Försättsblad till skriftlig tentamen vid Linköpings universitet

(fylls i av ansvarig)

Datum för tentamen	101220
Sal	KÅRA, T1, T2, U1
Tid	8-12
Kurskod	TDTS08
Provkod	TEN1
Kursnamn/benämning	Datorarkitektur
Institution	*IDA*
Antal uppgifter som ingår i tentamen	14
Antal sidor på tentamen (inkl. försättsbladet)	5
Jour/Kursansvarig	Zebo Peng
Telefon under skrivtid	0702/582067
Besöker salen ca kl.	0930
Kursadministratör (namn + tfnnr + mailadress)	G. Mellheden 282297
Tillåtna hjälpmedel	English diotonary
Övrigt (exempel när resultat kan ses på webben, betygsgränser, visning, övriga salar tentan går i m.m.)	
Vilken typ av papper ska användas, rutigt eller linjerat	
Antal exemplar i påsen	

TEKNISKA HÖGSKOLAN I LINKÖPING
Institutionen för datavetenskap
Zebo Peng

Tentamen i kursen

TDTS08 Datorarkitektur

(Examination on TDTS 08 Advanced Computer Architecture)

2010-12-20, 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. 28 2067 / 070 258 2067

Exam for TDTS 08, Advanced Computer Architecture, 2010-12-20
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.
   (3p)

2. A computer has a cache, a main memory, and a hard disk used for virtual memory. If a referenced word is in the cache, 10 *ns* are required to access it. If it is in the main memory but not in the cache, 80 *ns* are needed to load it into the cache, and then the reference is started again. If the word is not in the main memory, 12 *ms* are required to fetch the word from the disk to the main memory, followed by 80 *ns* to copy it to the cache, and then the reference is started again. The cache hit ratio is 0.95 and the main memory hit ratio is 0.8. What is the average time in *ns* required to access a referenced word in this memory system?
   (2p)

3. A register file severs as a small and fast buffer for holding the variables that are currently manipulated by the CPU. From this point of view, a register file acts like a cache memory.
   a) Can we then draw the conclusion that we can replace the cache with a large register file? Why?
   b) Discuss the different characteristics of a large register file and a cache, including the way they are accessed and their usual contents, respectively.
   (3p)

4. What does it mean by static branch prediction? Give an example of a static branch prediction method. What are the advantages and disadvantages of static branch prediction, as compared to dynamic branch prediction?
   (2p)

5. A microprogrammed control unit is commonly used in a CISC computer to control the execution of the complex instructions. The execution of a complex instruction consisting of several operations in microcode is much faster than the execution of these operations as several machine instructions. Why? Give a simple example to support your argument.
   (3p)

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

6. a) In a superscalar architecture, the window of execution plays an important role. Why?
b) The window of execution is usually extended over the basic block borders by speculative execution. What is speculative execution?
c) What are the possible disadvantages associated with speculative execution?
(3p)

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) Why is the placement of the "load from memory" operations an important issue?
b) Describe the speculative loading technique. What are the advantages of this technique?
c) Illustrate the speculative loading technique with a simple example.
(3p)

9. What are the two most important characteristics of a program that limit the speedup of its execution on a parallel computer? Discuss the impact of these two characteristics on the achievable speedup on a parallel computer with a number of processors, respectively.
(3p)

10. a) Why does a vector processor provide very high performance for vector computations, even though it doesn't contain several CPUs working in parallel?
b) What is the role of the vector length register in a vector processor? Why is it important to have it?
(3p)

11. 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)

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

12. a) What is a MESI protocol? What are the conditions for a MESI protocol to work efficiently?

    b) Draw the MESI state transition diagrams corresponding to the situation when a read-miss occurs. Explain the different scenarios with a read-miss.

    (3p)

13. 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 only a single scalar processor in your computer?

    (3p)

14. a) One argument for using a graphics processing unit (GPU) is that it is power efficient. Describe all features of a GPU architecture that contribute to the reduction of power consumption.

    b) Can we use GPUs for non-graphics computation? Support your answer with some good arguments.

    (3p)