

General Aptitude (GA)

Q.1 – Q.5 Carry ONE mark Each

Q.1 If '→' denotes increasing order of intensity, then the meaning of the words [walk → jog → sprint] is analogous to [bothered → _____ → daunted]. Which one of the given options is appropriate to fill the blank?

- (A) phased
- (B) phrased
- (C) fazed
- (D) fused

Q.2 Two wizards try to create a spell using all the four elements, *water*, *air*, *fire*, and *earth*. For this, they decide to mix all these elements in all possible orders. They also decide to work independently. After trying all possible combination of elements, they conclude that the spell does not work.

How many attempts does each wizard make before coming to this conclusion, independently?

- (A) 24
- (B) 48
- (C) 16
- (D) 12

Q.3 In an engineering college of 10,000 students, 1,500 like neither their core branches nor other branches. The number of students who like their core branches is $1/4^{\text{th}}$ of the number of students who like other branches. The number of students who like both their core and other branches is 500.

The number of students who like their core branches is

- (A) 1,800
- (B) 3,500
- (C) 1,600
- (D) 1,500

Q.4 For positive non-zero real variables x and y , if

$$\ln\left(\frac{x+y}{2}\right) = \frac{1}{2}[\ln(x) + \ln(y)]$$

then, the value of $\frac{x}{y} + \frac{y}{x}$ is

- (A) 1
- (B) $1/2$
- (C) 2
- (D) 4

Q.5 In the sequence 6, 9, 14, x , 30, 41, a possible value of x is

(A) 25

(B) 21

(C) 18

(D) 20

GATE 2024

Q.6 – Q.10 Carry TWO marks Each

Q.6 Sequence the following sentences in a coherent passage.

P: This fortuitous geological event generated a colossal amount of energy and heat that resulted in the rocks rising to an average height of 4 km across the contact zone.

Q: Thus, the geophysicists tend to think of the Himalayas as an active geological event rather than as a static geological feature.

R: The natural process of the cooling of this massive edifice absorbed large quantities of atmospheric carbon dioxide, altering the earth's atmosphere and making it better suited for life.

S: Many millennia ago, a breakaway chunk of bedrock from the Antarctic Plate collided with the massive Eurasian Plate.

(A) QPSR

(B) QSPR

(C) SPRQ


(D) SRPQ

Q.7 A person sold two different items at the same price. He made 10% profit in one item, and 10% loss in the other item. In selling these two items, the person made a total of

- (A) 1% profit
- (B) 2% profit
- (C) 1% loss
- (D) 2% loss

GATE 2024

Q.8 The pie charts depict the shares of various power generation technologies in the total electricity generation of a country for the years 2007 and 2023.



The renewable sources of electricity generation consist of Hydro, Solar and Wind. Assuming that the total electricity generated remains the same from 2007 to 2023, what is the percentage increase in the share of the renewable sources of electricity generation over this period?

- (A) 25%
- (B) 50%
- (C) 77.5%
- (D) 62.5%

Q.9 A cube is to be cut into 8 pieces of equal size and shape. Here, each cut should be straight and it should not stop till it reaches the other end of the cube.

The minimum number of such cuts required is

- (A) 3
- (B) 4
- (C) 7
- (D) 8

GATE 2024

- Q.10 In the 4×4 array shown below, each cell of the first three rows has either a cross (X) or a number.

| | | | |
|---|---|---|---|
| 1 | X | 4 | 3 |
| X | 5 | 5 | 4 |
| 3 | X | 6 | X |
| | | | |

The number in a cell represents the count of the immediate neighboring cells (left, right, top, bottom, diagonals) NOT having a cross (X). Given that the last row has no crosses (X), the sum of the four numbers to be filled in the last row is

- (A) 11
(B) 10
(C) 12
(D) 9

Q.11 – Q.35 Carry ONE mark Each

Q.11 Consider a computer with a 4 MHz processor. Its DMA controller can transfer 8 bytes in 1 cycle from a device to main memory through cycle stealing at regular intervals. Which one of the following is the data transfer rate (*in bits per second*) of the DMA controller if 1% of the processor cycles are used for DMA?

- (A) 2,56,000
- (B) 3,200
- (C) 25,60,000
- (D) 32,000

Q.12 Let p and q be the following propositions:

p : Fail grade can be given.

q : Student scores more than 50% marks.

Consider the statement: “Fail grade cannot be given when student scores more than 50% marks.”

Which one of the following is the CORRECT representation of the above statement in propositional logic?

- (A) $q \rightarrow \neg p$
- (B) $q \rightarrow p$
- (C) $p \rightarrow q$
- (D) $\neg p \rightarrow q$

- Q.13 Consider the following C program. Assume parameters to a function are evaluated from *right to left*.

```
#include <stdio.h>

int g(int p) { printf("%d", p); return p; }
int h(int q) { printf("%d", q); return q; }

void f(int x, int y) {
    g(x);
    h(y);
}

int main() {
    f(g(10), h(20));
}
```

Which one of the following options is the **CORRECT** output of the above C program?

- (A) 20101020
- (B) 10202010
- (C) 20102010
- (D) 10201020

Q.14 The format of a single-precision floating-point number as per the IEEE 754 standard is:

| Sign (1bit) | Exponent (8 bits) | Mantissa (23 bits) |
|----------------|----------------------|-----------------------|
|----------------|----------------------|-----------------------|

Choose the largest floating-point number among the following options.

(A)

| Sign | Exponent | Mantissa |
|------|-----------|------------------------------|
| 0 | 0111 1111 | 1111 1111 1111 1111 1111 111 |

(B)

| Sign | Exponent | Mantissa |
|------|-----------|------------------------------|
| 0 | 1111 1110 | 1111 1111 1111 1111 1111 111 |

(C)

| Sign | Exponent | Mantissa |
|------|-----------|------------------------------|
| 0 | 1111 1111 | 1111 1111 1111 1111 1111 111 |

(D)

| Sign | Exponent | Mantissa |
|------|-----------|------------------------------|
| 0 | 0111 1111 | 0000 0000 0000 0000 0000 000 |

Q.15 Let $T(n)$ be the recurrence relation defined as follows:

$$T(0) = 1,$$

$$T(1) = 2, \text{ and}$$

$$T(n) = 5T(n-1) - 6T(n-2) \quad \text{for } n \geq 2$$

Which one of the following statements is TRUE?

(A) $T(n) = \Theta(2^n)$

(B) $T(n) = \Theta(n2^n)$

(C) $T(n) = \Theta(3^n)$

(D) $T(n) = \Theta(n3^n)$

Q.16 Let $f(x)$ be a continuous function from \mathbb{R} to \mathbb{R} such that

$$f(x) = 1 - f(2 - x)$$

Which one of the following options is the CORRECT value of $\int_0^2 f(x) dx$?

(A) 0

(B) 1

(C) 2

(D) -1

Q.17 Let A be the adjacency matrix of a simple undirected graph G . Suppose A is its own inverse. Which one of the following statements is *always* TRUE?

- (A) G is a cycle
- (B) G is a perfect matching
- (C) G is a complete graph
- (D) There is no such graph G

Q.18 When six unbiased dice are rolled simultaneously, the probability of getting all distinct numbers (*i.e.*, 1, 2, 3, 4, 5, and 6) is

- (A) $\frac{1}{324}$
- (B) $\frac{5}{324}$
- (C) $\frac{7}{324}$
- (D) $\frac{11}{324}$

Q.19 Once the DBMS informs the user that a transaction has been successfully completed, its effect should persist even if the system crashes before all its changes are reflected on disk. This property is called

- (A) durability
- (B) atomicity
- (C) consistency
- (D) isolation

Q.20 In the context of owner and weak entity sets in the ER (Entity-Relationship) data model, which one of the following statements is TRUE?

- (A) The weak entity set **MUST** have total participation in the identifying relationship
- (B) The owner entity set **MUST** have total participation in the identifying relationship
- (C) Both weak and owner entity sets **MUST** have total participation in the identifying relationship
- (D) Neither weak entity set nor owner entity set **MUST** have total participation in the identifying relationship

Q.21 Consider the following two sets:

Set X

- P. Lexical Analyzer
- Q. Syntax Analyzer
- R. Intermediate Code Generator
- S. Code Optimizer


Set Y

- 1. Abstract Syntax Tree
- 2. Token
- 3. Parse Tree
- 4. Constant Folding

Which one of the following options is the CORRECT match from **Set X** to **Set Y** ?

- (A) P – 4; Q – 1; R – 3; S – 2
- (B) P – 2; Q – 3; R – 1; S – 4
- (C) P – 2; Q – 1; R – 3; S – 4
- (D) P – 4; Q – 3; R – 2; S – 1

Q.22 Which one of the following regular expressions is equivalent to the language accepted by the DFA given below?



- (A) $0^*1(0 + 10^*1)^*$
- (B) $0^*(10^*11)^*0^*$
- (C) $0^*1(010^*1)^*0^*$
- (D) $0(1 + 0^*10^*1)^*0^*$

Q.23 Node X has a TCP connection open to node Y. The packets from X to Y go through an intermediate IP router R. Ethernet switch S is the first switch on the network path between X and R. Consider a packet sent from X to Y over this connection.

Which of the following statements is/are TRUE about the destination IP and MAC addresses on this packet at the time it leaves X?

- (A) The destination IP address is the IP address of R
- (B) The destination IP address is the IP address of Y
- (C) The destination MAC address is the MAC address of S
- (D) The destination MAC address is the MAC address of Y

Q.24 Which of the following tasks is/are the responsibility/responsibilities of the memory management unit (MMU) in a system with paging-based memory management?

- (A) Allocate a new page table for a newly created process
- (B) Translate a virtual address to a physical address using the page table
- (C) Raise a trap when a virtual address is not found in the page table
- (D) Raise a trap when a process tries to write to a page marked with read-only permission in the page table

Q.25 Consider a process P running on a CPU. Which one or more of the following events will *always* trigger a context switch by the OS that results in process P moving to a non-running state (e.g., ready, blocked)?

- (A) P makes a blocking system call to read a block of data from the disk
- (B) P tries to access a page that is in the swap space, triggering a page fault
- (C) An interrupt is raised by the disk to deliver data requested by some other process
- (D) A timer interrupt is raised by the hardware

Q.26 Which of the following file organizations is/are I/O efficient for the scan operation in DBMS?

- (A) Sorted
- (B) Heap
- (C) Unclustered tree index
- (D) Unclustered hash index

Q.27 Which of the following statements about the Two Phase Locking (2PL) protocol is/are TRUE?

- (A) 2PL permits only serializable schedules
- (B) With 2PL, a transaction always locks the data item being read or written just before every operation and always releases the lock just after the operation
- (C) With 2PL, once a lock is released on any data item inside a transaction, no more locks on any data item can be obtained inside that transaction
- (D) A deadlock is possible with 2PL

Q.28 Which of the following statements about IPv4 fragmentation is/are TRUE?

- (A) The fragmentation of an IP datagram is performed *only* at the source of the datagram
- (B) The fragmentation of an IP datagram is performed at any IP router which finds that the size of the datagram to be transmitted exceeds the MTU
- (C) The reassembly of fragments is performed *only* at the destination of the datagram
- (D) The reassembly of fragments is performed at all intermediate routers along the path from the source to the destination

Q.29 Which of the following statements is/are FALSE?

- (A) An attribute grammar is a syntax-directed definition (SDD) in which the functions in the semantic rules have no side effects
- (B) The attributes in a L-attributed definition cannot always be evaluated in a depth-first order
- (C) Synthesized attributes can be evaluated by a bottom-up parser as the input is parsed
- (D) All L-attributed definitions based on LR(1) grammar can be evaluated using a bottom-up parsing strategy

Q.30 For a Boolean variable x , which of the following statements is/are FALSE?

- (A) $x \cdot 1 = x$
- (B) $x + 1 = x$
- (C) $x \cdot x = 0$
- (D) $x + \bar{x} = 1$

Q.31 An instruction format has the following structure:

Instruction Number: *Opcode destination reg, source reg-1, source reg-2*

Consider the following sequence of instructions to be executed in a pipelined processor:

I1: DIV R3, R1, R2

I2: SUB R5, R3, R4

I3: ADD R3, R5, R6

I4: MUL R7, R3, R8

Which of the following statements is/are TRUE?

- (A) There is a RAW dependency on R3 between I1 and I2
- (B) There is a WAR dependency on R3 between I1 and I3
- (C) There is a RAW dependency on R3 between I2 and I3
- (D) There is a WAW dependency on R3 between I3 and I4

Q.32 Which of the following fields of an IP header is/are *always* modified by any router before it forwards the IP packet?

- (A) Source IP Address
- (B) Protocol
- (C) Time to Live (TTL)
- (D) Header Checksum

Q.33 Consider the following C function definition.

```
int fX(char *a) {  
    char *b = a;  
    while(*b)  
        b++;  
    return b - a; }
```

Which of the following statements is/are TRUE?

- (A) The function call `fX("abcd")` will always return a value
- (B) Assuming a character array `c` is declared as `char c[] = "abcd"` in `main()`, the function call `fX(c)` will always return a value
- (C) The code of the function will not compile
- (D) Assuming a character pointer `c` is declared as `char *c = "abcd"` in `main()`, the function call `fX(c)` will always return a value

Q.34 Let P be the partial order defined on the set $\{1,2,3,4\}$ as follows

$$P = \{(x, x) \mid x \in \{1,2,3,4\}\} \cup \{(1,2), (3,2), (3,4)\}$$

The number of total orders on $\{1,2,3,4\}$ that contain P is _____

Q.35 Let A be an array containing integer values. The distance of A is defined as the minimum number of elements in A that must be replaced with another integer so that the resulting array is sorted in non-decreasing order. The distance of the array $[2, 5, 3, 1, 4, 2, 6]$ is _____

GATE 2024

Q.36 – Q.65 Carry TWO marks Each

Q.36 What is the output of the following C program?

```
#include <stdio.h>

int main() {

    double a[2]={20.0, 25.0}, *p, *q;

    p = a;

    q = p + 1;

    printf("%d,%d", (int)(q - p), (int)(*q - *p));

    return 0;}
```

- (A) 4,8
- (B) 1,5
- (C) 8,5
- (D) 1,8

Q.37 Consider a single processor system with four processes A, B, C, and D, represented as given below, where for each process the first value is its arrival time, and the second value is its CPU burst time.

A (0, 10), B (2, 6), C (4, 3), and D (6, 7).

