TIFR GS Computer Science

Sample Paper – 4

Duration: 180 Minutes **Maximum Marks: 120**

Instructions

- This paper contains **30 Multiple Choice Questions** (single correct answer), divided into **Part A: Common (Q1–Q15)** and **Part B: Computer Science (Q16–Q30)**.
- **Part A** must be attempted by all candidates. **Part B** must be attempted for Computer Science (CS) programme admission.
- Each correct answer carries **+4 marks**. **1 mark is deducted** for each wrong answer. Unattempted questions carry **0 marks**.
- Only **one** option is correct per question.
- Use of calculators, mobile phones, laptops, tablets, or any electronic device is strictly **NOT permitted**. Rough sheets are provided at the examination centre.
- Syllabus level: **Graduate-level Computer Science** — Discrete Mathematics, Algorithms & Data Structures, Theory of Computation, Programming Languages, Operating Systems, Digital Circuits, Databases.

Part A: Common (Q1–Q15)

Q1. [Boolean Algebra – Absorption] Which of the following is the simplified form of the Boolean expression $A \cdot B + A \cdot B'$?

(A) $A \cdot B$

(B) B

(C) A

(D) $A + B$

Q2. [Combinatorics – Burnside's Lemma] How many distinct binary necklaces (each bead is either black or white) with exactly 6 beads exist, considering only rotational symmetry?

(A) 14

(B) 16

(C) 12

(D) 18

Q3. [Graph Theory – Graph Colouring] Consider the cycle graph C_5 shown below.

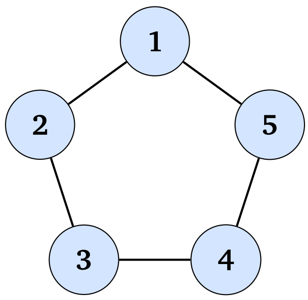

What is the chromatic number $\chi(C_5)$?

(A) 2

(B) 4

(C) 5

(D) 3

Q4. [Number Theory – Euler's Totient] What is $\varphi(36)$, the number of positive integers up to 36 that are coprime to 36?

(A) 10

(B) 12

(C) 18

(D) 24

Q5. [Combinatorics – Stirling Numbers] The Stirling number of the second kind $S(4,2)$ gives the number of ways to partition the set $\{1,2,3,4\}$ into exactly 2 non-empty, non-ordered subsets. What is $S(4,2)$?

(A) 4

(B) 6

(C) 7

(D) 9

Q6. [Probability – Bernoulli Variance] A random variable X follows a Bernoulli distribution with parameter $p = 0.3$. What is $\mathrm{Var}(X)$?

(A) 0.21

(B) 0.30

(C) 0.09

(D) 0.49

Q7. [Linear Algebra – Matrix Rank] What is the rank of the matrix

$$M = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}?$$

(A) 1

(B) 3

(C) 0

(D) 2

Q8. [Predicate Logic – Resolution] Using the resolution rule, what is the resolvent of the clauses $\{P, Q\}$ and $\{\neg P, R\}$?

(A) $\{P, Q, R\}$

(B) $\{Q, R\}$

(C) $\{\neg P, Q, R\}$

(D) $\{P, R\}$

Q9. [Graph Theory – Bipartite Matching] Consider the complete bipartite graph $K_{3,3}$ shown below, with vertex sets $X = \{x_1, x_2, x_3\}$ and $Y = \{y_1, y_2, y_3\}$.

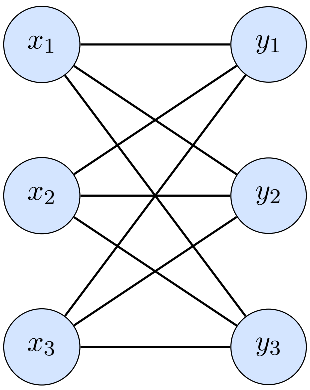

Does $K_{3,3}$ have a perfect matching?

(A) No, because not every vertex has degree 3

(B) No, Hall's condition fails for some subset

(C) Yes, because Hall's condition holds for all subsets of X

(D) Yes, but only one perfect matching exists

Q10. [Combinatorics – Multinomial Coefficient] What is the coefficient of a^2b^3c in the expansion of $(a+b+c)^6$?

(A) 15

(B) 20

(C) 30

(D) 60

Q11. [Probability – Markov's Inequality] Let X be a non-negative random variable with $E[X] = 5$. Using Markov's inequality, what is the best upper bound for $P(X \geq 10)$?

(A) 0.5

(B) 0.1

(C) 0.25

(D) 1

Q12. [Number Theory – Sum of Proper Divisors] The sum of all proper divisors of 12 (i.e., all divisors except 12 itself) is:

(A) 12

(B) 16

(C) 28

(D) 14

Q13. [Relations – Equivalence Relations] How many distinct equivalence relations can be defined on the set $\{1, 2, 3\}$?

(A) 3

(B) 4

(C) 5

(D) 6

Q14. [Logic – Semantic Tableaux] Given the premises $P \vee Q$ and $\neg P$, a semantic tableau is constructed. After branching on $P \vee Q$, the branch containing P closes because it contradicts $\neg P$. Which formula is provable from these premises?

