國立成功大學

114學年度碩士班招生考試試題

編　號：126

系　所：電機工程學系

科　目：計算機組織與作業系統

日　期：0210

節　次：第 1 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) Which of the following statements about **MIPS** (Millions of Instructions Per Second) is/are **TRUE**?
   (a) MIPS is a completely reliable performance metric that remains consistent across different instruction sets.
   (b) MIPS does not depend on the mix of instructions a program executes.
   (c) MIPS can vary among different programs running on the same hardware.
   (d) MIPS is always higher when a program is compiled with more optimizations.
   (e) Increasing the clock frequency always guarantees an increase in MIPS.

2. (10pts, no partial point, no penalty) Consider a **5-stage pipelined CPU** (IF, ID, EX, MEM, WB) implementing **full forwarding** (from MEM→EX and WB→EX) but **no branch prediction**. Which of the following statements about pipeline hazards and stalls is/are **correct**?
   (a) With full MEM→EX and WB→EX forwarding, a load-use hazard (when a load is immediately followed by an instruction that uses the loaded register) still requires at least one stall cycle.
   (b) An R-type instruction (e.g., add x5, x1, x2) immediately followed by another R-type instruction dependent on its result (e.g., sub x6, x5, x3) always requires a stall, even with forwarding.
   (c) A branch resolved in the EX stage (with no prediction) typically incurs a control hazard penalty (flush) of 2 pipeline stages in a classic 5-stage pipeline.
   (d) Storing a value (sw x5, 0(x6)) does not introduce a stall if the next instruction reads from x5, because that read can receive x5 from the register file directly (the store doesn't modify x5).
   (e) If two instructions both write to the *same* register back-to-back, a data hazard occurs that forwarding cannot fix, requiring an extra stall.

3. (10pts, no partial point, no penalty) In a **32-bit 2's complement integer** representation, which of the following statements is/are **FALSE**?
   (a) The most negative value is represented by the bit pattern 0x80000000.
   (b) The maximum positive value is 0x7FFFFFFF.
   (c) Adding 1 to 0x7FFFFFFF (the largest positive integer) results in 0x80000000, which is the most negative integer.
   (d) The range of representable 32-bit 2's complement integers is $[-2^{31}, 2^{31} - 1]$.
   (e) When adding two 32-bit 2's complement numbers, signed overflow can be detected simply by checking if the carry out of the MSB is 1.

4. (10pts, no partial point, no penalty) Which of the following is/are **NOT** a **memory management scheme** typically used in operating systems?
   (a) Paging
   (b) Segmentation
   (c) Demand Paging
   (d) Virtual Memory
   (e) Pipelining

5. (10pts, no partial point, no penalty) Which of the following statements about **semaphores** is/are **TRUE**?
   (a) A semaphore is an integer variable that is accessed only through two standard atomic operations, commonly called wait() (or P()) and signal() (or V()).
   (b) Semaphores can be used to implement both mutual-exclusion (mutex) locks and signaling/counting mechanisms (e.g., for producer-consumer buffers).
   (c) Binary semaphores (mutex semaphores) and counting semaphores are functionally identical in all concurrency scenarios.
   (d) If multiple processes (or threads) are waiting on a semaphore's queue, a single signal() operation typically unblocks exactly one waiting entity.
   (e) If used incorrectly (e.g., forgetting to call signal() after wait()), semaphores can lead to deadlocks or process starvation.

6. (10pts, no partial point, no penalty) A file system can use **contiguous allocation**, **linked allocation**, or **indexed allocation** to store files. Which of the following statements about these allocation methods is/are **TRUE**?
   (a) Contiguous allocation greatly simplifies file expansion at run time, since extending a file only requires adding a new block to any free location on the disk.
   (b) Linked allocation maintains a table in memory with direct pointers to all the blocks of the file, allowing random access to any block in constant time.
   (c) Indexed allocation uses a special block (or blocks) that hold a list of pointers to each data block of the file, enabling relatively easy random access but possibly requiring multiple index blocks for very large files.
   (d) A FAT (File Allocation Table)-based scheme is identical to contiguous allocation, since each file is stored in one single run of consecutive disk blocks.
   (e) Contiguous allocation requires no external fragmentation because files can be non-contiguous on disk.

7. (10pts, no partial point, no penalty) Which of the following statements describes a key difference between **paging** and **segmentation** is/are **TRUE**?

(a) Paging divides memory into variable-sized blocks, while segmentation divides memory into fixed-sized blocks.

(b) Paging can cause external fragmentation, while segmentation cannot.

(c) In paging, the logical address is split into page number and offset, whereas in segmentation, it is split into segment number and offset.

(d) Paging tables are always stored in virtual memory, but segmentation tables must be stored in physical memory.

(e) Segmentation typically uses multilevel page tables, while paging relies on a single-level translation scheme.

8. (10pts, no partial point, no penalty) Assume a **5-stage pipeline** (IF → ID → EX → MEM → WB) with **full forwarding** from:

- MEM → EX
- WB → EX

Which of the following back-to-back instruction pairs definitely **require inserting at least one stall** (a pipeline bubble) to avoid an incorrect result? (Select all that apply.)

(a) lw x5, 0(x6)
add x7, x5, x8

(b) lw x5, 0(x6)
beq x5, x0, L1

(c) lw x5, 0(x6)
sw x5, 0(x7)

(d) add x5, x1, x2
sub x6, x5, x3

(e) sw x6, 0(x7)
add x8, x6, x2

9. (10pts, no partial point) A system implements a two-level cache hierarchy with the following characteristics:

- L1 Cache: 1 cycle access time, 95% hit rate
- L2 Cache: 10 cycle access time, 90% hit rate (of those requests that miss in L1)
- Main Memory: 100 cycle access time (accessed only if an L2 miss occurs)

Assume there are no other stalls or overheads besides cache/memory accesses.

What is the **Average Memory Access Time (AMAT)** in cycles?

AMAT = _____ ?

10. (10pts) You have the following set of 4 processes, each with an arrival time (in milliseconds) and a CPU burst time (in milliseconds). You want to schedule them on a single CPU using FCFS, Non-Preemptive SJF, Round Robin (time quantum = 2 ms), and Shortest Remaining Time First (SRTF).
For each algorithm, you must compute the average waiting time.
What is the **average waiting time** for each scheduling method?

Process Table

Process	Arrival Time	Burst Time
P1	0	5
P2	1	3
P3	2	6
P4	4	2

FCFS = _____

Non-preemptive SJF = _____

Round Robin(q=2) = _____

SRTF = _____