CS/ECE 552: Introduction to Computer Architecture
Department of Computer Sciences
University of Wisconsin-Madison

Prof. Karthikeyan Sankaralingam

Final Examination

May 13, 2011
Approximate Weight: 27.5%

2 hour 00 minutes

CLOSED BOOK. You can bring two-cheat sheets (two-sides 8.5 x 11 page).

Exam is **one-sided, total of 14 numbered pages**. Plan your time carefully.
One blank page included in the end for rough work. Also use left-side blank pages for rough work.

NAME: _____

Email: _____

Problem number	Maximum points	Actual points
1	**12**	
2	**12**	
3	**24**	
4	**24**	
5	**36**	
6	**20**	
7	**20**	
Total	**148**	

Problem 1: Virtual Memory (4x3 = 12 points)

Consider a memory system with the following parameters. The virtual address is 64 bits (v63-v0 where v0 is least significant). The physical address is 48 bits (p47-p0). The page size is 16 kilobytes. The translation lookaside buffer (TLB) is 64 entries, two-way set-associative and accessed *before* the cache. The cache is unified (i.e., instructions and data), 32 kilobytes, four-way set-associative, LRU replacement, and 128-byte blocks (lines). All addresses are byte addresses.

a) Consider a naïve page table that stores an entry for every page in the virtual address space. How many entries would the page table store? In each page table entry, which bits of the physical address would be stored?

b) What are techniques for making page tables much smaller than the naïve page table of part (a)?

c) For the TLB, which bits will be stored in tag and which will select the set (the index)?

Problem 2: Caches (4x3 = 12 points)

Consider the *same* memory system as Problem 1. The virtual address is 64 bits (v63-v0 where v0 is least significant). The physical address is 48 bits (p47-p0). The page size is 16 kilobytes. The translation lookaside buffer (TLB) is 64 entries, two-way set-associative and accessed *before* the cache. The cache is unified (i.e., instructions and data), 32 kilobytes, four-way set-associative, LRU replacement, and 128-byte blocks (lines). All addresses are byte addresses.

(a) For the cache, which bits will be stored in tag, which will select the set (the index), and which bit select the data within the block?

(b) Increasing cache block size usually improves a cache's miss ratio, but does not always improve performance (e.g., the effective access time for memory access). How is this possible?

(c) A write-buffer for the cached designed for our single-cycle in-order processor in the homework can never make any difference to performance. Justify or explain why this statement is wrong.

Problem 3: Caches (4x6 = 24 points)

Consider a cache with the following characteristics:

1. 64-byte blocks
2. 6-way set associative
3. 1024 sets total
4. 50-bit addresses
5. Byte-addressable cache

a) How many bytes of data storage are there?

b) Show the break-down of the 50-bit address in terms of word-select (or offset) within the cache block, index bits and tag bits.

c) What is the total number of bits needed to implement the cache?

d) How is the index to the cache computed i.e. what bits are used and what logical or arithmetic operation is applied on the bits?

e) Draw a high-level cache organization diagram for this cache. Use the reference diagram shown alongside for a 4-way associate cache with 256 sets.
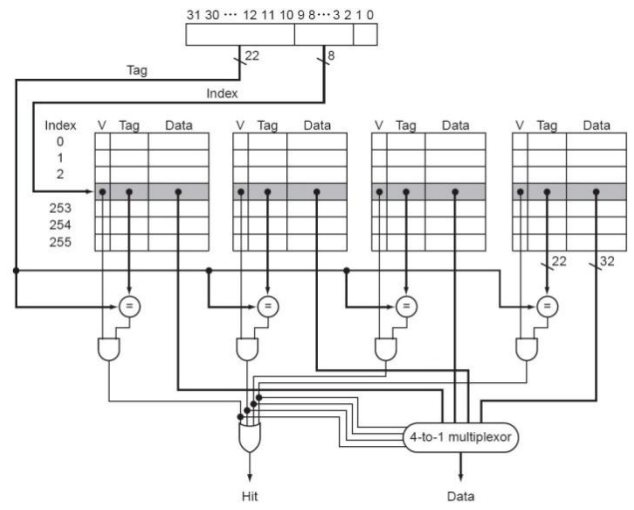
Problem 4: Short answers (4*6 = 24 points)

a) What are *denormalized* floating point numbers? What is their purpose?

b) Define Hamming distance. How many errors can be detected and corrected with a code with Hamming distance of 4?

c) What functions does the operating system provide in accessing I/O?

d) What is the difference between an interrupt and exception?

e) What are the advantages provided by a virtual machine?

f) In a RAID-10 array with 20 disks, how many failures can be tolerated and under what conditions?

