Computer Architecture

December 13th, 2004, 14:00-17:00 (or whatever the normal exam times are at the VU)

- This exam consists of 5 questions with several sub questions. Each subquestion has the same weight.
- This version is in English only; the real exam will have a version in both Dutch and in English. Feel free to answer either in English or Dutch.
- Each question should start on a *new page*. Make sure to write your name and student number on each page.
- The use of books, notes or papers is not allowed, but **a calculator is permitted** (and, indeed, seems advisable).
- Make your answers brief and to the point, do *not* use pencil or red ink.

Question 1

A processor has a clock-rate of 1GHz and the following timing information:

Type	**CPI**
ALU operations	1
FP operations	20
Branches	3
Loads	3
Stores	2

On this processor a benchmark is executed which has the following instruction profile:

Type	**Number of instructions** ($\times 10^6$)
ALU operations	90
FP operations	100
Branches	60
Loads	100
Stores	50

(a) What is the overall CPI?

(b) What is the total CPU time that is needed to execute the program?

Suppose we could change the design using one of the following alternatives:

- introduce a new instruction 'XFP' that is able to perform a FP operation on a vector of operands. XFP reduces the number of FP operations by 40%. The CPI for XFP is 30 and it is used 10×10^6 times in the benchmark.
- reduce the CPI for all FP operations to 18.

(c) Which is the better alternative and why?

Question 2

(a) In many of the architectures discussed in this course, the same registers could be read and written in the same clock cycle. How is this possible?

Consider the microinstruction in Figure 1.

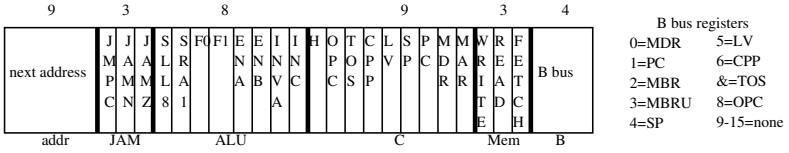

Figure 1: Microinstruction format

(b) Explain with an example the use of the JMPC field.

Question 3

Consider the following applications:

- playing video and audio on a computer.
- recording 'raw' video and audio and transforming them to a compressed format (e.g. MPEG) which is subsequently stored on disk.

The audio and video are the only data you need to consider.

(a) Argue, using only the concepts of 'spatial' and 'temporal' locality, whether (and if so, why) you think a data cache would be useful in each of these cases.

Consider a byte-addressable machine with 32 bits addresses, 8 bytes in a word and 32 words in a cache line. The size of the cache is 256 KByte and the cache is direct mapped. Initially, it is empty.

(b) Given the hexadecimal address 1234 5678. In what cache line will the byte at this address be stored? Explain. Also give the tag corresponding to this entry.

(c) Explain in detail what happens if subsequently the following (hexadecimal) addresses are referenced: (i) 1234 8765, (ii) 1234 5688, (iii) 4936 5678, (iv) 1234 5680

For the next subquestion, consider the following code fragment that occurs in a loop:

```
if (a == 0) a = 9;
if (a == 9) WriteOutput ("a is equal to 9");
```

(d) A '(2,2) branch predictor' is an example of a two-level (or 'correlating') branch predictor. Explain how a '(2,2) branch predictor' works and argue *for the example code* why it is more useful than a simple 1-bit branch predictor.

Question 4

(a) What is the difference between RISC and CISC?

(b) What is the advantage of having a fixed-length instruction? What might be a disadvantage?

(c) What is the difference between segments and pages? Also explain what each is used for.

Question 5

(a) Give a reason why the number of nodes that can be used in a Uniform Memory Access (UMA) machine is limited.

(b) Given an example of behaviour that *is* processor consistent and *not* sequentially consistent.