

EECS 370
Exam 2
Winter 2004
March 25, 2004

Name:

University of Michigan username:
(NOT your student ID number!)

Open book, open notes. Calculators are permitted, but no laptops, PDAs, cell phones, etc. This exam has 8 pages, 8 questions, and 100 points. Questions vary in difficulty; it is strongly recommended that you do not spend too much time on any one question.

The rules of the Honor Code of the University of Michigan College of Engineering apply for this exam.

Honor code pledge: I have neither given nor received aid on this examination, nor have I concealed any violations of the Honor Code.

Signature:

(examinations without a signed honor pledge will not be graded)

1. (12 points): Please circle TRUE or FALSE for each statement.

(a) TRUE FALSE

RANDOM cache line replacement is a great strategy to compare against other replacement algorithms, because it is the worst performing replacement strategy that exists.

(b) TRUE FALSE

For a given instruction set, a MULTI-CYCLE processor implementation requires less hardware than a PIPELINED processor implementation.

(c) TRUE FALSE

It is not possible to eliminate all data hazard stalls in the LC2K4 pipeline, even if the design is FULLY BYPASSED (has full forwarding).

(d) TRUE FALSE

Modifying a processor instruction set to reduce average CPI will always produce a faster design, assuming that clock speed is unaffected.

2. (17 points):

Please compute the accuracy (i.e., % of correct predictions on average) of the following branch predictors, assuming that branches are taken 60% of the time, and 40% of the branches have targets in the forward direction (this statistic holds true for both taken and not taken branches):

(a) Predict direction of branch randomly.

(b) Predict all branches TAKEN.

(c) Predict all branches NOT TAKEN.

(d) Predict TAKEN if offset negative, NOT TAKEN otherwise.

(e) Why is it more expensive (e.g., in performance or hardware) to build a **predict TAKEN** branch predictor, compared to a **predict NOT TAKEN** branch predictor?

3. (12 points):

Consider a normal 5-stage LC-2K4 pipeline as discussed in class. Branches are resolved in the MEM stage and branches are predicted as not taken. 1/3 of load instructions incur a one cycle stall due to data hazards that can not be resolved via forwarding. Assume the instruction cache has a perfect hit rate, but the data cache has a 10% miss rate for both reads and writes. For a memory read, a data cache miss causes a 20 cycle stall. For a memory write, a data cache miss does not cause a stall. Analysis has shown the following dynamic instruction breakdown for program X: 15% not taken branches, 10% taken branches, 10% loads, 15% stores, 50% R-type instructions (add/nand).

(a) What is the CPI of this machine when running program X?

(b) Now assume the pipeline was stretched out to 10 stages (numbered 1 to 10, with instruction fetch beginning in stage 1, and instructions finally completing in stage 10). Branches are resolved in stage 7. The clock frequency is increased to twice that of the original processor with the 5 stage pipeline. Data cache misses cause a 40 cycle stall. 1/3 of load instructions still incur a one cycle stall due to data hazards that can not be resolved via forwarding. What is the CPI for this new pipeline design?

(c) Which of these machines performs better, the original 5 stage pipeline or the new 10 stage pipeline? How many times faster is it than the other one?

4. (12 points):

Consider the LC2 processor with the 5 stage pipeline as described in the lecture slides. Assume it runs programs with the following instruction mix: 40% R-type (add/nand), 20% LW, 20% SW, 20% BEQ. All load and store operations access memory in one cycle. Data hazards are resolved by forwarding whenever possible. 50% of LW instructions are followed by a dependent instruction. 50% of BEQ instructions are taken. The processor clock speed is 100 MHz.

(a) How many MIPS (millions of instructions per second) does this processor run at with this instruction mix?

(b) The original LC2 processor is modified as follows: The branch hardware has been streamlined so that branches are resolved in the EX stage rather than the MEM stage. Nothing else is changed in the pipeline. The processor clock speed must be decreased to 80 MHz to accommodate this change. How many MIPS does this processor run at?

(c) The original LC2 processor is modified to shorten the pipeline from five stages to four as follows: Data memory access is moved into the EX stage. A special dedicated adder (separate from the ALU) is put into this stage to calculate the **base register plus offset** value and send it as an address to the data memory. This effectively merges EX and MEM into one stage and removes the load-use stall, but with a slower clock as a result. The old MEM stage goes away entirely; the WB stage immediately follows the EX stage. Branches are resolved in the WB stage. Data hazards are resolved by forwarding whenever possible. The processor clock speed must be decreased to 90 MHz to accommodate this change. How many MIPS does this processor run at?

