

# Advanced Computer Architecture CMSC 611

## Homework 6

Due in class at 1.05pm, Dec 10<sup>th</sup>, 2012

(If you wish to **go green**, then you can submit the entire Homework electronically. Make sure you include the string “**CMSC 611 Homework**” in your subject line. You will get a canned response immediately if your message is filtered correctly! I use an exact substring match for the filter so make sure you include the exact string in your subject line. Deadline remains the same)

Please **DO NOT** email your homework to Dr. Olano!! **DO NOT** include him in the CC either!! There is a strong chance it won't be graded if you do!! **Send it only to <abhay1@umbc.edu>**

1) (20 points)

- a) RAID 0 does not provide any redundancy to improve reliability. Instead it uses a technique called striping to improve disk performance. What exactly is striping and how does it improve I/O performance? Why might a person prefer to have 2 (or more) independent drives rather than have them in a RAID 0 setup?
- b) RAID3 and RAID 4 both provide extra disks for data redundancy. In what situations might one prefer to use RAID 4 over RAID 3? What is the bottleneck of RAID 4?
- c) What is the main difference between snooping protocols and directory protocols?
- d) Implementations of directory protocol where the whole directory is centralized in one location have problems scaling to more than a couple hundred processors. What is the bottleneck in basic directory protocols and why does it occur? How might this bottleneck be removed?
- e) In snooping protocols, what is the difference between invalidate schemes and update schemes? For each scheme, discuss a situation where the scheme performs poorly and explain why.

2) (20 points)

(Textbook 4<sup>th</sup> edition Exercise 4.17, p.278) Directory protocols are more scalable than snooping protocols because they send explicit request and invalidate messages to those nodes that have copies of a block, while snooping protocols broadcast all requests and invalidates to all nodes. Consider the 16-processor system illustrated in Figure 4.42 and assume that all caches not shown have invalid blocks. For each of the sequences below, identify which nodes receive each request and invalidate.

- a. P0: write 110 <-- 80
- b. P0: write 108 <-- 88
- c. P0: write 118 <-- 90
- d. P0: write 128 <-- 98

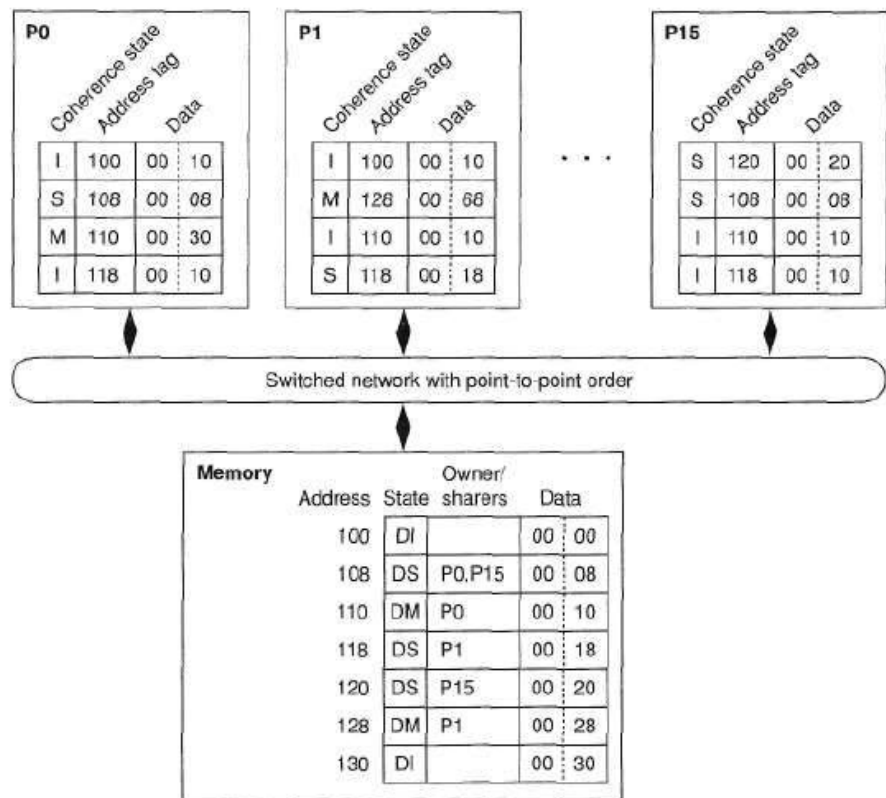
### Case Study 3: Simple Directory-Based Coherence

