

I. COMPUTER SCIENCE

- C1. Consider n given points in the 2D plane where each point is specified by its coordinates and is colored either RED or BLUE. Design an efficient algorithm for testing whether a straight line $L : y = ax + b$ exists in the plane such that all the RED points lie on one side of the line L and all the BLUE points lie on the other side of the line L . If such a line exists, then report the constants a and b that define the line L ; otherwise report “*No such line exists*”.

[12]

- C2. (a) A k -stage pipelined processor is executing a program consisting of three types of instructions namely, memory reference instruction (MRI), register reference instruction (RRI) and input-output instruction (IOI). In the program, 40% of the instructions are of MRI type, 40% are of RRI type and the rest are of IOI type. Assume that 10% of MRIs cause stalls of 10 clock cycles each and 50% of IOIs cause stalls of 2 clock cycles each, whereas no RRI causes any stalls. Calculate the average number of clock cycles required to execute an instruction by the processor in steady state.
- (b) Consider a processor for which instructions are encoded using 16 bytes. Each instruction operates on 16-byte operands. Assume that the processor implements a 32 KB cache that is 8-way set associative and contains 1024 sets. Identify the problem, if any, in this cache design and propose a modification to rectify the problem.

[6+(2+4)=12]

C3. Let X and Y be two 2-bit inputs to a circuit. The circuit works as follows:

if $(X \geq Y)$ **then** output $(X + Y)$ **else** output $(X \oplus Y)$.

Here ‘+’ denotes the integer addition operation and ‘ \oplus ’ denotes the bitwise XOR operation.

- (a) How many bits are required for the output?
- (b) Identify the major blocks of the circuit and the connections among them.
- (c) Provide a gate level design for each block.

[2+4+6=12]

C4. Consider a relational scheme $R(A, B, C, D)$, where A is the primary index. The objective is to choose a secondary index on the basis of the following.

- A search query specifies only one attribute: B in 75% of the cases, C in 20% of the cases, and D in 5% of the cases.
- The cost of an insertion/deletion is 1 unit of time and an additional $\frac{1}{2}$ unit of time per secondary index.
- The cost of a search query is 1 unit of time if there is a secondary index on the search attribute, and 10 units of time if there is no secondary index on the specified attribute.

Suppose 40% of the operations are insertion/deletion and the rest are search operations. Determine the attribute on which the secondary index should be created.

[12]

C5. Consider the generator polynomial $g = 110011$ and the data bits $m = 11100011$.

- (a) Find the CRC (Cyclic Redundancy Check) bits and the transmitted string T .
- (b) Suppose errors occur only in the data bits (CRC bits unchanged). Provide an example where such error(s) cannot be detected.
- (c) Instead of the above, assume that the error(s) may occur only in the CRC bits (data bits unchanged). Prove that such error(s) can always be detected.

[3+4+5=12]

C6. Consider the following information about a magnetic hard disk, with cylinders numbered 0 (innermost) to 511 (outermost):

- at time $t = 0$ ms, the disk head is at cylinder **120**, moving outwards towards cylinder **511**;
- the disk arm takes **1** ms to seek from cylinder i to cylinder $i + 1$ or cylinder $i - 1$;
- the time taken to read sectors may be neglected.

Let $T_{LOOK}(S)$ and $T_{SSTF}(S)$ denote the total time taken by the disk controller to service a sequence S of requests, using LOOK and SSTF algorithms respectively (standard definitions are given at the end of this question*).

- (a) Two requests arrive at almost the same time for cylinder numbers n_1 and n_2 . Show that, for any $n_1, n_2 \in [0, 511]$, $T_{LOOK}(S) \geq T_{SSTF}(S)$ for the sequence $S = \langle n_1, n_2 \rangle$.

- (b) Construct a sequence of length 3 for which $T_{SSTF}(S) > T_{LOOK}(S)$. (You should show that $T_{SSTF}(S) > T_{LOOK}(S)$ for the sequence that you have constructed.)

