1. (10%) What are the two properties critical to program correctness that are normally preserved by control dependency? Explain them briefly and clearly with examples.

2. (10%) Assume we apply an optimization for the load-store machine for which the measurements in the following table have been made. The optimization reduces 40% of Store instructions, although it cannot reduce ALU, loads, or branches. Ignoring systems issues and assuming a 0.5-ns clock cycle time. Compare the performance before and after the optimization. The optimized machine will be how much faster?

Instruction Type	Frequency	Clock Cycle Count
ALU ops	20%	1
Loads	20%	2
Stores	50%	2
Branches	10%	2

3. (10%) Which of the following statements are true for load-store (register-to-register) architecture?
   (A) The hardware implementation of control units is typically simpler than that of Complex Instruction Set Computer (CISC) architectures.
   (B) A compiler can directly perform an addition operation between a data value in a register and a data value at a memory address.
   (C) It generally requires more registers than a Register-Memory architecture to avoid frequent memory access bottlenecks.
   (D) Instruction memory traffic is reduced because instructions are usually encoded in variable lengths.
   (E) The execution time (in clock cycles) of most instructions tends to be uniform and predictable.

4. (10%) In the context of advanced processor design, which of the following statements regarding hardware-based speculation and Reorder Buffers are TRUE?
   (A) The introduction of a Reorder Buffer allows instructions to execute out-of-order but ensures they commit in-order.
   (B) When a branch prediction is incorrect, the speculative results from the Reorder Buffers to other waiting instructions are discarded.
   (C) Prior to the instruction commit stage, speculative results are held in the main architectural register file.
   (D) Once a mis-prediction is detected, speculative instructions are committed to maintain program order.
   (E) In the hardware-based speculation scheme, the Reorder Buffer provides temporary storage for speculative results before they are permanently committed.

5. (10%) If modules have exponentially distributed lifetimes (age of module does not affect probability of failure), the overall failure rate is the sum of failure rates of the all modules in a system. A disk subsystem contains the following components: 16 disks (1M hour MTTF per disk), 1 disk controller (1M hour MTTF), and 1 power supply (1M hour MTTF). 1 fan (1M hour MTTF), 1 SCSI cable (1M hour MTTF)
   (A) Calculate MTTF for the disk subsystem
   (B) Assume it takes on average 24 hours for an operator to notice that a power supply has failed and replace it. After adding one redundant power supply, what is the reliability improvement of the entire disk subsystem?

$$MTTF_{\text{power supply pair}} = \frac{MTTF_{\text{power supply}} / 2}{MTTR_{\text{power supply}} / MTTF_{\text{power supply}}}$$

6. (16%) A system has a TLB, a page table in main memory, and a physically indexed, physically tagged cache. Each row gives whether the TLB hits, whether the referenced page is in main memory, and whether the cache hits.

   For each combination, write P (possible) or I (impossible).

TLB	Page in main memory?	Cache	P/I
Hit	Yes	Hit	(1)
Hit	Yes	Miss	(2)
Hit	No	Hit	(3)
Hit	No	Miss	(4)
Miss	Yes	Hit	(5)
Miss	Yes	Miss	(6)
Miss	No	Hit	(7)
Miss	No	Miss	(8)

7. (10%) Given a 128-byte cache with 16-byte blocks and 32 bits addresses, answer the following questions.
   a. Offset bits: (1)
   b. 2-way set associative: tag bits (2); index bits (3)
   c. Direct-mapped: tag bits (4); index bits (5)

8. (12%) A 4 GHz processor has a base CPI of 1.2 with a perfect cache. Every instruction makes one instruction fetch, and 30% of instructions are loads/stores. L1 hits are included in the base CPI. L1 miss rates are 2% (I-cache) and 5% (D-cache). The unified L2 has a 12-cycle hit time and a 25% local miss rate. Main memory access time is 80 ns (i.e., 80x10^-9 s).
    a. The main memory access takes (1) cycles, and the number of L1 misses per instruction is (2).
    b. The effective CPI is (3).
    c. Without L2, the CPI would be (4).

9. (12%) A TLB lookup takes 2 ns and is performed on every reference. A memory access takes 100 ns. Page-table entries are never cached. The TLB hit rate is 98%. One in every 10^6 references causes a page fault, which costs 8 ms (i.e., 8x10^-6 s) to service (add this on top of the normal access).
    a. With a two-level page table, the effective memory access time is (1) ns.
    b. To keep the effective access time at or below 110 ns (two-level table), the page-fault rate must be at most (2) per reference.
    c. With a four-level page table (all else unchanged, original fault rate), the effective access time is (3) ns.