CS 433 Homework 5

Assigned on 11/7/2017
Due in class on 11/30/2017

Instructions:

1. Please write your name and NetID clearly on the first page.
2. Refer to the course fact sheet for policies on collaboration.
3. Due **IN CLASS** on 11/30/2017.

Problem 1 [4 Points]

Consider a system with the following processor components and policies:

- A direct-mapped L1 data cache of size 4KB (4×2^{10} bytes) and block size of 16 bytes, indexed and tagged using physical addresses, and using a write-allocate, write-back policy
- A fully-associative data TLB with 4 entries and an LRU replacement policy
- Physical addresses of 32 bits, and virtual addresses of 40 bits
- Byte addressable memory
- Page size of 1MB

Part A [2 points]

Which bits (counting from 0 at the LSB) of the virtual address are used to obtain a virtual to physical translation from the TLB? Explain exactly how these bits are used to make the translation, assuming there is a TLB hit.

Part B [2 points]

Which bits of the virtual or physical address are used as the tag, index, and block offset bits for accessing the L1 data cache? Explicitly specify which of these bits can be used directly from the virtual address without any translation.

Problem 2 [12 points]

Consider a tiny system with virtual memory. Physical addresses are 8 bits long, but only $2^7 = 128$ bytes of physical memory is installed, at physical addresses 0 up to 127. Pages are $2^4 = 16$ bytes long. Virtual addresses are 10 bits long. An exception is raised if a program accesses a virtual address whose virtual page has no mapping in the page table, or is mapped to a physical page outside of installed physical memory.

The contents of main memory are shown in Table . To find the physical address of a byte, read the least significant digit from the column label and the most significant digit from the row label. For example, the shaded byte in the second row is at physical address 0x12. All entries are in hexadecimal.

Table 1: Main Memory

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0x0_	4e	65	76	65	72	20	67	6f	6e	6e	61	20	67	69	76	65
0x1_	20	79	**6f**	75	20	75	70	0a	4e	65	76	65	72	20	67	6f
0x2_	6e	6e	61	20	**6c**	**65**	**74**	**20**	79	6f	75	20	64	6f	77	6e
0x3_	0a	4e	65	76	65	72	20	67	6f	6e	6e	61	20	72	75	6e
0x4_	20	61	72	6f	75	6e	64	20	**61**	**6e**	**64**	**20**	64	65	73	65
0x5_	72	74	20	79	6f	75	0a	4e	65	76	65	72	20	67	6f	6e
0x6_	6e	61	20	6d	61	6b	65	20	79	6f	75	20	**63**	**72**	**79**	**0a**
0x7_	4e	65	76	65	72	20	67	6f	6e	6e	61	20	73	61	79	20

The page table is shown in Table 2. The virtual page number in the left column is mapped to the physical page number in the second column. Virtual page numbers are listed in binary.

Table 2: Page Table

Virtual page	Physical page
0	0x2
1	0x4
10	0x1
11	0x5
100	0x4
101	0x7
110	0x9

Part A[2 points]

List the four bytes in the word beginning at physical address 0x34.

Part B[2 points]

How many virtual addresses refer to the first byte of the shaded word in row 0x2_? List them.

Part C[2 points]

How many virtual addresses refer to the first byte of the shaded word in row 0x4_? List them.

Part D[2 points]

How many virtual addresses refer to the first byte of the shaded word in row 0x6_? List them.

Part E[2 points]

What data is returned if the program loads a word from virtual address 0x5C (01011100)?

Part F[2 points]

What is the result if the program loads a word from virtual address 0x64 (01100010)?

Problem 3 [11 points]

Assume the code below executes concurrently on two processors on data f in shared memory as shown. Assume s, i, and j are allocated in registers within each processor. Assume each processor has its own cache which is kept coherent through an invalidation based protocol. Assume a cache block size of 64 bytes with one valid bit per block and an int datatype size of 32 bytes. Assume f is aligned on a cache block boundary; i.e., $f.x$ is the first word in a cache block and $f.y$ is the second word in the cache block.

```
struct foo {
    int x;
    int y;
};
foo f;

//----Processor 1
for (i=0; i<100,000,000; i++)
    s += f.x;

//----Processor 2
for (j=0; j<100,000,000; j++)
    ++f.y;
```

Part A [3 points]

What can you say about the possible cache performance of this code? Explain your answer.

Part B [4 points]

Explain how the code will perform on an update based cache coherent system?

Part C [4 points]

Describe a software-only technique that could eliminate virtually all the misses in the invalidation protocol and perform better than the update protocol scenarios above. Please explain the technique and show the changed code. Your solution should not add synchronization.

Problem 4 [21 points]

This problem concerns MOESI, an invalidation based snooping cache coherence protocol, for bus-based shared-memory multiprocessors with a single level of cache per processor. The MOESI protocol has five states. A block starting at address Addr can be in one of the following states in cache C:

- Modified: The block is present only in cache C and the data in the cache is dirty or modified (i.e., it reflects a more recent version than the copy in memory).
- Owned: The block is present in cache C and may also be present in other caches. The memory may not have an up-to-date copy of this block. The block is said to be owned by cache C and C must service the requests of other caches to this block since memory may not have an up-to-date copy.
- Exclusive: The block is present in a single cache (C) but is clean (i.e., memory has an up-to-date copy of the block).
- Shared: The block is present in cache C and possibly present in other caches.
- Invalid: The block is not valid in cache C (space for the block may or may not be currently allocated in this cache).