Which one of the following options gives the average waiting times when preemptive Shortest Remaining Time First (SRTF) and Non-Preemptive Shortest Job First (NP-SJF) CPU scheduling algorithms are applied to the processes?

- (A) SRTF = 6, NP-SJF = 7
- (B) SRTF = 6, NP-SJF = 7.5
- (C) SRTF = 7, NP-SJF = 7.5
- (D) SRTF = 7, NP-SJF = 8.5

Q.38 Which one of the following CIDR prefixes exactly represents the range of IP addresses 10.12.2.0 to 10.12.3.255?

- (A) 10.12.2.0/23
- (B) 10.12.2.0/24
- (C) 10.12.0.0/22
- (D) 10.12.2.0/22

- Q.39 You are given a set V of distinct integers. A binary search tree T is created by inserting all elements of V one by one, starting with an empty tree. The tree T follows the convention that, at each node, all values stored in the left subtree of the node are smaller than the value stored at the node. You are not aware of the sequence in which these values were inserted into T , and you do not have access to T .

Which one of the following statements is TRUE?

- (A) Inorder traversal of T can be determined from V
- (B) Root node of T can be determined from V
- (C) Preorder traversal of T can be determined from V
- (D) Postorder traversal of T can be determined from V

Q.40 Consider the following context-free grammar where the start symbol is S and the set of terminals is $\{a,b,c,d\}$.

$$\begin{aligned}
 S &\rightarrow AaAb \mid BbBa \\
 A &\rightarrow cS \mid \epsilon \\
 B &\rightarrow dS \mid \epsilon
 \end{aligned}$$

The following is a partially-filled LL(1) parsing table.


| | a | b | c | d | $\$$ |
|-----|--------------------------|--------------------------|--------------------|--------------------|------|
| S | $S \rightarrow AaAb$ | $S \rightarrow BbBa$ | (1) | (2) | |
| A | $A \rightarrow \epsilon$ | (3) | $A \rightarrow cS$ | | |
| B | (4) | $B \rightarrow \epsilon$ | | $B \rightarrow dS$ | |

Which one of the following options represents the CORRECT combination for the numbered cells in the parsing table?

Note: In the options, "blank" denotes that the corresponding cell is empty.

- (A) (1) $S \rightarrow AaAb$ (2) $S \rightarrow BbBa$ (3) $A \rightarrow \epsilon$ (4) $B \rightarrow \epsilon$
- (B) (1) $S \rightarrow BbBa$ (2) $S \rightarrow AaAb$ (3) $A \rightarrow \epsilon$ (4) $B \rightarrow \epsilon$
- (C) (1) $S \rightarrow AaAb$ (2) $S \rightarrow BbBa$ (3) blank (4) blank
- (D) (1) $S \rightarrow BbBa$ (2) $S \rightarrow AaAb$ (3) blank (4) blank

Q.41 Let M be the 5-state NFA with ϵ -transitions shown in the diagram below.



Which one of the following regular expressions represents the language accepted by M ?

- (A) $(00)^* + 1(11)^*$
- (B) $0^* + (1 + 0(00)^*)(11)^*$
- (C) $(00)^* + (1 + (00)^*)(11)^*$
- (D) $0^+ + 1(11)^* + 0(11)^*$

Q.42 Consider an array X that contains n positive integers. A subarray of X is defined to be a sequence of array locations with consecutive indices.

The C code snippet given below has been written to compute the length of the longest subarray of X that contains at most two distinct integers. The code has two missing expressions labelled (P) and (Q) .

```
int first=0, second=0, len1=0, len2=0, maxlen=0;
for (int i=0; i < n; i++) {
    if (X[i] == first) {
        len2++; len1++;
    } else if (X[i] == second) {
        len2++;
        len1 = (P) ;
        second = first;
    } else {
        len2 = (Q) ;
        len1 = 1; second = first;
    }
    if (len2 > maxlen) {
        maxlen = len2;
    }
    first = X[i];
}
```

Which one of the following options gives the CORRECT missing expressions?

(Hint: At the end of the i -th iteration, the value of $len1$ is the length of the longest subarray ending with $X[i]$ that contains all equal values, and $len2$ is the length of the longest subarray ending with $X[i]$ that contains at most two distinct values.)

- (A) (P) $len1+1$ (Q) $len2+1$
- (B) (P) 1 (Q) $len1+1$
- (C) (P) 1 (Q) $len2+1$
- (D) (P) $len2+1$ (Q) $len1+1$

- Q.43 Consider the following expression: $x[i] = (p + r) * -s[i] + u/w$. The following sequence shows the list of triples representing the given expression, with entries missing for triples (1), (3), and (6).

| | | | |
|-----|--------|-----|-----|
| (0) | + | p | r |
| (1) | | | |
| (2) | uminus | (1) | |
| (3) | | | |
| (4) | / | u | w |
| (5) | + | (3) | (4) |
| (6) | | | |
| (7) | = | (6) | (5) |

Which one of the following options fills in the missing entries CORRECTLY?

- (A) (1) = [] $s\ i$ (3) * (0) (2) (6) [] = $x\ i$
- (B) (1) [] = $s\ i$ (3) - (0) (2) (6) = [] $x\ (5)$
- (C) (1) = [] $s\ i$ (3) * (0) (2) (6) [] = $x\ (5)$
- (D) (1) [] = $s\ i$ (3) - (0) (2) (6) = [] $x\ i$

Q.44 Let x and y be random variables, not necessarily independent, that take real values in the interval $[0,1]$. Let $z = xy$ and let the mean values of x, y, z be $\bar{x}, \bar{y}, \bar{z}$, respectively. Which one of the following statements is TRUE?

(A) $\bar{z} = \bar{x}\bar{y}$

(B) $\bar{z} \leq \bar{x}\bar{y}$

(C) $\bar{z} \geq \bar{x}\bar{y}$

(D) $\bar{z} \leq \bar{x}$

Q.45 The relation schema, $\text{Person}(\underline{\text{pid}}, \text{city})$, describes the city of residence for every person uniquely identified by pid . The following relational algebra operators are available: selection, projection, cross product, and rename.

To find the list of cities where at least 3 persons reside, using the above operators, the minimum number of cross product operations that must be used is

(A) 1

(B) 2

(C) 3

(D) 4

Q.46

Consider a multi-threaded program with two threads T1 and T2. The threads share two semaphores: s_1 (initialized to 1) and s_2 (initialized to 0). The threads also share a global variable x (initialized to 0). The threads execute the code shown below.

| | |
|---|---|
| <pre>// code of T1 wait(s1); x = x+1; print(x); wait(s2); signal(s1);</pre> | <pre>// code of T2 wait(s1); x = x+1; print(x); signal(s2); signal(s1);</pre> |
|---|---|

Which of the following outcomes is/are possible when threads T1 and T2 execute concurrently?


- (A) T1 runs first and prints 1, T2 runs next and prints 2
- (B) T2 runs first and prints 1, T1 runs next and prints 2
- (C) T1 runs first and prints 1, T2 does not print anything (deadlock)
- (D) T2 runs first and prints 1, T1 does not print anything (deadlock)

Q.47

Let A be an $n \times n$ matrix over the set of all real numbers \mathbb{R} . Let B be a matrix obtained from A by swapping two rows. Which of the following statements is/are TRUE?

- (A) The determinant of B is the negative of the determinant of A
- (B) If A is invertible, then B is also invertible
- (C) If A is symmetric, then B is also symmetric
- (D) If the trace of A is zero, then the trace of B is also zero

- Q.48 Let S1 and S2 be two stacks. S1 has capacity of 4 elements. S2 has capacity of 2 elements. S1 already has 4 elements: 100, 200, 300, and 400, whereas S2 is empty, as shown below.



Only the following three operations are available:

PushToS2: Pop the top element from S1 and push it on S2.

PushToS1: Pop the top element from S2 and push it on S1.

GenerateOutput: Pop the top element from S1 and output it to the user.

Note that the pop operation is not allowed on an empty stack and the push operation is not allowed on a full stack.

Which of the following output sequences can be generated by using the above operations?

- (A) 100, 200, 400, 300
- (B) 200, 300, 400, 100
- (C) 400, 200, 100, 300
- (D) 300, 200, 400, 100

Q.49 Which of the following is/are EQUAL to 224 in radix-5 (*i.e.*, base-5) notation?

(A) 64 in radix-10

(B) 100 in radix-8

(C) 50 in radix-16

(D) 121 in radix-7

GATE 2024


Q.50 Consider 4-variable functions f_1, f_2, f_3, f_4 expressed in sum-of-minterms form as given below.

$$f_1 = \sum(0,2,3,5,7,8,11,13)$$

$$f_2 = \sum(1,3,5,7,11,13,15)$$

$$f_3 = \sum(0,1,4,11)$$

$$f_4 = \sum(0,2,6,13)$$



With respect to the circuit given above, which of the following options is/are CORRECT?

- (A) $Y = \sum(0,1,2,11,13)$
- (B) $Y = \Pi(3,4,5,6,7,8,9,10,12,14,15)$
- (C) $Y = \sum(0,1,2,3,4,5,6,7)$
- (D) $Y = \Pi(8,9,10,11,12,13,14,15)$

Q.51 Let G be an undirected connected graph in which every edge has a positive integer weight. Suppose that every spanning tree in G has *even* weight. Which of the following statements is/are TRUE for *every* such graph G ?

- (A) All edges in G have even weight
- (B) All edges in G have even weight **OR** all edges in G have odd weight
- (C) In each cycle C in G , all edges in C have even weight
- (D) In each cycle C in G , either all edges in C have even weight **OR** all edges in C have odd weight

Q.52 Consider a context-free grammar G with the following 3 rules.

$$S \rightarrow aS, S \rightarrow aSbS, S \rightarrow c$$


Let $w \in L(G)$. Let $n_a(w)$, $n_b(w)$, $n_c(w)$ denote the number of times a, b, c occur in w , respectively. Which of the following statements is/are TRUE?

- (A) $n_a(w) > n_b(w)$
- (B) $n_a(w) > n_c(w) - 2$
- (C) $n_c(w) = n_b(w) + 1$
- (D) $n_c(w) = n_b(w) * 2$


- Q.53 Consider a disk with the following specifications: rotation speed of 6000 RPM, average seek time of 5 milliseconds, 500 sectors/track, 512-byte sectors. A file has content stored in 3000 sectors located *randomly* on the disk. Assuming average rotational latency, the total time (*in seconds, rounded off to 2 decimal places*) to read the entire file from the disk is _____
- Q.54 Consider a TCP connection operating at a point of time with the congestion window of size 12 MSS (Maximum Segment Size), when a timeout occurs due to packet loss. Assuming that all the segments transmitted in the next two RTTs (Round Trip Time) are acknowledged correctly, the congestion window size (*in MSS*) during the third RTT will be _____
- Q.55 Consider an Ethernet segment with a transmission speed of 10^8 bits/sec and a maximum segment length of 500 meters. If the speed of propagation of the signal in the medium is 2×10^8 meters/sec, then the minimum frame size (*in bits*) required for collision detection is _____
- Q.56 A functional dependency $F: X \rightarrow Y$ is termed as a *useful functional dependency* if and only if it satisfies all the following three conditions:
- X is not the empty set.
 - Y is not the empty set.
 - Intersection of X and Y is the empty set.
- For a relation R with 4 attributes, the total number of possible useful functional dependencies is _____
- Q.57 A processor with 16 general purpose registers uses a 32-bit instruction format. The instruction format consists of an opcode field, an addressing mode field, two register operand fields, and a 16-bit scalar field. If 8 addressing modes are to be supported, the maximum number of unique opcodes possible for every addressing mode is _____

Q.58 A non-pipelined instruction execution unit operating at 2 GHz takes an average of 6 cycles to execute an instruction of a program P. The unit is then redesigned to operate on a 5-stage pipeline at 2 GHz. Assume that the ideal throughput of the pipelined unit is 1 instruction per cycle. In the execution of program P, 20% instructions incur an average of 2 cycles stall due to data hazards and 20% instructions incur an average of 3 cycles stall due to control hazards. The speedup (*rounded off to one decimal place*) obtained by the pipelined design over the non-pipelined design is _____

Q.59 The number of distinct minimum-weight spanning trees of the following graph is _____



Q.60 The chromatic number of a graph is the minimum number of colours used in a *proper* colouring of the graph. The chromatic number of the following graph is _____



- Q.61 A processor uses a 32-bit instruction format and supports byte-addressable memory access. The ISA of the processor has 150 distinct instructions. The instructions are equally divided into two types, namely R-type and I-type, whose formats are shown below.

R-type Instruction Format:

| | | | | |
|--------|--------|--------------|---------------|----------------|
| OPCODE | UNUSED | DST Register | SRC Register1 | SRC Register 2 |
|--------|--------|--------------|---------------|----------------|

I-type Instruction Format:

| | | | |
|--------|--------------|--------------|---------------------------|
| OPCODE | DST Register | SRC Register | # Immediate value/address |
|--------|--------------|--------------|---------------------------|

In the OPCODE, 1 bit is used to distinguish between I-type and R-type instructions and the remaining bits indicate the operation. The processor has 50 architectural registers, and all register fields in the instructions are of equal size.

Let X be the number of bits used to encode the UNUSED field, Y be the number of bits used to encode the OPCODE field, and Z be the number of bits used to encode the immediate value/address field. The value of $X + 2Y + Z$ is _____

- Q.62 Let L_1 be the language represented by the regular expression $b^*ab^*(ab^*ab^*)^*$ and $L_2 = \{w \in (a+b)^* \mid |w| \leq 4\}$, where $|w|$ denotes the length of string w . The number of strings in L_2 which are also in L_1 is _____

- Q.63 Let Z_n be the group of integers $\{0, 1, 2, \dots, n-1\}$ with addition modulo n as the group operation. The number of elements in the group $Z_2 \times Z_3 \times Z_4$ that are their own inverses is _____

- Q.64 Consider a 32-bit system with 4 KB page size and page table entries of size 4 bytes each. Assume $1 \text{ KB} = 2^{10}$ bytes. The OS uses a 2-level page table for memory management, with the page table containing an outer page directory and an inner page table. The OS allocates a page for the outer page directory upon process creation. The OS uses demand paging when allocating memory for the inner page table, i.e., a page of the inner page table is allocated only if it contains at least one valid page table entry.

An active process in this system accesses 2000 unique pages during its execution, and none of the pages are swapped out to disk. After it completes the page accesses, let X denote the minimum and Y denote the maximum number of pages across the two levels of the page table of the process.