(A) $P \vee Q$

(B) $\neg Q$

(C) P

(D) Q

Q15. [Recurrences – Characteristic Equation] Consider the homogeneous linear recurrence $T(n) = 5T(n-1) - 6T(n-2)$ for $n \geq 2$. Solving its characteristic equation gives roots $r_1 = 2$ and $r_2 = 3$. What is the general solution?

(A) $\alpha \cdot 2^n + \beta \cdot 3^n$

(B) $\alpha \cdot 3^n + \beta \cdot 5^n$

(C) $\alpha \cdot (5/2)^n$

(D) $\alpha \cdot 6^n$

Part B: Computer Science (Q16–Q30)

Q16. [Algorithms – Prim's MST] Consider the weighted undirected graph below with 4 nodes and 5 edges.

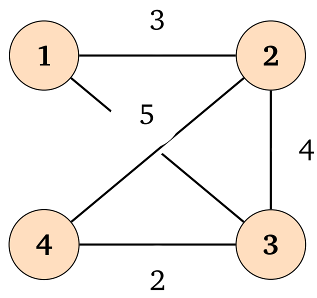

Starting from node 1, Prim's algorithm selects edges greedily to build a Minimum Spanning Tree. What is the total weight of the MST?

(A) 8

(B) 9

(C) 10

(D) 11

Q17. [Data Structures – B-tree] In a B-tree with minimum degree $t = 2$ (also known as a 2-3-4 tree), what is the minimum number of keys that a non-root, non-leaf internal node must contain?

(A) 2

(B) 0

(C) 1

(D) 3

Q18. [Theory of Computation – Chomsky Hierarchy] Which of the following correctly classifies the language $L = \{a^n b^n c^n \mid n \geq 1\}$?

(A) Regular

(B) Context-free but not regular

(C) Context-free

(D) Decidable but not context-free

Q19. [Algorithms – Longest Increasing Subsequence] What is the length of the longest increasing subsequence (LIS) of the sequence

$$[2,\ 1,\ 5,\ 3,\ 6,\ 4,\ 8,\ 9,\ 7]\ ?$$

(A) 5

(B) 4

(C) 6

(D) 3

Q20. [Data Structures – Hash Table] For a hash table using *linear probing* with load factor $\alpha = 0.5$, the expected number of probes for an **unsuccessful** search is approximately $\frac{1}{2}\left(1 + \frac{1}{(1-\alpha)^2}\right)$. What is this value?

(A) 2.0

(B) 2.5

(C) 3.0

(D) 4.0

Q21. [Theory of Computation – Rice's Theorem] Is the language $L = \{\langle M\rangle \mid \text{Turing machine } M \text{ accepts at least one string}\}$ decidable?

(A) Yes, decidable by checking all possible inputs one by one

(B) Yes, decidable using a dovetailing simulation

(C) No, undecidable by Rice's theorem since it is a non-trivial semantic property

(D) Yes, decidable only when the Turing machine's language is finite

Q22. [Operating Systems – TLB Effective Access Time] A memory system has TLB hit ratio $h = 0.9$, TLB access time $t_{\text{TLB}} = 2\,\text{ns}$, and main memory access time $t_m = 10\,\text{ns}$. On a TLB miss, one extra main memory access is required for the page table. The Effective Memory Access Time (EMAT) is:

$$\text{EMAT} = h(t_{\text{TLB}} + t_m) + (1 - h)(t_{\text{TLB}} + 2t_m)$$

What is the EMAT?

(A) 13 ns

(B) 10 ns

(C) 12 ns

(D) 22 ns

Q23. [Compilers – Syntax-Directed Definition] An arithmetic expression grammar uses synthesized attributes: $E \to E_1 + T$ with $E.val = E_1.val + T.val$, and $T \to T_1 * F$ with $T.val =$

$T_1.val \times F.val$, with operator precedence respected. What is the synthesized value of E for the input `3 + 4 * 2`?

(A) 14

(B) 20

(C) 7

(D) 11

Q24. [Theory of Computation – Many-One Reduction] Suppose $A \leq_m B$ (language A is many-one reducible to language B) and B is decidable. What can we conclude about A?

(A) A is undecidable

(B) A is decidable

(C) A is Turing-recognizable but not decidable

(D) Nothing definitive can be concluded

Q25. [Algorithms – Max-Flow] Consider the flow network below, where each edge is labelled with its capacity.

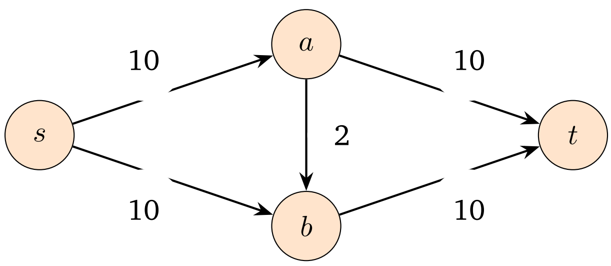

What is the maximum flow from s to t?

(A) 10

(B) 15

(C) 20

(D) 22

Q26. [Operating Systems – Peterson's Algorithm] Peterson's mutual exclusion algorithm for two processes uses two shared variables to coordinate access to the critical section. Which pair of shared variables does it use?

