LINKÖPINGS TEKNISKA HÖGSKOLA
Institutionen för datavetenskap
Petru Eles

Tentamen i kursen

Datorarkitektur - TDDI03

2014-01-13, kl. 08-12

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 !!!

Tentamen i kursen Datorarkitektur - TDDI03, 2014-01-13, kl. 08-12
Du kan skriva på svenska eller engelska!

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

(2p)

2.
a) Why do we need special *write strategies* for cache memories?
b). We have discussed three write strategies: *write-through*, *write through with buffered write*, and *copy back*. How do they work? Which are their advantages and disadvantages?

(3p)

3. Data hazards in pipelines can sometimes be avoided by a technique called *forwarding*. How does this technique work? Give an example in which forwarding produces an acceleration (draw a figure which illustrates the corresponding pipelined execution).

(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. Delayed load. Why do we need it with RISC architectures? How does it work?
   Give an example.

(2p)

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.

(3p)

9.

Consider the following sequence of machine instructions:

1.	R3	<-	R10 + R1
2.	R7	<-	R5 * 10
3.	R11	<-	R3 - R0
4.	R4	<-	R11 - R16
5.	R10	<-	R7 + 101
6.	R9	<-	R4 + R12
7.	R8	<-	R6 * R0
8.	R2	<-	R4 * R7
9.	R13	<-	R17 + 2
10.	R15	<-	R12 * R22
11.	R14	<-	R15 + 3

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 two tables, according to the model below, showing how instructions are executed in consecutive cycles, if we assume:

b1) in order execution

b2) out of order execution

In the table cells indicate the sequence number (1, 2, .., 11) 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.
a) What is hardware multithreading?
b) Why do multithreaded processors provide higher performance?
c) We have described three approaches to multithreading: interleaved, blocked, and simultaneous; what is the main characteristic of each of them?

(3p)