國立成功大學

115學年度碩士班招生考試試題

編　號：123

系　所：電機工程學系

科　目：計算機組織

日　期：0203

節　次：第 2 節

注　意：1.不可使用計算機
2.請於答案卷(卡)作答，於試題上作答，不予計分。

1. (10pts, no partial point, no penalty) Computer design often relies on "Great Ideas in Computer Architecture." Which of the following statements is/are TRUE?
   (a) "Make the common case fast" suggests optimizing the instructions or events that occur most frequently, rather than the rare cases.
   (b) Amdahl's Law states that the performance improvement to be gained from using some faster mode of execution is limited by the fraction of the time the faster mode can be used.
   (c) Pipelining improves performance by decreasing the execution time of an individual instruction (latency) rather than increasing the throughput of the system.
   (d) Spatial locality refers to the tendency of a program to access data items that are close in storage locations to the items recently accessed.
   (e) Predicting branches is an example of "Guessing" (Speculation) to overcome control hazards and improve parallelism.

2. (10pts, no partial point, no penalty)In a shared-memory multiprocessor system using a Snooping Cache Coherence Protocol (like MSI or MESI), which of the following statements is/are TRUE?
   (a) In a "Write-Invalidate" protocol, a writing processor commands all other processors to update their copies of the data to the new value immediately.
   (b) False sharing occurs when two unrelated variables reside in the same cache block, causing unnecessary invalidations between cores even though they are not sharing data.
   (c) The "Snooping" mechanism relies on a shared communication medium (like a bus) where all cache controllers monitor (snoop) for requests to memory addresses they hold.
   (d) In the MESI protocol, if a block is in the Shared (S) state, the processor can write to it immediately without notifying other caches.
   (e) Directory-based protocols are generally more scalable than snooping protocols because they avoid broadcasting every memory request to all processors.

3. (10pts, no partial point, no penalty)Regarding RISC-V Instruction Set Architecture and Synchronization, which of the following statements is/are TRUE?
   (a) RISC-V uses special "Load Reserved" (lr.d) and "Store Conditional" (sc.d) instructions to implement atomic operations for synchronization (like locks).
   (b) A "Store Conditional" instruction will always succeed and update memory, regardless of whether the value at the address has changed since the "Load Reserved".
   (c) In the RISC-V ISA, there are no condition codes (flags); instead, conditional branches compare two registers directly (e.g., beq x1, x2, Label).
   (d) The jal (Jump and Link) instruction saves the return address (PC + 4) into a destination register, usually x1 (ra).
   (e) RISC-V is a "Load-Store" architecture, meaning arithmetic operations (like add, sub) can operate directly on operands located in main memory.

4. (10pts, no partial point, no penalty) Modern processors utilize Pipelining and Hazard Detection. Which of the following statements is/are TRUE?
   (a) A "Structural Hazard" occurs when the hardware cannot support the combination of instructions that we want to execute in the same clock cycle (e.g., single memory port for both instruction fetch and data access).
   (b) "Data Hazards" occur when an instruction depends on the result of a previous instruction that has not yet completed.
   (c) Forwarding (Bypassing) can resolve data hazards, including the "Load-Use" hazard.
   (d) A "Control Hazard" (or branch hazard) arises from the need to make a decision based on the results of one instruction while others are executing.
   (e) Using a "Branch Prediction Buffer" (or Branch History Table) guarantees that the CPU will not stall due to a branch instruction.

5. (10pts, no partial point, no penalty) A processor runs at a clock frequency of 3 GHz. A benchmark program executes 2 billion instructions 2×10^9 with the following mix:
   40% ALU instructions, 3 cycles each.
   20% Load instructions, 6 cycles each.
   10% Store instructions, 5 cycles each.
   30% Branch instructions, 2 cycles each.
   Which of the following correctly gives the Effective CPI and Total CPU Time?
   (a) Effective CPI = 3.5; CPU Time = 2.33 seconds
   (b) Effective CPI = 4.0; CPU Time = 2.67 seconds
   (c) Effective CPI = 3.5; CPU Time = 2.00 seconds
   (d) Effective CPI = 3.0; CPU Time = 2.00 seconds
   (e) Effective CPI = 3.8; CPU Time = 2.53 seconds

6. (10pts, no partial point, no penalty) Regarding Memory Hierarchy and Caches (L1, L2, Main Memory), which of the following statements is/are FALSE?
   (a) "Temporal Locality" implies that if a data location is referenced, it will tend to be referenced again soon.
   (b) A "Write-Through" cache updates both the cache and the main memory (or next level cache) simultaneously on every write.
   (c) In a Set-Associative cache, a memory block can be placed in any cache line within a specific set.
   (d) Increasing the cache associativity (e.g., from 2-way to 4-way) generally increases the miss rate but reduces the hit time.
   (e) The "Miss Penalty" is the time required to fetch a block from a lower level of the memory hierarchy (e.g., main memory) to the upper level (e.g., cache).

7. (10pts, no partial point, no penalty) Regarding Data-Level Parallelism, SIMD, and GPUs, which of the following statements is/are FALSE?

(a) SIMD (Single Instruction, Multiple Data) allows one instruction to operate on multiple data elements simultaneously, which is efficient for matrix operations.

(b) GPUs (Graphics Processing Units) typically rely on massive multithreading to hide memory latency rather than using large caches to reduce latency.

(c) In Vector Architectures, a single instruction can launch a set of operations on large vectors of data stored in vector registers.

(d) SIMD extensions in x86 (like AVX) or RISC-V (Vector Extension) are strictly handled by the operating system and are invisible to the compiler.

(e) GPUs use a "Single Instruction, Multiple Thread" (SIMT) model where many threads execute the same instruction stream on different data.

8. (10pts, no partial point) Assume a standard 5-stage RISC-V pipeline (IF → ID → EX → MEM → WB) with:
Full Forwarding logic implemented (to solve RAW hazards where possible).
No Branch Prediction: Branches are resolved in the EX stage.
The pipeline always predicts "Not Taken" (fetches the next sequential instruction). If the branch is actually taken, the pipeline must flush instructions.
Consider the following RISC-V assembly snippet:

```
ld    x10, 0(x2)
add   x11, x10, x3
sub   x12, x11, x4
beq   x12, x0, Label
addi x5, x5, 1
...
Label:
sub   x5, x5, 1
```

Assume the branch is TAKEN. How many stall cycles (bubbles) and flush cycles are introduced in total by this sequence?

9. (10pts, no partial point) A system uses Virtual Memory with the following parameters:
Virtual Address: 40 bits
Page Size: 8 KB
Page Table Entry (PTE) size: 4 bytes
TLB: Fully Associative, 64 entries.
System Cache: Physically indexed, Physically tagged.
Answer the following:

(a) (5pts)How many bits are used for the Virtual Page Number (VPN)?

(b) (5pts)If the system uses a flat (single-level) page table, what is the total size of the Page Table in bytes? (Assume one valid entry per possible VPN).

10. (10pts, no partial point) A processor has a 64-bit physical address space and a 256 KB L2 cache. The cache is 8-way set-associative. The block (line) size is 64 bytes.
    Answer the following questions regarding the address bit breakdown:
    (a) (2pts)Number of Offset bits = _____
    (b) (2pts)Number of Sets = _____
    (c) (3pts)Number of Index bits = _____
    (d) (3pts)Number of Tag bits = _____