TEKNISKA HÖGSKOLAN I LINKÖPING
Institutionen för datavetenskap (IDA)
Zebo Peng

Tentamen i kursen
TDTS08 Datorarkitektur

Examination of the course
TDTS08 Advanced Computer Architecture

2018-01-08, 8:00-12:00

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. 013-28 2067

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

1. a) What is the basic idea of associative mapping for cache organization? What are the advantages and disadvantages of the associative mapping organization, as compared with the direct mapping organization?

   b) Why is the fully associative cache organization seldom used in practical computers? Which cache organization is commonly used? Why?

   (3p)

2. a) What does it by mean by a memory of sequential access type?

   b) Give an example of a sequential access memory.

   c) Can a memory of sequential access type be used as the main memory of a computer system? Why?

   (3p)

3. a) Discuss how the bimodal prediction technique work for branch prediction.

   b) Why does the bimodal prediction technique give better performance than the one-bit prediction method? Give a concrete example to support your argument.

   (3p)

4. A computer has an instruction pipeline with four pipeline stages. As a designer, you are asked to consider the possibility of increasing the number of pipeline stages of this computer in order to improve its performance. It turns out that one of the stages can't be divided into two or several shorter stages. Does it make sense to divide the other three stages so that the number of stages increased? Why?

   (3p)

5. The design of RISC architectures is based on certain characteristics of program execution.

   a) What are the characteristics related to procedure calls and returns?

   b) What is the mechanism used in a RISC architecture to make procedure calls and returns efficient? How does this mechanism work?

   (3p)

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

6. a) What are the most essential characteristics of a superscalar architecture?

   b) Why is the window of execution an important mechanism in a superscalar architecture?

   c) What does it means by register renaming? What is the main purpose of using such a technique?

   (3p)

7. 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?

   c) What is the main problem of a traditional VLIW computer? How is this problem addressed by the IA-64 architecture?

   (4p)

8. a) Define the Amdahl's law for parallel processing.

   b) What are the main parameters that determine the efficiency of a parallel computer when running an application?

   (2p)

9. How are general-purpose microprocessors usually extended to handle multimedia applications with little extra hardware cost?

   (2p)

10. a) Describe the different multithreading approaches and discuss how they are applied in the context of a superscalar architecture. 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)

11. a) What is a symmetric multiprocessor (SMP) system?

    b) What are the advantages of an SMP?

    (2p)

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

12. When we have an L1 cache that does not connect to the bus, and would still like to use the MESI cache coherence protocol, we can use the write-though policy for the L1 cache. This forces any modification to an L1 cache line out to the L2 cache, and therefore makes it visible to other L2 caches.

    a) In order for this scheme to work, it is required that the L1 content must be a subset of the L2 content. This means that the associativity of the L2 cache should be equal to or greater than that of the L1 associativity. Why?

    b) Give a concrete example (i.e., a cache access sequence) to support your answer to (a).

    (3p)

13. a) What are the main features of a graphics processing unit (GPU)?

    b) Which of the GPU features have contributed to its high performance? How?

    c) Discuss the concept of divergent execution in a GPU processor. What is the main impact of such divergent execution?

    (3p)

14. a) The dynamic power consumption for CMOS circuits is determined by three parameters, and therefore we have three degrees of freedom in the space for low-power design. What are these three parameters?

    b) For each of these parameters, describe one technique which is most efficient, in your opinion, to achieve low-power design. Explain why you think the described technique is most efficient. (Note: You should present only one technique for each parameter; discussing more than one technique for a give parameter will receive no point).

    (3p)