The value of $X+Y$ is _____

- Q.65 Consider the following augmented grammar, which is to be parsed with a SLR parser. The set of terminals is $\{a,b,c,d,\#,@\}$

$$\begin{aligned} S' &\rightarrow S \\ S &\rightarrow SS \mid Aa \mid bAc \mid Bc \mid bBa \\ A &\rightarrow d\# \\ B &\rightarrow @ \end{aligned}$$

Let $I_0 = \text{CLOSURE}(\{S' \rightarrow \bullet S\})$. The number of items in the set $\text{GOTO}(I_0, S)$ is _____

GATE 2024 CS and IT Set 2 Question Paper with Solution

Time Allowed :3 Hours

Maximum Marks :100

Total Questions :65

General Instructions

Read the following instructions very carefully and strictly follow them:

1. The GATE Exam will be structured with a total of 100 marks.
2. The exam mode is Online CBT (Computer Based Test)
3. The total duration of Exam is 3 Hours.
4. It will include 65 questions , divided in 3 sections.
5. Section 1 : General Aptitude.
6. Section 2 : Engineering Mathematics.
7. Section 3 : Subject Based Questions.
8. The marking scheme is as such : 1 and 2 marks Questions. Each correct answer will carry marks as specified in the question paper. Incorrect answers may carry negative marks, as indicated in the question paper.
9. Question Types: The exam will include a mix of Multiple Choice Questions (MCQs), Multiple Select Questions (MSQs), and Numerical Answer Type (NAT) questions.

GENERAL APTITUDE

Question 1-5 carry one mark each

1. If → denotes increasing order of intensity, then the meaning of the words [walk → jog → sprint] is analogous to [bothered → ----- → daunted].

- (A) phased
- (B) phrased
- (C) fazed
- (D) fused

Correct Answer: (C) fazed

Solution: The relationship described in the first analogy is based on increasing intensity:

walk → jog → sprint.

This progression moves from low intensity (walk) to high intensity (sprint). Similarly, in the second analogy:

bothered → ___ → daunted.

Step 1: Analyze the terms.

”Bothered” indicates mild disturbance, and ”daunted” represents a high level of emotional disturbance. The intermediate level is best represented by ”fazed.”

Step 2: Eliminate incorrect options.

- (1) ”phased” does not relate to emotional intensity.
- (2) ”phrased” is unrelated to emotions.
- (4) ”fused” does not fit the context.

Hence, (3) ”fazed” aligns with the progression.

Final Answer:

fazed

Quick Tip

When solving analogy questions, focus on the nature and progression of relationships between the terms in the first set to identify the correct answer in the second set.

2. Two wizards try to create a spell using all the four elements, water, air, fire, and earth. For this, they decide to mix all these elements in all possible orders. They also decide to work independently. After trying all possible combinations of elements, they conclude that the spell does not work.

How many attempts does each wizard make before coming to this conclusion, independently?

- (A) 24
- (B) 48
- (C) 16
- (D) 12

Correct Answer: (A) 24

Solution: The total number of possible orders for n elements is given by the factorial of n , denoted by $n!$.

Step 1: Calculate the number of orders for 4 elements.

The number of permutations of 4 elements is:

$$4! = 4 \times 3 \times 2 \times 1 = 24.$$

Step 2: Verify independent attempts.

Since both wizards work independently and test all possible permutations, each wizard makes 24 attempts.

Final Answer:

24

Quick Tip

The number of arrangements (permutations) of n elements is always $n!$. Use this for questions involving ordering.

3. In an engineering college of 10,000 students, 1,500 like neither their core branches nor other branches. The number of students who like their core branches is $\frac{1}{4}$ of the number of students who like other branches. The number of students who like both their core and other branches is 500.

The number of students who like their core branches is:

- (A) 1,800
- (B) 3,500
- (C) 1,600
- (D) 1,500

Correct Answer: (A) 1,800

Solution: Let Ω , C , and O be the set of total students (universal set), students who like their core branch, and students who like other branches, respectively. We are given:

$$|\Omega| = 10,000, \quad |C^c \cap O^c| = 1,500, \quad |O| = 4|C|, \quad |C \cap O| = 500.$$

We need to find $|C|$.

Step 1: Calculate $|C \cup O|$.

Using De Morgan's law:

$$|C \cup O| = |\Omega| - |C^c \cap O^c| = 10,000 - 1,500 = 8,500.$$

Step 2: Relation between $|C|$ and $|O|$.

From set theory:

$$|C \cup O| = |C| + |O| - |C \cap O|.$$

Substitute $|C \cup O| = 8,500$, $|O| = 4|C|$, and $|C \cap O| = 500$:

$$8,500 = |C| + 4|C| - 500.$$

Simplify:

$$5|C| = 8,500 + 500 = 9,000 \implies |C| = \frac{9,000}{5} = 1,800.$$

Final Answer:

1,800

Quick Tip

In set theory problems, use union and intersection relations systematically, and leverage De Morgan's laws for complements.

4. For positive non-zero real variables x and y , if

$$\ln\left(\frac{x+y}{2}\right) = \frac{1}{2}[\ln(x) + \ln(y)],$$

then the value of $\frac{x}{y} + \frac{y}{x}$ is:

- (A) 1
- (B) $\frac{1}{2}$
- (C) 2
- (D) 4

Correct Answer: (C) 2

Solution: The given equation is:

$$\ln\left(\frac{x+y}{2}\right) = \frac{1}{2}[\ln(x) + \ln(y)].$$

Step 1: Simplify using logarithmic properties.

Using the property of logarithms, $\ln(ab) = \ln(a) + \ln(b)$, the right-hand side can be rewritten as:

$$\frac{1}{2}[\ln(x) + \ln(y)] = \ln(\sqrt{xy}).$$

Equating both sides, we get:

$$\ln\left(\frac{x+y}{2}\right) = \ln(\sqrt{xy}).$$

Step 2: Exponentiate both sides.

By removing the logarithm, we have:

$$\frac{x + y}{2} = \sqrt{xy}.$$

Step 3: Solve for $\frac{x}{y} + \frac{y}{x}$.

Square both sides:

$$\left(\frac{x + y}{2}\right)^2 = xy \implies \frac{(x + y)^2}{4} = xy \implies x^2 + y^2 + 2xy = 4xy.$$

Simplify to find $x^2 + y^2$:

$$x^2 + y^2 = 2xy.$$

Thus:

$$\frac{x}{y} + \frac{y}{x} = \frac{x^2 + y^2}{xy} = \frac{2xy}{xy} = 2.$$

Final Answer:

2

Quick Tip

When dealing with logarithmic equations, always simplify using properties such as $\ln(ab) = \ln(a) + \ln(b)$ and exponentiate to eliminate the logarithm.

5. In the sequence 6, 9, 14, x , 30, 41, a possible value of x is:

- (A) 25
- (B) 21
- (C) 18
- (D) 20

Correct Answer: (B) 21

Solution: Step 1: Identify the pattern in the sequence.

The given sequence is 6, 9, 14, x , 30, 41. To find the missing value x , observe the differences

between consecutive terms:

$$9 - 6 = 3, \quad 14 - 9 = 5.$$

The differences are increasing by 2: 3, 5, ... The next difference should be 7:

$$x - 14 = 7 \implies x = 14 + 7 = 21.$$

Verify the pattern for the next terms:

$$30 - 21 = 9, \quad 41 - 30 = 11.$$

The differences 7, 9, 11 confirm the pattern.

Final Answer:

21

Quick Tip

When solving sequence problems, analyze the differences or ratios between consecutive terms to identify patterns.

Question 6-10 carry two marks each

6. Sequence the following sentences in a coherent passage.

P: This fortuitous geological event generated a colossal amount of energy and heat that resulted in the rocks rising to an average height of 4 km across the contact zone.

Q: Thus, the geophysicists tend to think of the Himalayas as an active geological event rather than as a static geological feature.

R: The natural process of the cooling of this massive edifice absorbed large quantities of atmospheric carbon dioxide, altering the earth's atmosphere and making it better suited for life.

S: Many millennia ago, a breakaway chunk of bedrock from the Antarctic Plate collided with the massive Eurasian Plate.

(A) QPSR

(B) QSPR

(C) SPRQ

(D) SRPQ

Correct Answer: (C) SPRQ

Solution: To form a coherent passage, the sentences need to be ordered logically and chronologically.

Step 1: Identify the starting point.

The passage begins with a historical event. Sentence **S** describes the initial collision between the Antarctic Plate and the Eurasian Plate.

Step 2: Follow with the resulting impact.

Sentence **P** explains the geological consequences of the collision described in **S**.

Step 3: Continue with further outcomes.

Sentence **R** details the natural cooling process and its environmental impact, logically following **P**.

Step 4: Conclude with the modern interpretation.

Sentence **Q** provides a current perspective on the geological activity, making it a suitable ending.

The coherent sequence is:

$$S \rightarrow P \rightarrow R \rightarrow Q$$

Final Answer:

SPRQ

Quick Tip

When sequencing sentences, look for chronological order or cause-effect relationships to form a logical passage.

7. A person sold two different items at the same price. He made 10% profit in one item, and 10% loss in the other item. In selling these two items, the person made a total of:

(A) 1% profit

(B) 2% profit

(C) 1% loss

(D) 2% loss

Correct Answer: (C) 1% loss

Solution: Let the selling price (SP) of each item be 100 units.

Step 1: Calculate the cost price (CP) for each item.

For the first item, there is a 10% profit:

$$CP_1 = \frac{SP_1}{1.1} = \frac{100}{1.1} \approx 90.91 \text{ units.}$$

For the second item, there is a 10% loss:

$$CP_2 = \frac{SP_2}{0.9} = \frac{100}{0.9} \approx 111.11 \text{ units.}$$

Step 2: Compute the total cost price and total selling price.

$$\text{Total CP} = CP_1 + CP_2 = 90.91 + 111.11 = 202.02 \text{ units.}$$

$$\text{Total SP} = 100 + 100 = 200 \text{ units.}$$

Step 3: Determine the overall profit or loss.

$$\text{Loss} = \text{Total CP} - \text{Total SP} = 202.02 - 200 = 2.02 \text{ units.}$$

$$\text{Loss percentage} = \frac{\text{Loss}}{\text{Total CP}} \times 100 = \frac{2.02}{202.02} \times 100 \approx 1\%.$$


Final Answer:

1% loss

Quick Tip

When equal selling prices are given, the overall result depends on the cost price difference and the percentage profit or loss for each item.

8. The pie charts depict the shares of various power generation technologies in the total electricity generation of a country for the years 2007 and 2023.



The renewable sources of electricity generation consist of Hydro, Solar, and Wind. Assuming that the total electricity generated remains the same from 2007 to 2023, what is the percentage increase in the share of the renewable sources of electricity generation over this period?

- (1) 25%
- (2) 50%
- (3) 77.5%
- (4) 62.5%

Correct Answer: (4) 62.5%

Solution: Step 1: Calculate the total share of renewable sources in 2007.

In 2007, the shares of renewable sources (Hydro, Solar, and Wind) are:

$$\text{Hydro: } 30\%, \quad \text{Solar: } 5\%, \quad \text{Wind: } 5\%.$$

The total share of renewable sources in 2007 is:

$$30\% + 5\% + 5\% = 40\%.$$

Step 2: Calculate the total share of renewable sources in 2023.

In 2023, the shares of renewable sources (Hydro, Solar, and Wind) are:

$$\text{Hydro: } 35\%, \quad \text{Solar: } 20\%, \quad \text{Wind: } 10\%.$$

The total share of renewable sources in 2023 is:

$$35\% + 20\% + 10\% = 65\%.$$

Step 3: Calculate the percentage increase in the share of renewable sources.

The increase in share is:

$$65\% - 40\% = 25\%.$$

The percentage increase relative to 2007 is:

$$\frac{25\%}{40\%} \times 100 = 62.5\%.$$

Final Answer:

(4) 62.5%

Quick Tip

To calculate percentage increases, always compare the absolute increase to the original value. Use this approach for analyzing changes in shares or percentages.

9. A cube is to be cut into 8 pieces of equal size and shape. Here, each cut should be straight, and it should not stop till it reaches the other end of the cube. The minimum number of such cuts required is:

- (A) 3
- (B) 4
- (C) 7
- (D) 8

Correct Answer: (A) 3

Solution: To divide a cube into 8 equal parts (smaller cubes), the cube must be cut along its three dimensions: length, breadth, and height.

Step 1: First Cut (Lengthwise)

The first cut divides the cube into 2 equal parts along its length.

Step 2: Second Cut (Breadthwise)

The second cut divides each of the 2 parts into 2 smaller parts, making a total of 4 parts.

Step 3: Third Cut (Heightwise)

The third cut divides each of the 4 parts into 2, resulting in 8 smaller cubes.

Thus, a total of 3 cuts are required to divide the cube into 8 equal pieces.

Final Answer:

3

Quick Tip

For dividing a cube into 2^n smaller cubes, the minimum number of cuts required is n , one cut along each dimension (length, breadth, and height).

10. In the 4×4 array shown below, each cell of the first three rows has either a cross (X) or a number.

| | | | |
|---|---|---|---|
| 1 | X | 4 | 3 |
| X | 5 | 5 | 4 |
| 3 | X | 6 | X |
| ? | ? | ? | ? |

The number in a cell represents the count of the immediate neighboring cells (left, right, top, bottom, diagonals) NOT having a cross (X). Given that the last row has no crosses (X), the sum of the four numbers to be filled in the last row is:

- (A) 11
- (B) 10
- (C) 12
- (D) 9

Correct Answer: (A) 11

Solution: The task is to determine the numbers in the last row such that each represents the count of immediate neighboring cells that do not contain a cross (X). The last row contains

no crosses (X), and its neighboring cells are the third row.

For the first cell in the last row (?), the neighbors are: 3, X. Only 3 is not a cross, so this cell gets 1.

For the second cell in the last row (?), the neighbors are: 3, X, 6. Both 3 and 6 are not crosses, so this cell gets 2.

For the third cell in the last row (?), the neighbors are: X, 6, X. Only 6 is not a cross, so this cell gets 1.

For the fourth cell in the last row (?), the neighbors are: 6, X. Only 6 is not a cross, so this cell gets 1.

The numbers in the last row are: 2, 4, 3, 2. Their sum is:

$$2 + 4 + 3 + 2 = 11.$$

Final Answer:

11

Quick Tip

For such array-based questions, carefully examine the neighboring cells for each cell and count only those that meet the conditions given in the question.