If the cache has a block in Owned state, then it services any requests to that block from other processors. Assume that the memory does NOT update its copy even if the request is a read by some other cache and the Owner cache has put the block on the bus. Hence, the owner cache remains in Owned state and continues to service other requests, until the block is replaced from its cache. Also assume that the only way to reach the Owned state is from the Modified or Exclusive state, when some other cache issues a read request for that block.

If a cache C has a block in exclusive or modified state, then it is responsible for servicing any requests to that block from other processors. If the request is a read, then the cache C transitions to the Owned state, and memory does NOT update its copy.

On replacement of a block in Owned or Modified state, the block is sent to memory, and memory resumes responsibility for servicing subsequent requests to that block. Replacement of a block in Exclusive state is similar, except that the block need not be sent to memory (since memory already has a copy).

Assume that after a cache performs a transaction on the bus, there is a mechanism for it to know whether other caches have a copy of the requested block or not at that time. This enables the cache to determine whether to transition to exclusive state.

Part A [2 points]

Consider a Block B in Owned state in the cache of processor P. Can B be in a non-invalid state in any other processors cache? If yes, then what are the possible (non-invalid) states in which B could be in any of the other caches? If no, then explain why not.

Part B [9 points]

This part concerns the response of the cache of processor i to bus transactions initiated by the cache of processor j for a block that starts at address B (referred to as block B below). Fill out the following rows for the state transition table for the cache of processor i, showing the next state for block B in the cache and any action taken by the cache. Each entry should be filled out as: Next State/Action (e.g., S/Send block to memory) where

Next State = M, O, E, S, or I

Action = Send block to memory, Send block to cache, Send block to cache and memory, or No action

Note: If an entry is not possible (i.e., the system cannot be in such a state), write Not Possible:

Current state of block B in cache of processor i	Read of block B by cache of processor j	Invalidate of block B by cache of processor j (with no read request for the block)	Read +Invalidate of block B by cache of processor j
M			
O			
E			

Part C [10 points]

Consider the following sequence of operations by two processors for a block that starts at address B. Determine the state of that block in the caches of both the processors after each operation in the sequence for the MOESI protocol. Both caches are initially empty. The table below is provided to help organize your answer.

No.	Operation	MOESI	
		P1	P2
1	P1 reads B		I
2	P1 writes B		I
3	P2 writes B		
4	P1 reads B		
5	P1 writes B		
6	P2 reads B		

Problem 5: Graduate students only [12 points]

Consider the following graph obtained by varying the amount of data accessed by a certain benchmark. The only thing you know of the experiments is that the system uses virtual memory, a data TLB, only one level of data cache, and the data TLB maps a much smaller amount of data than can be contained in the data cache. You may assume that there are no conflict misses in the caches and TLB. Further assume that instructions always fit in the instruction TLB and an L1 instruction cache.

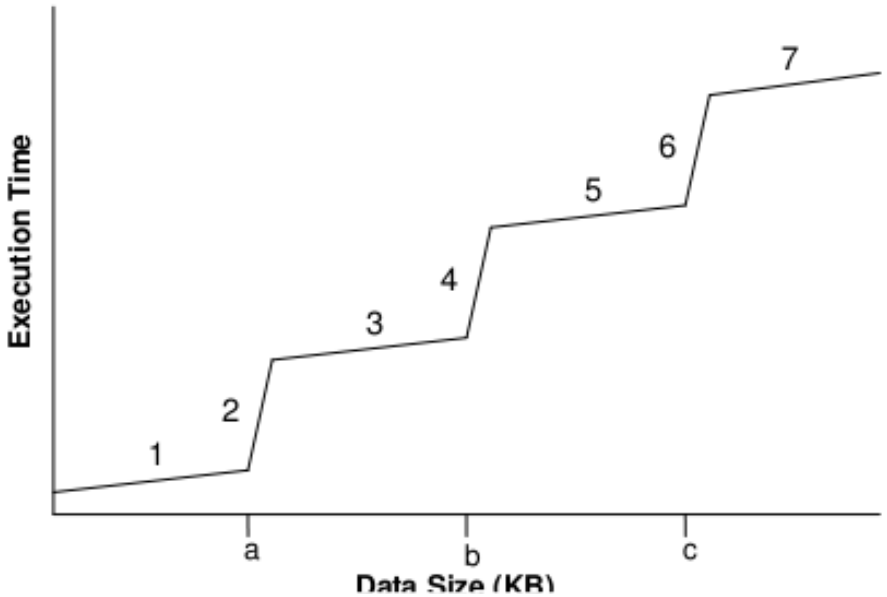

Part A [7 points]

Give an explanation for the shape of the curve in each of the regions numbered 1 through 7.

Region on graph	Explanation
1	
2	
3	
4	
5	
6	
7	

Part B [5 points]

From the graph, can you make a reasonable guess at any of the following system properties? If so, what are they? If not, why not? Explain your answers. (Note: your answers can be in terms of a, b, and c).

1. Number of TLB entries
2. Page size
3. Physical memory size
4. Virtual memory size
5. Cache size