NTNU – Trondheim
Norwegian University of
Science and Technology

Department of Computer and Information Science

Examination paper for TDT4260 Computer Architecture

Academic contact during examination: Lasse Natvig

Phone: 906 44 580

Examination date: 31. May 2014

Examination time (from-to): 09:00 – 13:00

Permitted examination support material: code D; No written or handwritten examination support materials are permitted. A specified, simple calculator is permitted.

Other information:

The exam accounts for 80% of the final grade, and the provided points show the maximal number of points that can be achieved on each assignment. Read the problem texts thoroughly. You can answer the questions in English or Norwegian.

For all multiple choice questions: Answer by writing the question-ID and one alternative, like this: "X1 b" where X1 is the question ID and b is your answer. You are awarded 3.0 points for a correct answer and 0 points if you do not answer. If your answer is wrong or you give more than one alternative, you will get -1.5 points.

Language: English

Number of pages: 7

Checked by:

……..

20/5/2014 *Magnus Jahre* (sign)

Date Signature

Problem A: Architecture and Systems, multiple choice (Max 15 points)

A1: Computer architectures categorized as MISD according to Flynn's taxonomy are sometimes called *software pipeline*. The reason for this is that

a) the same program code is transferred to several processors (stages) through a linear pipeline.
b) data flows through a sequence of processors where each processor executes its own software (program code) on the data.
c) the instruction stream is executed by one single instruction pipeline (fetch, decode, execute, write back).

A2: A central element in *virtual memory* is page replacement. Which of the following statements is most correct in this context?

a) All pages in memory belonging to a process are automatically written out to disk when the processor does a context switch to another process.
b) A page fault is time consuming since it loads a page from disk and will normally lead to a context switch.
c) A page fault means that the page is lost and a file must be repaired from a backup.

A3: A *distributed shared memory* (DSM) multiprocessor

a) usually has non-uniform memory access.
b) has multiple address spaces on a set of distributed nodes.
c) cannot offer message-passing as a communication mechanism for parallel programs.

A4: The term "*dark silicon*" is used to denote the fact that in the future processors

a) will be made in silicon with a feature size that is so small that impurities will lead to several parts of the chip to be malfunctioning, or dark. These can be tolerated using redundancy.
b) will be made in a new material composed of both silicon and dark (black) carbon nanotubes.
c) will have so many transistors that only a subset of them can be powered at the same time to stay within the power budget of the chip.

A5: In the paper by Borkar et. al. about The Future of Microprocessors the authors describe how *customization* as a general technique can lead to much more energy efficient multicores. Which of the following claims is most correct?

a) Customization means that a customer buying a PC will get a multicore processor that is customized to his or her applications since it contains self adaptive on-chip FPGA logic.
b) Customization means the inclusion of several special purpose accelerators on the multicore, and some parts of the application are executed on those for more energy efficient execution.
c) Customization is used to develop heterogeneous multicores where different cores have different word-size and every function call in an application is mapped to its optimal word-size or a combination of cores forming a pareto optimal space of energy efficient hybrid designs.

Problem B: Cache systems (Max 15 points)

B1: Two main alternatives for *write strategy in a cache system* are "write through" and "write back". One argument for using "write back" is that

a) the system can reduce the energy consumption when you have multiple writes to the same data element.
b) false sharing will never happen since the processor only writes data.
c) both the cache and the memory will always store the right value of the data, so the memory system remains consistent.

B2: Assume a cache of fixed storage capacity, i.e. the total amount of bytes that can be stored in the cache at the same time. We *observe the miss-rate when the cache block size (cache line size) is varied* from small to large block sizes. What alternative will most often describe the typical behaviour of the cache system?

a) Compulsory misses stays constant.
b) Conflict misses is slightly decreasing.
c) Compulsory misses is decreasing while conflict misses increases.

B3: The 2:1 *cache rule of thumb* states that

a) the instruction set cache should be twice the size of the data cache to achieve a balanced hit rate.
b) a direct-mapped cache of size N has about the same miss rate as a two way set-associative cache of size N/2.
c) the miss rate in a cache is roughly doubled when its size is reduced from N to N/2.