5. (14 points):

The alpha version of the LC2K4 processor did not support using immediate values in the main ALU as operands. Hence, the LC2K4-Alpha had an extra stage in the pipeline (with its own dedicated adder separate from the main ALU). This stage is called **AC** (Address Calculation); the **base register + immediate offset** address calculation is done in this stage. Given that a program has the following properties:

- 20% of all instructions are loads and stores that use an address with a base register whose value is set in the immediately-preceding instruction, which is an **r-type** instruction
- 10% of all instructions are loads and stores that use an address with a base register whose value is set in the immediately-preceding instruction, which is an **lw** instruction
- 20% of all instructions are **nand** and **add** instructions that immediately follow a **lw** and depend on the result of that load.
- 20% of all instructions are **nand** and **add** instructions that follow two instructions after a **lw** and depend on the result of that load (e.g. **lw**, something, **add**)

These four groups listed above are totally disjoint. More specifically, there are no **lw** instructions that are followed by two dependent instructions in a row; something like the following never occurs: **lw 0 3 one, add 1 3 4, add 2 3 5**. There are no instruction combinations other than those specifically enumerated above that could cause any stalling. There are no **BEQ** instructions; the program runs once through in consecutive address order from beginning to end. There is full forwarding to the extent physically (temporally) possible. Calculate the CPI of a program with the preceding statistics on the following two pipeline implementations of the LC2K4.

(a) IF ID AC EX MEM WB

(b) IF ID EX AC MEM WB

6. (9 points):

A certain new, high-performance computer has a cache with the following characteristics:

- 48 bit addresses
- Byte addressable memory
- 8 megabyte cache size (1 megabyte = 2^{20} bytes = 1,048,576 bytes)
- 4 kilobyte cache block size (1 kilobyte = 2^{10} bytes = 1,024 bytes)
- 16-way set associative cache
- Address bit numbering: most significant bit is bit 47, least significant bit is bit 0.

(a) What range of address bit numbers are used for each of the following?

Set index:

Block offset:

Tag:

(b) For $0x3f2c84de7b29$, show the set index, block offset, and tag. You may state your answers in either binary or hexadecimal.

Set index:

Block offset:

Tag:

7. (9 points):

You are designing a data cache with a single objective: minimize the total number of cache blocks transferred between the cache and main memory. Each copy of a block from the main memory to the cache or from the cache to the main memory counts as one transfer (regardless of the block size). For situations when a CPU write operation goes directly through to main memory, each such write operation counts as one transfer. You are limited to a specific total cache data size, and this size is significantly smaller than the size of the program that you are going to run on this system, but you have complete freedom on all other cache parameters. You know that the program you are going to run has a high degree of both spatial and temporal locality, for both reads and writes. For each of the three design decisions below, circle the BEST answer:

5.1. Write policies:

- (a) You should use a write back, allocate on write policy
- (b) You should use a write back, no allocate on write policy
- (c) You should use a write through, allocate on write policy
- (d) You should use a write through, no allocate on write policy
- (e) The write policy is not likely to have much effect on the total number of blocks transferred.

5.2. Block size:

- (a) You should use a 1 byte block size
- (b) You should make the block size as large as possible.
- (c) For your given program, there is likely to be some intermediate block size that is optimal.
- (d) The block size is not likely to have much effect on the total number of blocks transferred.

5.3. Associativity:

- (a) You should use a direct mapped cache
- (b) You should use a 2-way set associative cache
- (c) You should use a 4-way set associative cache
- (d) You should use a fully associative cache
- (e) There is likely to be some intermediate degree of associativity that is optimal.
- (f) The associativity is not likely to have much effect on the total number of blocks transferred.

8. (15 points):

For each of the following cache configurations, how many cache misses will result for the following series of word addressed references (given in decimal)? The cache is initially empty and **least recently used** (LRU) is used to evict blocks from the cache. All caches are fully associative. Show your work for partial credit.

1, 4, 8, 5, 1, 3, 2, 10, 9, 3, 4, 7, 6, 1, 8, 0, 6, 5, 7, 8, 2, 10, 1

(a) Cache contains eight one-word blocks.

(b) Cache contains four two-word blocks.

(c) Cache contains two four-word blocks.