EECS 470 *Midterm Exam*

Winter 2008

Name: _____ unique name: _____

Sign the honor code:

I have neither given nor received aid on this exam nor observed anyone else doing so.

_____

Scores:

#Page	Points
2	**/10**
3	**/12**
4	**/15**
5	**/13**
6	**/17**
7	**/13**
8&9	**/20**
Total	***/100***

NOTES:

- Open book and Open notes
- Calculators are allowed, but no PDAs, Portables, Cell phones, etc.
- Don't spend too much time on any one problem.
- You have about 120 minutes for the exam.
- There is a 3 point extra credit problem on page 6. It is insanely hard, so don't even start it until you are done with everything else and have checked everything over.
- There are ***9*** pages including this one.
- Be sure to show work and explain what you've done when asked to do so.
- **There are two "answer areas" for problem 8 (pages 8&9). Clearly mark which one you want graded or we will grade the first one.**

1. **Written short answer (Your answers must be one or two sentences.) [10 points]**
    a. In the algorithm we're calling T3 (using an RRAT) explain exactly when a PRF entry is freed assuming there is no mis-speculation. **[3]**

    b. In the algorithm we're calling T3 (using an RRAT) explain exactly when a PRF entry is freed by a mis-predicted branch. **[3]**

    c. In the algorithm we're calling T1, explain exactly when the ARF is updated. **[4]**

2. **Short answer (*Clearly* show your work). [12 points]**
    a. If you have a 32-bit address space (addresses are 32 bits) and you have a 4-way associative cache that uses 5 bits as the set index and 4 bits as the byte offset, how large is the data portion of the cache (in bytes)? **[4]**

    b. In T2 if you have 16 architected registers, 8 Reservation stations and 32 Reorder Buffer entries, how many bits of memory will you need to store the rename table? **[3]**

    c. Consider the pipelined processor you did as part of project 3 *but with the structural hazards removed* (so you can fetch and load/store in the same cycle). Say 20% of all instructions were loads, 30% were branches, and 10% were stores. Assume the program is quite long (millions and millions of instructions). In addition, 25% of all instructions are data dependent on the instruction in front of them, and branches are taken 25% of the time. What would you expect the CPI to be? **[5]**

3. **Mr. Amdahl (*Clearly* show your work) [7 points]**
   a. Consider a workload where 50% of the execution time consists of multimedia processing for which a multi-media instruction set extension might be helpful. According to Amdahl's law, what is the maximum speedup that can be achieved by implementing these extensions? **[3]**

   b. Now, say that you work at AMD and the designers there claim that multimedia code sequences will see a 3.5 times (3.5X) speedup by using special multi-media extensions. What is the fraction of the execution time that must be multi-media code in order to achieve an overall speedup of 1.5X? **[4]**

4. **Physical Register File [8 points]**
   a. Consider the algorithm we've named "T3". What is the size of the physical register file you should use if you want to *guarantee* there will never be a stall because the free list for the physical register file is empty? Assume that there are 32 architected registers, 48 RoB entries, 12 reservation stations, and that we are using an RRAT. Clearly *explain* your work. **[4]**

   b. Explain why you should probably not make your PRF as big as your calculation above would suggest. **[4]**

5. **CMOV [13 points]**

a. Consider the following pseudo-assembly code segment that computes the absolute value of a number in R1 and puts it into R2.

```
        If(R1>0) goto DONE
        R1=(-1)*R1
DONE:   R2=R1
```

Re-write this code using a conditional move so that does the same thing but without using a branch. Your conditional move instruction should be of the form: `If(Rx==0) Ry=Rz` where the "==" can be replaced by ">" or "<" as you desire (but the 0 must be a 0) and the registers can be specified to be any given register. After this code segment, all registers other than R1 are potentially live (meaning they may be read before they are written to). Your solution should be no more than 6 lines of assembly. **[5]**

b. Now re-write the following code using a conditional move to again eliminate a branch while preserving functionality (including faults). You must follow the restrictions described above on the form of CMOV you can use.
In this case, all registers other than R3, R5, R6 and R7 are potentially live after this code segment. Your solution should have no more than 12 lines of assembly for full credit. **[8]**

```
        R4=0
TOP:    R3=R3+1
        If(R4>2000) goto BOT
        R5=R2/R3
        R4=R4+R5
BOT:    IF(R3<50) goto TOP
```