#### Concepts illustrated by this case study

*m* Directory Coherence Protocol Transitions

- Coherence Protocol Performance
- Coherence Protocol Optimizations

Consider the distributed shared-memory system illustrated in Figure 4.42. Each processor has a single direct-mapped cache that holds four blocks each holding two words. To simplify the illustration, the cache address tag contains the full address and each word shows only two hex characters, with the least significant word on the right. The cache states are denoted M, S, and I for Modified, Shared, and Invalid. The directory states are denoted DM, DS, and DI for Directory Modified,



**Figure 4.42** Multiprocessor with directory cache coherence.

Directory Shared, and Directory Invalid. The simple directory protocol is described in Figures 4.21 and 4.22.

3) (20 points)

Consider the following hard drive:

Seek Time - 8 ms

Rotation Speed - 7200 rpm

Transfer Rate - 80 MB/s

Controller Overhead - 0.1 ms

Sector Size – 512 bytes

- a. We have a 1 MB binary file stored on a single track. The file contains sorted key/value pairs. Each pair consists of two integers. The first integer is the key and the second is its corresponding value. Integers are 4 bytes long. Assuming the integer key we are searching for is contained within the file, calculate the average time needed to find its corresponding value using sequential search. Assume we must do the search on disk, that is, we can only keep one sector (512 bytes) of the file in memory at any one time (e.g. we read 1 sector, process it, then read the next sector). Also, assume the only significant factor in the search time is the time to transfer sectors from disk (once transferred to memory, the time needed for to process each record is negligible).
- b. Consider the same setup as part a, except we can process the file as we read (instead of reading 1 sector, processing it, then going on to the next). On average, how long does it take to find a key/value pair on the hard disk?
- c. What would be the worst case time needed if we used binary search?

4) (40 points)

Consider the SNOOPING cache coherence protocol as follows.

Direct mapping is used. This scheme combines the advantages of both write-through and write-back invalidation. Only the very first write of a cache block uses a write-through policy. If a cache block is written more than once, it applies a write-back policy. There are four states in this protocol:

- **Invalid (Stage 00):** The block is not found in the cache.
  - **Valid (Stage 01):** The cache block, which is consistent with the memory copy, has been read from shared memory and has not been modified.
  - **Reserved (Stage 10):** Data has been written exactly once since it has been read from shared memory. The cache copy is consistent with the memory copy, which is the only other copy. (Write-through policy applies.)
  - **Dirty (Stage 11):** The cache block has been modified (written) more than once, and the cache copy is the only one in the system (thus inconsistent with all other copies). (Write-back policy applies).
- I. Assume that a multiprocessor has 2 nodes, P0 and P1. There is a 32-word centrally shared memory module. Each node has an 8-word local cache. Each memory block and each cache block have **2 words**. Draw the multiprocessor architecture. You have to specify the size of each block and the length of the tag field.

II. Trace the following events. Show all of the changes of the local cache blocks and the shared memory. You can directly mark the changes in the figure you have drawn in part 2). Please indicate the event identifier along with each change. Also, please determine if each of the events is a hit or miss. Note that Mem[x] refers to the access of the centrally shared memory, where x is the hexadecimal number of the WORD address. Initially, Mem[x]=0 for x can be any number between 00 and 1F. All of the local caches are in the invalid state initially.

Event id	Proc Id	Read/ Write	Word address (hexadecimal)	Value (decimal)	<b>Hit/miss? (you have to fill in your answers here)</b>
a	P1	Write	1B	16	
b	P0	Read	1B		
c	P0	Write	13	5	
d	P0	Write	12	6	
e	P1	Read	12		