

I. COMPUTER SCIENCE

Notation

- \mathbb{N} : the set of all positive integers.
- For all $n \in \mathbb{N}$, $[n]$ denotes the set $\{1, 2, \dots, n\}$.

C1. A *perfect binary tree* is a rooted binary tree where each node stores an integer, each internal node has two children, and all the leaf nodes are at the same level. Consider a perfect binary tree \mathbf{B} with the following property: the value stored in each internal node is greater than or equal to the values stored in all the nodes in its left subtree and strictly less than the values stored in all the nodes in its right subtree. Figure 1 shows an example.

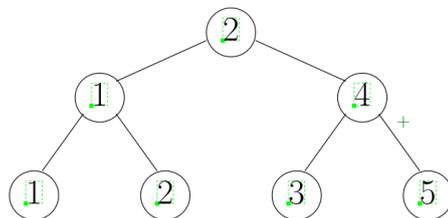


Figure 1: The above figure shows an example of a perfect binary tree of depth 2. Here $S = \{1, 2, 3, 4, 5\}$ and $|S| = 5$. Note that $|S|$ could have been even less in this case.

For any $k \in \mathbb{N}$, if \mathbf{B} has a depth of k , then \mathbf{B} has $2^{k+1} - 1$ nodes. Consider the set S of all the values stored in \mathbf{B} . Find the minimum value of $|S|$ in terms of k . Prove your result.

[12]

C2. (a) A *closed interval* on the real line is defined as

$$[a, b] = \{x \in \mathbb{R} \mid a \leq x \leq b\}$$

where $a < b$. Let $S = \{I_1, I_2, \dots, I_n\}$ be a collection of n non-empty closed intervals with $I_j = [a_j, b_j]$ for $1 \leq j \leq n$. Design an algorithm that checks if all pairs of intervals in S intersect, i.e., $I_i \cap I_j \neq \emptyset$ for $1 \leq i < j \leq n$. The algorithm should have time complexity $O(n)$ to receive full credit.

- (b) Let G be a simple, connected, edge-weighted graph on $n \geq 6$ vertices with no negative cycles. The minimum edge weight in G is ω , with $\omega < 0$. Let s and t be two distinct vertices in G .

Dijkstra's shortest-path algorithm is used to find a shortest path between two given vertices in a graph when all the edge-weights are positive. To ensure that all the edge-weights are positive, we obtain G' from G by adding $2|\omega|$ to the weight of every edge in G . The resulting new graph G' has all edge weights strictly greater than zero.

A shortest path between s and t in G' found by Dijkstra's algorithm is also a shortest path between s and t in G .

Prove or disprove the above claim.

[6 + 6 = 12]

- C3.** (a) Consider a floating point representation in the following format:

$$x = \pm 0.1m \times 2^e,$$

where m and e stand for mantissa and exponent, respectively, as given in Figure 2.

Explain the overflow and underflow regions of this floating-point number system, clearly stating the assumptions made.

- (b) Consider a pipeline with the following 5 stages: IF (Instruction Fetch), ID (Instruction Decode), OF

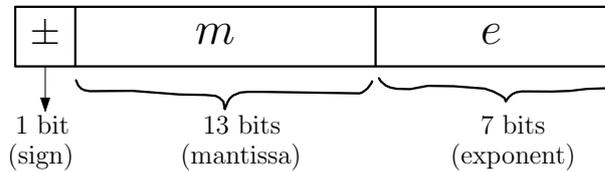


Figure 2: Normalized Floating Point Representation

(Operand Fetch), EX (Execute), and WB (Write Back). Each of the IF, ID, OF, and WB stages takes one clock cycle to complete. The number of clock cycles required by the EX stage depends on the instruction. The ADD and SUB instructions require 1 clock cycle each, while the MUL instruction needs 3 clock cycles.

All (memory or register) reads take place in the second half of a clock cycle and all writes occur in the first half.

Now consider the following instruction sequence:

MUL	$R_3, R_1, R_2;$	$\#R_3 \leftarrow R_1 \times R_2$
ADD	$R_5, R_3, R_4;$	$\#R_5 \leftarrow R_3 + R_4$
SUB	$R_8, R_6, R_7;$	$\#R_8 \leftarrow R_6 - R_7$

- (i) Identify and explain the different hazards that the above code can generate when executed on the five-stage pipeline.
- (ii) Assume that STALLs are allowed. Determine an execution of the above instruction sequence with stalls inserted in the pipeline at appropriate stages that would avoid all the hazards you identified. How many clock cycles does your solution take to complete?
- (iii) Demonstrate how the same result can be achieved in fewer clock cycles and show the corresponding execution.

[4+(2+3+3)=12]

C4. A sender S and a receiver R are connected through a path of length k , consisting of k bidirectional links. Link i ($1 \leq i \leq k$) has a packet loss probability of p_i in the direction from S to R and q_i in the direction from R to S . Assume that each packet on a link is received or lost independently of the other packets, and that each packet's loss probability on a link is the same as any other's.