6. **Predictors [<u>17 points</u>]**
Consider the following types of predictors:
A) A PC indexed branch-history table with eight 1-bit entries
B) A PC indexed branch-history table with four 2-bit entries (done as a standard counter)
C) A local history predictor where the branch-history table has eight 3-bit entries and points to a pattern-history table that has 1-bit entries.

a. *Assuming the entire program had only one branch,* what would be the expected hit rate of each of these branch predictors on a branch that has a repeating pattern of NNNNT (it does that pattern forever, so NNNNTNNNNT etc.).
[5 points, -2 for each wrong or blank answer, min 0]

A)	B)	C)

b. Now consider the following code segment that is 5 assembly instructions in length:

```
X:    If(R1>0) goto BRANCH X
      ALU or MEMORY OP
Y:    if(R1%2==0) goto Z
      ALU or MEMORY OP
Z:    if(R1!=1000000) goto X
```

The branch at X is taken 50% of the time (coin-flip each time). The branch at Y follows the pattern NNNNT repeating forever and the branch at Z is always taken. For each given predictor, compute the expected rate of *<u>correctly</u>* predicting the branch. (The one with an "X" is extra credit, if you get it right *and* have a correct detailed explanation you get +3 points. It's <u>much</u> harder than it's worth so don't waste time on it until you have everything else done and checked over.)
[12 points, -2 per box wrong or blank, min 0]

	A)	**B)**	**C)**
Branch X			
Branch Y			X
Branch Z			

Extra credit (with detailed explanation)**[+3]**.

7. **Performance of Tomasulo's [<u>13 points</u>]**
Consider the following pseudo-assembly which *<u>loops forever</u>*:

```
TOP:  R1=MEM[R2]
      R2=R2+4
      R3=R3+1
      R4=R1+R4
      R4=R4-2
      R5=R1+R5
      R5=R5/R3
      If(R5<100000) goto TOP
```

Say you have a machine which can issue one instruction per cycle, finish execution of one instruction per cycle, and retire one instruction per cycle. Branch mis-predictions are resolved when the branch hits the head of the RoB *and nothing is done about mis-predicted branches before that.* This machine implements what we have called "Tomasulo's 3", has a RS size of 3 and a RoB size of 5. Loads take 10 cycles to execute and all other instructions take 1 cycle to execute. There is no value prediction or other ways to break RAW hazards, but you can assume perfect branch prediction. Instructions don't release an RS until they finish execution.

a. What is the best CPI that could be achieved? To receive credit you must clearly explain your answer. **[6]**

b. What is the minimum size of the RS and RoB to get this program to get a CPI of one (ignoring warm-up issues). To receive credit, you must clearly explain your answer. **[7]**

8. Consider the following state of a machine implementing what we've called Tomasulo's third algorithm.

RAT

Arch Reg #	Phy. Reg #
0	7
1	9
2	2
3	3
4	8

ROB

Buffer Number	PC	Executed?	Dest. PRN	Dest ARN	
0	12	Y	5	1	← HEAD
1	16	N	6	1	
2	20	N	7	0	
3	24	Y	--	--	
4	28	N	8	4	
5	32	Y	9	1	← TAIL
6					
7					
8					

RRAT

Arch Reg #	Phy. Reg #
0	0
1	1
2	2
3	3
4	4

RS

RS#	Op Type	Op1 Ready?	Op1 PRN/value	Op2 Ready?	Op2 PRN/value	Dest PRN	ROB
0	ADD	N	6	Y	8	7	2
1	LD	Y	13	Y	8	6	1
2	ADD	N	7	Y	6	8	4
3							
4							