(A) A `mutex` lock and a boolean array `flag[]`

(B) A semaphore S and a boolean array `flag[]`

(C) Only a `turn` variable (round-robin)

(D) A boolean array `flag[]` and an integer `turn`

Q27. [Digital Circuits – Multiplexer] The figure below shows a 2:1 multiplexer (MUX). We wish to implement $f(A, B) = A \oplus B$ (XOR) using this MUX, with A as the SELECT input.

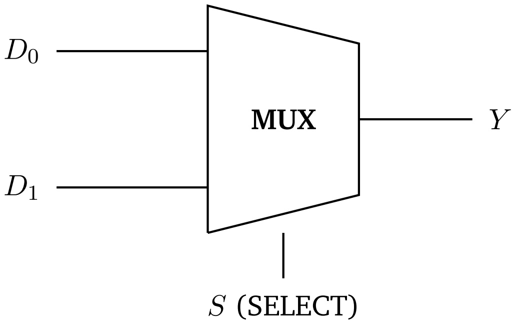

When $A = 0$: $f = B$, and when $A = 1$: $f = \overline{B}$. Which connection correctly implements $f = A \oplus B$?

(A) Select $= A,\ D_0 = B,\ D_1 = \overline{B}$

(B) Select $= B,\ D_0 = A,\ D_1 = \overline{A}$

(C) Select $= A,\ D_0 = 0,\ D_1 = B$

(D) Select $= A,\ D_0 = \overline{B},\ D_1 = B$

Q28. [Databases – Conflict Serializability] A schedule S over a set of transactions is conflict-serializable if and only if which of the following holds for its conflict (precedence) graph?

(A) The graph is bipartite

(B) The graph is acyclic (a DAG)

(C) The graph is strongly connected

(D) The graph is complete

Q29. [Programming Languages – System F Polymorphism] In System F (the polymorphic lambda calculus), the type $\forall\alpha.\,\alpha \to \alpha$ is the type of the polymorphic identity function. Which of the following terms has this type?

(A) $\lambda x.\ x\,x$

(B) $\lambda x.\,\lambda y.\ x$

(C) $\lambda\alpha.\,\lambda x.\ x+1$

(D) $\Lambda\alpha.\,\lambda x{:}\alpha.\ x$

Q30. [Algorithms – FPTAS for Knapsack] A Fully Polynomial-Time Approximation Scheme (FPTAS) for the 0/1 Knapsack problem with n items and maximum profit P works by scaling down item profits by a factor $K = \lfloor \varepsilon P/n \rfloor$ and solving the resulting DP. What is the running time of this FPTAS?

(A) $O(n/\varepsilon)$

(B) $O(n \cdot P/\varepsilon)$

(C) $O(n^2/\varepsilon)$

(D) $O(2^n)$

Detailed Solutions

Solution

Concept: Boolean Algebra – Distributive and Complement laws.

Step 1: Factor out A from both terms:

$$A \cdot B + A \cdot B' = A(B + B').$$

Step 2: Apply the complement law $B + B' = 1$:

$$A(B + B') = A \cdot 1 = A.$$

Why other options are wrong:

Q1.
- (A) $A \cdot B$: This is only one term, not the full simplification.
- (B) B: Incorrect; A is factored out, not B.
- (D) $A + B$: This over-expands; the expression collapses to just A.

Answer: (C) ↑ Go Back to Q1

Solution

Concept: Burnside's Lemma for counting distinct necklaces under rotational symmetry.

Setup: The cyclic group $\mathbb{Z}_6$ acts on 6-bead binary necklaces. For rotation by k positions, the number of fixed colourings is $2^{\gcd(k,6)}$.

Step 1: Compute fixed points for each rotation $k = 0, 1, \ldots, 5$:

k	0	1	2	3	4	5
$\gcd(k,6)$	6	1	2	3	2	1
$2^{\gcd(k,6)}$	64	2	4	8	4	2

Step 2: Apply Burnside's lemma:

$$\text{Distinct necklaces} = \frac{1}{6}(64 + 2 + 4 + 8 + 4 + 2) = \frac{84}{6} = 14.$$

Why other options are wrong:

Q2.
- (B) 16, (C) 12, (D) 18: Common miscounts from wrong $\gcd$ values or incorrect group size.

Answer: (A) ↑ Go Back to Q2

Q3.

Solution

Concept: Chromatic number of cycle graphs.

Key fact: For a cycle C_n:

$$\chi(C_n) = \begin{cases} 2 & \text{if } n \text{ is even,} \\ 3 & \text{if } n \text{ is odd.} \end{cases}$$

Reasoning for C_5: C_5 is an odd cycle ($n = 5$). At least 3 colours are needed because any two adjacent vertices must differ, and the odd cycle forces a colour conflict if only 2 are used. Three colours suffice: colour the vertices $1, 2, 3, 4, 5$ with colours R, G, R, G, B (going around the pentagon).

Why other options are wrong:

- (A) 2: Only works for even cycles (bipartite graphs).
- (B) 4 and (C) 5: More colours than necessary; 3 already suffice.

