國立中山大學資訊工程學系

101 學年度第 1 學期博士班資格考試　計算機結構

1. Consider two different machines, with two different instruction sets, both of which have a clock rate of 200 MHz. The following measurements are recorded on the two machines running a given set of benchmark programs:

Machine	Instruction Type	Instruction Count (millions)	Cycles per Instruction
A	Arithmetic and logic	8	1
	Load and store	4	3
	Branch	2	4
	Others	4	3
B	Arithmetic and logic	10	1
	Load and store	8	2
	Branch	2	4
	Others	4	3

(1) (12%) Determine the effective CPI, MIPS rate, and execution time for each machine.

(2) (6%) Comment on the results.

2. (1) (6%) Consider a single-level cache with an access time of 2.5 ns, a line size of 64 bytes, and a hit ratio of $H = 0.95$. Main memory uses a block transfer capability that has a firstword (4 bytes) access time of 50 ns and an access time of 5 ns for each word thereafter. What is the access time when there is a cache miss? Assume that the cache waits until the line has been fetched from main memory and then re-executes for a hit.

(2) (6%) Continued with the previous problem, suppose that increasing the line size to 128 bytes increases the H to 0.97. Does this reduce the average memory access time?

(3) (8%) A computer has a cache, main memory, and a disk used for virtual memory. If a referenced word is in the cache, 20 ns are required to access it. If it is in main memory but not in the cache, 60 ns are needed to load it into the cache, and then the reference is started again. If the word is not in main memory, 12 ms are required to fetch the word from disk, followed by 60 ns to copy it to the cache, and then the reference is started again. The cache hit ratio is 0.9 and the main memory hit ratio is 0.6. What is the average time in nanoseconds required to access a referenced word on this system?

3. (1) (4%) In virtually all systems that include DMA modules, DMA access to main memory is given higher priority than CPU access to main memory. Why?

(2) (8%) A DMA module is transferring characters to memory using cycle stealing, from a device transmitting at 14400 bps. The processor is fetching instructions at the rate of 1 million instructions per second (1 MIPS). By how much will the processor be slowed down due to the DMA activity?

4. A nonpipelined processor has a clock rate of 3 GHz and an average CPI (cycles per instruction) of 4. An upgrade to the processor introduces a six-stage pipeline. However, due to internal pipeline delays, such as latch delay, the clock rate of the new processor has to be reduced to 2.5 GHz.

(1) (6%) What is the speedup achieved for a typical program?

(2) (6%) What is the MIPS rate for each processor?

5. Consider the following sequence of instructions, where the syntax consists of an opcode followed by the destination register followed by one or two source registers:

```
0     ADD     R3,  R1,  R2
1     LOAD    R6,  [R3]
2     AND     R7,  R5,  3
3     ADD     R1,  R6,  R0
4     SRL     R7,  R0,  8
5     OR      R2,  R4,  R7
6     SUB     R5,  R3,  R4
7     ADD     R0,  R1,  R10
8     LOAD    R6,  [R5]
9     SUB     R2,  R1,  R6
10    AND     R3,  R7,  15
```

Assume the use of a four-stage pipeline: fetch, decode/issue, execute, write back. Assume that all pipeline stages take one clock cycle except for the execute stage. For simple integer arithmetic and logical instructions, the execute stage takes one cycle, but for a LOAD from memory, five cycles are consumed in the execute stage.

If we have a simple scalar pipeline but allow out-of-order execution, we can construct the following table for the execution of the first seven instructions:

Instruction	Fetch	Decode	Execute	Write Back
0	0	1	2	3
1	1	2	4	9
2	2	3	5	6
3	3	4	10	11
4	4	5	6	7
5	5	6	8	10
6	6	7	9	12

The entries under the four pipeline stages indicate the clock cycle at which each instruction begins each phase. In this program, the second ADD instruction (instruction 3) depends on the LOAD instruction (instruction 1) for one of its operands, R6. Because the LOAD instruction takes five clock cycles, and the issue logic encounters the dependent ADD instruction after two clocks, the issue logic must delay the ADD instruction for three clock cycles. With an out-of-order capability, the processor can stall instruction 3 at clock cycle 4, and then move on to issue the following three independent instructions, which enter execution at clocks 6, 8, and 9. The LOAD finishes execution at clock 9, and so the dependent ADD can be launched into execution on clock 10.

(1) (6%) Complete the preceding table.

(2) (6%) Redo the table assuming no out-of-order capability. What is the savings using the capability?

(3) (6%) Redo the table assuming a superscalar implementation that can handle two instructions at a time at each stage.

6. Please answer the following questions.

(1) (4%) For a system with two levels of cache, define T_{c1} = first-level cache access time; T_{c2} = second-level cache access time; T_m = memory access time; H_1 = first-level cache hit ratio; H_2 = combined first/second level cache hit ratio. Provide an equation for T_a for a read operation.

(2) (4%) What is the purpose of a translation lookaside buffer?

(3) (4%) Suppose that two registers contain the following hexadecimal values: AB0890C2, 4598EE50. What is the result of adding them using MMX instructions for packed byte. Assume saturation arithmetic is used.

(4) (8%) Briefly explain the following cache coherence protocols: directory protocols and snoopy protocols.