LIACS - Computerarchitectuur - 2012 - BSc 2de jaar

15 januari 2013, 14:00-17:00

U moet deel 1 beantwoorden, dan één van de 2 andere delen. Boeken en aantekeningen zijn niet toegestaan. Tentamen telt voor 30% van het eindcijfer (dan practica 50%, huiswerk 20%). Elk deel in het tentamen telt voor 15% van het eincijfer. Iedere vraag wordt aangegeven met het % van het eindcijfer tussen haakjes (bvb [x%]).

U mag vragen beantwoorden in het Engels of het Nederlands.

Tip: begin door uw tijd te verdelen tussen vragen. Bepaal dan pas de lengte van een uitleg op basis van beschikbare tijd.

Part 1 (15%)

a. List the names of the 4 stages of the original MIPS pipeline, OR of the 6 stages of MGSim's RISC core. Next to each name, describe briefly what the role of the stage is. [3%]

b. Explain in your own words, if possible using examples, the difference between each of the following pairs of words [3%]:

   i. bandwidth vs. throughput

   ii. Uniform Memory Access (UMA) vs. Non-Uniform Memory Access (NUMA)

   iii. memory-mapped I/O vs. dedicated I/O

   iv. single-issue vs. multiple-issue (for processor pipelines)

c. Which concept in processor architecture is called "Pollack's rule"? What is its relationship with out-of-order execution? [3%]

d. Explain in your own words the factors that have caused increased use of multi-core architectures since 2000. [6%]

Part 2 (15%)

a. Give the name of and explain each of the following pipeline hazards:

   i. Empty (unused) remaining stages after a branch is taken [1%]

   ii. Stalling stages due to contention on a functional unit [1%]

   iii. Invalid result value due to completion reordering in out-of-order execution [1%]

   iv. Invalid result value to issue reordering in out-of-order execution [1%]

   (Assembly format for examples: op rega <- regb, regc; rega is the target register.)

b. The Intel Nehalem (i7) processor described in the diagram below uses ITLB and DTLB components. State what these acronyms stand for and explain what role these components play. [3%].

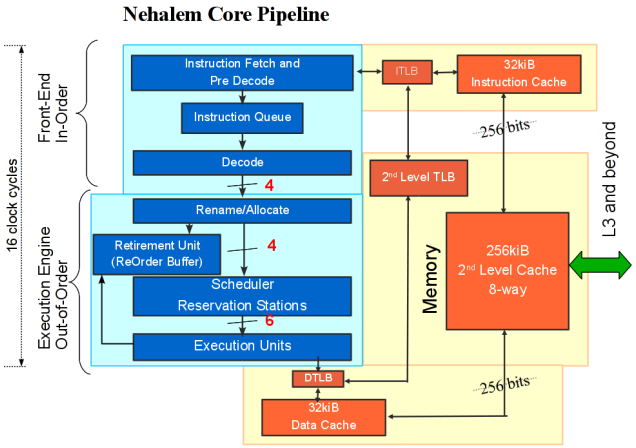

c. The Nehalem core above uses a 256KiB, 8-way set-associative L2 cache with 256-bit data paths to the L1 I- and D-caches. Determine the number of cache lines and the number of sets in the L2 cache. [3%]

d. For a new embedded application, you are in charge of selecting a low-power, in-order core architecture. You are thinking at a choice between a 2-way VLIW core at 100MHz or a single-issue pipeline with branch predictor at 150MHz. However you do not know yet what the software will do. What do you need to ask and how would the answer influence your choice? [5%]

Part 3 (15%)

a. A given D-RAM controller can answer a single load request in x second if the row for the address is already loaded in the row buffer; it needs to wait y seconds between subsequent requests, and it also needs z seconds to change rows. Express the minimum read latency, maximum read latency, and maximum read bandwidth as functions of x, y and z. [4%]

b. Define and separate the following concepts: throughput, bandwidth, latency, access time, performance and efficiency in your own words. Explain the units in which these values are expressed, with examples. [3%]

c. Increasing the L1 cache size increases the access time from 1 to 2ns. To preserve the cycle time at 1ns, another memory stage is added to the pipeline. How does this affect the startup time and half-performance vector length? [3%]

d. You are in charge of selecting a memory configuration for a new video streaming server. You have a choice between two configurations: either 1 DRAM module at 1600MHz or 2 modules at 800MHz. Both configurations use the same DDR3 protocol (1 controller per module), are connected to the same processors, have nearly the same price. How do you choose? Why? [5%]