

國立成功大學

115學年度碩士班招生考試試題

編 號：125

系 所：電機工程學系

科 目：計算機組織與作業系統

日 期：0203

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 CPU Performance and the "Iron Law" of performance is/are TRUE?
 - (a) If you increase the clock rate of a CPU, the execution time of a program will always decrease, regardless of the instruction count or CPI.
 - (b) Improving the compiler to reduce the number of instructions (Instruction Count) will always improve performance, assuming the CPI and Clock Cycle Time remain constant.
 - (c) Moving from a single-cycle implementation to a pipelined implementation typically increases the instruction latency but improves the overall instruction throughput.
 - (d) The CPI (Cycles Per Instruction) of a program is a fixed value determined solely by the hardware design and never changes based on the program's instruction mix.
 - (e) MIPS (Millions of Instructions Per Second) is generally considered a better metric than Execution Time for comparing computers with different instruction sets (ISAs).

2. (10pts, no partial point, no penalty) Consider a standard 5-stage RISC-V pipeline (IF, ID, EX, MEM, WB) with full forwarding. Which of the following statements about Branch Prediction and Hazards is/are correct?
 - (a) In a pipeline that assumes "branch not taken," if a branch is actually taken, the instructions fetched from the fall-through path must be flushed.
 - (b) A "control hazard" occurs because the decision to branch (and the target address) is not known until the MEM stage in a standard optimized RISC-V pipeline.
 - (c) Dynamic branch prediction uses a Branch History Table (BHT) to store the history of recent branches to predict the outcome of future branches.
 - (d) Moving the branch execution decision from the EX stage to the ID stage reduces the branch penalty (flush cost) from 2 cycles to 1 cycle.
 - (e) Branch prediction is only useful for conditional branches (beq, bne); unconditional jumps (jal) never cause control hazards.

3. (10pts, no partial point, no penalty) Which of the following statements about RISC-V Instruction Formats and Addressing is/are TRUE?
 - (a) RISC-V uses fixed-length 32-bit instructions, which simplifies the instruction fetch and decode stages compared to x86.
 - (b) The jal (Jump and Link) instruction uses the J-type format and saves the return address (PC+4) into register x1 (ra).
 - (c) In the I-type instruction format (used for addi, lw), the immediate field is 12 bits, allowing signed values between -2048 and +2047.
 - (d) The Store instruction (S-type) splits the immediate value into two parts (imm[11:5] and imm[4:0]) because the source register rs2 occupies the spot where the immediate bits usually sit in other formats.
 - (e) PC-relative addressing is calculated by adding the immediate value directly to the current value in the Stack Pointer (sp/x2).

4. (10pts, no partial point, no penalty) Assume a Direct Mapped Cache with the following parameters: 32-bit addresses, cache size of 16 KiB, and a block size of 16 bytes. Which of the following statements is/are correct?
- (a) The "Byte Offset" field requires 4 bits.
 - (b) The "Index" field requires 10 bits.
 - (c) The "Tag" field requires 18 bits.
 - (d) If the cache were changed to 2-way Set Associative (keeping total size and block size the same), the Index field size would decrease by 1 bit.
 - (e) A conflict miss occurs when the cache is full and the processor attempts to access a block that is not currently in the cache.
5. (10pts, no partial point, no penalty) Which of the following statements about Process Scheduling and Context Switching is/are TRUE?
- (a) Context switching is purely overhead; the system does no useful work while switching between processes.
 - (b) In a Round Robin scheduler, if the time quantum is extremely large, the algorithm behaves exactly like First-Come, First-Served (FCFS).
 - (c) "Hard affinity" guarantees that a process will never migrate to a different processor, whereas "soft affinity" only attempts to keep it on the same processor.
 - (d) A process in the "Ready" state is waiting for an I/O event to complete.
 - (e) Throughput is defined as the amount of time it takes to execute a particular process from submission to completion.
6. (10pts, no partial point, no penalty) Which of the following statements regarding Deadlocks and the Banker's Algorithm is/are TRUE?
- (a) A deadlock can only occur if the four necessary conditions hold simultaneously: Mutual Exclusion, Hold and Wait, No Preemption, and Circular Wait.
 - (b) The "Hold and Wait" condition can be prevented by requiring a process to request and be allocated all its resources before execution begins.
 - (c) The Banker's Algorithm is used for Deadlock Detection, not Deadlock Avoidance.
 - (d) An "Unsafe State" implies that a deadlock has already occurred.
 - (e) A Resource Allocation Graph with a cycle always indicates a deadlock if there is only single instance of each resource type.

7. (10pts, no partial point, no penalty) Which of the following statements about Page Replacement Algorithms and Virtual Memory is/are TRUE?
- (a) The Optimal Page Replacement algorithm (OPT) is impossible to implement in a real general-purpose OS because it requires knowledge of future memory accesses.
 - (b) The LRU (Least Recently Used) algorithm suffers from Belady's Anomaly (increasing frames increases page faults).
 - (c) Thrashing occurs when the sum of the sizes of the working sets of all active processes exceeds the total size of physical memory.
 - (d) A "dirty bit" in a page table entry indicates that the page is read-only and cannot be modified.
 - (e) Demand paging brings a page into memory only when it is needed (accessed) by the running process.
8. (10pts, no partial point, no penalty) Which of the following statements about Mass-Storage Structure and RAID is/are TRUE?
- (a) RAID 0 provides high performance via striping but offers no redundancy; if one drive fails, data is lost.
 - (b) RAID 1 implements disk mirroring, which effectively doubles the storage cost for the same capacity but provides 100% redundancy.
 - (c) RAID 5 requires a minimum of 3 disks and uses distributed parity to recover data if a single disk fails.
 - (d) Low-level formatting (physical formatting) divides a disk into sectors that the disk controller can read and write.
 - (e) In HDD scheduling, the SSTF (Shortest Seek Time First) algorithm guarantees that starvation will never occur.
9. (10pts, no partial point) A program runs on a CPU with the following instruction mix and CPI (Cycles Per Instruction):
- ALU ops: 50% of instructions, CPI = 1
 - Loads: 20% of instructions, CPI = 5
 - Stores: 10% of instructions, CPI = 4
 - Branches: 20% of instructions, CPI = 2
- If the clock frequency is 2 GHz and the program executes 10^9 instructions, what is the Total Execution Time in seconds?
- Execution Time = _____ seconds

10. (10pts, no partial point, no penalty) A system uses a Translation Look-aside Buffer (TLB) and Paging.

Assume the following timing characteristics:

TLB Hit Ratio: 90%

TLB Access Time: 10 ns

Main Memory Access Time: 100 ns

What is the Effective Memory Access Time (EMAT) in nanoseconds? (Assume a single-level page table and that the TLB lookup happens sequentially before memory access. If TLB miss, you must access memory to get the page table entry, then access memory again for the data).

EMAT = _____ ns