(a) Suppose the probability that a data packet sent by S does not reach R is P , and the probability that an acknowledgment packet sent by R does not reach S is Q . Deduce expressions for P and Q in terms of the p_i, q_i ($1 \leq i \leq k$).

(b) Suppose S and R use a stop-and-wait ARQ protocol to communicate. What is the expected number of transmissions of a packet before S can send the next packet in sequence? Deduce your answer in terms of P and Q .

(c) Let RTT be the round-trip time of the connection between S and R when neither packet loss nor acknowledgment loss occurs over any link. Let RTO be the retransmission timeout. What is the expected time taken to send a packet from S to R and receive its acknowledgment back at S ? Express your answer in terms of RTT, RTO, P , and Q .

[2 + 4 + 6 = 12]

C5. For an attribute set A with a set of underlying functional dependencies, the closure A^+ is defined as the set of attributes that can be derived from A . Given a relation R with attribute sets X, Y, Z , prove or disprove each the following statements.

$$(a) (X \cup Y)^+ \cap (X \cup Z)^+ = X^+ \cup (Y \cap Z)^+.$$

$$(b) (X \cap Y)^+ \cup (X \cap Z)^+ = X^+ \cap (Y \cup Z)^+.$$

$$[6 + 6 = 12]$$

C6. Suppose the following code is compiled, and run for different values of N .

```
int main() {
    for (int i = 1; i <= N; i++) {
        fprintf(stderr, "ISI\n");
        fork();
        fork();
    }
    fprintf(stderr, "JRF\n");
    return 0;
}
```

Assume that all fork calls are successful.

- (a) How many times will the strings **ISI** and **JRF** be printed if the value of N is (i) 1, (ii) 2?
- (b) For a given N , let I_N and J_N represent, respectively, the number of times the strings **ISI** and **JRF** are printed. Derive recurrence relations for I_N and J_N .

$$[(1 + 3) + 8 = 12]$$

II. MATHEMATICS FOR COMPUTER SCIENCE

Notation

- \mathbb{N} : the set of all positive integers.
- For all $n \in \mathbb{N}$, $[n]$ denotes the set $\{1, 2, \dots, n\}$.

M1. Let $\Sigma = \{0, 1\}$, and Σ^* be the set of all finite strings over Σ . Suppose that for any string $x \in \Sigma^*$, let $f_0(x)$ and $f_1(x)$ denote the number of occurrences of 0 and 1 in x , respectively. Let $L \subseteq \Sigma^*$ be a language defined as follows

$$L = \{x \in \Sigma^* : 2f_0(x) + f_1(x) \equiv 0 \pmod{3}\}.$$

Design a deterministic finite automaton (DFA) recognizing the language L . For full credit, your DFA should have three states.

[12]

M2. Consider the following two definitions from graph theory.

- A *tournament* is a directed graph with *exactly* one edge between any pair of vertices, in one of the two possible directions.
- A *king* in a directed graph is a vertex from which every other vertex in the graph can be reached through a directed path of length at most 2.

Show that every tournament on at least 2 vertices contains a king.

[12]

M3. For $n \in \mathbb{N}$, define $\sin_n(x)$ as follows:

$$\sin_n(x) = \begin{cases} \sin(x) & \text{if } n = 1, \\ \sin(\sin_{n-1}(x)) & \text{if } n \geq 2. \end{cases}$$

(a) If $\sin(x) > 0$, then calculate

$$\lim_{n \rightarrow \infty} \sin_n(x).$$

(b) Given $n \in \mathbb{N}$, calculate

$$\lim_{x \rightarrow 0} \frac{\sin_n(x)}{x}.$$

[8 + 4 = 12]

M4. (a) Let A be an $n \times n$ matrix with $A_{ij} = \min\{i, j\}$. Compute the determinant of A .

(b) Let B be an $n \times n$ matrix with $B_{ij} = \frac{1}{\min\{i, j\}}$. Prove or disprove the following claim:

The determinant of B is zero.

[4 + 8 = 12]

M5. For all $n \in \mathbb{N}$, let S_n denote the set of all the bijective maps from $[n]$ to $[n]$. Let π be chosen uniformly at random from the set S_n . For a given $i \in [n]$, compute the following probabilities:

(a) $\Pr[\pi(i) \neq i]$

(b) $\Pr[\pi(i) = \min\{\pi(1), \dots, \pi(i)\}]$.

[4 + 8 = 12]

M6. Consider the following recursion

$$F_n = \begin{cases} 0 & \text{if } n = 0, \\ 2 & \text{if } n = 1, \\ 6F_{n-1} - 4F_{n-2} & \text{if } n \geq 2. \end{cases}$$

Answer the following questions.

-
- (a) For what values of n do we have $F_n = 2028$?
- (b) Does there exist infinitely many n such that $F_n = 6n^6 - 4n^4$?
Justify your answer.

$$[4 + 8 = 12]$$