. Lamport's.
- c. Token ring.
- d. All of the above.

Lamport's requires reliable group multicast.  
Token ring requires the ability to construct a logical ring of processes based on the group.  
Centralized doesn't require you to know about any other process.

Page 47

16. How does IP multicast achieve reliable packet delivery?

- a. It doesn't.
- b. Via best-effort multicast.
- c. Via atomic multicast.
- d. Via IGMP (Internet Group Management Protocol).

Page 48

17. A *copyset* is used:

- a. In a DSM system that employs page caching.
- b. In a DSM system as an alternative to page caching.
- c. To implement lazy release consistency.
- d. To implement entry consistency.

A copyset is nothing more than a list of processors that have a copy of a particular shared memory page.

Page 49

The End.

Page 50