Answer: (D) ↑ Go Back to Q3

Q4.

Solution

Concept: Euler's totient function $\varphi(n)$.

Step 1: Factorise $36 = 2^2 \cdot 3^2$.

Step 2: Apply the product formula:

$$\varphi(36) = 36 \cdot \left(1 - \frac{1}{2}\right)\left(1 - \frac{1}{3}\right) = 36 \cdot \frac{1}{2} \cdot \frac{2}{3} = 12.$$

Why other options are wrong:

- (A) 10: Incorrect arithmetic.
- (C) 18: This equals $\varphi(27)$, not $\varphi(36)$.
- (D) 24: This equals $\varphi(35)$ or $\varphi(25 \cdot \ldots)$; not $\varphi(36)$.

Answer: (B) ↑ Go Back to Q4

Solution

Concept: Stirling numbers of the second kind $S(n, k)$ count partitions of an n-element set into exactly k non-empty subsets.

Step 1: Use the recurrence $S(n, k) = k \cdot S(n - 1, k) + S(n - 1, k - 1)$, or list all partitions of $\{1,2,3,4\}$ into exactly 2 non-empty parts:

$$\{1\}|\{2,3,4\}, \quad \{2\}|\{1,3,4\}, \quad \{3\}|\{1,2,4\}, \quad \{4\}|\{1,2,3\},$$

$$\{1,2\}|\{3,4\}, \quad \{1,3\}|\{2,4\}, \quad \{1,4\}|\{2,3\}.$$

Step 2: Count: 7 partitions, so $S(4,2) = 7$.

Why other options are wrong:

Q5.

- (A) 4: Counts only the "one singleton" partitions.
- (B) 6: Off by one; misses one partition.
- (D) 9: Overcounts, possibly confusing with ordered pairs.

Answer: (C) ↑ Go Back to Q5

Solution

Concept: Variance of a Bernoulli random variable.

Formula: If $X \sim \text{Bernoulli}(p)$, then $\text{Var}(X) = p(1-p)$.

Calculation:

$$\text{Var}(X) = 0.3 \times (1 - 0.3) = 0.3 \times 0.7 = 0.21.$$

Why other options are wrong:

Q6.

- (B) 0.30: This is $E[X] = p$, not the variance.
- (C) 0.09: This is p^2, not $p(1-p)$.
- (D) 0.49: This is $(1-p)^2$, not $p(1-p)$.

Answer: (A) ↑ Go Back to Q6

Solution

Concept: Rank of a matrix via row reduction.

Step 1: Observe the rows: $R_1 = (1,2,3)$, $R_2 = (4,5,6)$, $R_3 = (7,8,9)$.

Step 2: Check linear dependence: $R_3 - R_1 = (6,6,6) = 2(R_2 - R_1)$, so $R_3 = 2R_2 - R_1$. Hence all three rows are linearly dependent and span only a 2-dimensional space.

Step 3: Verify that R_1 and R_2 are linearly independent (not proportional). Therefore $\text{rank}(M) = 2$.

Why other options are wrong:

Q7.

- (A) 1: Would require all rows to be proportional; $(1, 2, 3)$ and $(4, 5, 6)$ are not.
- (B) 3: Would require all rows to be linearly independent, but they are not.
- (C) 0: Only the zero matrix has rank 0.

Answer: (D) ↑ Go Back to Q7

Solution

Concept: Resolution rule in propositional logic.

Rule: Given clauses $C_1 = \{P, Q\}$ and $C_2 = \{\neg P, R\}$, the resolvent on the complementary literal P is obtained by removing P from C_1 and $\neg P$ from C_2, then taking the union of remaining literals:

$$\{P, Q\} \cup \{\neg P, R\} \xrightarrow{\text{resolve on } P} \{Q\} \cup \{R\} = \{Q, R\}.$$

Why other options are wrong:

Q8.

- (A) $\{P, Q, R\}$: Keeps P, which is incorrect; resolution eliminates P.
- (C) $\{\neg P, Q, R\}$: Keeps $\neg P$, which is incorrect.
- (D) $\{P, R\}$: Drops Q, which should remain.

Answer: (B) ↑ Go Back to Q8

Solution

Concept: Hall's Marriage Theorem for bipartite perfect matching.

Hall's condition: A bipartite graph $G = (X \cup Y, E)$ has a perfect matching iff for every subset $S \subseteq X$, $|N(S)| \geq |S|$ (where $N(S)$ is the neighbourhood of S in Y).

For $K_{3,3}$: Every vertex in X is connected to all 3 vertices in Y. For any $S \subseteq X$ with $|S| = k$, $|N(S)| = 3 \geq k$. Hall's condition holds for all subsets, so a perfect matching exists. (In fact, $K_{3,3}$ has $3! = 6$ perfect matchings.)

Why other options are wrong:

Q9.

- (A) Every vertex has degree 3, so this is factually false.
- (B) Hall's condition is satisfied, not violated.
- (D) $K_{3,3}$ has 6 distinct perfect matchings, not just one.

Answer: (C) ↑ Go Back to Q9

Solution

Concept: Multinomial coefficient for expansion of $(a+b+c)^6$.

