國立成功大學

114學年度碩士班招生考試試題

編　號：131

系　所：電機工程學系

科　目：計算機組織

日　期：0210

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) **Abstractions** is an important technique for complex design. Which of the following statements is/are **TRUE**?
   (a) Abstraction hides the complexities of hardware from software, allowing programmers to focus on high-level algorithms without needing to understand the underlying hardware specifics.
   (b) Abstraction refers to the process of directly controlling hardware components such as memory, registers, and buses in a low-level programming environment.
   (c) Abstraction in computer architecture is about optimizing hardware performance by reducing the number of components in the processor.
   (d) Abstraction eliminates the need for an instruction set architecture by directly controlling the execution of machine-level instructions.
   (e) Abstraction provides a clear separation between the hardware and software layers by defining an instruction set architecture (ISA) that allows software to execute on any hardware.

2. (10pts, no partial point, no penalty) In a shared-memory multiprocessor system that implements the **MESI cache-coherence protocol**, which of the following statements is/are **TRUE**?
   (a) The Modified (M) state indicates that the cache line is dirty and present only in the current cache.
   (b) The Shared (S) state indicates that the cache line may be present in multiple caches, but none of those copies is dirty.
   (c) The Exclusive (E) state indicates that the cache line is present in only one cache and is clean (not dirty).
   (d) The Invalid (I) state indicates that the cache line is not valid and cannot be used.
   (e) A cache line transitions directly from Modified (M) to Exclusive (E) when another cache attempts to read the same data.

3. (10pts, no partial point, no penalty) In a multiprocessor system, the memory consistency model defines the order in which operations (reads and writes) appear to execute from the perspective of different processors. Consider the **Sequential Consistency (SC) model**. Which of the following statements about Sequential Consistency is/are **TRUE**?
   (a) Under SC, all memory accesses appear to execute in a single global order that respects each processor's program order.
   (b) Under SC, a processor is free to reorder its loads and stores arbitrarily to improve performance.
   (c) SC is stronger (i.e., imposes more ordering constraints) than weaker consistency models, such as Release Consistency (which allowing certain memory operations to be reordered or overlapped).
   (d) Implementing SC often requires hardware or compiler mechanisms (such as memory fences) to ensure the correct ordering of memory operations.
   (e) Under SC, a read on one processor may return a value written by another processor before that write is visible to all other processors.

4. (10pts, no partial point, no penalty) Modern **superscalar processors** use out-of-order execution, branch prediction, and speculative execution to improve performance. Which of the following statements is/are **TRUE**?
   (a) Branch prediction reduces pipeline stalls by anticipating the outcome of a branch.
   (b) Out-of-order execution allows the processor to execute independent instructions while waiting for data dependencies to resolve.
   (c) Out-of-order execution can reorder instructions with true data dependencies (RAW) without any extra hardware or compiler support.
   (d) Precise exceptions require instructions to complete in-order at the commit stage, even if they are executed out-of-order.
   (e) When a branch is mispredicted, only the instructions that have already committed must be flushed from the pipeline.

5. (10pts, no partial point, no penalty) A processor runs at a clock frequency of 2 GHz. A particular program executes 1 billion instructions (10^9) with the following instruction mix and clock-cycle counts per instruction type:
   - 20% R-type instructions, each requiring 4 cycles
   - 50% I-type instructions, each requiring 5 cycles
   - 30% Memory (load/store) instructions, each requiring 8 cycles

   Which of the following statements **correctly** gives both the **effective CPI** and the **total CPU execution time** for this program?
   (a) Effective CPI = 5.2 and CPU time = 2.60 seconds.
   (b) Effective CPI = 5.7 and CPU time = 2.85 seconds
   (c) Effective CPI = 5.7 and CPU time = 1.70 seconds
   (d) Effective CPI = 6.2 and CPU time = 3.10 seconds
   (e) Effective CPI = 5.0 and CPU time = 2.50 seconds

6. (10pts, no partial point, no penalty) In a modern computer system with **multiple levels of caches** (L1, L2, and possibly L3), which of the following statements about cache design and performance is/are **FALSE**?
   (a) Increasing the cache block size (line size) beyond a certain point can reduce the miss rate, but it may also increase the miss penalty due to more data transfer on each miss.
   (b) Write-back caches always require a write to main memory on every store instruction, ensuring data consistency at the expense of performance.
   (c) The inclusion property states that everything in the L1 cache must also be present in the L2 cache, ensuring consistent copies of data across levels.
   (d) Associative caches (with more ways) generally reduce the conflict miss rate but can have higher access latency and more complex hardware.
   (e) A fully associative cache of the same size as a direct-mapped cache will always have a lower miss rate, but its hardware complexity and access time might be higher.

7. (10pts, no partial point, no penalty) A modern processor may implement multi-core designs (e.g., 4 cores), and each core may support hardware-based multithreading (e.g., 2 hardware threads per core). Which of the following statements about **multicore** and **hardware multithreading** is/are **FALSE**?

   (a) Hardware-based multithreading (like Simultaneous Multithreading, SMT) allows each core to issue instructions from multiple threads in the same clock cycle, when resources (functional units) are available.

   (b) When the CPU supports 2 hardware threads per core on a 4-core processor, the operating system sees a total of 4 logical CPUs in the system.

   (c) Multicore processors place multiple independent processing units (cores) on a single chip, each capable of running its own thread of execution.

   (d) A multicore design can achieve parallel execution even if hardware multithreading is disabled or not implemented, because each core can run an independent process or thread.

   (e) Hardware-based multithreading helps hide pipeline stalls (e.g., due to memory latency) by quickly switching to another hardware thread, utilizing idle execution resources.

8. (10pts, no partial point) Assume a simple 5-stage RISC-V pipeline (IF → ID → EX → MEM → WB) with:

   - Forwarding from MEM to EX (to reduce some data hazards).
   - No branch prediction, and branches are resolved in EX (meaning the pipeline does not know if the branch is taken or not until the EX stage of the branch instruction).

   Consider the following snippet of RISC-V assembly:

```
   lw x5, 0(x10)
   add x6, x5, x7
   sub x8, x6, x5
   beq x8, x0, L1
   or x9, x8, x0
L1:
   add x9, x0, x0
```

   How many **stall cycles (pipeline bubbles)** are introduced in total by these instructions due to hazards, under the given assumptions?

9. (10pts, no partial point) A processor has:
   - A 32-bit virtual address space with a two-level page table.
   - A page size of 4 KB (4096 bytes).
   - A single TLB (Translation Lookaside Buffer) with:
     - TLB hit rate = 98% (so TLB miss rate = 2%).
     - Access time for a TLB hit is already included (hidden) in the base pipeline.
     - Miss penalty: On a TLB miss, the hardware must walk the two-level page table in memory. Each level of the page table is stored in main memory, which takes 50 cycles per memory access.

   Other system parameters:
   - The base CPI (assuming no TLB misses) is 1.0.
   - 40% of the instructions perform a data access (load/store) in addition to the instruction fetch.
   - Each instruction fetch needs a TLB translation, and each load/store also needs a TLB translation.
   - Ignore any other stall sources (cache misses, etc.). Focus only on the extra stalls from TLB misses.

   What is the **overall (effective) CPI** once TLB misses and their penalties are included?

   Overall CPU = ____.

10. (10pts) A system uses 32-bit addresses and has an 8 KB (8192 bytes) 2-way set-associative cache. The block (line) size is 16 bytes. Answer the following question
    - number of sets = _____ .
    - Number of offset bits = ____.
    - The number of index bits = ____.
    - The number of tag bits = ____ .