Problem 5: Qualitative analysis (4x9 = 36 points)

On the left are 10 descriptions of graphs describing different trends labeled a-j. On the right are 10 graphs. Match each description with a graph that shows the trend and fill in the table below. You can reuse graphs and not all graphs are used at least once.

1. Cache miss penalty vs. block size (on X-axis)

2. Cache hit rate vs. block size

3. AMAT vs. block size

4. Cache access time vs. cache size

5. Execution time of program vs. cache size (single level of cache)

6. Pipeline latency (in cycles) vs. pipeline depth (# of stages)

7. Probability of reference vs. Address range (over entire execution of program)

8. Maximum speedup vs. fraction of program that can be speedup using a given technique

9. MTTF for a disk array vs. number of disks (no redundancy or RAID)

DESCRIPTION	GRAPH NUMBER	DESCRIPTION	GRAPH NUMBER
1		6	
2		7	
3		8	
4		9	
5			
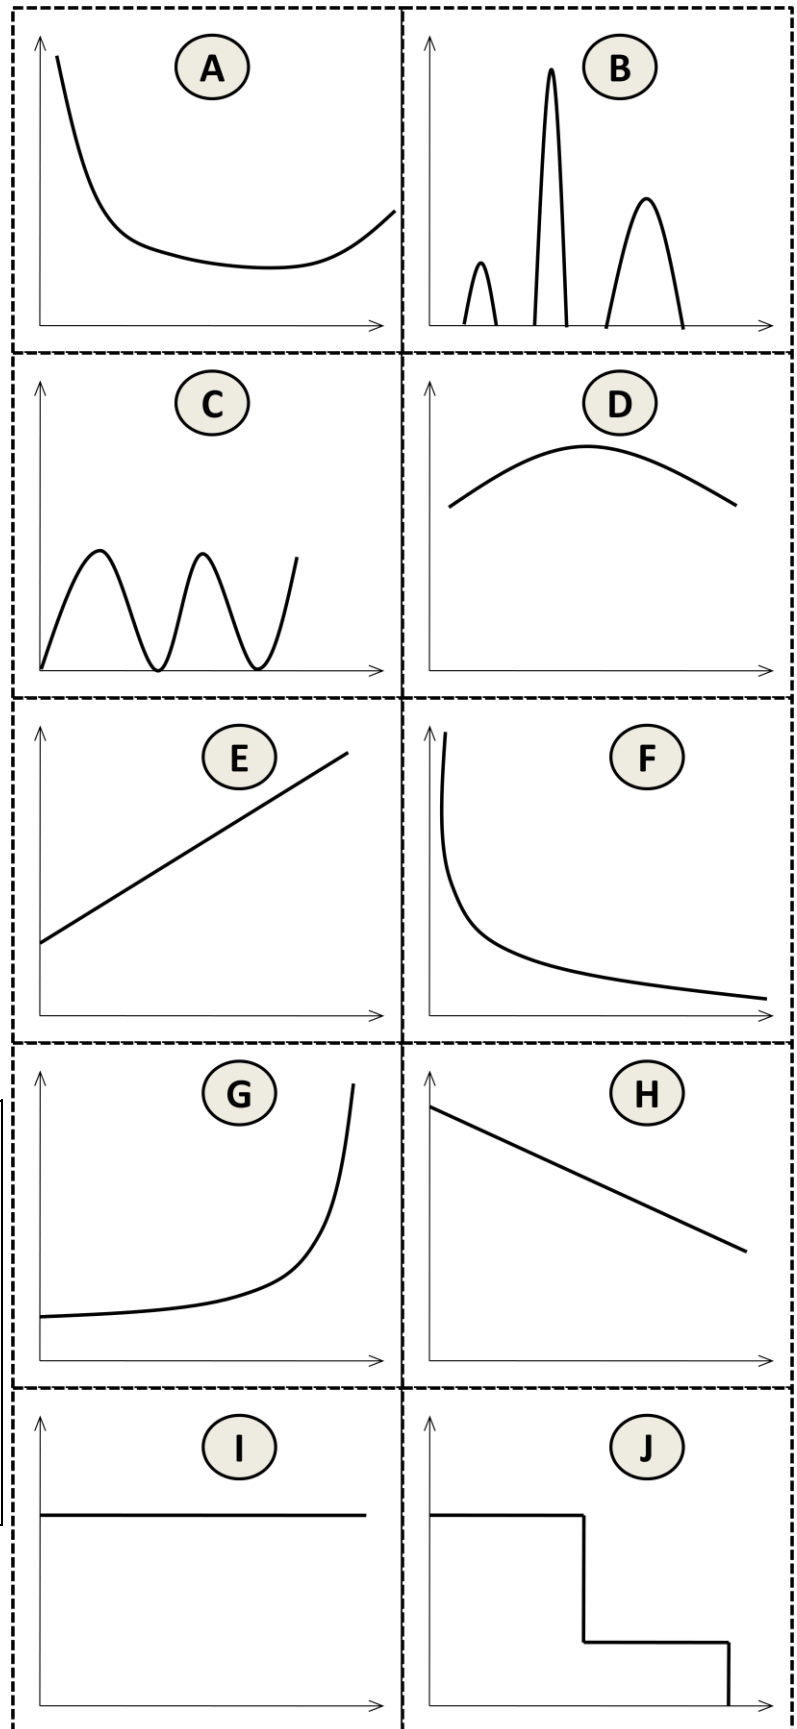
Problem 6: Processor Implementation (20 points)