[5+7=12]

*Recall that in LOOK, the disk head alternately moves in the outward and inward directions, servicing any currently pending request as it reaches each cylinder. While travelling in a particular direction, the head goes only as far as the last request currently pending in that direction, before reversing the direction. In SSTF, for a sequence of requests, the disk head first moves to the cylinder nearest to its current position to service the corresponding request, and proceeds similarly to service the remaining requests.

II. MATHEMATICS FOR COMPUTER SCIENCE

- MC1. (a) Find all the positive integers a for which $a^{10} + 1$ is divisible by 10.
- (b) If k is odd, then prove that 2^{n+2} divides $k^{2^n} - 1$ for any natural number n .

[6+6=12]

- MC2. Consider the infinite sequence w_1, w_2, \dots , where $w_1 = 0$ and w_{i+1} , $i \geq 1$, is obtained by replacing each 0 in w_i by 001 and each 1 in w_i by 0. For example, $w_2 = 001$ and $w_3 = 0010010$.

- (a) For $k \geq 2$, write the recurrences for (i) a_k , the number of 0s in w_k , and (ii) b_k , the number of 1s in w_k .
- (b) Find the 23-rd bit (from the left) in w_{100} .

[(4+2)+6=12]

- MC3. Shares of company A are priced at Rs. 20 per unit today. Tomorrow it will be priced at either Rs. 10 per unit or Rs. 40 per unit, with both possibilities being equally likely. You have a budget of Rs. 1000 to buy shares of A .

What should be your strategy for trading (buying and selling) the shares of A so that by the end of tomorrow

- (a) the expected amount of money at hand is maximum?
- (b) the expected number of units you have is maximum?

[6+6=12]

- MC4. (a) Consider the regular language L consisting of all strings ending with 010, defined over the alphabet $\Sigma = \{0, 1\}$. Further, consider the DFA D shown in Figure 1. Let \hat{L} be the language recognized by D . What is the relation between L and \hat{L} ? Justify your answer.

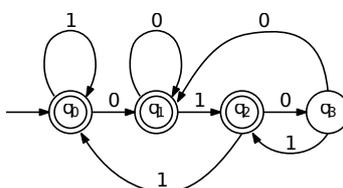


Figure 1: DFA D

- (b) Consider a Turing machine M with exactly one accept state and exactly one reject state. Construct a Turing machine N by swapping the accept and reject states of M . Prove or disprove (with a counterexample) the following statement about the languages $L(M)$ and $L(N)$ accepted by the machines M and N respectively:

$$L(N) = \overline{L(M)}.$$

[5+7=12]

- MC5. Let A and B be two $n \times n$ real symmetric and non-negative definite matrices. Define $tr(A)$ as the sum of all the diagonal elements of A .
- (a) Prove that $0 \leq tr(AB) \leq tr(A)tr(B)$ when A is a diagonal matrix.

- (b) Prove that the above result holds even if A is not a diagonal matrix.

[4+8=12]

MC6. Let V be a vector space over \mathbb{R} (the set of real numbers) and $g_0, g_1, \dots, g_k : V \rightarrow \mathbb{R}$ be linear maps such that $\ker(g_0) \supset \bigcap_{j=1}^k \ker(g_j)$, where $\ker(\cdot)$ denotes the kernel of the corresponding linear map.

- (a) If $G : V \rightarrow \mathbb{R}^k$ is the map $G(v) = (g_1(v), g_2(v), \dots, g_k(v))$, then show that there exists a linear map $F : \mathbb{R}^k \rightarrow \mathbb{R}$ satisfying $F \circ G = g_0$.
- (b) Show that there exists $a_1, a_2, \dots, a_k \in \mathbb{R}$ such that

$$g_0 = \sum_{j=1}^k a_j g_j.$$

[8+4=12]