

|                  |                       |                      |
|------------------|-----------------------|----------------------|
| <b>COSC 6360</b> | <b>SECOND MIDTERM</b> | <b>APRIL 3, 2008</b> |
|------------------|-----------------------|----------------------|

This exam is **closed book**. You can have **one sheet** (that is, **two pages**) of notes. UH expels cheaters  
Please answer every part of every question.

1. A system of physical clocks consists of two clocks, namely, one that is slow and loses 10 minutes every hour and another that is fast and advances by 10 minutes every hour. Assuming that the clocks are managed by Lamport's physical clock protocol, what will be the time marked by each clock at 2:00 PM given that
  - (a) both clocks indicated the correct time at noon;
  - (b) the processors on which the clock reside continuously exchanged messages between themselves until 1:00 PM then ceased to communicate; and
  - (c) the message transmission delays are negligible. (2×5 points)

The fast clock will indicate   2:20 PM   plus or minus a few seconds at 2:00 PM.

The slow clock will indicate   2:00 PM   plus or minus a few seconds at 2:00 PM.

2. In Kerberos, which keys are used to encrypt:
  - a) The tickets issued by the *Kerberos server*? (5 points)  
The secret key of the ticket granting service  $K_{tgs}$  \_\_\_\_\_
  - b) The tickets issued by the *ticket granting service*? (5 points)  
The secret key of the server for which the ticket was issued  $K_s$  \_\_\_\_\_

3. What is the function of Totem *guaranteed vector messages*? (10 *easy* points)

Totem *guaranteed vector messages* let processes in a multiple ring protocol keep delivering the messages they receive when one of the rings does not forward any messages.

4. Consider a *single-ring* Totem system comprising two processors A and B. Assuming that each of these two processors has received the messages with the following sequence numbers:

| Processor | Messages      |
|-----------|---------------|
| A         | 3, 4, 6, 8    |
| B         | 3, 4, 5, 7, 8 |

Which messages will be delivered by each processor if all messages are *safe delivery messages*? (2×5 points and no partial credit)

A will deliver messages numbered 3 and 4 \_\_\_\_\_

B will deliver messages numbered 3 and 4 \_\_\_\_\_

5. Which are the main advantage and the main disadvantage of using journaling with *asynchronous log updates* compared to using journaling with *synchronous log updates*?

a) *Main advantage:* (5 points)

It is much faster as there are less physical writes to the log. \_\_\_\_\_

b) *Main disadvantage:* (5 points)

It does not guarantee the durability of metadata updates. \_\_\_\_\_

6. How would you implement *close-to-open consistency* in Sun NFS without modifying the protocol? (10 points) You should assume that the server can send to the server *validate(block)* and *write(block)* requests among other requests. (*Hint: The answer is fairly simple.*)

When a client closes a file, it will forward to the server all the blocks of that file it has modified since it opened the file. (Note that the operation will be blocking because NFS uses blocking write requests.)

When a client opens a file, it will validate with the server all the blocks of that file it has in its cache

7. Which are the main advantage and the main disadvantage of Coda/AFS *callbacks*?

a) *Main advantage:* (5 points).

The server has to process fewer client validation requests. \_\_\_\_\_

b) *Main disadvantage:* (5 points)

Lost callbacks can result in inconsistent file updates. \_\_\_\_\_

8. What is the main advantage of *RAID level 5 disk arrays* over

a) *RAID level 3 arrays:* (5 points)

Increased parallelism. \_\_\_\_\_

b) *RAID level 4 arrays:* (5 points)

Distributing parity data among all disks removes a write bottleneck. \_\_\_\_\_

9. What is the purpose of *file hoarding* in Coda? (5 points) Which files does Coda hoard? (2×5 points) How would you implement file hoarding on a modern laptop? (5 points)

In addition to improving the overall file system performance, file hoarding allows Coda clients to keep accessing the files hoarded on their workstation while operating in disconnected mode.

By default, Coda file bases its file hoarding decisions on its normal LRU caching policy. In addition, it lets users specify a prioritized list of files and directories that it should attempt to keep in the cache of their workstations.

Since today's laptops have much larger disks than the laptops of the late eighties, we should strive to hoard on each laptop a complete copy of all the files and directories that its owner keeps on the server.