a) Explain the general verification process you used for the processor you implemented in the project. You may use a diagram if you like. (5 points)

b) Do you believe there may still be some functional bugs in your processor? (0 points)

c) Why? What is missing in the current test strategy? (5 points)

d) What additional steps or changes to the testing strategy will you take to confirm the absence of bugs before deciding to manufacture the processor and ship it to customers? (10 points)

Problem 7: (20 points): Pipeline similar to what you saw in Exam 1.

High performance datapaths use bypass paths (also known as data forwarding logic) to reduce pipeline stalls. However, bypass paths are relatively expensive especially in some wire constrained technologies. To reduce the cost (and potential cycle time impact), some architecture have explored omitting some of the possible bypass paths. Consider the datapath illustrated below (note that the PC update logic and all control logic is intentionally omitted).

This pipelined datapath is similar to the one in the book, but only has bypass paths on one side of the ALU.

Also it has one memory only, which has one port, so that any cycle either an instruction can be fetched or a ld/store can access memory. Assume the mux enforces fairness, so that instruction-fetch and data memory access are never starved and get equal priority by alternating.

Assume that the register file internally bypasses the value, so that if register $i is read and written in the same cycle, then the read returns the new value. Assume that the control logic bypasses the data as soon as possible using the given forwarding data paths, and stalls decode otherwise. You may NOT add additional data paths.