Formula: The coefficient of $a^i b^j c^k$ in $(a+b+c)^n$ where $i+j+k=n$ is $\frac{n!}{i!\,j!\,k!}$.

Calculation: For $a^2b^3c^1$ (so $i=2, j=3, k=1, n=6$):

$$\binom{6}{2,3,1} = \frac{6!}{2!\cdot 3!\cdot 1!} = \frac{720}{2\cdot 6\cdot 1} = \frac{720}{12} = 60.$$

Why other options are wrong:

Q10.

- (A) 15: This is $\binom{6}{2}$, ignoring the multinomial structure.
- (B) 20: This is $\binom{6}{3}$.
- (C) 30: Off by a factor of 2; possibly using $3!$ instead of $2!\cdot 3!$.

Answer: (D) ↑ Go Back to Q10

Solution

Concept: Markov's Inequality.

Statement: For a non-negative random variable X and $a>0$:

$$P(X \geq a) \leq \frac{E[X]}{a}.$$

Calculation: With $E[X]=5$ and $a=10$:

$$P(X \geq 10) \leq \frac{5}{10} = 0.5.$$

Why other options are wrong:

Q11.

- (B) 0.1: Would require $E[X]=1$; Markov gives 0.5, not 0.1.
- (C) 0.25: Would require $E[X]=2.5$ with $a=10$.
- (D) 1: Trivially true but not the best bound Markov gives.

Answer: (A) ↑ Go Back to Q11

Solution

Concept: Proper divisors and perfect numbers.

Step 1: Find all divisors of 12: $1, 2, 3, 4, 6, 12$.

Step 2: Proper divisors (all divisors except 12 itself): $\{1, 2, 3, 4, 6\}$.

Step 3: Sum: $1 + 2 + 3 + 4 + 6 = 16$.

Note: Since $16 \neq 12$, the number 12 is *not* perfect (compare with 6: $1 + 2 + 3 = 6$, which is perfect).

Why other options are wrong:

Q12.

- (A) 12: Would make 12 a perfect number, but it is not.
- (C) 28: This is the sum of proper divisors of 28 ($1 + 2 + 4 + 7 + 14 = 28$), a perfect number.
- (D) 14: Misses the divisor 2 or another; arithmetic error.

Answer: (B) ↑ Go Back to Q12

Solution

Concept: Equivalence relations and Bell numbers.

Key fact: The number of distinct equivalence relations on a set of n elements equals the n-th Bell number B_n. This is because each equivalence relation corresponds to a unique partition of the set.

For $n = 3$: The partitions of $\{1, 2, 3\}$ are:

$$\{\{1, 2, 3\}\}, \quad \{\{1\}, \{2, 3\}\}, \quad \{\{2\}, \{1, 3\}\}, \quad \{\{3\}, \{1, 2\}\}, \quad \{\{1\}, \{2\}, \{3\}\}.$$

Count: 5 partitions, so $B_3 = 5$.

Why other options are wrong:

Q13.

- (A) 3: Counts only partitions with one singleton block.
- (B) 4: Off by one; misses either the discrete or the indiscrete partition.
- (D) 6: One too many; only 5 distinct partitions exist.

Answer: (C) ↑ Go Back to Q13

Solution

Concept: Semantic tableaux (analytic tableau) for propositional logic.

Setup: We have premises $P \vee Q$ and $\neg P$; we wish to derive the conclusion.

Step 1: Begin the tableau with both premises as assumptions.

Step 2: Branch on $P \vee Q$:

Q14.

- **Left branch**: Add P. This contradicts $\neg P$ — branch closes ($\times$).
- **Right branch**: Add Q. No contradiction — branch remains open.

Step 3: Since all branches that don't contain Q are closed, the formula Q is provable from the premises.

Why other options are wrong:

- (A) $P \vee Q$: This is a premise, not the conclusion derived.
- (B) $\neg Q$: Contradicts the open branch containing Q.
- (C) P: Contradicted by $\neg P$ (closed branch).

Answer: (D) ↑ Go Back to Q14

Q15.

Solution

Concept: Solving linear homogeneous recurrences via characteristic equations.

Step 1: Write the characteristic equation for $T(n) = 5T(n-1) - 6T(n-2)$:

$$r^2 - 5r + 6 = 0.$$

Step 2: Factor: $(r-2)(r-3) = 0$, giving distinct roots $r_1 = 2$ and $r_2 = 3$.

Step 3: Since the roots are distinct, the general solution is:

$$T(n) = \alpha \cdot 2^n + \beta \cdot 3^n,$$

where α and β are determined by initial conditions.

Why other options are wrong:

- (B) $\alpha \cdot 3^n + \beta \cdot 5^n$: 5 is a coefficient, not a root.
- (C) $\alpha \cdot (5/2)^n$: No such root exists from the characteristic equation.
- (D) $\alpha \cdot 6^n$: 6 is the constant term's product $r_1 \cdot r_2$, not a root itself.

Answer: (A) ↑ Go Back to Q15

Solution

Concept: Prim's Minimum Spanning Tree (MST) algorithm.

Graph edges: 1–2: 3, 1–3: 6, 2–3: 4, 2–4: 5, 3–4: 2.

