LINKÖPINGS TEKNISKA HÖGSKOLA
Institutionen för datavetenskap
Petru Eles

Tentamen i kursen

Datorarkitektur - TDDI03

2015-01-12, kl. 8-12

Hjälpmedel:

Engelsk ordbok.

Supporting material:

English dictionary.

Poänggränser:

Maximal poäng är 40.
För godkänt krävs sammanlagt
21 poäng.

Points:

Maximum points: 40.
In order to pass the exam you need a
total of minimum 21 points.

Jourhavande lärare:

Petru Eles, tel. 0703681396

Good luck !!!

1. Unified caches and separate data and instruction caches: draw a picture for each of the two alternatives and comment on advantages and disadvantages.

    (3p)

2. The Pentium 4 has an L1 instruction cache which is particular in several regards.
    In what consists the particularity and what is the reason behind it?

    (2p)

3. Consider a pipelined processor with k pipeline stages.
a) What is the theoretical acceleration (ignoring overheads) for a sequence of n instructions, compared to a similar but non-pipelined processor? Show how you obtain the formula!
b) What is the acceleration of a sequence of 20 instructions if the number of pipeline stages is 5?
c) What is the acceleration for an infinitely long sequence if the number of pipeline stages is 5?

    (3p)

4. Branch history table: what does it contain and how is it used?

    (2p)

5. The design of RISC architectures is based on certain characteristics of typical programs which are frequently used. Enumerate at least five such characteristics of programs.

    (2p)

6. Dynamic branch prediction with a two-bit scheme. How does it work?
    Illustrate with the case of a loop like the one below. Compare with one-bit prediction.

    ```
            - - - - - - - - - - - - - - - -
    LOOP    - - - - - - - - - - - - - - - -
            - - - - - - - - - - - - - - - -
    BNZ     LOOP
            - - - - - - - - - - - - - - - -
    ```

    (3p)

7.
a) What is the role of the page table in a virtual memory system? What data does it store?
b) The page table is very large, usually too large to be stored in main memory. Such a large size, at the same time, makes access to the page table very slow. How is this solved in current microprocessor architectures.

    (3p)

8.

a)   What is a superscalar architecture?

b)   Draw a block-diagram of a superscalar unit.

(2p)

9.

Consider the following sequence of machine instructions:

1:	R1	←	100
2:	R5	←	R1 + R2
3:	R7	←	R5 + 1
4:	R1	←	R2 * R4
5:	R5	←	0
6:	R2	←	R4 - 25
7:	R3	←	R7 * 2
8:	R4	←	R1 + R3
9:	R10	←	0
10:	R1	←	R1 * 30

a) Indicate the data dependencies among instructions.

b) Consider a superscalar computer on which the execution of each instruction takes one cycle; the computer has two units that execute addition&subtraction and one unit for multiplication. Produce a table, according to the model below, showing how instructions are executed in consecutive cycles, if we assume in order execution.

c) Rename the registers in the above sequence to prevent, where possible, dependency problems. Produce a second table, according to the model below, showing how instructions (after register renaming) are executed in consecutive cycles, if we assume out of order execution.

In the table cells indicate the sequence number (1, 2, .., 10) of the instruction executed in the corresponding cycle on the respective unit.

	ADD/SUB	ADD/SUB	MUL
Cycle 1			
Cycle 2			
Cycle 3			
. . .			

(4p)

10. Compare VLIW architectures with superscalar architectures:
a) Show similarities and differences.
b) Show the advantages and disadvantages of the two approaches.
c) Why is a superscalar consuming more power, compared to a VLIW computer?

(4p)

11. What is loop unrolling? How does it work? Why is it important with VLIW architectures? Illustrate by an example.

(3p)

12.
a) What is branch predication (like in the Itanium architecture)?
b) Compare with ordinary branch prediction.

(3p)

13. What is the role of the mask register in a vector unit? Give an example.

(3p)

14. Formulate Amdahl's law and comment.

(3p)