Försättsblad till skriftlig tentamen vid Linköpings universitet

(fylls i av ansvarig)

Datum för tentamen	2010-04-08
Sal	U14, R42, R44, R41
Tid	8-12
Kurskod	TDTS08
Provkod	TEN1
Kursnamn/benämning	Datorarkitektur
Institution	IDA
Antal uppgifter som ingår i tentamen	13
Antal sidor på tentamen (inkl. försättsbladet)	4
Jour/Kursansvarig	Zebo Peng
Telefon under skrivtid	0702582067
Besöker salen ca kl.	9:40
Kursadministratör (namn + tfnnr + mailadress)	J. Mellheden 2297
Tillåtna hjälpmedel	Engelsk ordbok
Övrigt (exempel när resultat kan ses på webben, betygsgränser, visning, övriga salar tentan går i m.m.)	

TEKNISKA HÖGSKOLAN I LINKÖPING
Institutionen för datavetenskap
Zebo Peng

Tentamen i kursen

TDTS08 Datorarkitektur

(Examination on TDTS 08 Advanced Computer Architecture)

2010-04-08, kl. 8 - 12

Hjälpmedel:

Engelsk ordbok.

Supporting material:

English dictionary.

Poänggränser:

Maximal poäng är 40.
För godkänt krävs 21 poäng.

Points:

Maximum points: 40.
You need 21 points to pass the exam.

Jourhavande lärare (Teacher on duty):

Zebo Peng, tel. 070 258 2067

Note: You can give the answers in English or Swedish.

1. a) There are several write policies that are used to keep the cache contents and the contents of the main memory consistent. Describe briefly each of these policies and discuss the advantages and disadvantages of each of them, respectively.

   b) Describe the additional problems we have when applying these policies in a multiprocessor system.

   (4p)

2. Consider a 32-bit computer that has an on-chip 16-kbyte four-way set associative cache. Assume that the cache has a line size of eight 32-bit words. Draw a block diagram of this cache, showing its organization and how the different address fields are used to determine a cache hit/miss. (3p)

3. a) What is branch prediction? Why is it important? Illustrate your answer by showing how a certain instruction sequence, of your choice, passes a pipelined unit, in the case of a correct and of an incorrect prediction, respectively.

   b) Dynamic branch prediction: How does it work? What is the basic principle that is used to implement dynamic branch prediction?

   (4p)

4. a) Discuss the main features of a typical RISC architecture.

   b) What are the arguments for the RISC computers?

   c) In your opinion, what are the main features of a RISC computer that have contributed to its high performance?

   (3p)

5. a) What is the most essential characteristics of a superscalar architecture?

   b) What are the differences between the superscalar and superpipelining approaches? Compare these two approaches to each other, and discuss their advantages and disadvantages respectively.

   c) Why is data dependence a bigger issue in a superscalar architecture than in an ordinary computer?

   (3p)

6. a) Draw a picture to show a typical VLIW architecture.

   b) A VLIW architecture is said to support explicit parallel instruction execution. Define the concept of explicit parallelism. What are the advantages of exploiting explicit parallelism?

   (3p)

Note: You can give the answers in English or Swedish.

7. a) Define the concept of loop unrolling. Why is loop unrolling very useful in the context of a VLIW processor?

   b) Use an example to illustrate how loop unrolling is used in a VLIW processor.

   (3p)

8. a) How are the parallel architectures classified according to Flynn?

   b) Describe and draw a diagram for each of the architectures which have been implemented.

   (3p)

9. a) Discuss the concept of sub-word execution. In which context is sub-word execution usually used?

   b) If we use a single ALU to perform several arithmetic operations of sub-words, what is required in order to make sure that the results are always correct?

   (3p)

10. a) What is a snoopy protocol? What is it used for? How does it work?

    b) What is a MESI protocol? Draw the MESI state transition diagram corresponding to the situation when a write-miss occurs. Explain the different scenarios with a write-miss.

    (4p)

11. Describe all the mechanisms the Crusoe processor uses to reduce power consumption.

    (2p)

12. a) Describe the different multithreading approaches and discuss how they are applied in the context of superscalar architectures. What are the advantages and disadvantages of these different approaches, respectively?

    b) Why does multithreading improve system performance even in the case when there is a single scalar processor in your computer?

    (3p)

13. It is very common that the cores in a multi-core processor architecture share the same L2 (Level 2) cache. What are the advantages of sharing the L2 cache that is implemented on the same chip together with the processor cores? (2p)