16.482 / 16.561: Computer Architecture and Design

Spring 2014

Homework #8: EXTRA CREDIT

Due **5:00 PM on Monday, 4/28/14**

Notes:

- This assignment is strictly for extra credit—if you complete this assignment, your grade will replace your lowest homework score from HW 1-7, unless your score on HW 8 is lower than that on all previous homework assignments.
- **<u>All assignments must be submitted by 5:00 PM on Monday, 4/28, and no late submissions will be accepted.</u>** Once the solution to this assignment is posted on Monday evening, no additional submissions will be allowed.
- While typed submissions are preferred, handwritten submissions are acceptable.
- Any handwritten solutions that are scanned and submitted electronically <u>must</u> be clearly legible and combined into a single file—<u>simply sending a picture of each scanned page is not an acceptable form of submission</u>.

1. <u>RAID</u> (40 points) In this problem, you will deal with reliability issues in a RAID system. Assume we have a system in which each disk contains 60 GB, with a peak sequential read rate of 150 MB/sec and a peak sequential write rate of 90 MB/sec. The disks share a bus that can transfer 320 MB/sec. The mean time to failure for each disk is 2.4M hours.

a. (10 points) Assume this system uses RAID 4 and contains 6 disks. Draw a diagram showing the basic layout of blocks across disks in this system.

b. (5 points) If a single disk fails, the RAID system will perform *reconstruction*—the process of determining what data was on a disk when it failed and restoring that data. What is the expected time until reconstruction is required?

c. (5 points) What read and write operations are required to perform reconstruction?

d. (10 points) In *offline reconstruction*, the RAID system devotes all resources to performing reconstruction and services no other requests until reconstruction is complete. How long will it take offline reconstruction process to complete in the RAID 4 system discussed in this problem? Assume that the reads and writes described in part (c) can be executed in parallel and that the bus is the limiting factor.

e. (10 points) In *online reconstruction*, the RAID system continues to service requests while performing reconstruction. This technique removes service interruptions, but the reconstruction process is limited to a fraction of the total system bandwidth. Online reconstruction therefore takes longer than offline reconstruction, leaving the system more vulnerable to a second disk failure.

   Assuming the reconstruction process is limited to 20 MB/sec, how long will online reconstruction take in this system?

2. <u>Snooping protocols</u> (30 points) You are given a four-processor system that uses a write-invalidate, snooping coherence protocol. Each direct-mapped, write-back cache has four lines, each of which holds eight bytes; in the diagram below, only the least-significant byte of each word is shown. The cache states are I (invalid), S (shared), and M (modified/exclusive). The caches and memory have the following initial state:

P0

	State	Tag	Data	
B0	I	0x100	01	23
B1	S	0x108	00	88
B2	M	0x110	00	30
B3	I	0x118	00	10

P1

	State	Tag	Data	
B0	I	0x100	01	23
B1	M	0x128	AB	CD
B2	S	0x130	14	12
B3	S	0x118	14	92

P2

	State	Tag	Data	
B0	S	0x120	13	31
B1	S	0x108	00	88
B2	I	0x130	51	55
B3	I	0x138	01	38

P3

	State	Tag	Data	
B0	S	0x120	13	31
B1	S	0x108	00	88
B2	I	0x110	00	30
B3	S	0x118	14	92

Memory

Address	Data	
0x100	00	00
0x108	00	88
0x110	20	08
0x118	14	92
0x120	13	31
0x128	FF	FE
0x130	14	12
0x138	AB	BA

For each of the transactions listed below, list all cache blocks modified and their final state, as well as all memory blocks modified and their final state. Assume each set of transactions starts with the same initial state—in other words, your answer to part (b) does not depend on your answer to part (a). However, you should track the state transitions of each block throughout the problem.

a. P1: read 0x110
P2: read 0x110
P3: read 0x110

b. P2: write 0x128 ← 14
P2: write 0x12C ← DF
P2: write 0x108 ← 20

c. P3: write 0x110 ← 45
P0: write 0x110 ← 67
P3: read 0x114

3. <u>Directory protocols</u> (30 points) Say we have a four-processor system that uses a write-invalidate, directory coherence protocol. The system contains a total of 8 memory blocks, as shown in the initial directory state below:

Block #	P0	P1	P2	P3	Dirty
0	1	0	0	0	1
1	0	0	0	0	0
2	1	1	0	0	0
3	0	0	1	0	1
4	1	1	1	1	0
5	0	1	1	0	0
6	0	0	0	1	1
7	0	0	0	0	0
8	0	1	0	0	0

a. (5 points) What state is each block in?

b. (25 points) What messages are sent between the nodes and directory for each of the following transactions to ensure the processor has the most up-to-date block copy, the appropriate invalidations are made, and the directory holds the appropriate state?

i. P1: write block 0
ii. P2: read block 7
iii. P3: write block 4
iv. P0: write block 2
v. P1: read block 6