Question 11-35 carry one mark each

11. Consider a computer with a 4 MHz processor. Its DMA controller can transfer 8 bytes in 1 cycle from a device to main memory through cycle stealing at regular intervals. Which one of the following is the data transfer rate (in *bits per second*) of the DMA controller if 1% of the processor cycles are used for DMA?

- (A) 2, 56, 000
- (B) 3, 200
- (C) 25, 60, 000
- (D) 32, 000

Correct Answer: (C) 25,60,000

Solution: Processor speed = 4 MHz = 4×10^6 cycles per second.

DMA uses 1% of processor cycles, so the number of cycles used for DMA:

$$DMA_{cycles} = 1\% \times 4 \times 10^6 = 4 \times 10^4 \text{ cycles per second.}$$

Data transferred per DMA cycle = 8 bytes = $8 \times 8 = 64$ bits.

Total data transfer rate:

$$Rate = DMA \text{ cycles} \times \text{Data per cycle} = 4 \times 10^4 \times 64 = 25,60,000 \text{ bits per second.}$$

Final Answer:

25,60,000 bits per second

Quick Tip

For DMA transfer rate calculations, multiply the number of cycles allocated for DMA by the data transferred per cycle.

12. Let p and q be the following propositions:

p : Fail grade can be given.

q : Student scores more than 50% marks.

Consider the statement: “Fail grade cannot be given when student scores more than 50% marks.”

Which one of the following is the CORRECT representation of the above statement in propositional logic?

(A) $q \rightarrow \neg p$

(B) $q \rightarrow p$

(C) $p \rightarrow q$

(D) $\neg p \rightarrow q$

Correct Answer: (A) $q \rightarrow \neg p$

Solution: The given statement can be interpreted as:

If a student scores more than 50% marks (q), then a fail grade (p) cannot be given ($\neg p$).

This can be directly represented in propositional logic as:

$$q \rightarrow \neg p$$

Validation of options:

(A) $q \rightarrow \neg p$: Correct. Matches the interpretation of the given statement.

(B) $q \rightarrow p$: Incorrect. This implies that scoring $> 50\%$ leads to a fail grade.

(C) $p \rightarrow q$: Incorrect. This implies that failing implies scoring $> 50\%$, which is incorrect.

(D) $\neg p \rightarrow q$: Incorrect. This implies that not failing implies scoring $> 50\%$, which is not equivalent to the given statement.

Final Answer:

$$q \rightarrow \neg p$$

Quick Tip

To convert statements into propositional logic, carefully identify the conditional relationships and apply logical operators accordingly.

13. Consider the following C program. Assume parameters to a function are evaluated from right to left.

```
#include <stdio.h>

int g(int p) { printf("%d", p); return p; }
int h(int q) { printf("%d", q); return q; }

void f(int x, int y) {
    g(x);
    h(y);
}
```

```
int main() {  
    f(g(10), h(20));  
}
```

Which one of the following options is the CORRECT output of the above C program?

- (A) 20101020
- (B) 10202010
- (C) 20102010
- (D) 10201020

Correct Answer: (A) 20101020

Solution: The execution of the given C program proceeds as follows:

In the `main()` function, the function `f(g(10), h(20))` is called.

Parameters to the function `f()` are evaluated **from right to left**.

`h(20)` is evaluated first, printing 20 and returning 20.

`g(10)` is evaluated next, printing 10 and returning 10.

Therefore, the first part of the output is 20 followed by 10.

The function `f()` is then executed:

`g(x)` is called with $x = 10$, printing 10.

`h(y)` is called with $y = 20$, printing 20.

The second part of the output is 10 followed by 20.

Combining these parts, the total output is:

20101020

Final Answer:

20101020

Quick Tip

When function parameters are evaluated **right to left**, start with the last parameter, and carefully track the order of execution and the sequence of prints.

14. The format of a single-precision floating-point number as per the IEEE 754 standard is:

| | | |
|--------------|-------------------|--------------------|
| Sign (1 bit) | Exponent (8 bits) | Mantissa (23 bits) |
|--------------|-------------------|--------------------|

Choose the largest floating-point number among the following options:

- (A) **Sign:** 0, **Exponent:** 0111 1111, **Mantissa:** 1111 1111 1111 1111 1111 111
- (B) **Sign:** 0, **Exponent:** 1111 1110, **Mantissa:** 1111 1111 1111 1111 1111 111
- (C) **Sign:** 0, **Exponent:** 1111 1111, **Mantissa:** 1111 1111 1111 1111 1111 111
- (D) **Sign:** 0, **Exponent:** 0111 1111, **Mantissa:** 0000 0000 0000 0000 0000 000

Correct Answer: (B)

Solution: To determine the largest floating-point number, the following points are considered based on the IEEE 754 standard:

The value of a floating-point number is given by:

$$(-1)^{\text{Sign}} \times (1.\text{Mantissa}) \times 2^{\text{Exponent}-127}.$$

The larger the exponent and the closer the mantissa is to 1, the larger the number.

In the given options:

Option (A): Exponent = 127, Mantissa = All 1's.

Option (B): Exponent = 254, Mantissa = All 1's.

Option (C): Exponent = 255, Mantissa = All 1's (this represents *NaN*, not a valid number).

Option (D): Exponent = 127, Mantissa = All 0's.

Analysis:

Option (B) has the largest exponent (254) among the valid numbers, making it the largest floating-point number.

Option (C) represents *NaN* (Not a Number) due to an exponent of 255 and a non-zero mantissa.

Option (D) has the smallest value due to its mantissa being all 0's.

Final Answer:

(B)

Quick Tip

For IEEE 754 floating-point numbers, larger exponents and mantissas closer to 1 result in larger values. Be cautious about special cases like *NaN* and *Infinity*.

15. Let $T(n)$ be the recurrence relation defined as follows:

$$T(0) = 1, \quad T(1) = 2, \quad T(n) = 5T(n-1) - 6T(n-2) \quad \text{for } n \geq 2$$

Which one of the following statements is TRUE?

- (A) $T(n) = \Theta(2^n)$
- (B) $T(n) = \Theta(n2^n)$
- (C) $T(n) = \Theta(3^n)$
- (D) $T(n) = \Theta(n3^n)$

Correct Answer: (A) $T(n) = \Theta(2^n)$

Solution: The recurrence relation is:

$$T(n) = 5T(n-1) - 6T(n-2)$$

Step 1: Find the characteristic equation.

The characteristic equation is:

$$x^2 - 5x + 6 = 0$$

Factoring gives:

$$(x-3)(x-2) = 0$$

Thus, the roots are $r_1 = 3$ and $r_2 = 2$.

Step 2: General solution for the recurrence relation.

The general solution is:

$$T(n) = C_1(3^n) + C_2(2^n)$$

Using the initial conditions $T(0) = 1$ and $T(1) = 2$, we find:

$$C_1 = 0, \quad C_2 \neq 0$$

Step 3: Asymptotic behavior.

Since $C_1 = 0$, the solution reduces to:

$$T(n) = C_2(2^n)$$

Thus, the asymptotic complexity is:

$$T(n) = \Theta(2^n)$$

Final Answer:

$$\boxed{\Theta(2^n)}$$

Quick Tip

When solving recurrence relations, always find the characteristic equation and analyze the roots to determine the asymptotic behavior.

16. Let $f(x)$ be a continuous function from \mathbb{R} to \mathbb{R} such that:

$$f(x) = 1 - f(2 - x)$$

Which one of the following options is the CORRECT value of $\int_0^2 f(x) dx$?

- (A) 0
- (B) 1
- (C) 2
- (D) -1

Correct Answer: (B) 1

Solution: The given functional equation is:

$$f(x) + f(2 - x) = 1$$

Step 1: Split the integral.

The integral can be split as:

$$\int_0^2 f(x) dx = \int_0^1 f(x) dx + \int_1^2 f(x) dx$$

Step 2: Use substitution for symmetry.

In the second term, substitute $u = 2 - x$:

$$\int_1^2 f(x) dx = \int_0^1 f(2 - u) du$$

Using $f(2 - u) = 1 - f(u)$, this becomes:

$$\int_1^2 f(x) dx = \int_0^1 (1 - f(u)) du$$

Step 3: Simplify the integral.

$$\begin{aligned} \int_0^2 f(x) dx &= \int_0^1 f(x) dx + \int_0^1 1 du - \int_0^1 f(u) du \\ \int_0^2 f(x) dx &= \int_0^1 1 du = 1 \end{aligned}$$

Final Answer:

1

Quick Tip

Use symmetry properties of functions and substitution to simplify definite integrals.

17. Let A be the adjacency matrix of a simple undirected graph G . Suppose A is its own inverse. Which one of the following statements is always TRUE?

- (A) G is a cycle
- (B) G is a perfect matching
- (C) G is a complete graph
- (D) There is no such graph G

Correct Answer: (B) G is a perfect matching

Solution: If A is the adjacency matrix of an undirected graph G , and $A \cdot A = I$ (where I is the identity matrix), it means that the graph G satisfies the following conditions: 1. Each node in G is connected to exactly one other node because $A \cdot A = I$ implies no node connects back to itself (diagonal is 1 in I). 2. There are no paths of length 2 between any two nodes in G .

This describes a perfect matching, where each node is paired with exactly one other node.

Step 1: Analyze other options.

- (A) G is a cycle: A cycle would not satisfy $A \cdot A = I$ because there are paths of length 2.
- (C) G is a complete graph: A complete graph does not satisfy $A \cdot A = I$ because all nodes are connected.
- (D) There is no such graph G : This is incorrect because graphs with perfect matchings exist.

Final Answer:

Perfect Matching

Quick Tip

When the adjacency matrix of a graph satisfies $A \cdot A = I$, the graph represents a perfect matching.

18. When six unbiased dice are rolled simultaneously, the probability of getting all distinct numbers (i.e., 1, 2, 3, 4, 5, and 6) is

- (A) $\frac{1}{324}$
- (B) $\frac{5}{324}$
- (C) $\frac{7}{324}$
- (D) $\frac{11}{324}$

Correct Answer: (B) $\frac{5}{324}$

Solution: To find the probability of rolling six distinct numbers, calculate the total number of favorable outcomes and divide by the total number of possible outcomes.

Step 1: Total possible outcomes.

$$6^6 = 46656$$

Step 2: Favorable outcomes. The first die can take any of the 6 numbers, the second can take any of the remaining 5, and so on:

$$6 \cdot 5 \cdot 4 \cdot 3 \cdot 2 \cdot 1 = 720$$

Step 3: Calculate the probability.

$$P = \frac{\text{Favorable Outcomes}}{\text{Total Outcomes}} = \frac{720}{46656} = \frac{5}{324}$$

Final Answer:

$$\frac{5}{324}$$

Quick Tip

When calculating probabilities involving permutations, always account for the order of outcomes when the arrangement matters.

19. Once the DBMS informs the user that a transaction has been successfully completed, its effect should persist even if the system crashes before all its changes are reflected on disk. This property is called

- (A) durability
- (B) atomicity
- (C) consistency
- (D) isolation

Correct Answer: (A) durability

Solution: The durability property ensures that once the DBMS confirms the successful completion of a transaction, its changes to the database are permanent and will persist even if there is a system failure. This property is one of the ACID properties of transactions.

Step 1: Analyze the other options.

- (B) Atomicity ensures that a transaction is completed entirely or not at all, but it does not guarantee persistence after system crashes.
- (C) Consistency ensures that a database remains in a valid state before and after a transaction, but it does not address persistence.
- (D) Isolation ensures that transactions do not interfere with each other, but it does not guarantee durability.

Conclusion: Durability is the correct answer as it guarantees the effect of a committed transaction persists after system crashes.

Final Answer:

durability

Quick Tip

The ACID properties—Atomicity, Consistency, Isolation, and Durability—are fundamental to understanding transaction management in databases.

20. In the context of owner and weak entity sets in the ER (Entity-Relationship) data model, which one of the following statements is TRUE?

- (A) The weak entity set **MUST** have total participation in the identifying relationship
- (B) The owner entity set **MUST** have total participation in the identifying relationship
- (C) Both weak and owner entity sets **MUST** have total participation in the identifying relationship
- (D) Neither weak entity set nor owner entity set **MUST** have total participation in the identifying relationship

Correct Answer: (A) The weak entity set **MUST** have total participation in the identifying relationship

Solution: In the ER model, a weak entity set cannot exist without being associated with an owner entity set. This dependency requires a weak entity set to have total participation in the

identifying relationship, meaning every instance of a weak entity must be related to at least one instance of the owner entity set.

Step 1: Analyze the other options.

(B) The owner entity set does not require total participation in the identifying relationship.

(C) Both sets do not need total participation; only the weak entity set does.

(D) This is incorrect because the weak entity set must have total participation.

Conclusion: The weak entity set must have total participation in the identifying relationship, making (A) the correct answer.

Final Answer:

The weak entity set **MUST** have total participation

Quick Tip

In the ER model, always associate a weak entity set with an identifying relationship involving its owner entity set for complete design.

21. Consider the following two sets:

| Set X | | Set Y | |
|-------|-----------------------------|-------|----------------------|
| P. | Lexical Analyzer | 1. | Abstract Syntax Tree |
| Q. | Syntax Analyzer | 2. | Token |
| R. | Intermediate Code Generator | 3. | Parse Tree |
| S. | Code Optimizer | 4. | Constant Folding |

Which one of the following options is the CORRECT match from Set X to Set Y?

(A) $P - 4; Q - 1; R - 3; S - 2$

(B) $P - 2; Q - 3; R - 1; S - 4$

(C) $P - 2; Q - 1; R - 3; S - 4$

(D) $P - 4; Q - 3; R - 2; S - 1$

Correct Answer: (B)

Solution: To correctly match Set X with Set Y, we analyze each component:

P. Lexical Analyzer: The lexical analyzer is responsible for generating **tokens** from the input stream. Hence, $P - 2$.

Q. Syntax Analyzer: The syntax analyzer generates a **parse tree** by analyzing the syntactical structure of tokens. Hence, $Q - 3$.

R. Intermediate Code Generator: The intermediate code generator creates an **abstract syntax tree (AST)** as part of its processing. Hence, $R - 1$.

S. Code Optimizer: The code optimizer applies techniques like **constant folding** to optimize the intermediate code. Hence, $S - 4$.