Step 1: Start at node 1. Cheapest edge from $\{1\}$: edge 1–2 (weight 3). Add it. MST set: $\{1, 2\}$.

Step 2: Cheapest edge from $\{1, 2\}$ to outside: 2–3 (weight 4), 2–4 (weight 5), 1–3 (weight 6). Choose 2–3 (weight 4). MST set: $\{1, 2, 3\}$.

Step 3: Cheapest edge from $\{1, 2, 3\}$ to outside: 3–4 (weight 2), 2–4 (weight 5). Choose 3–4 (weight 2). MST set: $\{1, 2, 3, 4\}$.

Total MST weight: $3 + 4 + 2 =$ **9**.

Why other options are wrong:

Q16.

- (A) 8: No combination of 3 spanning edges sums to 8.
- (C) 10: Would include a heavier edge like 2–4 (5) instead of 2–3 (4).
- (D) 11: Even heavier edge selection.

Answer: (B) ↑ Go Back to Q16

Solution

Concept: B-tree (minimum degree t) node capacity constraints.

Definition: In a B-tree with minimum degree t:

Q17.

- Every non-root node has **at least** $t - 1$ keys and at most $2t - 1$ keys.
- The root has at least 1 key (and at most $2t - 1$ keys).

For $t = 2$: Minimum number of keys in any non-root node $= t - 1 = 2 - 1 = 1$.

This means each non-root node (leaf or internal) holds at least 1 and at most 3 keys, giving rise to nodes with 2, 3, or 4 children – hence the name "2-3-4 tree."

Why other options are wrong:

- (A) 2: This equals t, not $t - 1$; it would be the minimum number of *children*.
- (B) 0: An empty non-root node violates the B-tree property.
- (D) 3: This is $2t - 1$, the *maximum* number of keys.

Answer: (C) ↑ Go Back to Q17

Solution

Concept: Chomsky hierarchy classification of $L = \{a^n b^n c^n \mid n \geq 1\}$.

Not Regular: A DFA cannot count unbounded n; regular languages cannot enforce the triple equality.

Not Context-Free: By the Context-Free Pumping Lemma, pumping any substring cannot simultaneously maintain the equal counts of a's, b's, and c's. (Any single pumped block can touch at most two of the three symbols, breaking the third count.)

Decidable: A multi-tape Turing machine can count a's, then decrement for b's, then decrement for c's, and accept iff all counters reach zero simultaneously. Hence L is decidable (recursive).

Classification: L is in the class of Context-Sensitive Languages (and hence decidable), but not context-free.

Why other options are wrong:

Q18.

- (A) Regular: Fails; no DFA suffices.
- (B) and (C) Context-free: Fails the CFL pumping lemma.

Answer: (D) ↑ Go Back to Q18

Solution

Concept: Longest Increasing Subsequence (LIS) via patience sorting.

Sequence: $[2, 1, 5, 3, 6, 4, 8, 9, 7]$.

Patience sorting trace (each pile shows its current top):

Element	Piles (tops)
2	$[2]$
1	$[1]$
5	$[1][5]$
3	$[1][3]$
6	$[1][3][6]$
4	$[1][3][4]$
8	$[1][3][4][8]$
9	$[1][3][4][8][9]$
7	$[1][3][4][7][9]$

(7 replaces 8 since $7 < 8$ and $7 > 4$.)

LIS length: Number of piles $=$ **5**.

One such LIS: $1, 3, 4, 8, 9$ (or $1, 3, 6, 8, 9$, or $1, 3, 4, 7, 9$).

Why other options are wrong:

Q19.

- (B) 4: Undercounts; a subsequence of length 5 clearly exists.
- (C) 6: Overcounts; no 6-element increasing subsequence exists.
- (D) 3: Far too short.

Answer: (A) ↑ Go Back to Q19

Q20.

Solution

Concept: Expected probes in linear-probing hash tables.

Formula for unsuccessful search (Knuth's analysis):

$$E[\text{unsuccessful probes}] = \frac{1}{2}\left(1 + \frac{1}{(1-\alpha)^2}\right).$$

Calculation with $\alpha = 0.5$:

$$= \frac{1}{2}\left(1 + \frac{1}{(0.5)^2}\right) = \frac{1}{2}(1+4) = \frac{5}{2} = 2.5.$$

Why other options are wrong:

- (A) 2.0: Corresponds to successful search with $\alpha = 0.5$: $\frac{1}{2}(1 + \frac{1}{1-\alpha}) = \frac{1}{2}(1+2) = 1.5$; not 2.
- (C) 3.0: Requires $\alpha \approx 0.577$.
- (D) 4.0: Requires $\alpha \approx 0.667$.

Answer: (B) ↑ Go Back to Q20

Q21.

Solution

Concept: Rice's theorem for undecidability of semantic TM properties.

Rice's Theorem: Any non-trivial property of the *language* recognized by a Turing machine is undecidable. A property is non-trivial if it holds for some TM languages and fails for others.

Application: The property "M accepts at least one string" is equivalent to asking whether $L(M) \neq \emptyset$ (the non-emptiness property). This is:

- **Non-trivial**: The empty TM has $L(M) = \emptyset$ (property fails), while any TM that accepts a single string has $L(M) \neq \emptyset$ (property holds).
- **Semantic**: It depends only on $L(M)$, not the structure of M.