In this problem, you will look at how a program snippet performs on this pipeline. Recall that R-format instructions have the form: `opcode rd, rs, rt`
and I-format instructions have the form `opcode rt, imm(rs)`
or `opcode rt, rs, imm`
Use the table on the next page to show how the given instruction sequence flows through the pipeline and where stalls are necessary to resolve hazards.
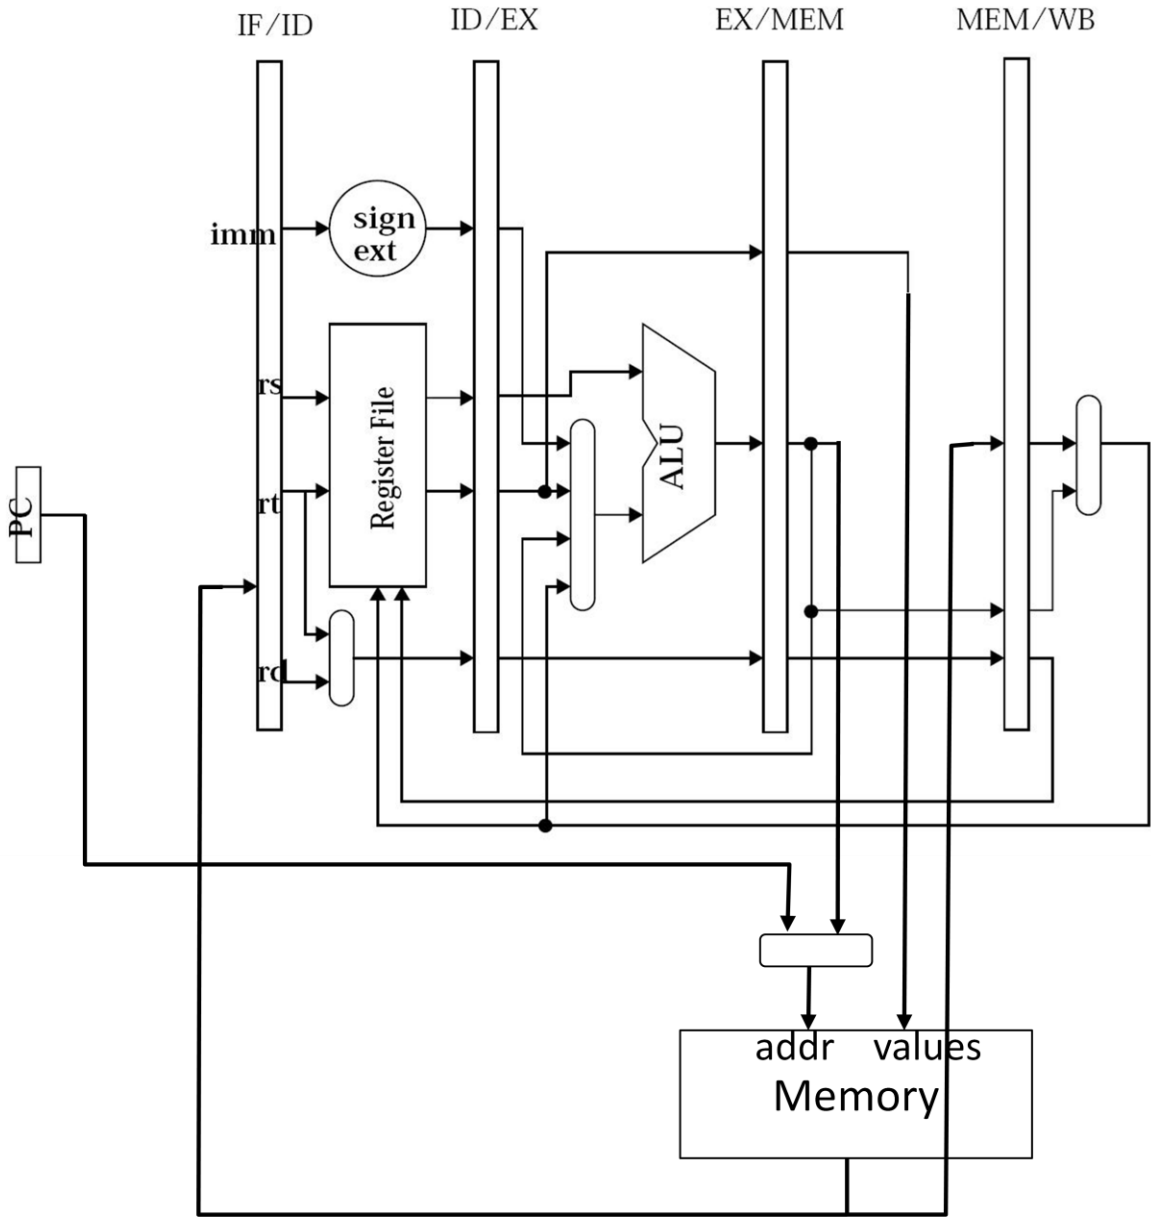
Instruction	Cycle																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
add $1, $2, $3	F	D	X	M	W															
sub $4, $5, $1		F	D																	
or $6, $1, $4																				
and $7, $4, $6																				
lw $9, 4($7)																				
add $1, $9, $2																				
sw $1, 4($7)																				

Consider the code and pipeline above. Show the execution of this code on the pipeline above.
Use the letters, F, D, X, M, and W.

For each cycle where a stall occurs explain why.

Scratch sheet