

COE 308 – Computer Architecture

Exam II – Fall 2008

Monday, January 12, 2009

7:00 – 9:00 pm

Computer Engineering Department
College of Computer Sciences & Engineering
King Fahd University of Petroleum & Minerals

Student Name: _____

Student ID: _____

Q1	/ 15	Q2	/ 15
Q3	/ 15	Q4	/ 25
Q5	/ 15	Q6	/ 15
Total	/ 100		

Important Reminder on Academic Honesty

Using unauthorized information on an exam, peeking at others work, or altering graded exams to claim more credit are severe violations of academic honesty. Detected cases will receive a failing grade in the course.

Q1. (15 pts) Using the refined division hardware, show the **unsigned** division of:

Dividend = **11011001** by Divisor = **00001010**

The result of the division should be stored in the Remainder and Quotient registers.
Eight iterations are required. Show your steps.

Q2. (15 pts) A program, being executed on a processor, has the following instructions mix:

Operation	Frequency	Clock cycles per instruction
ALU	40 %	2
Load	20 %	10
Store	15 %	4
Branches	25 %	3

- a) (3 pts) Compute the average clock cycles per instruction
- b) (5 pts) Compute the percent of execution time spent by each class of instructions
- c) (5 pts) A designer wants to improve the performance. He designs a new execution unit that makes 80% of ALU operations take only **1** cycle to execute. The other 20% of ALU operations will still take **2** cycles to execute. The designer also wants to improve the execution of the memory access instructions. He does it in a way that **95%** of the **load** instructions take only **2** cycles to execute, while the remaining **5%** of the **load** instructions take **10** cycles to execute per **load**. He also improves the store instructions in such a way that each **store** instruction takes **2** cycles to execute.
- Compute the new average CPI
- d) (2 pts) What is the speedup factor by which the performance has improved in part c?

Q3. (15 pts) The following code fragment processes two double-precision floating-point arrays *A* and *B*, and produces an important result in register **\$f0**. Each array consists of **10000** double words. The base addresses of the arrays *A* and *B* are stored in **\$a0** and **\$a1** respectively.

```

        ori    $t0, $zero, 10000
        sub.d  $f0, $f0, $f0
loop:   ldc1   $f2, 0($a0)
        ldc1   $f4, 0($a1)
        mul.d  $f6, $f2, $f4
        add.d  $f0, $f0, $f6
        addi   $a0, $a0, 8
        addi   $a1, $a1, 8
        addi   $t0, $t0, -1
        bne    $t0, $zero, loop

```

- a) (3 pts) Count the total number of instructions executed by all the iterations (including those executed outside the loop).

Assume that the code is run on a machine with a **2 GHz** clock that requires the following number of cycles for each instruction:

Instruction	Cycles
addi, ori	1
ldc1	3
add.d, sub.d	5
mul.d	6
bne	2

- b) (5 pts) How many cycles does it take to execute the above code?

- c) (3 pts) What is the execution time in nanoseconds?

- d) (2 pts) What is the average CPI for the above code?

- e) (2 pts) What is the MIPS rate for the above code?

Q4. (25 pts) Consider the following idea: we want to modify all load and store instructions in the instruction set such that the offset is always 0 (the addressing mode is register indirect only). This means that all load and store instructions will have the following format, where **Rs** is the register that contains the memory address.

```
LW  Rt, (Rs)    # No immediate constant used
SW  Rt, (Rs)    # No immediate constant used
```

a) (10 pts) Draw the modified single-cycle datapath (rotate the page for wider drawing). Identify the changes that you are making to the single-cycle (non-pipelined) datapath.

- b) (7 pts) Assume that the operation delays for the major components are as follows:

Instruction Memory: 200 ps

Data Memory: 200 ps

ALU: 100 ps

Register file (read or write): 50 ps

Ignore the delays in the multiplexers, control, PC access, extension logic, and wires.

What is the cycle time for the single-cycle non-pipelined datapath BEFORE and AFTER making the modification?

- c) (8 pts) Because we have reduced the offset to zero in all load and store instructions, all original load-store instructions with non-zero offsets would now require an additional **ADDI** instruction to compute the address. This will increase the instruction count.

Suppose we have a program in which 20% of the instructions are load-store instructions. Assume further that only 10% of the original load-store instructions have a non-zero offset and would require an additional **ADDI** instruction to compute the address.

What is the percent increase in the instruction count when additional **ADDI** instructions are used?

Which design is better, the original one that allowed non-zero offsets, or the modified one with zero offsets, and why?

What is the speedup factor?

Q5. (15 pts) Consider the following MIPS code sequence:

lw \$5, 100(\$2)

add \$2, \$3, \$5

sub \$5, \$5, \$2

sw \$5, 100(\$2)

- a) (6 pts) Identify all pairs of instructions that have RAW dependencies
- b) (3 pts) Identify all pairs of instructions that have WAR dependencies
- c) (2 pts) Identify all pairs of instructions that have WAW dependencies
- d) (4 pts) Rewrite the above instructions to eliminate all WAR and WAW dependencies

Q6. (15 pts) Consider the following MIPS assembly language code:

```

I1: ADD $4, $1, $0
I2: SUB $9, $3, $4
I3: ADD $4, $5, $6
I4: LW  $2, 100($3)
I5: LW  $2, 0($2)
I6: SW  $2, 100($4)
I7: AND $2, $2, $1
    
```

a) (8 pts) Show the timing of one loop iteration on a 5-stage (IF, ID, EX, MEM, WB) pipeline **without forwarding hardware**. Complete the timing table, showing all the stall cycles that are caused by data hazards. Label all stall cycles (Draw an X in the box). Compute the average CPI for the above code fragment.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
I1: ADD	IF	ID	EX	M	WB															
I2: SUB																				
I3: ADD																				
I4: LW																				
I5: LW																				
I6: SW																				
I7: AND																				

b) (7 pts) Repeat part (a) on a pipeline that **supports forwarding**. Label all data forwards that the forwarding unit detects with an arrow between the stage handing off the data and the stage receiving the data. Compute the average CPI for the above code fragment.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
I1: ADD	IF	ID	EX	M	WB															
I2: SUB																				
I3: ADD																				
I4: LW																				
I5: LW																				
I6: SW																				
I7: AND																				

Additional page if needed