

COSC 6360	SECOND MIDTERM	APRIL 6, 2009
------------------	-----------------------	----------------------

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. Consider a RAID 5 array with four data blocks (b_0, b_1, b_2, b_3) and one parity block p per stripe.
 - a) How much of the total disk space is used by parity blocks? (5 *easy* points) 20 percent
 - b) How can we reconstitute the contents of block b_1 after the disk holding that block *has failed*? (5 points)

$$b_1 = b_0 \oplus b_2 \oplus b_3 \oplus p$$

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

Processor	Messages
A	2, 3, 4, 6, 8
B	2, 3, 4, 5, 7, 8

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

A will deliver messages numbered 2, 3, 4

B will deliver messages numbered 2, 3, 4, 5

3. Consider the following CSP program:

```

ch :char; count : integer; count := 0;
*[one ? ch → count := count + 1];
    
```

- a) What does the program do? (5 points)
 It counts the number of characters sent by process one. _____
- b) When will it terminate? (5 points)
 When process "one" terminates. _____

4. A system of physical clocks consists of two clocks, namely, one that is slow and loses one minute every hour and another that is fast and advances by one minute every hour. Assuming that the clocks are managed by Lamport's physical clock protocol, what will be the time marked by each clock at noon given that:

- (a) both clocks indicated the correct time at midnight;
- (b) the processors on which the clock reside continuously exchanged messages between themselves from midnight until nine o'clock ; and
- (c) message transmission delays are negligible. (2×5 points)

The fast clock will indicate 12:12 pm plus or minus a few seconds at noon.
 The slow clock will indicate 12:06 pm plus or minus a few seconds at noon.

5. In the ARC cache replacement policy,

a) When would a block be moved from the **T1** list to the **T2** list? (5 points)

When it is referenced a second time. _____

b) Why do we say that ARC is *self-tuning*? (5 points)

Because it has no user-tunable parameter. _____

6. Explain BitTorrent *rarest first policy*. (5 points) Show how the policy benefits all downloaders. (5 points)

See paper.

7. Consider a log-structured file system (LFS) that is being accessed immediately after the system has been rebooted. Assuming that a final checkpoint was taken at the moment the system was powered down, which steps must be taken to access a specific i-node. (20 points)

- Fetch address of i-node map block _____
located in checkpoint area _____
- Fetch i-node map block _____
whose address is given by address found in checkpoint area _____
- Fetch i-node block _____
whose address is given by i-node map block _____
- Fetch _____
whose address is given by _____

8. Which are the main advantage and the main disadvantage of using *journaling with synchronous log updates* compared to using *soft updates*? (2×5 points)

a) *Main advantage:*

 It guarantees the durability of metadata updates. _____

b) *Main disadvantage:*

 It is much slower than journaling file systems using asynchronous log updates. _____

9. Which are the two ways that *MapReduce* applications can benefit from *Corey address ranges*. (2×5 points)

a) <See paper > _____

b) < See paper > _____
