

MCS-284 Final Exam

Serial #:

This exam is closed-book and mostly closed-notes. You may, however, use a single 8 1/2 by 11 sheet of paper with *hand-written* notes for reference. (Both sides of the sheet are OK.)

Please write your name only on this page. Do not turn the page until instructed, in order that everyone may have the same time. Then, be sure to look at all problems before deciding which one to do first. Some problems are easier than others, so plan your time accordingly. You have 120 minutes to work.

Write the answer to each problem on the page on which that problem appears. You may also request additional paper, which should be labeled with your test number and the problem number.

Printed name: _____

On my honor, I pledge that I have not given, received, nor tolerated others' use of unauthorized aid in completing any work for this course.

Signature for above honor pledge: _____

Problem	Page	Possible	Score
1	2	10	
2	3	10	
3	4	10	
4	5	15	
5	6	10	
6	7	10	
7	8	12	
8	9	11	
9	10	12	
Total		100	

1. [**10 Points**] Consider a direct-mapped cache with a total capacity of 16 words and a block size of one word. Assume the cache is initially empty and that we are using word addresses, so that words have addresses 0, 1, 2, etc. For the following sequence of addresses, indicate which are misses and which are hits. Also show which address (if any) is contained in each location of the cache at the end. To obtain partial credit, you will need to show your work by showing not just the final contents of each cache location, but also all prior contents, crossing them out when they are replaced.

0 10 8 5 26 10 8 4 0 5

2. [**10 Points**] Consider a direct-mapped cache with a total capacity of 16 words and a block size of two words. Assume the cache is initially empty and that we are using word addresses, so that words have addresses 0, 1, 2, etc. For the following sequence of addresses, indicate which are misses and which are hits. Also show which addresses (if any) are contained in each location of the cache at the end. To obtain partial credit, you will need to show your work by showing not just the final contents of each cache location, but also all prior contents, crossing them out when they are replaced.

0 10 8 5 26 10 8 4 0 5

3. [**10 Points**] Consider a two-way set associative cache with a total capacity of 16 words, a block size of one word, and LRU replacement. Assume the cache is initially empty and that we are using word addresses, so that words have addresses 0, 1, 2, etc. For the following sequence of addresses, indicate which are misses and which are hits. Also show which addresses (if any) are contained in each location of the cache at the end. To obtain partial credit, you will need to show your work by showing not just the final contents of each cache location, but also all prior contents, crossing them out when they are replaced.

0 10 8 5 26 10 8 4 0 5

4. [15 Points]

- (a) Why do you think the designers of the Cell processor decided to have half the ring networks carry messages clockwise and half counterclockwise, rather than all in the same direction?
- (b) Why do you think the designers of the Cell processor decided to have two ring networks oriented in each direction, rather than only one in each direction?
- (c) Why do you think the designers of the Oak Ridge National Laboratory's Jaguar supercomputer (a Cray XT) decided to use a 3D hypertorus network to connect the 182,000 nodes rather than using a ring network?
- (d) A program needs to compute the sum of all the elements of a large one-dimensional array. In order to fully utilize a Core 2 Duo, the program runs two threads of execution, one on each core, with each thread adding the values from half the array locations; the two partial sums are added together later. There are two alternatives for how the work could be partitioned. (1) One thread could process the first half of the array and the other the second half. (2) One thread could process the even numbered elements and the other the odd numbered ones. Why might alternative (1) be significantly faster than alternative (2)?
- (e) Consider a situation similar to problem 4d, but now the program wants to increment each array element rather than computing their sum. Why might the difference between options (1) and (2) be even more pronounced for this program?