Correct match:

$$P - 2, Q - 3, R - 1, S - 4.$$


Final Answer:

(B)

Quick Tip

Understand the role of each compiler phase and its output to accurately match components.

22. Which one of the following regular expressions is equivalent to the language accepted by the DFA given below?



- (A) $0^*1(0 + 10^*1)^*$
- (B) $0^*(10^*11)^*0^*$
- (C) $0^*1(010^*1)^*0^*$
- (D) $0(1 + 0^*10^*1)^*0^*$

Correct Answer: (A)

Solution: The DFA has three states:

q_0 : Initial state.

q_1 : Intermediate state.

q_2 : Accepting state.

Analysis:

From q_0 , a 1 transitions to q_1 .

From q_1 , a 0 transitions to q_2 , and another 1 transitions back to q_0 .

From q_2 , 0 loops back to q_2 .

This DFA accepts strings that:

Start with 0^* (any number of 0s at the beginning).

Have at least one 1 to reach q_1 .

Follow by $0 + 10^*1$ repeatedly (to loop between q_1 and q_0).

Optionally end in 0^* (to stay in q_2).

The equivalent regular expression is:

$$0^*1(0 + 10^*1)^*$$

Final Answer:

$$0^*1(0 + 10^*1)^*$$

Quick Tip

To find the regular expression for a DFA, analyze transitions for each state and consider all possible loops and paths to the accepting state(s).

23. Node X has a TCP connection open to node Y. The packets from X to Y go through an intermediate IP router R. Ethernet switch S is the first switch on the network path between X and R. Consider a packet sent from X to Y over this connection.

Which of the following statements is/are TRUE about the destination IP and MAC addresses on this packet at the time it leaves X?

(A) The destination IP address is the IP address of R.

- (B) The destination IP address is the IP address of Y.
- (C) The destination MAC address is the MAC address of S.
- (D) The destination MAC address is the MAC address of Y.

Correct Answer: (B)

Solution: In a TCP/IP packet, the destination **IP address** remains constant and always points to the final destination node, which in this case is *Y*. Hence, the destination IP address is the IP address of *Y*, making **Option (B)** correct.

The **MAC address**, however, is updated at every hop and points to the next device in the path. When the packet leaves *X*, the destination MAC address is set to the MAC address of *R* (the next-hop router), not *S* or *Y*.

Therefore, **Option (B)** is correct.

Final Answer:

(B)

Quick Tip

In a TCP/IP network, the IP address indicates the final destination, while the MAC address is updated for each hop to represent the next device on the path.

24. Which of the following tasks is/are the responsibility/responsibilities of the memory management unit (MMU) in a system with paging-based memory management?

- (A) Allocate a new page table for a newly created process
- (B) Translate a virtual address to a physical address using the page table
- (C) Raise a trap when a virtual address is not found in the page table
- (D) Raise a trap when a process tries to write to a page marked with read-only permission in the page table

Correct Answer: (B), (C), (D)

Solution: **(A) Allocate a new page table for a newly created process:** This task is typically handled by the operating system, not the MMU. Hence, **Option (A)** is incorrect.

(B) Translate a virtual address to a physical address using the page table: This is one of the core functions of the MMU. The MMU performs the translation of virtual addresses to physical addresses based on the page table. Hence, **Option (B)** is correct.

(C) Raise a trap when a virtual address is not found in the page table: The MMU raises a page fault trap when a virtual address is not found in the page table. This prompts the operating system to handle the fault. Hence, **Option (C)** is correct.

(D) Raise a trap when a process tries to write to a page marked with read-only permission in the page table: The MMU enforces memory protection by checking permissions in the page table. If a process violates these permissions, such as writing to a read-only page, the MMU raises a trap. Hence, **Option (D)** is correct.

Final Answer:

(B), (C), (D)

Quick Tip

The MMU handles address translation, permission checks, and traps for invalid memory accesses, while the operating system handles tasks like page table allocation.

25. Consider a process P running on a CPU. Which one or more of the following events will always trigger a context switch by the OS that results in process P moving to a non-running state (e.g., ready, blocked)?

- (A) P makes a blocking system call to read a block of data from the disk.
- (B) P tries to access a page that is in the swap space, triggering a page fault.
- (C) An interrupt is raised by the disk to deliver data requested by some other process.
- (D) A timer interrupt is raised by the hardware.

Correct Answer: (A), (B)

Solution: When process P makes a blocking system call (e.g., to read a block of data from the disk), it cannot proceed until the I/O operation completes. The OS places P in the *blocked* state, triggering a context switch. Therefore, **Option (A)** is correct.

If P tries to access a page that is not in memory but in swap space, a page fault occurs. The

OS retrieves the page from disk while P is placed in the *blocked* state. This also triggers a context switch. Hence, **Option (B)** is correct.

An interrupt raised by the disk to deliver data requested by another process does not directly involve P . The OS handles the interrupt but does not always switch context for P . Thus, **Option (C)** is incorrect.

A timer interrupt raised by hardware does not always trigger a context switch. The OS scheduler may decide to continue running P or preempt it. Hence, **Option (D)** is incorrect.

Final Answer:

(A), (B)

Quick Tip

Blocking system calls and page faults always cause a context switch, as the process must wait for the event to complete.

26. Which of the following file organizations is/are I/O efficient for the scan operation in DBMS?

- (A) Sorted
- (B) Heap
- (C) Unclustered tree index
- (D) Unclustered hash index

Correct Answer: (A), (B)

Solution: In a scan operation, sequential access to data minimizes disk I/O. Sorted file organizations provide sequential access and are efficient for scan operations. Thus, **Option (A)** is correct.

Heap files store data without any particular order, but for a full scan, every block is accessed sequentially. This makes heap files I/O efficient for scan operations. Hence, **Option (B)** is correct.

Unclustered tree indexes, however, are not I/O efficient for scans because data blocks may not be accessed sequentially, leading to multiple random I/O operations. Hence, **Option (C)** is

incorrect.

Unclustered hash indexes are designed for equality searches and not optimized for scan operations. Random I/O makes them inefficient for scans. Thus, **Option (D)** is incorrect.

Final Answer:

(A), (B)

Quick Tip

Heap and sorted file organizations are efficient for scans due to their sequential access patterns, reducing random I/O.

27. Which of the following statements about the Two Phase Locking (2PL) protocol is/are TRUE?

- (A) 2PL permits only serializable schedules
- (B) With 2PL, a transaction always locks the data item being read or written just before every operation and always releases the lock just after the operation
- (C) With 2PL, once a lock is released on any data item inside a transaction, no more locks on any data item can be obtained inside that transaction
- (D) A deadlock is possible with 2PL

Correct Answer: (A), (C), (D)

Solution: The Two Phase Locking (2PL) protocol ensures two distinct phases: the growing phase (acquiring locks) and the shrinking phase (releasing locks).

Option (A): 2PL permits only serializable schedules by ensuring conflict-serializability. Hence, (A) is correct.

Option (B): This statement is incorrect because locks are not acquired and released immediately for every operation. Locks are held until the shrinking phase.

Option (C): In 2PL, once a lock is released during the shrinking phase, no new locks can be acquired. This is a key property of 2PL. Hence, (C) is correct.

Option (D): Deadlocks can occur in 2PL due to circular wait conditions. Hence, (D) is correct.

Final Answer:

(A), (C), (D)

Quick Tip

2PL ensures serializability but does not prevent deadlocks. Use techniques like timeout or wait-die to handle deadlocks in 2PL.

28. Which of the following statements about IPv4 fragmentation is/are TRUE?

- (A) The fragmentation of an IP datagram is performed only at the source of the datagram.
- (B) The fragmentation of an IP datagram is performed at any IP router which finds that the size of the datagram to be transmitted exceeds the MTU.
- (C) The reassembly of fragments is performed only at the destination of the datagram.
- (D) The reassembly of fragments is performed at all intermediate routers along the path from the source to the destination.

Correct Answer: (B), (C)

Solution: IPv4 fragmentation occurs when the size of a datagram exceeds the Maximum Transmission Unit (MTU) of a network link.

Option (A): Incorrect. Fragmentation can occur at any intermediate router along the path, not just at the source.

Option (B): Correct. Any IP router that encounters an MTU smaller than the datagram size performs fragmentation.

Option (C): Correct. Reassembly of fragments is performed only at the destination to reconstruct the original datagram.

Option (D): Incorrect. Intermediate routers do not reassemble fragments; they forward them as they are.

Final Answer:

(B), (C)

Quick Tip

Fragmentation can occur at any router where the MTU is exceeded, but reassembly is always performed only at the destination.

29. Which of the following statements is/are FALSE?

- (A) An attribute grammar is a syntax-directed definition (SDD) in which the functions in the semantic rules have no side effects

- (B) The attributes in an L-attributed definition cannot always be evaluated in a depth-first order
- (C) Synthesized attributes can be evaluated by a bottom-up parser as the input is parsed
- (D) All L-attributed definitions based on LR(1) grammar can be evaluated using a bottom-up parsing strategy

Correct Answer: (B), (D)

Solution: Option (A): This is true. In attribute grammars, functions in semantic rules are required to have no side effects, as they must not interfere with other computations.

Option (B): This is false. L-attributed definitions allow evaluation in a depth-first order since inherited attributes are passed only along the parse tree edges and are computable in one depth-first pass.

Option (C): This is true. Synthesized attributes are computed from the values of the children nodes and can be evaluated in a bottom-up parsing strategy.

Option (D): This is false. While L-attributed definitions can be evaluated using top-down parsing, they are not inherently suitable for bottom-up parsing strategies.

Final Answer:

(B), (D)

Quick Tip

Remember that L-attributed grammars are designed for top-down parsing, while synthesized attributes are naturally evaluated in bottom-up parsing.

30. For a Boolean variable x , which of the following statements is/are FALSE?

- (A) $x \cdot 1 = x$
- (B) $x + 1 = x$
- (C) $x \cdot x = 0$
- (D) $x + \bar{x} = 1$

Correct Answer: (B), (C)

Solution: Option (A): This is true. The Boolean AND operation with 1 always yields the variable itself: $x \cdot 1 = x$.

Option (B): This is false. The Boolean OR operation with 1 always yields 1, regardless of the value of x : $x + 1 = 1$.

Option (C): This is false. The Boolean AND operation with itself yields the variable itself: $x \cdot x = x$, not 0.

Option (D): This is true. The Boolean OR operation with the complement of a variable always yields 1: $x + \bar{x} = 1$.

Final Answer:

(B), (C)

Quick Tip

Boolean algebra follows basic logic rules: $x + 1 = 1$, $x \cdot x = x$, and $x + \bar{x} = 1$. Review these rules for accuracy.

31. An instruction format has the following structure:

Instruction Number: *Opcode destination reg, source reg-1, source reg-2*

Consider the following sequence of instructions to be executed in a pipelined processor:

I₁: DIV R3, R1, R2

I₂: SUB R5, R3, R4

I₃: ADD R3, R5, R6

I₄: MUL R7, R3, R8

Which of the following statements is/are TRUE?

(A) There is a RAW dependency on R3 between I₁ and I₂

(B) There is a WAR dependency on R3 between I₁ and I₃

(C) There is a RAW dependency on R3 between I₂ and I₃

(D) There is a WAW dependency on R3 between I₃ and I₄

Correct Answer: (A)

Solution: Option (A): True. Instruction I_1 writes to R3, and instruction I_2 reads from R3, creating a Read After Write (RAW) dependency between I_1 and I_2 .

Option (B): False. Instruction I_3 writes to R3, but I_1 also writes to R3. This would be a Write After Write (WAW) dependency, not a WAR dependency.

Option (C): False. Instruction I_3 writes to R3, but I_2 does not read R3. Hence, there is no RAW dependency between I_2 and I_3 .

Option (D): False. Both I_3 and I_4 write to R3, creating a Write After Write (WAW) dependency, but this dependency does not exist in the given instruction sequence.

Final Answer:

(A)

Quick Tip

To identify dependencies, carefully analyze the source and destination registers used by consecutive instructions.

32. Which of the following fields of an IP header is/are always modified by any router before it forwards the IP packet?

- (A) Source IP Address
- (B) Protocol
- (C) Time to Live (TTL)
- (D) Header Checksum

Correct Answer: (C), (D)

Solution: Option (A): The source IP address is not modified by routers. It remains the same throughout the packet's journey. Hence, **Option (A)** is incorrect.

Option (B): The protocol field indicates the higher-level protocol (e.g., TCP, UDP) and is not modified by routers. Hence, **Option (B)** is incorrect.

Option (C): The Time to Live (TTL) field is decremented by 1 by each router to prevent

infinite loops in the network. This field is always modified by routers. Hence, **Option (C)** is correct.

Option (D): Since the TTL field is modified, the Header Checksum is recomputed by routers to ensure integrity. Hence, **Option (D)** is correct.

Final Answer:

(C), (D)

Quick Tip

The TTL field prevents infinite loops, and any modification to the header requires re-computing the checksum.

33. Consider the following C function definition.

```
int fX(char *a){
    char *b = a;
    while(*b)
        b++;
    return b - a;
}
```

Which of the following statements is/are TRUE?

- (A) The function call `fX("abcd")` will always return a value
- (B) Assuming a character array `c` is declared as `char c[] = "abcd"` in `main()`, the function call `fX(c)` will always return a value
- (C) The code of the function will not compile
- (D) Assuming a character pointer `c` is declared as `char *c = "abcd"` in `main()`, the function call `fX(c)` will always return a value

Correct Answer: (A), (B), (D)

Solution: Option (A): The string literal "abcd" is valid, and the function calculates the length of the string by iterating until the null character. Hence, **Option (A)** is correct.

Option (B): The character array `c` is a valid input for the function, and the function correctly calculates the length of the string. Hence, **Option (B)** is correct.

Option (C): This is false. The code is valid and will compile without any issues. Hence, **Option (C)** is incorrect.

Option (D): The character pointer `c` pointing to the string literal "abcd" is valid, and the function calculates the length of the string correctly. Hence, **Option (D)** is correct.

Final Answer:

(A), (B), (D)

Quick Tip

String literals and character arrays can be used as inputs to functions that calculate string lengths, provided they are properly null-terminated.

34. Let P be the partial order defined on the set $\{1, 2, 3, 4\}$ as follows:

$$P = \{(x, x) \mid x \in \{1, 2, 3, 4\}\} \cup \{(1, 2), (3, 2), (3, 4)\}$$

The number of total orders on $\{1, 2, 3, 4\}$ that contain P is _____.