B4: The *MESI protocol* for multiprocessor cache coherence adds an additional state E (Exclusive) to the simpler MSI (Modified, Shared, Invalid) protocol. State E indicates when a cache block is resident only in a single cache but is clean. An advantage of adding the state E to the protocol is that

a) a write to a block in state E by the same core need not acquire bus access or generate an invalidate.
b) two cores might communicate very efficiently by using one shared cache block marked for exclusive buffer-communication.
c) A memory block can be marked as exclusive to ensure that only one processor core can access its data.

B5: *A directory based cache coherence protocol* has the advantage that

a) invalid copies are stored only in directories so that broadcast operations for update becomes very infrequent.
b) the file directories can be hierarchical as assumed by most operating systems.
c) updates and invalidation messages can be sent to only those caches that are registered in the directory for a given cache block.

Problem C: Performance and software (Max 15 points)

C1: A general trend in computer architecture is called "*latency lags bandwidth*". This is also true for the communication of data between the cache system and the processor. A consequence of this is that

a) the memory bandwidth is reduced so that more accesses can be handled in parallel.
b) the number of cache levels should be kept as low as possible.
c) memory level parallelism will typically increase, i.e. we will often have more uncompleted memory requests in the system that have not reached the processor.

C2: *Instruction scheduling* may be done by the compiler (static) or the processor (dynamic). Which of the following statements is not true?

a) The compiler has more time to do complex scheduling algorithms than the processor.
b) The compiler must be conservative and assume more conflicts than what actually will occur.
c) Profiling gives the compiler more accurate information about conflicts than what is available to the dynamic scheduler in the processor.

C3 : Embedded computers often run applications with a real-time (RT) performance requirement. The term *soft real time* is used in cases where the RT requirement

a) is achieved by special-purpose software running on standard hardware.
b) is achieved by a slow software solution.
c) must be achieved in most cases, but not always.

C4: Suppose you have a multiprocessor with 100 processors and you want to achieve a *speedup* of 60 or larger. What fraction of the program can be sequential? Use the same assumptions as in *Amdahl's law*.

a) Not more than 0.67 % of the computation can be sequential.
b) Not more than 60 % of the computation can be sequential.
c) Not more than 6.7 % of the computation can be sequential.

C5: Two kinds of multiple-issue processors are *superscalar* processors and *VLIW* (very long instruction word) processors. Which of the following claims is most correct?

a) VLIW processors have so wide instructions that the decode stage is split into two stages, one at the front of the pipeline and one just before the write-back stage of the pipeline.
b) Binary code compatibility is in general lower for VLIW than for superscalar processors since VLIW processors with different number of functional units requires different versions of the code, while superscalar processors can execute the same code on processor variants with different number of functional units.
c) VLIW processors are in general faster than superscalar processors since they to a larger extent can adapt the execution of instructions to the dynamic data dependencies found during execution time.

Problem D: Parallel processing (Max 15 points)

D1: A central technique for GPUs to speed up the processing is the use of "*tile based rendering*". This technique will

a) always lead to better load balancing between the multiple cores in the GPU.
b) increase the cache locality in the multiple cores by splitting up the frame buffer into appropriate pieces.
c) render an image sequentially, tile by tile.

D2: *Loop unrolling* is a technique that might have effects both on cache miss rate and register usage. Which of the following three claims is most correct?

a) Loop unrolling might increase the instruction cache miss rate and the "register pressure" (the need for many register).
b) Loop unrolling will remove all jump instructions used to control the loop and hence reduce the number of instructions so that instruction miss rate is reduced.
c) Loop unrolling will reduce the register usage since the need for storing the loop iteration number is removed.

D3: A *spin lock* is a good choice

a) to avoid deadlock in unidirectional interconnection networks.
b) to protect shared data in cases where the lock is held for a long time.
c) to protect shared data in cases where the lock is held for a very short amount of time.

D4: The main feature of *simultaneous multithreading* is that it exploits both

a) thread-level parallelism and instruction-level parallelism.
b) coarse-grained multithreading and speculation.
c) fine-grained multithreading and multilevel cache hierarchies.

D5: The use of multimedia SIMD extensions like SSE and AVX can give significant performance improvement for some applications. A concept called *partitioned adders* has been used to achieve efficient implementation of such instructions. The advantage of partitioned adders is that

a) the processor efficiently can operate on short vectors of operands of variable length (measured in bits).
b) the number of pipeline-stages for doing an addition is increased.
c) addition and subtraction is performed in different logic, so that more parallelism can be exploited.