5. [10 Points] Consider the following 8-bit numerals:

$$A = 01010101$$

$$B = 00101010$$

$$C = 10101010$$

$$D = 11010101$$

- (a) Fill the letters A , B , C , and D into the appropriate blanks to make this statement true, provided that the numerals are interpreted using the twos-complement notation for signed integers:

$$\underline{\quad} > \underline{\quad} > \underline{\quad} > \underline{\quad}$$

- (b) Fill the letters A , B , C , and D into the appropriate blanks to make this statement true, provided that the numerals are interpreted as unsigned integers:

$$\underline{\quad} > \underline{\quad} > \underline{\quad} > \underline{\quad}$$

- (c) Which two of the numerals represent even integers? (This doesn't depend on twos-complement versus unsigned.)
- (d) Which two of the numerals, if added using twos-complement arithmetic, would result in an overflow?
- (e) Taking B as an unsigned integer, what is its value in decimal?


6. [10 Points] A state machine with three states has the following next-state function:

S_1	S_0	B	NS_1	NS_0
0	0	0	0	1
0	0	1	0	0
0	1	0	1	0
0	1	1	0	0
1	0	0	1	0
1	0	1	0	0
1	1	X	X	X

In this table, S_1 and S_0 encode the current state, NS_1 and NS_0 encode the next state, and B is the single input to the state machine.

- (a) The last row in the table has two output don't cares as well as an input don't care. Briefly explain why.

- (b) Label the three states in the following incomplete diagram with their S_1 and S_0 values. For example, write 01 next to the state that has $S_1 = 0$ and $S_0 = 1$.



- (c) Complete the diagram by labeling the unlabeled arrow and by adding the remaining four arrows and labeling them.

7. [**12 Points**] Identify which networking layer (application, transport, network, link, or physical) each of the following pertains to:
- (a) IP
 - (b) HTTP
 - (c) TCP
 - (d) WiFi's use of particular radio frequencies
 - (e) Addressing data to another computer on the same network
 - (f) Ensuring a stream of bytes is delivered in order
 - (g) Addressing data to a computer on another network
 - (h) Verifying that a cached web page is up to date
 - (i) Ethernet hubs
 - (j) Ethernet switches
 - (k) Web proxy caches
 - (l) Routers

8. [11 Points] Recall that the MIPS pipeline we studied contained five stages: IF (instruction fetch), ID (instruction decode and register fetch), EX (execute), MEM (memory), and WB (write back). Suppose that the ALU contained in the EX stage is found to be the slowest unit, and so in order to support a faster clock rate, it is divided into two halves. In place of the EX stage, there are now two pipeline stages, EX1 and EX2. The first half of the ALU is in EX1 and produces some intermediate results into the EX1/EX2 pipeline register. The second half of the ALU is in EX2 and uses those intermediate results to produce the final ALU result. The following questions concern this modified pipeline.

- (a) For an R-format instruction such as

```
add $t0, $t1, $t2
```

which pipeline stage are the operand values (in this case, \$t1 and \$t2) needed in?

- (b) For an R-format instruction such as

```
add $t0, $t1, $t2
```

which pipeline stage is the result value (in this case, the sum) computed in?

- (c) Suppose a program contains the following two instructions:

```
add $t0, $t1, $t2
```

```
add $t3, $t1, $t4
```

Give an example of a third instruction that could follow these first two such that forwarding would be needed to allow it to be correctly computed without delay.

- (d) Give an example of an `add` instruction such that if a program contained a whole bunch of copies of that same instruction in a row, they would run with a CPI of 2 rather than 1. What accounts for the extra clock cycles?

9. [**12 Points**] A two-processor bus-based multiprocessor uses the MESI coherence protocol we discussed in class. Initially address A is in the invalid state in both processors' caches. Processor 1 reads from address A, then processor 1 writes to it, then processor 2 writes to it, then processor 1 reads from it, then processor 2 reads from it, then processor 2 writes to it. These actions are shown in the following table, along with the initial states of address A in each cache. Fill in the remaining states. You may use the letters M, E, S, and I to denote the states modified, exclusive, shared, and invalid.

cache 1 state	cache 2 state
I	I
processor 1 reads	
processor 1 writes	
processor 2 writes	
processor 1 reads	
processor 2 reads	
processor 2 writes	