Correct Answer: 5

Solution: A total order is an extension of a partial order in which every pair of elements is comparable. Based on the given partial order P , we know the following relations:

$$1 \leq 2, \quad 3 \leq 2, \quad 3 \leq 4.$$

Using these relations, the valid total orders containing P are:

1. $1 \leq 3 \leq 4 \leq 2$
2. $1 \leq 3 \leq 2 \leq 4$
3. $1 \leq 4 \leq 3 \leq 2$
4. $3 \leq 1 \leq 4 \leq 2$
5. $3 \leq 1 \leq 2 \leq 4$

Thus, there are 5 total orders that extend P .

Final Answer:

5

Quick Tip

To determine the number of total orders from a partial order, analyze the transitive closure and enumerate all valid permutations that satisfy the given relations.

35. Let A be an array containing integer values. The distance of A is defined as the minimum number of elements in A that must be replaced with another integer so that the resulting array is sorted in non-decreasing order. The distance of the array $[2, 5, 3, 1, 4, 2, 6]$ is

Correct Answer: 3

Solution: To calculate the distance, we need to determine the Longest Increasing Subsequence (LIS) of the array. The length of the LIS represents the largest subset of elements already in sorted order.

The given array is $[2, 5, 3, 1, 4, 2, 6]$. The LIS for this array is:

$[2, 3, 4, 6]$ (length = 4).

The minimum number of elements to replace is given by:

$$\text{Distance} = \text{Length of array} - \text{Length of LIS} = 7 - 4 = 3.$$

Final Answer:

3

Quick Tip

The distance of an array is calculated as the total number of elements minus the length of its Longest Increasing Subsequence (LIS).

Question 36-65 carry two mark each

36. What is the output of the following C program?

```
#include <stdio.h>
int main() {
    double a[2] = {20.0, 25.0}, *p, *q;
    p = a;
    q = p + 1;
    printf("%d,%d", (int)(q - p), (int)(*q - *p));
    return 0; }
```

- (A) 4, 8
- (B) 1, 5
- (C) 8, 5
- (D) 1, 8

Correct Answer: (B)

Solution: The program initializes an array $a[2] = \{20.0, 25.0\}$, two pointers p and q , and performs the following steps:

1. $p = a$: Pointer p now points to the first element of the array, i.e., $a[0] = 20.0$.
2. $q = p + 1$: Pointer q now points to the second element of the array, i.e., $a[1] = 25.0$.
3. The statement $q - p$ calculates the difference in pointer positions. Since q points to the second element and p points to the first element, $q - p = 1$.
4. The statement $*q - *p$ calculates the difference between the values pointed to by q and p . This gives:

$$*q - *p = a[1] - a[0] = 25.0 - 20.0 = 5.0.$$

5. The `printf` statement prints these values as integers:

$$\text{Output: } (int)(q - p) = 1, \quad (int)(*q - *p) = 5.$$

Final Answer:

(B)

Quick Tip

Pointer arithmetic on arrays calculates the difference in element positions, while dereferencing pointers gives the actual value stored in the array.

37. Consider a single processor system with four processes A, B, C, and D, represented as given below, where for each process the first value is its arrival time, and the second value is its CPU burst time.

$$A(0, 10), B(2, 6), C(4, 3), D(6, 7).$$

Which one of the following options gives the average waiting times when preemptive Shortest Remaining Time First (SRTF) and Non-Preemptive Shortest Job First (NP-SJF) CPU scheduling algorithms are applied to the processes?

- (A) SRTF = 6, NP-SJF = 7
- (B) SRTF = 6, NP-SJF = 7.5
- (C) SRTF = 7, NP-SJF = 7.5
- (D) SRTF = 7, NP-SJF = 8.5

Correct Answer: (B)

Solution:

Non-Preemptive SJF Scheduling:

| Process | Arrival Time | Burst Time | Completion Time | Turnaround Time | Waiting Time |
|---------|--------------|------------|-----------------|-----------------|--------------|
| A | 0 | 10 | 10 | 10 | 0 |
| B | 2 | 6 | 19 | 17 | 11 |
| C | 4 | 3 | 13 | 9 | 6 |
| D | 6 | 7 | 26 | 20 | 13 |

The total average waiting time for NP-SJF is:

$$\text{Average Waiting Time} = \frac{0 + 11 + 6 + 13}{4} = \frac{30}{4} = 7.5.$$

Preemptive SRTF Scheduling:

| Process | Arrival Time | Burst Time | Completion Time | Turnaround Time | Waiting Time |
|---------|--------------|------------|-----------------|-----------------|--------------|
| A | 0 | 10 | 26 | 26 | 16 |
| B | 2 | 6 | 11 | 9 | 3 |
| C | 4 | 3 | 7 | 3 | 0 |
| D | 6 | 7 | 18 | 12 | 5 |

The total average waiting time for SRTF is:

$$\text{Average Waiting Time} = \frac{16 + 3 + 0 + 5}{4} = \frac{24}{4} = 6.$$

Note:

- **Turnaround Time** is calculated as: Completion Time – Arrival Time.
- **Waiting Time** is calculated as: Turnaround Time – Burst Time.

Final Answer:

$$(2) \text{ SRTF} = 6, \text{ NP-SJF} = 7.5$$

Quick Tip

When solving CPU scheduling problems, use tables to calculate Completion Time, Turnaround Time, and Waiting Time for each process. Apply the appropriate algorithm rules, such as preemption in SRTF and non-preemption in NP-SJF.

38. Which one of the following CIDR prefixes exactly represents the range of IP addresses 10.12.2.0 to 10.12.3.255?

- (A) 10.12.2.0/23
- (B) 10.12.2.0/24
- (C) 10.12.0.0/22
- (D) 10.12.2.0/22

Correct Answer: (A)

Solution: The given range is 10.12.2.0 to 10.12.3.255.

The CIDR notation 10.12.2.0/23 has a subnet mask of 255.255.254.0, which covers the following range of IPs:

10.12.2.0 (start) to 10.12.3.255 (end).

Other options do not correctly represent this range:

(B) 10.12.2.0/24: Only covers 10.12.2.0 to 10.12.2.255.

(C) 10.12.0.0/22: Covers a larger range from 10.12.0.0 to 10.12.3.255.

(D) 10.12.2.0/22: Invalid as it overlaps incorrectly with the range 10.12.4.x.

Final Answer:

(A)

Quick Tip

To find the correct CIDR prefix, calculate the subnet mask and verify the range of IPs it covers.

39. You are given a set V of distinct integers. A binary search tree T is created by inserting all elements of V one by one, starting with an empty tree. The tree T follows the convention that, at each node, all values stored in the left subtree of the node are smaller than the value stored at the node. You are not aware of the sequence in which these values were inserted into T , and you do not have access to T .

Which one of the following statements is TRUE?

- (A) Inorder traversal of T can be determined from V
- (B) Root node of T can be determined from V
- (C) Preorder traversal of T can be determined from V
- (D) Postorder traversal of T can be determined from V

Correct Answer: (A)

Solution: In a Binary Search Tree (BST), the inorder traversal of the tree produces the elements in sorted order, regardless of the sequence in which they were inserted. Since V contains all the elements of T as a set, sorting V will give the inorder traversal of T . Hence,

Option (A) is correct.

Option (B): The root node of T depends on the sequence of insertion and cannot be determined solely from V . Hence, **Option (B)** is incorrect.

Option (C): The preorder traversal depends on the structure of T , which in turn depends on the insertion order. Therefore, it cannot be determined from V . Hence, **Option (C)** is incorrect.

Option (D): The postorder traversal also depends on the structure of T and the insertion sequence, and cannot be determined from V . Hence, **Option (D)** is incorrect.

Final Answer:

(A)

Quick Tip

In a Binary Search Tree (BST), the inorder traversal always gives the sorted order of elements, regardless of the insertion sequence.

40. Consider the following context-free grammar where the start symbol is S and the set of terminals is $\{a, b, c, d\}$:

$$S \rightarrow AaAb \mid BbBa \quad A \rightarrow cS \mid \epsilon \quad B \rightarrow dS \mid \epsilon$$

The following is a partially-filled LL(1) parsing table.

| | a | b | c | d | $\$$ |
|-----|--------------------------|--------------------------|--------------------|--------------------|------|
| S | $S \rightarrow AaAb$ | $S \rightarrow BbBa$ | (1) | (2) | |
| A | $A \rightarrow \epsilon$ | (3) | $A \rightarrow cS$ | | |
| B | (4) | $B \rightarrow \epsilon$ | | $B \rightarrow dS$ | |

Which one of the following options represents the CORRECT combination for the numbered cells in the parsing table?

Note: In the options, "blank" denotes that the corresponding cell is empty.

(A) (1) $S \rightarrow AaAb$, (2) $S \rightarrow BbBa$, (3) $A \rightarrow \epsilon$, (4) $B \rightarrow \epsilon$

(B) (1) $S \rightarrow BbBa$, (2) $S \rightarrow AaAb$, (3) $A \rightarrow \epsilon$, (4) $B \rightarrow \epsilon$

(C) (1) $S \rightarrow AaAb$, (2) $S \rightarrow BbBa$, (3) blank, (4) blank

(D) (1) $S \rightarrow BbBa$, (2) $S \rightarrow AaAb$, (3) blank, (4) blank

Correct Answer: (A)

Solution: For the given grammar and the parsing table:

Cell (1): This corresponds to the production for S when the input starts with c . From the grammar, $A \rightarrow cS$ and $S \rightarrow AaAb$, leading to $S \rightarrow AaAb$ in this cell. Hence, (1) is $S \rightarrow AaAb$.

Cell (2): This corresponds to S when the input starts with d . From $B \rightarrow dS$ and $S \rightarrow BbBa$, this cell contains $S \rightarrow BbBa$. Hence, (2) is $S \rightarrow BbBa$.

Cell (3): This corresponds to A when the input starts with b . From the grammar, $A \rightarrow \epsilon$, this cell contains $A \rightarrow \epsilon$. Hence, (3) is $A \rightarrow \epsilon$.

Cell (4): This corresponds to B when the input starts with a . From the grammar, $B \rightarrow \epsilon$, this cell contains $B \rightarrow \epsilon$. Hence, (4) is $B \rightarrow \epsilon$.

The correct configuration for the cells is:

$$(1)S \rightarrow AaAb, (2)S \rightarrow BbBa, (3)A \rightarrow \epsilon, (4)B \rightarrow \epsilon.$$


Final Answer:

(A)

Quick Tip

To fill an LL(1) parsing table, use the FIRST and FOLLOW sets of the grammar rules for each non-terminal.

41. Let M be the 5-state NFA with ϵ -transitions shown in the diagram below.



Which one of the following regular expressions represents the language accepted by M ?

- (A) $(00)^* + 1(11)^*$
- (B) $0^* + (1 + 0(00)^*)(11)^*$
- (C) $(00)^* + (1 + 0(00)^*)(11)^*$
- (D) $0^+ + 1(11)^* + 0(11)^*$

Correct Answer: (B)

Solution: The given NFA M has five states (1, 2, 3, 4, 5) and includes ϵ -transitions. To determine the language accepted by M , consider the paths and transitions:

1. From state 1 to 2: The ϵ -transition allows state 2 to be reached directly without consuming any input.
2. From state 2 to 3: Input 0 can be consumed repeatedly (loop on 2) or proceed to 3 with a single 0. This corresponds to $(00)^*$.
3. From state 1 to 4: The ϵ -transition allows state 4 to be reached directly without consuming any input.
4. From state 4 to 5: Input 1 can be consumed repeatedly (loop on 5) after consuming 1. This corresponds to $(11)^*$.
5. Combining paths: The overall language includes all possible combinations:
 - 0^* , representing sequences of 0's from state 2.
 - $(1 + 0(00)^*)(11)^*$, combining sequences starting with 1 or sequences starting with 0 followed by $(00)^*$, and ending with $(11)^*$.

The regular expression representing the language is:

$$0^* + (1 + 0(00)^*)(11)^*.$$

Final Answer:

(B)

Quick Tip

For NFAs with ϵ -transitions, trace all paths from the start to accept states, combining loops and concatenations to construct the regular expression.

42. Consider an array X that contains n positive integers. A subarray of X is defined to be a sequence of array locations with consecutive indices.

The C code snippet given below has been written to compute the length of the longest subarray of X that contains at most two distinct integers. The code has two missing expressions labelled (P) and (Q) .

```
int first=0, second=0, len1=0, len2=0, maxlen=0;
for (int i=0; i < n; i++) {
    if (X[i] == first) {
        len2++; len1++;
    } else if (X[i] == second) {
        len2++;
        len1 = (P);
        second = first;
    } else {
        len2 = (Q);
        len1 = 1; second = first;
    }
    if (len2 > maxlen) {
        maxlen = len2;
    }
    first = X[i];
}
```

Which one of the following options gives the CORRECT missing expressions?

- (A) (P) $len1 + 1$, (Q) $len2 + 1$
- (B) (P) 1 , (Q) $len1 + 1$
- (C) (P) 1 , (Q) $len2 + 1$
- (D) (P) $len2 + 1$, (Q) $len1 + 1$

Correct Answer: (B)

Solution: The goal is to maintain the length of the longest subarray ending at the current

index i with at most two distinct integers. The variables have the following roles:

1. $len1$: Length of the subarray ending at $X[i]$ where all elements are equal to $X[i]$.
2. $len2$: Length of the subarray ending at $X[i]$ with at most two distinct integers.

Analysis of (P): When $X[i] == \text{second}$, $len1$ must be reset to 1 because this marks the start of a new sequence of identical integers. Thus, $(P) = 1$.

Analysis of (Q): When $X[i]$ is neither *first* nor *second*, $len2$ must be updated to include the length of the subarray containing the new two distinct integers. Since the previous $len1$ represents the length of the contiguous subarray of *first*, the new $len2 = len1 + 1$ (the current element $X[i]$ is added). Thus, $(Q) = len1 + 1$.

The correct replacements for (P) and (Q) are:

$$(P) = 1, \quad (Q) = len1 + 1.$$

Final Answer:

(B)

Quick Tip

Carefully track the roles of variables in maintaining properties of subarrays. $len1$ tracks identical elements, while $len2$ tracks subarrays with at most two distinct integers.

