Practice Exam

CS 3853

Adopted from previous offerings of the course

Name (please print): _____

Student ID: _____

Total points: 90

1. (*6 points*) Given the function $F(A, B, C) = \Sigma m(1, 2, 5, 6, 7)$.

- Draw the Karnaugh map for the function

- Find all prime implicants. Circle them on the map above and give their boolean expressions.

- List all essential prime implicants

- Find a minimal SOP

- Is the minimal solution unique? If it is, prove it, otherwise show an alternative.

2. (*8 points*) Given the function $F(A, B, C, D) = \Sigma m(1, 2, 3, 4, 5, 7, 10, 12, 13) + d(8, 11)$.

- Draw the Karnaugh map for the function

- Find all prime implicants. Circle them on the map above and give their boolean expressions.

- List all essential prime implicants.

- Find a minimal SOP

3. (*8 points*) The following branch prediction algorithm has been suggested:

- start by predicting "taken"
- if you predicting "taken", then after two consecutive incorrect predictions predict "not taken"
- if you predicting "not taken", then after two consecutive incorrect predictions predict "taken"

The system has one input Z. If $Z = 1$, the current branch is taken, not taken otherwise. The output $X = 1$ predicts that the branch should be taken, not taken otherwise.

a. (*4 pts*) Draw the state diagram.

b. (*1 pt*) How many flip flops are required to implement the finite state machine?

c. (*1 pt*) How many state assignment are possible for the finite state machine described above? (Considering the number of flip flops given in part (b).)

d. (*2 pts*) Show a sequence of branches where this prediction will perform very poorly. You answer should look like: `TTTNNNNTNTTT.....` How many branches will be correctly predicted in your example?

4. (*5 points*) Show how a T Flip Flop can be used to implement a JK Flip Flop. The characteristic table of a JK Flip Flop is shown below. Show your work: draw the K-map(s), derive the equation(s), and draw the circuit.

J	K	q^*
0	0	q
0	1	0
1	0	1
1	1	$\overline{q}$

5. (*2 points*) Give two reason why a split cache is preferred over a unified one.

6. (*8 points*) Given a processor with 8K data cache and lines with 64 bytes. The following matrix has been declared `int mat[1024][1024]` in a C program. Two dimensional arrays are stored continuously one row after the other. Assumptions: an `int` is 4 bytes, `r` and `c` are stored in registers.
Consider the two C program snippets below
(a)

```
for(int r = 0; r < 1024; r++) {
    for(int c = 0;  c < 1024; c++) {
        mat[r][c] = mat[r][c] * 2;
    }
}
```

(b)

```
for(int c = 0;  c < 1024; c++) {
     for(int r = 0; r < 1024; r++){
        mat[r][c] = mat[r][c] * 2;
    }
}
```

Please answer the following questions (if the exact values is too difficult to calculate, then give reasonable approximations.)

a. (*2 pts*) What is the data cache hit ratio for snippet (a) (with a direct cache)?

b. (*2 pts*) What is the data cache hit ratio for snippet (b) (with a direct cache)?

c. (*2 pts*) What is the data cache hit ratio for snippet (a) (with a 4-way associate cache)? Explain how you obtained the result.

d. (*2 pts*) What is the data cache hit ratio for snippet (b) (with a 4-way associate cache)? Explain how you obtained the result.

7. (*6 points*) Consider the following scenarios:

1. A single 2 TB disk.
2. Eight 250 GB disks set up as RAID 5 (block level distributed parity.)

Describe the advantages and disadvantages of the first setup versus the second one. You may assume that the access time and the data transfer time is the same for the 2 TB disk as for each 250 GB disk. You may also assume that each system has the same cost. Choose your criteria of comparison carefully (I can think of 6 important ones.)

8. (*11 points*) For this question, one memory access will read/write 4 bytes, registers hold 4 bytes of data, the size of each instruction is determined as follows:

- Opcode – 1 byte
- Memory Address – 4 bytes
- Register – 4 bits

The calculation $Q = (X - Y)/(X + Y)$ can be accomplished on a CISC computer with the following statements:

		Size	Mem. Access
`ADD`	`W,X,Y`	____	____
`SUB`	`Q,X,Y`	____	____
`DIV`	`Q,Q,W`	____	____

The same calculation can be accomplished on RISC computer with the following statements:

		Size	Mem. Access
`LOAD`	`R1,X`	____	____
`LOAD`	`R2,Y`	____	____
`LOAD`	`R3,R1`	____	____
`SUB`	`R1,R2`	____	____
`ADD`	`R3,R2`	____	____
`DIV`	`R1,R2`	____	____
`STORE`	`R1,Q`	____	____

a. (*2 pts*) Give the size of each instruction in the spaces provided above (in bytes).

b. (*2 pts*) Give the number of memory accesses for each instruction in the spaces provided above.

c. (*4 pts*) Assume that the CICS computer also has some registers. In the instructions above any memory reference can be replaced with a register. For example, `ADD R1,X,R2` would be valid. Additionally, `LOAD` and `STORE` instructions are supported. Write the code for the calculation $Q = (X - Y)/(X + Y)$ for such a processor. Compare its size and number of memory access to the other two versions.

d. (*3 pts*) If you had to design a pipeline for the CISC (with registers) or the RISC computer (as sketched above), which would you prefer? Give three good arguments.

9. (*10 points*)

a. (*2 pts*) Explain the principle of predicated instructions.

b. (*8 pts*) Give an example where predicated instructions would be useful. Consider following for your answer:

1. Show the C-like code that you use for your illustration.

2. Show the assembly code without predication.

3. Show the assembly code with predication.

4. Explain the benefits of predication as shown in your example.

10. (*3 points*) Name three charactristics of a superscalar processor.

11. (*4 points*) Write a piece of C-like code which would have at least three branches when translated to assembly code. Which branches could you predict? What would your expected success rate be?

12. (*6 points*) Consider a processor with 4 cores. Draw **two** diagrams on how a 3 level cache could be organized. What are the main issues with caches on multi-core processors in general and on your designs in particular? Briefly contrast the advantages/disadvantages of the two designs.

13. (*9 points*) Consider the data paths of a CPU shown in the figure below. Give the micro instructions for the following (indicate which micro instructions can be done in parallel):

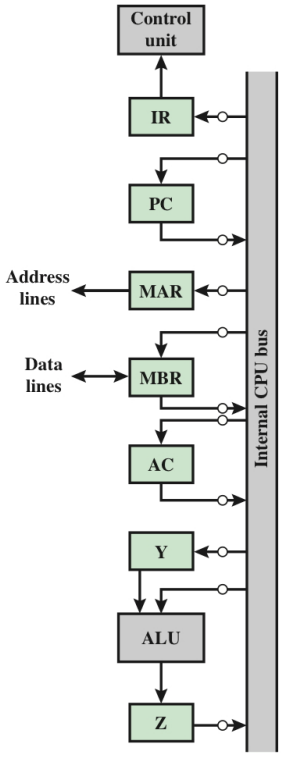

- Fetch an instruction
- Execute the instruction `ADD M(X)` ie AC = AC+M(X)
- Execute the instruction `JUMP X`

You can assume that the address of an instruction `X`, is still in the MBR after the instruction has been fetched.

14. (*4 points*) You are asked to implement pipelining for the computer structure above. Briefly explain the complexity of this endeavour. Consider the effort of adding a "fetch stage."