By Rice's theorem, $L = \{\langle M \rangle \mid L(M) \neq \emptyset\}$ is **undecidable**.

Why other options are wrong:

- (A) and (B): No finite or dovetailing procedure can decide this for all TMs.
- (D): Undecidability applies regardless of whether the language is finite.

Answer: (C) ↑ Go Back to Q21

Q22.

Solution

Concept: Effective Memory Access Time (EMAT) with TLB.

Parameters: $h = 0.9$, $t_{\text{TLB}} = 2\,\text{ns}$, $t_m = 10\,\text{ns}$.

Formula:

$$\text{EMAT} = h\,(t_{\text{TLB}} + t_m) + (1-h)\,(t_{\text{TLB}} + 2t_m).$$

Calculation:

$$= 0.9 \times (2+10) + 0.1 \times (2+20) = 0.9 \times 12 + 0.1 \times 22 = 10.8 + 2.2 = 13\,\text{ns}.$$

Why other options are wrong:

- (B) 10 ns: Ignores TLB access time entirely.
- (C) 12 ns: Assumes 100% hit rate ($h = 1$): $2 + 10 = 12$.
- (D) 22 ns: Only counts the miss case ($t_{\text{TLB}} + 2t_m = 22$), ignoring hits.

Answer: (A) ↑ Go Back to Q22

Q23.

Solution

Concept: Syntax-Directed Definitions (SDD) with synthesized attributes for arithmetic.

Input: `3 + 4 * 2`.

Parse tree structure (with standard precedence, $*$ binds tighter than $+$):

$$E \to E_1 + T, \quad T \to T_1 * F, \quad F \to \text{digit}.$$

Bottom-up evaluation:

(a) $F_1.val = 3$ (the digit 3).

(b) $T_2.val = 4 * 2 = 8$ (from $T_1 * F$ where $T_1.val = 4$, $F.val = 2$).

(c) $E.val = 3 + 8 = 11$.

Why other options are wrong:

- (A) 14: Evaluates left-to-right ignoring precedence: $(3+4) * 2 = 14$.
- (B) 20: Arithmetic error.
- (C) 7: Adds all digits: $3+4+2 = 9$? Still wrong; no standard grammar gives 7.

Answer: (D) ↑ Go Back to Q23

Solution

Concept: Many-one (mapping) reducibility and decidability.

Definition: $A \leq_m B$ means there exists a computable total function f such that for all x: $x \in A \iff f(x) \in B$.

Key theorem: If $A \leq_m B$ and B is decidable, then A is also decidable.

Proof sketch: To decide membership in A on input x: (1) compute $f(x)$ (always terminates since f is total computable), then (2) run the decider for B on $f(x)$. The composed procedure halts and correctly decides A.

Why other options are wrong:

Q24.

- (A) Undecidable: Wrong direction; decidability of B "pulls down" to A via $\leq_m$.
- (C) Recognizable but not decidable: The reduction to a decidable set ensures full decidability.
- (D) Nothing concluded: The theorem gives a definitive conclusion.

Answer: (B) ↑ Go Back to Q24

Solution

Concept: Maximum flow via Ford-Fulkerson / augmenting paths.

Network: $s \to a$: 10, $s \to b$: 10, $a \to b$: 2, $a \to t$: 10, $b \to t$: 10.

Step 1: Augmenting path $s \to a \to t$: bottleneck = $\min(10, 10) = 10$. Send 10 units.

Step 2: Augmenting path $s \to b \to t$: bottleneck = $\min(10, 10) = 10$. Send 10 units.

Step 3: No more augmenting paths (all paths from s to t are saturated). Total flow = $10 + 10 =$ **20**.

Min-cut verification: The cut $\{s\} \mid \{a, b, t\}$ has capacity $10 + 10 = 20$. By the max-flow min-cut theorem, max flow $= 20$.

Why other options are wrong:

Q25.

- (A) 10: Only one path used.
- (B) 15: No such bottleneck exists in this graph.
- (D) 22: Exceeds the minimum cut capacity of 20.

Answer: (C) ↑ Go Back to Q25

Solution

Concept: Peterson's algorithm for mutual exclusion.

Algorithm overview: Peterson's solution for 2 processes (P_0 and P_1) uses exactly two shared variables:

Q26.

- `flag[2]`: A boolean array where `flag[i] = true` signals that process P_i wants to enter its critical section.
- `turn`: An integer (0 or 1) indicating whose turn it is to enter the critical section when there is contention.

Each process P_i sets `flag[i] = true` and `turn = 1-i` (yields to the other), then busy-waits while `flag[1-i] && turn == 1-i`.

Why other options are wrong:

- (A) `mutex` and `flag[]`: Peterson uses no mutex primitive; it constructs mutual exclusion from scratch.
- (B) Semaphore and `flag[]`: Peterson requires no OS-level semaphore.
- (C) `turn` only: `turn` alone gives strict alternation, not the full Peterson guarantee.