43. Consider the following expression: $x[i] = (p+r) * -s[i] + u/w$. **The following sequence shows the list of triples representing the given expression, with entries missing for triples (1), (3), and (6).**

| | | | |
|------------|--------|-----|--------|
| (0) | + | p | r |
| (1) | | | |
| (2) | uminus | (1) | |
| (3) | | | |
| (4) | / | u | w |
| (5) | + | (3) | (4) |
| (6) | | | |
| (7) | = | (6) | $x[i]$ |

Which one of the following options fills in the missing entries **CORRECTLY**?

- (A) (1) = $s[i]$, (3) * (0)(2), (6) = $x[i]$
- (B) (1) = $s[i]$, (3) - (0)(2), (6) = $x(5)$
- (C) (1) = $s[i]$, (3) * (0)(2), (6) = $x(5)$
- (D) (1) = $s[i]$, (3) - (0)(2), (6) = $x[i]$

Correct Answer: (A)

Solution: To break down the expression $x[i] = (p + r) * -s[i] + u/w$ into intermediate triples:

1. Triple (0): + operation for p and r . Already given as:

$$(0) : + p r.$$

2. Triple (1): Assignment of $-s[i]$. Using the unary minus (uminus) operator:

$$(1) := s[i].$$

3. Triple (2): Unary minus applied to (1). Already given as:

$$(2) : \text{uminus } (1).$$

4. Triple (3): Multiplication of (0) and (2):

$$(3) : * (0) (2).$$

5. Triple (4): Division of u and w . Already given as:

$$(4) : / u w.$$

6. Triple (5): Addition of (3) and (4):

$$(5) : + (3) (4).$$

7. Triple (6): Assignment of (5) to $x[i]$:

$$(6) := [] x[i].$$

The correct missing entries are:

$$(1) := [] s[i], (3) : * (0) (2), (6) := [] x[i].$$

Final Answer:

(A)

Quick Tip

For three-address code generation, process the operations in order of precedence and create intermediate triples for subexpressions.

44. Let x and y be random variables, not necessarily independent, that take real values in the interval $[0, 1]$. Let $z = xy$ and let the mean values of x, y, z be $\bar{x}, \bar{y}, \bar{z}$, respectively. Which one of the following statements is TRUE?

(A) $\bar{z} = \bar{x}\bar{y}$

(B) $\bar{z} \leq \bar{x}\bar{y}$

(C) $\bar{z} \geq \bar{x}\bar{y}$

(D) $\bar{z} \leq \bar{x}$

Correct Answer: (D)

Solution: The random variables x and y are in the interval $[0, 1]$, and $z = xy$. For any pair of real values $x, y \in [0, 1]$, the value of xy is always less than or equal to x (since $y \leq 1$). Therefore, the mean of z , \bar{z} , satisfies:

$$\bar{z} \leq \bar{x}.$$

This is always true, regardless of the dependence or independence of x and y .

Analysis of the other options:

Option (A): $\bar{z} = \bar{x}\bar{y}$ is true only when x and y are independent. This is not guaranteed here. Hence, (A) is incorrect.

Option (B): $\bar{z} \leq \bar{x}\bar{y}$ is true due to the general inequality $E[xy] \leq E[x]E[y]$ (Cauchy-Schwarz inequality), but this is weaker than the stricter constraint $\bar{z} \leq \bar{x}$. Hence, (B) is not the strongest statement.

Option (C): $\bar{z} \geq \bar{x}\bar{y}$ is false because the inequality $E[xy] \leq E[x]E[y]$ holds.

Option (D): $\bar{z} \leq \bar{x}$ is always true and the strongest statement in this case. Hence, (D) is correct.

Final Answer:

(D)

Quick Tip

When working with random variables in bounded intervals, remember that products of variables are constrained by their individual ranges.

45. The relation schema, Person(pid, city), describes the city of residence for every person uniquely identified by pid. The following relational algebra operators are available: selection, projection, cross product, and rename.

To find the list of cities where at least 3 persons reside, using the above operators, the minimum number of cross product operations that must be used is

(A) 1

(B) 2

(C) 3

(D) 4

Correct Answer: (B)

Solution: The problem requires finding cities where at least 3 people reside. This can be accomplished as follows:

1. Use `selection` and `projection` to create a list of city values.
2. Perform `cross product` operations to create all combinations of rows for a given city, ensuring at least 3 unique entries for each city.

The minimum number of `cross product` operations required to verify at least 3 people in each city is 2.

Final Answer:

(B)

Quick Tip

Cross product is used in relational algebra to combine rows from two tables. Use it judiciously for solving cardinality constraints.

46. Consider a multi-threaded program with two threads T_1 and T_2 . The threads share two semaphores: s_1 (initialized to 1) and s_2 (initialized to 0). The threads also share a global variable x (initialized to 0). The threads execute the code shown below.

```
// code of T1           // code of T2
wait(s1);               wait(s1);
x = x+1;                x = x+1;
print(x);                print(x);
wait(s2);                signal(s2);
signal(s1);              signal(s1);
```

Which of the following outcomes is/are possible when threads T_1 and T_2 execute concurrently?

- (A) $T1$ runs first and prints 1, $T2$ runs next and prints 2
- (B) $T2$ runs first and prints 1, $T1$ runs next and prints 2
- (C) $T1$ runs first and prints 1, $T2$ does not print anything (deadlock)
- (D) $T2$ runs first and prints 1, $T1$ does not print anything (deadlock)

Correct Answer: (B), (C)

Solution: Outcome (A): This is not possible because $T2$ cannot proceed after $T1$ without acquiring $s1$, but $s1$ is released only when $T1$ completes. Thus, (A) is incorrect.

Outcome (B): $T2$ can run first, acquire $s1$, increment x , and print 1. Then $T1$ acquires $s1$, increments x , and prints 2. This outcome is possible.

Outcome (C): $T1$ acquires $s1$, increments x , and prints 1. However, it waits on $s2$, which is never signaled by $T2$, causing a deadlock. This outcome is possible.

Outcome (D): This is not possible because $T2$ signals $s2$ after printing, allowing $T1$ to proceed. Thus, (D) is incorrect.

Final Answer:

(B), (C)

Quick Tip

Analyze semaphore dependencies to identify potential deadlocks and permissible thread execution orders.

47. Let A be an $n \times n$ matrix over the set of all real numbers \mathbb{R} . Let B be a matrix obtained from A by swapping two rows. Which of the following statements is/are TRUE?

- (A) The determinant of B is the negative of the determinant of A
- (B) If A is invertible, then B is also invertible
- (C) If A is symmetric, then B is also symmetric
- (D) If the trace of A is zero, then the trace of B is also zero

Correct Answer: (A), (B)

Solution:

Option (A): When two rows of a matrix are swapped, the determinant of the new matrix is

the negative of the determinant of the original matrix. Hence, (A) is correct.

Option (B): Swapping two rows of an invertible matrix does not affect its invertibility because the determinant remains non-zero. Hence, (B) is correct.

Option (C): Swapping rows of a symmetric matrix does not guarantee symmetry, as the symmetry condition $A[i][j] = A[j][i]$ may be violated. Hence, (C) is incorrect.

Option (D): The trace of a matrix is the sum of its diagonal elements, which is independent of row swaps. Hence, if the trace of A is zero, the trace of B remains zero. This makes (D) incorrect.


Final Answer:

(A), (B)

Quick Tip

Understand how row operations affect matrix properties such as determinant, symmetry, and invertibility.

48. Let S_1 and S_2 be two stacks. S_1 has a capacity of 4 elements, and S_2 has a capacity of 2 elements. S_1 already has 4 elements: 100, 200, 300, and 400, whereas S_2 is empty, as shown below.



Only the following three operations are available:

PushToS2: Pop the top element from S_1 and push it on S_2 .

PushToS1: Pop the top element from S_2 and push it on S_1 .

GenerateOutput: Pop the top element from S_1 and output it to the user.

Note: - The pop operation is not allowed on an empty stack.

- The push operation is not allowed on a full stack.

Which of the following output sequences can be generated by using the above operations?

- (A) 100, 200, 400, 300
- (B) 200, 300, 400, 100
- (C) 400, 200, 100, 300
- (D) 300, 200, 400, 100

Correct Answer: (B), (C), (D)

Solution: The available operations allow for flexible reordering of elements in S_1 and S_2 . By analyzing the given sequences:

1. Sequence (A): The sequence 100, 200, 400, 300 cannot be generated because 400 is above 300 in S_1 , and moving 400 before 300 without outputting 300 is not possible. Hence, (A) is not feasible.
2. Sequence (B): The sequence 200, 300, 400, 100 can be generated as follows:
 - Use `GenerateOutput` to output 200 and 300.
 - Push 400 to S_2 and output 100.
 - Retrieve 400 from S_2 and output it. Hence, (B) is feasible.
3. Sequence (C): The sequence 400, 200, 100, 300 can be generated by:
 - Use `GenerateOutput` to output 400.
 - Push 300 and 200 to S_2 .
 - Output 100 from S_1 .
 - Retrieve 300 from S_2 and output it. Hence, (C) is feasible.
4. Sequence (D): The sequence 300, 200, 400, 100 can be generated by:
 - Push 400 to S_2 .
 - Use `GenerateOutput` to output 300 and 200.
 - Retrieve 400 from S_2 and output it.
 - Finally, output 100. Hence, (D) is feasible.

Final Answer:

(B), (C), (D)

Quick Tip

When solving stack-based problems, simulate the operations step by step to verify the feasibility of each sequence.

49. Which of the following is/are EQUAL to 224 in radix-5 (i.e., base-5) notation?

- (A) 64 in radix-10
- (B) 100 in radix-8
- (C) 50 in radix-16
- (D) 121 in radix-7

Correct Answer: (A), (B), (D)

Solution: The given number 224 in radix-5 can be converted to decimal (radix-10) as follows:

$$224_5 = 2 \cdot 5^2 + 2 \cdot 5^1 + 4 \cdot 5^0 = 50 + 10 + 4 = 64.$$

Now, check each option:

Option (A): 64 in radix-10 matches the decimal value of 224_5 . Hence, (A) is correct.

Option (B): Convert 100_8 (radix-8) to decimal:

$$100_8 = 1 \cdot 8^2 + 0 \cdot 8^1 + 0 \cdot 8^0 = 64.$$

This matches 224_5 . Hence, (B) is correct.

Option (C): Convert 50_{16} (radix-16) to decimal:

$$50_{16} = 5 \cdot 16^1 + 0 \cdot 16^0 = 80.$$

This does not match 224_5 . Hence, (C) is incorrect.

Option (D): Convert 121_7 (radix-7) to decimal:

$$121_7 = 1 \cdot 7^2 + 2 \cdot 7^1 + 1 \cdot 7^0 = 49 + 14 + 1 = 64.$$

This matches 224_5 . Hence, (D) is correct.

Final Answer:

(A), (B), (D)


Quick Tip

To convert a number from any radix to decimal, multiply each digit by the base raised to its positional power and sum up the results.

50. Consider 4-variable functions f_1, f_2, f_3, f_4 expressed in sum-of-minterms form as given below.

$$f_1 = \Sigma(0, 2, 3, 5, 7, 8, 11, 13) \quad f_2 = \Sigma(1, 3, 5, 7, 11, 13, 15) \quad f_3 = \Sigma(0, 1, 4, 11) \quad f_4 = \Sigma(0, 2, 6, 13)$$

The circuit is represented as:



With respect to the circuit given above, which of the following options is/are CORRECT?

- (A) $Y = \Sigma(0, 1, 2, 11, 13)$
- (B) $Y = \Pi(3, 4, 5, 6, 7, 8, 9, 10, 12, 14, 15)$
- (C) $Y = \Sigma(0, 1, 2, 3, 4, 5, 6, 7)$
- (D) $Y = \Pi(8, 9, 10, 11, 12, 13, 14, 15)$

Correct Answer: (C), (D)

Solution: Step 1: Calculate the result of $f_1 \wedge f_2$:

$$f_1 \wedge f_2 = \Sigma(3, 5, 7, 11, 13).$$

Step 2: Calculate the result of $f_3 \vee f_4$:

$$f_3 \vee f_4 = \Sigma(0, 1, 2, 4, 6, 11, 13).$$

Step 3: XOR the outputs of $f_1 \wedge f_2$ and $f_3 \vee f_4$: The XOR operation results in the minterms present in one function but not both:

$$Y = \Sigma(0, 1, 2, 3, 4, 5, 6, 7).$$

Analysis of Options:

- Option (A): Incorrect, as Y is not equal to $\Sigma(0, 1, 2, 11, 13)$.
- Option (B): Incorrect, as $\Pi(3, 4, 5, 6, 7, 8, 9, 10, 12, 14, 15)$ does not match Y .
- Option (C): Correct, as $Y = \Sigma(0, 1, 2, 3, 4, 5, 6, 7)$.
- Option (D): Correct, since $Y = \Sigma(0, 1, 2, 3, 4, 5, 6, 7)$ implies the complement is:

$$Y = \Pi(8, 9, 10, 11, 12, 13, 14, 15).$$

Final Answer:

(C), (D)

Quick Tip

When solving circuit problems involving logical operations, calculate intermediate results step by step and carefully apply XOR and other operations.

51. Let G be an undirected connected graph in which every edge has a positive integer weight. Suppose that every spanning tree in G has even weight. Which of the following statements is/are TRUE for every such graph G ?

- (A) All edges in G have even weight
- (B) All edges in G have even weight **OR** all edges in G have odd weight
- (C) In each cycle C in G , all edges in C have even weight
- (D) In each cycle C in G , either all edges in C have even weight **OR** all edges in C have odd weight

Correct Answer: (D)

Solution:

If every spanning tree in G has an even weight, it implies a specific structure for the graph G . Consider the properties:

Option (A): Incorrect. Not all edges in G need to have even weight; some edges can have odd weights as long as the total weight of spanning trees remains even.

Option (B): Incorrect. G may contain both even and odd weighted edges.

Option (C): Incorrect. It is not necessary for all edges in any cycle C to have even weight.

Option (D): Correct. For every cycle C in G , the parity of the edges must be consistent (all even or all odd) to ensure that every spanning tree has an even total weight.

Final Answer:

(D)

Quick Tip

To analyze spanning trees, focus on the parity of edge weights in cycles of the graph.

52. Consider a context-free grammar G with the following 3 rules:

$$S \rightarrow aS, \quad S \rightarrow aSbS, \quad S \rightarrow c$$

Let $w \in L(G)$. Let $n_a(w), n_b(w), n_c(w)$ denote the number of times a, b, c occur in w , respectively. Which of the following statements is/are TRUE?

