Department of Electrical and Computer Engineering

The University of Texas at Austin
ECE 460N, Spring 2025
Problem Set 4
Due: April 14th, before class
Yale N. Patt, Instructor
Luke Mason, Roy Mor, Jenna May, Margaret Lee, Rathna Sivakumar

Instructions

You are encouraged to work on the problem set in groups and turn in one problem set for the entire group. The problem sets are to be submitted on Gradescope. Only one student should submit the problem set on behalf of the group, but everyone should create a gradescope account and be tagged on the homework. **You will need to assign your questions to pages on gradescope. Each problem you fail to do this will receive a 3 point deduction.**

You will need to refer to the assembly language handout and the LC-3b ISA, microarchitecture, and state diagram documents on the course website.

Cache Memory Questions:

Problem 1

A particular system, similar to the one shown in class, has 256 bytes of memory (i.e. as defined by the ISA). The cache meanwhile has the capacity to store 64 bytes. Every cache line (or cache block) is made up of 8 bytes.

The cache has the following four design options :

1) Direct mapped, 8 sets, 1 way per set
2) 4 sets, Two-way set associative
3) 2 sets, Four-way set associative
4) Fully associative (1 set, 8 ways)

Given these designs, do the following :

1) Come up with an address access pattern (i.e. a list of addresses that are accessed sequentially) such that the direct mapped cache has more hits than the fully associative cache. The address access pattern should have at least 16 items for each of these problems.
2) Come up with an address access pattern such that the fully associative cache has more hits than the direct mapped cache.
3) When would the two way or four way set associative designs be beneficial? What are the tradeoffs between the two designs?

Problem 2

An LC-3b system with a 384B cache using an LRU (least recently used) replacement scheme has the following access pattern (the cache is initially empty):

```
x3000 MISS
x3006 HIT
x3008 MISS
x3104 MISS
x3180 MISS
x300A HIT
x3142 MISS
x3200 MISS
x3006 MISS
x3140 HIT
x3100 MISS
x3186 MISS - (try doing HIT instead!)
```

a. What is the block size of the cache?

b. What is the associativity of the cache? If set-associative, how many ways and sets are there?

c. What is the address mapping of the cache? (i.e. tag bits, index bits, offset bits)

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0

Problem 3

A computer has an 8KB write-through cache. Each cache block is 64 bits, the cache is 4-way set associative and uses a victim/next-victim pair of bits for each block for its replacement policy.
Assume a 24-bit address space and byte-addressable memory. How big (in bits) is the tag store (all metadata, not just tag)?

Problem 4

An LC-3b system ships with a two-way set associative, write back cache with perfect LRU replacement. The tag store requires a total of 4352 bits of storage. What is the block size of the cache? Please show all your work.

Hint 1 : For a 2-way set associative cache, only 1 LRU bit is needed per set.

Hint 2 : $4352 = 2^{12} + 2^{8}$

Problem 5

Based on Hamacher et al., p. 255, question 5.18. You are working with a computer that has a first level cache that we call L1 and a second level cache that we call L2. Use the following information to answer the questions.

● The L1 hit rate is 0.95 for instruction references and 0.90 for data references.
● The L2 hit rate is 0.85 for instruction references and 0.75 for data references.
● 30% of all instructions are loads and stores.
● The size of each cache block is 8 words.
● The time needed to access a cache block in L1 is 1 cycle and the time needed to access a cache block in L2 is 6 cycles.

● The accesses to the caches and memory are done sequentially. If there is a miss in the L1 and a hit in the L2 then the total latency is 7 cycles.
● Memory is accessed only if there is a miss in both caches.
● The width of the memory bus is one word.
● It takes one clock cycle to send an address to main memory.
● It takes 20 cycles to access the main memory.
● It takes one cycle to send one word from the memory to the processor. Thus the total latency to get a word from memory to the processor is 22 cycles.
● The bus allows sending a new address to memory in the same cycle that data is sent from memory to the processor.
● Assume the data is accessible to the processor only AFTER the whole cache block has been brought in from the memory, and buffered on the processor chip. The processor can then access the data independent of and during the cache fill.

1. What is the average access time per instruction (assume no interleaving)?
2. What is the average access time per instruction if the main memory is 4-way interleaved?
3. What is the average access time per instruction if the main memory is 8-way interleaved?
4. What is the improvement obtained with interleaving?

Problem 6

Hamacher, pg.255, question 5.13. A byte-addressable computer has a small data cache capable of holding eight 32-bit words. Each cache block consists of one 32-bit word. When a

given program is executed, the processor reads data from the following sequence of hex addresses:
200, 204, 208, 20C, 2F4, 2F0, 200, 204, 218, 21C, 24C, 2F4
This pattern is repeated four times.
1. Show the contents of the cache at the end of each pass throughout this loop if a direct-mapped cache is used. Compute the hit rate for this example. Assume that the cache is initially empty.
2. Repeat part (a) for a fully-associative cache that uses the LRU-replacement algorithm.
3. Repeat part (a) for a four-way set-associative cache that uses the LRU replacement algorithm.

Problem 7

Below, we have given you four different sequences of addresses generated by a program running on a processor with a data cache. Cache hit ratio for each sequence is also shown below. Assuming that the cache is initially empty at the beginning of each sequence, find out the following parameters of the processor's data cache:

- Associativity (1, 2, or 4 ways)
- Block size (1, 2, 4, 8, 16, or 32 bytes)
- Total cache size (256B, or 512B)
- Replacement policy (LRU or FIFO)

Assumptions: all memory accesses are one byte accesses. All addresses are byte addresses.

Number	**Address Sequence**	**Hit Ratio**
1	0, 2, 4, 8, 16 ,32	0.33
2	0, 512, 1024, 1536, 2048, 1536, 1024, 512, 0	0.33
3	0, 64, 128, 256, 512, 256, 128, 64, 0	0.33
4	0, 512, 1024, 0, 1536, 0, 2048, 512	0.25

ARITHMETIC

Problem 8

Binary-coded Decimal:

An integer, expressed as a 16-bit BCD is 0000 0100 0110 0000.

a. What is its decimal representation?

b. What is its 16-bit 2's complement representation?

c. A register filled with xXXXX is required to add two BCD values in hardware, where xXXXX represents a 16-bit value. What is xXXXX and why?

Problem 9

Using Booth's algorithm, we want to perform an integer multiplication, which is 79 x 30. Since both given values fit within 8 bits, we need at most 4 iterations to complete the multiplication.

1. Let multiplicand be 79 and multiplier be 30. Show the values of the registers MCAND and MULTIPLIER at the end of each iteration.

	MCAND	MULTIPLIER
1		
2		
3		
4		

2. Now let multiplicand be 30 and multiplier be 79. Show the values of the registers MCAND and MULTIPLIER at the end of each iteration.

	MCAND	MULTIPLIER
1		
2		
3		
4		

3. Express Booth's Algorithm for 79 × 30 in a single equation using decimal values. (mcand =79 and multiplier=30)