Problem E: Research papers and multiprocessors (Max 20 points)

E1: (Max 4 points) Figure 1 gives an overview of the architecture of the *UltraSPARC T1 processor*. Describe briefly what main resources/components are used in the architecture to implement simultaneous multithreading.

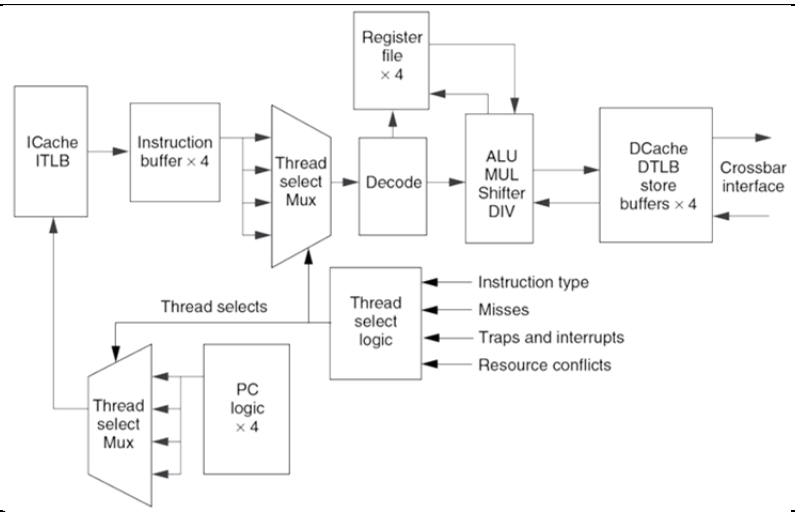

Figure 1.

E2: (Max 4 points) Based on Figure 2 (From the paper «*The Future of Multiprocessors*») explain very briefly what architectural advance gave (a) the best improvement in Floating Point performance, and (b) what gave the best improvement in energy-efficiency for integers.

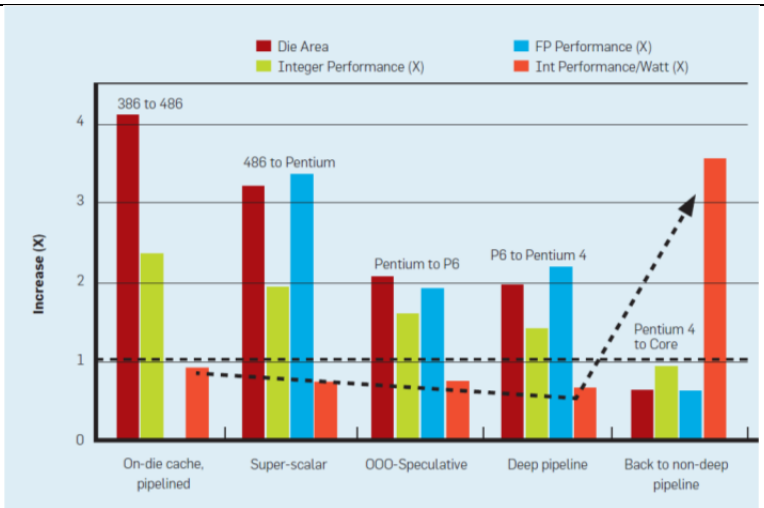

Figure 2.

E3: (Max 4 points) Describe the main architecture of the *matching unit* used in the Manchester Data-flow Machine. (Two keywords are hash-table and overflow unit)

E4: (Max 4 points) A popular framework for batch processing (E.g. creating search indexes from web crawls) in «warehouse-scale computing» (e.g. Google services) is called *MapReduce*. Explain briefly how this operation is divided in two main steps, and how MapReduce can be thought of as a generalization of the single-instruction multiple-data (SIMD) operation.

E5: (Max 4 points) In the paper «Exploring the Design Space of Future CMP's» the authors perform a *design space exploration* where several main architectural parameters are varied assuming a fixed total chip area of 400mm^2. Outline the approach by explaining Figure 3.

Gate length	CBE (Megabytes)	λ^2 area	P_{IN}	P_{OUT}
100nm	7.6	1.60e+11	68	24
70nm	15.5	3.26e+11	139	50
50nm	30.5	6.40e+11	273	99
35nm	61.9	1.30e+12	556	201

Figure 3.

…---oooOOOooo---…