Dr Kanu S Pandey 8/12/2022

205

END SEMESTER EXAMINATION

EC-455 Computer Architecture and Organization

Maximum Marks 50

Time allowed 180 minutes maximum

There are 5 questions in all. One additional question is for extra credit. All the questions are compulsory and carry 10 marks each. Questions vary considerably in difficulty and in how long each will take but they are all worth the same marks. This is an open-book exam. You may use your books and notes but not consult with anyone.

Que 1:

Compare the performance of the following 4 different RISC-V implementations:

1. a single cycle machine with a 10 nsec clock
2. a multi-cycle machine with a 2 nsec clock
3. a single issue, five stage pipelined machine with a 2 nsec clock
4. a 2-way in-order-issue in-order-complete, six stage pipelined machine (the extra pipeline stage is for result commit) with a 3 nsec clock

The performance needs to be compared for the following loop as workload. You may ignore the pipeline fill and pipeline drain cycles for the pipelined machines.

```
loop:   lw      x2,50(x3)
        addi    x1,x1,1
        addi    x3,x3,4
        add     x4,x2,x4
        bne     x1,x5,loop
```

Que 2:

Consider three processors with different cache (both i-cache as well as d-cache) configurations:

- Cache 1: Direct mapped with one word cache line size
- Cache 2: Direct mapped with four word cache line size
- Cache 3: Two way set associative with four word cache line size

The following miss rate measurements have been made:

- Cache 1: Instruction miss rate is 4%; data miss rate is 6%
- Cache 2: Instruction miss rate is 2%; data miss rate is 4%
- Cache 3: Instruction miss rate is 2%; data miss rate is 3%

For these processors, one half of the instructions contain a data reference. Assume that the cache miss penalty is 6 + cache line size in words. The CPI for this workload was measured on a processor with cache 1 and was found to be 2.0. Determine which processor spends the most cycles on cache misses.

Que 3:

Consider the loop code shown in Que 1 running on the five stage RISC-V pipeline discussed in class.

If x1 is initially 0 and x5 is 3 (i.e., the loop is executed three times), how many cycles are incurred if the following branch prediction schemes are used for both a branch penalty of *three* cycles and a branch penalty of *one* cycle. Assume for each, if the prediction is successful that the branch penalty is reduced to zero.

	Three cycle penalty	One cycle penalty
Static, not taken		
Static, taken		
Dynamic, 1-bit initial state is not taken		
Dynamic, 2-bit initial state is not taken and predict not taken		

Que 4:

Below is a list of 32 bit memory address references given as word addresses for two programs (a) and (b):

Program	Sequence of Words referenced
a	1, 134, 212, 1, 135, 213, 162, 161, 2, 44, 41, 221
b	6, 214, 175, 214, 6, 84, 65, 174, 64, 105, 85, 215

You are asked to optimize a cache design for the given references of both the programs. There are three direct mapped cache designs possible, all with a total of eight words of data: C1 has one word cache lines, C2 has two word cache lines and C3 has four word cache lines. In terms of miss rate, which cache design is the best? If the stall time is 25 cycles, and C1 has an access time of 2 cycles, C2 takes 3 cycles and C3 takes 5 cycles, which is the best cache design?

Que 5:

Suppose a computer's address size is k bits (using byte addressing), the cache size is S bytes, the cache line size is B bytes and the cache is A-way set associative. Assume that B is a power of two, so $B = 2^b$. Figure out what the following quantities are in terms of S, B, A, b and k:

a. The number of sets in the cache
b. The number of index bits in the address
c. The number of bits needed to implement the cache

Extra Credit:

Assume that, of all instructions executed in a processor, the following fraction of these instructions has a particular type of RAW data dependence. The type of RAW dependence is identified by the stage that produces the result (EX or MEM) and the instruction that consumes the result (1st instruction that follows the one that produces the result, 2nd instruction that follows or both). Also assume that the CPI of the processor is 1 if there are no data hazards.

	EX to 1st only	EX to 1st and 2nd	EX to 2nd only	MEM to 1st only
a	10%	10%	5%	25%
b	15%	5%	10%	20%

a. If we use no forwarding, what fraction of cycles are we stalling due to data hazards in both the cases a and b respectively?
b. If we use full forwarding, what fraction of cycles are we stalling due to data hazards in both the cases a and b respectively?

GOOD LUCK!