KEY:

- **Op1 PRN/value** is the value of the first argument if "Op1 ready?" is yes; otherwise it is the Physical Register Number that is being waited upon.
- **Op2 PRN/value** is the same as above but for the second argument.
- **Dest. PRN** is the destination Physical Register Number.
- **Dest. ARN** is the destination Architectural Register Number.
- **ROB** is the associated ROB entry for this instruction.
- **Free/Valid** indicates if the PRF entry is currently available for allocation and if the valid in it is valid. A free entry should be marked as invalid.

PRF

Phy Reg #	Value	Free	Valid
0	4	N	Y
1	5	N	Y
2	6	N	Y
3	7	N	Y
4	8	N	Y
5	13	N	Y
6	10	N	N
7	12	N	N
8	13	N	N
9	12	N	Y
10	15	Y	N
11	16	Y	N
12	17	Y	N

Say that the instruction in ROB #3 is a branch and it was mis-predicted: the next PC should have been 100. Also, assume the load in RS1 returns a value of 30. Now, say that the instruction in memory location 100 is R1=R1+R4 and in location 104 is R2=R3+R1. Update the machine to the state where the branch has left the RoB, and the instructions at location 100 and 104 have issued but not executed. When faced with an arbitrary decision, just be sure to make a legal choice. ***Be sure to update the head and tail pointers!*** **[20]**

On the following page is an extra copy of this state. You may use this one or the one on the next page but be sure to cross out (with a BIG X) the one you don't want graded.

EXTRA COPY. BE SURE TO CROSS OUT EITHER THIS ONE OR THE ONE ON THE PREVIOUS PAGE.

RAT

Arch Reg #	Phy. Reg #
0	7
1	9
2	2
3	3
4	8

ROB

Buffer Number	PC	Executed?	Dest. PRN	Dest ARN	
0	12	Y	5	1	← HEAD
1	16	N	6	1	
2	20	N	7	0	
3	24	Y	--	--	
4	28	N	8	4	
5	32	Y	9	1	← TAIL
6					
7					
8					

RRAT

Arch Reg #	Phy. Reg #
0	0
1	1
2	2
3	3
4	4

RS

RS#	Op Type	Op1 Ready?	Op1 PRN/value	Op2 Ready?	Op2 PRN/value	Dest PRN	ROB
0	ADD	N	6	Y	8	7	2
1	LD	Y	13	Y	8	6	1
2	ADD	N	7	Y	6	8	4
3							
4							

PRF

Phy Reg #	Value	Free	Valid
0	4	N	Y
1	5	N	Y
2	6	N	Y
3	7	N	Y
4	8	N	Y
5	13	N	Y
6	10	N	N
7	12	N	N
8	13	N	N
9	12	N	Y
10	15	Y	N
11	16	Y	N
12	17	Y	N

KEY:

- **Op1 PRN/value** is the value of the first argument if "Op1 ready?" is yes; otherwise it is the Physical Register Number that is being waited upon.
- **Op2 PRN/value** is the same as above but for the second argument.
- **Dest. PRN** is the destination Physical Register Number.
- **Dest. ARN** is the destination Architectural Register Number.
- **ROB** is the associated ROB entry for this instruction.
- **Free/Valid** indicates if the PRF entry is currently available for allocation and if the valid in it is valid. A free entry should be marked as invalid.

Say that the instruction in ROB #3 is a branch and it was mis-predicted: the next PC should have been 100. Also, assume the load in RS1 returns a value of 30. Now, say that the instruction in memory location 100 is R1=R1+R4 and in location 104 is R2=R3+R1. Update the machine to the state where the branch has left the RoB, and the instructions at location 100 and 104 have issued but not executed. When faced with an arbitrary decision, just be sure to make a legal choice. ***Be sure to update the head and tail pointers!***