Answer: (D) ↑ Go Back to Q26

Solution

Concept: Implementing XOR using a 2:1 multiplexer.

Truth table analysis for $f = A \oplus B$:

A	B	f
0	0	0
0	1	1
1	0	1
1	1	0

MUX with Select $= A$**:**

Q27.

- When $A = 0$: output $= D_0$. For $f = B$, set $D_0 = B$.
- When $A = 1$: output $= D_1$. For $f = \overline{B}$, set $D_1 = \overline{B}$.

Connection: **Select** $= A$, $D_0 = B$, $D_1 = \overline{B}$ correctly implements $A \oplus B$.

Why other options are wrong:

- (B) Select $= B$: Uses B as select; when $B = 0$ outputs $D_0 = A$, when $B = 1$

outputs $D_1 = \overline{A}$. This also gives XOR, but the question asks specifically with A as select.

- (C) $D_0 = 0$, $D_1 = B$: When $A = 0$, output $= 0 \neq B$; incorrect.
- (D) $D_0 = \overline{B}$, $D_1 = B$: When $A = 0$, output $= \overline{B}$, but $f(0, B) = B$; wrong.

Answer: (A) ↑ Go Back to Q27

Q28.

Solution

Concept: Conflict serializability and precedence graphs.

Definition: Two operations *conflict* if they belong to different transactions, access the same data item, and at least one is a write. A schedule S is *conflict-serializable* if it can be transformed into a serial schedule by swapping non-conflicting operations.

Theorem: A schedule S is conflict-serializable if and only if its conflict (precedence) graph is **acyclic** (a Directed Acyclic Graph, DAG). The topological sort of the DAG gives a valid serial order.

Why other options are wrong:

- (A) Bipartite: Bipartiteness is irrelevant to serializability; the graph is directed.
- (C) Strongly connected: A strongly connected graph has cycles; it would imply *non*-serializability.
- (D) Complete: Completeness is also irrelevant and generally inconsistent with acyclicity on ≥ 2 nodes.

Answer: (B) ↑ Go Back to Q28

Q29.

Solution

Concept: System F (polymorphic λ-calculus) and the polymorphic identity.

Type $\forall\alpha.\, \alpha \to \alpha$**:** This type says "for any type α, given an α, return an α." The only (canonical) inhabitant is the polymorphic identity function.

Term $\Lambda\alpha.\, \lambda x{:}\alpha.\, x$**:** This is the polymorphic identity. $\Lambda\alpha$ introduces a type abstraction (type-level λ), and $\lambda x{:}\alpha.\, x$ is the term-level identity at type α. Applying it to a type τ yields $\lambda x{:}\tau.\, x$, which has type $\tau \to \tau$.

Why other options are wrong:

- (A) $\lambda x.\, x\, x$: Self-application; has type $(\alpha \to \beta) \to \beta$ (or is untypeable in simple λ-calculus); not $\forall\alpha.\, \alpha \to \alpha$.

- (B) $\lambda x.\, \lambda y.\, x$: This is the **K** combinator with type $\alpha \to \beta \to \alpha$; not $\forall \alpha.\, \alpha \to \alpha$.
- (C) $\lambda \alpha.\, \lambda x.\, x + 1$: Uses arithmetic $(+1)$, so x is forced to be a number; not polymorphic.

Answer: (D) ↑ Go Back to Q29

Q30.

Solution

Concept: FPTAS for 0/1 Knapsack.

Idea: The standard DP for Knapsack runs in $O(n \cdot P)$ where P is the maximum total profit. To get an FPTAS with approximation ratio $(1 - \varepsilon)$:

(a) Set scaling factor $K = \lfloor \varepsilon P_{\max}/n \rfloor$ where $P_{\max}$ is the maximum item profit.

(b) Scale each profit: $p'_i = \lfloor p_i/K \rfloor$.

(c) The scaled profits are at most $\lfloor P_{\max}/K \rfloor \leq n/\varepsilon$.

(d) Run the DP on scaled profits: time $O(n \times n/\varepsilon) = O(n^2/\varepsilon)$.

Running time: $O(n^2/\varepsilon)$, which is polynomial in both n and $1/\varepsilon$ — confirming it is a fully polynomial-time approximation scheme.

Why other options are wrong:

- (A) $O(n/\varepsilon)$: Too optimistic; DP has a factor of n for the table size and another for items.
- (B) $O(n \cdot P/\varepsilon)$: This is the pseudo-polynomial DP time without scaling.
- (D) $O(2^n)$: Exponential; FPTAS avoids this entirely.

Answer: (C) ↑ Go Back to Q30

Answer Key

Part A: Common (Q1–Q15)

Q	Ans	Q	Ans	Q	Ans	Q	Ans	Q	Ans
1	C	2	A	3	D	4	B	5	C
6	A	7	D	8	B	9	C	10	D
11	A	12	B	13	C	14	D	15	A

Part B: Computer Science (Q16–Q30)

Q	Ans	Q	Ans
16	B	17	C
18	D	19	A
20	B	21	C
22	A	23	D
24	B	25	C
26	D	27	A
28	B	29	D
30	C		

Scan the QR code on each page footer to access more TIFR GS sample papers and solutions on Collegedunia.