(A) $n_a(w) > n_b(w)$

(B) $n_a(w) > n_c(w) - 2$

(C) $n_c(w) = n_b(w) + 1$

(D) $n_c(w) = n_b(w) * 2$

Correct Answer: (B), (C)

Solution: The grammar G generates strings of the form:

$$w = a^p cb^q,$$

where the derivation of c adds one b and one a . Hence, the relationships between $n_a(w), n_b(w), n_c(w)$ can be analyzed:

Option (A): Incorrect. While $n_a(w)$ is often greater than $n_b(w)$, there is no guarantee, as $n_a(w)$ and $n_b(w)$ depend on the specific derivation.

Option (B): Correct. Since the derivation adds at least two a 's for every c , the inequality $n_a(w) > n_c(w) - 2$ holds.

Option (C): Correct. Each derivation of $S \rightarrow aSbS$ adds one b for every c , making $n_c(w) = n_b(w) + 1$.

Option (D): Incorrect. The number of c 's is not necessarily twice the number of b 's.

Final Answer:

(B), (C)

Quick Tip

For context-free grammars, analyze the relationship between terminal symbols by considering the structure of derivations.

53. Consider a disk with the following specifications: rotation speed of 6000 RPM, average seek time of 5 milliseconds, 500 sectors/track, 512-byte sectors. A file has content stored in 3000 sectors located randomly on the disk. Assuming average rotational latency, the total time (in seconds, rounded off to 2 decimal places) to read the entire file from the disk is

Correct Answer: 29.5 to 30.5

Solution: The total time to read the file consists of:

- Seek time: Average seek time is 5 ms.
- Rotational latency: For a rotation speed of 6000 RPM, the time for one rotation is:

$$\text{Time per rotation} = \frac{60 \text{ sec}}{6000} = 0.01 \text{ sec.}$$

The average rotational latency is half the time per rotation:

$$\text{Average rotational latency} = \frac{0.01}{2} = 0.005 \text{ sec} = 5 \text{ ms.}$$

- Data transfer time per sector: Each track has 500 sectors. The time to transfer one sector is:

$$\text{Time per sector} = \frac{\text{Time per rotation}}{\text{Sectors per track}} = \frac{0.01}{500} = 0.00002 \text{ sec} = 20 \mu\text{s.}$$

- Total time per sector: Total time to access one sector is:

Time per sector = Seek time + Rotational latency + Transfer time.

$$\text{Time per sector} = 5 \text{ ms} + 5 \text{ ms} + 20 \mu\text{s} = 10.02 \text{ ms}.$$

- Total time for 3000 sectors:

$$\text{Total time} = 3000 \times 10.02 \text{ ms} = 30060 \text{ ms} = 30.06 \text{ sec}.$$

Final Answer:

29.5 to 30.5 sec

Quick Tip

To calculate disk access time, sum up the seek time, rotational latency, and data transfer time.

54. Consider a TCP connection operating at a point of time with the congestion window of size 12 MSS (Maximum Segment Size), when a timeout occurs due to packet loss. Assuming that all the segments transmitted in the next two RTTs (Round Trip Time) are acknowledged correctly, the congestion window size (in MSS) during the third RTT will be ____.

Correct Answer: 4

Solution: When a timeout occurs, TCP congestion control reduces the congestion window to 1 MSS and enters the slow-start phase. In slow start:

During the first RTT, the window size doubles: $1 \rightarrow 2$ MSS.

During the second RTT, the window size doubles again: $2 \rightarrow 4$ MSS.

Thus, at the start of the third RTT, the congestion window size will be 4 MSS.

Final Answer:

4 MSS

Quick Tip

After a timeout in TCP, the congestion window grows exponentially in the slow-start phase.

55. Consider an Ethernet segment with a transmission speed of 10^8 bits/sec and a maximum segment length of 500 meters. If the speed of propagation of the signal in the medium is 2×10^8 meters/sec, then the minimum frame size (in bits) required for collision detection is _____.

Correct Answer: 500

Solution: The round-trip time for a signal to propagate across the segment is:

$$\text{Time} = 2 \times \frac{\text{Distance}}{\text{Propagation speed}} = 2 \times \frac{500}{2 \times 10^8} = 5 \mu\text{s}.$$

To detect collisions, the frame transmission time must be at least equal to the round-trip time. The frame transmission time is:

$$\text{Frame transmission time} = \frac{\text{Frame size}}{\text{Transmission speed}}.$$

Equating the two:

$$\frac{\text{Frame size}}{10^8} = 5 \mu\text{s} \implies \text{Frame size} = 10^8 \times 5 \times 10^{-6} = 500 \text{ bits}.$$

Final Answer:

500 bits

Quick Tip

Collision detection in Ethernet requires the frame size to be large enough to cover the round-trip propagation delay.

56. A functional dependency $F : X \rightarrow Y$ is termed as a useful functional dependency if and only if it satisfies all the following three conditions:

- X is not the empty set.
- Y is not the empty set.
- Intersection of X and Y is the empty set.

For a relation R with 4 attributes, the total number of possible useful functional dependencies is _____.

Correct Answer: 50

Solution: For a relation R with 4 attributes $\{A, B, C, D\}$, we need to compute the total number of useful functional dependencies $X \rightarrow Y$, where:

- X is a non-empty subset of the attributes.
- Y is a non-empty subset of the attributes.
- X and Y are disjoint ($X \cap Y = \emptyset$).

Step 1: Total number of subsets of R The total number of subsets of $\{A, B, C, D\}$ is $2^4 = 16$. These subsets include:

$$\{\emptyset, \{A\}, \{B\}, \{C\}, \{D\}, \{A, B\}, \{A, C\}, \dots, \{A, B, C, D\}\}.$$

Excluding the empty set, there are 15 non-empty subsets.

Step 2: Counting valid combinations of X and Y Since $X \cap Y = \emptyset$, Y must be chosen from the attributes that are not in X . For a fixed X :

- The size of X determines the attributes available for Y .
- If X contains k attributes, then $4 - k$ attributes are available for Y .

For every non-empty subset X (15 options), the number of possible non-empty subsets of the remaining attributes (Y) is:

$$2^{4-|X|} - 1.$$

Step 3: Summing over all valid combinations To find the total number of useful functional

dependencies, sum over all non-empty subsets X :

$$\text{Total functional dependencies} = \sum_{k=1}^4 \binom{4}{k} \cdot (2^{4-k} - 1),$$

where $\binom{4}{k}$ is the number of ways to choose k attributes for X .

Step 4: Calculate each term For $k = 1, 2, 3, 4$:

- For $k = 1$:

$$\binom{4}{1} \cdot (2^{4-1} - 1) = 4 \cdot (8 - 1) = 4 \cdot 7 = 28.$$

- For $k = 2$:

$$\binom{4}{2} \cdot (2^{4-2} - 1) = 6 \cdot (4 - 1) = 6 \cdot 3 = 18.$$

- For $k = 3$:

$$\binom{4}{3} \cdot (2^{4-3} - 1) = 4 \cdot (2 - 1) = 4 \cdot 1 = 4.$$

- For $k = 4$:

$$\binom{4}{4} \cdot (2^{4-4} - 1) = 1 \cdot (1 - 1) = 0.$$

Step 5: Add the results The total number of useful functional dependencies is:

$$28 + 18 + 4 + 0 = 50.$$

Final Answer:

50

Quick Tip

Useful functional dependencies ensure X and Y are disjoint, non-empty subsets.

57. A processor with 16 general-purpose registers uses a 32-bit instruction format. The instruction format consists of an opcode field, an addressing mode field, two register operand fields, and a 16-bit scalar field. If 8 addressing modes are to be supported, the maximum number of unique opcodes possible for every addressing mode is ____.

Correct Answer: 32

Solution: The 32-bit instruction format is divided as follows:

16 general-purpose registers: Each register requires 4 bits ($2^4 = 16$).

Two register operand fields: Each operand field is 4 bits, for a total of 8 bits.

Addressing mode field: 3 bits to represent 8 addressing modes ($2^3 = 8$).

Scalar field: 16 bits.

The remaining bits are allocated to the opcode field:

$$\text{Opcode bits} = 32 - (8 + 3 + 16) = 5 \text{ bits.}$$

The total number of unique opcodes for every addressing mode is:

$$2^{\text{Opcode bits}} = 2^5 = 32.$$

Final Answer:

32

Quick Tip

When determining the opcode size, subtract the bits required for other fields from the total instruction size.

58. A non-pipelined instruction execution unit operating at 2 GHz takes an average of 6 cycles to execute an instruction of a program P . The unit is then redesigned to operate on a 5-stage pipeline at 2 GHz. Assume that the ideal throughput of the pipelined unit is 1 instruction per cycle. In the execution of program P , 20% instructions incur an average of 2 cycles stall due to data hazards and 20% instructions incur an average of 3 cycles stall due to control hazards. The speedup (rounded off to one decimal place) obtained by the pipelined design over the non-pipelined design is ____.

Correct Answer: 2.9 to 3.1

Solution: Step 1: Define throughput for a pipelined processor.

The throughput of a pipelined processor is the number of instructions coming out from the last stage of the pipeline per unit time. In this case, the time unit is 1 clock cycle, which

is 0.5 ns. This means that each stage in the pipelined processor takes only 1 clock cycle to operate (ignoring register delay and clock skew).

Step 2: Define CPI for pipelined and non-pipelined processors.

The CPI (Clock Cycles Per Instruction) is defined as the average clock cycles per instruction. For a non-pipelined processor, it takes 6 clock cycles to complete an instruction, whereas for a pipelined processor, it takes only 1 clock cycle on average to complete an instruction. Thus:

$$\text{CPI (non-pipelined)} = 6, \quad \text{CPI (pipelined)} = 1.$$

Step 3: Define the speedup of the pipelined processor.

The speedup of the pipelined processor compared to the non-pipelined processor is given by:

$$\text{Speedup} = \frac{\text{CPI (non-pipelined)}}{\text{Ideal CPI (pipelined)} + \text{Pipeline stall clock cycles}}.$$

Here, the pipeline stall clock cycles are added because some clock cycles are wasted due to stalls in the pipeline.

Step 4: Calculate the speedup.

For a program with a total instruction count of IC , the pipeline stall clock cycles are due to 20% of the instructions stalling for 2 cycles and another 20% of the instructions stalling for 3 cycles. The effective CPI for the pipelined processor becomes:

$$\text{Effective CPI (pipelined)} = 1 + (20\% \times 2) + (20\% \times 3).$$

Substitute the values:

$$\text{Effective CPI (pipelined)} = 1 + 0.4 + 0.6 = 2.$$

The speedup is then calculated as:

$$\text{Speedup} = \frac{\text{CPI (non-pipelined)}}{\text{Effective CPI (pipelined)}} = \frac{6}{2} = 3.$$

Step 5: Conclude the result.

The pipelined processor is 3.0 times faster than the non-pipelined processor.


Final Answer:

The pipelined processor is 3.0 times faster than the non-pipelined processor.

Quick Tip

When calculating speedup, always consider the impact of hazards and stalls in pipelined designs.

59. The number of distinct minimum-weight spanning trees of the following graph is



Correct Answer: 9

Solution: Step 1: Analyze the graph and find the minimum weight. The graph has 7 edges with weights: 2, 2, 2, 3, 3, 3, 1. To construct a minimum spanning tree (MST), the sum of edge weights must be minimized.

Step 2: Apply Kruskal's algorithm. Using Kruskal's algorithm:

- Select the edge with weight 1 (unique choice).
- Choose three edges of weight 2. These edges form a cycle, allowing multiple choices.
- Choose one edge of weight 3 to complete the MST.

Step 3: Count the distinct combinations.

- There are $\binom{3}{2} = 3$ ways to choose two edges of weight 2.
- For each choice, there are 3 ways to select one edge of weight 3.

The total number of distinct MSTs is:

$$3 \times 3 = 9.$$


Final Answer:

9

Quick Tip

When determining the number of distinct MSTs, consider the flexibility in selecting edges of the same weight without violating the MST property.

60. The chromatic number of a graph is the minimum number of colours used in a proper colouring of the graph. The chromatic number of the following graph is



Correct Answer: 2

Solution: Step 1: Definition of chromatic number. The chromatic number of a graph is the smallest number of colours required to colour its vertices such that no two adjacent vertices share the same colour.

Step 2: Analyze the graph structure. The given graph is a bipartite graph, as the vertices can be divided into two disjoint sets such that:

- All edges connect vertices from one set to the other.
- There are no edges between vertices within the same set.

A bipartite graph can always be coloured using exactly 2 colours.

Step 3: Verify bipartite property. The graph has two disjoint sets of vertices:

Set 1: Outer vertices. Set 2: Inner vertices.

All edges connect vertices from Set 1 to Set 2, satisfying the bipartite condition. Thus, the chromatic number of the graph is 2.

Final Answer:

2

Quick Tip

Bipartite graphs always have a chromatic number of 2, as their vertices can be divided into two independent sets.

61. A processor uses a 32-bit instruction format and supports byte-addressable memory access. The ISA of the processor has 150 distinct instructions. The instructions are equally divided into two types, namely R-type and I-type, whose formats are shown below.

R-type Instruction Format:

| | | | | |
|--------|--------|--------------|---------------|---------------|
| OPCODE | UNUSED | DST Register | SRC Register1 | SRC Register2 |
|--------|--------|--------------|---------------|---------------|

I-type Instruction Format:

| | | | |
|--------|--------------|--------------|---------------------------|
| OPCODE | DST Register | SRC Register | # Immediate Value/Address |
|--------|--------------|--------------|---------------------------|

In the OPCODE, 1 bit is used to distinguish between I-type and R-type instructions, and the remaining bits indicate the operation. The processor has 50 architectural registers, and all register fields in the instructions are of equal size.

Let X be the number of bits used to encode the UNUSED field, Y be the number of bits used to encode the OPCODE field, and Z be the number of bits used to encode the immediate value/address field. The value of $X + 2Y + Z$ is ____.

Correct Answer: 34

Solution: Step 1: Instruction format analysis. Each instruction is 32 bits in total. Based on the formats:

- **R-type instruction fields:**

$$\text{OPCODE} + \text{UNUSED} + \text{DST Register} + \text{SRC Register1} + \text{SRC Register2} = 32 \text{ bits.}$$

- **I-type instruction fields:**

$$\text{OPCODE} + \text{DST Register} + \text{SRC Register} + \text{Immediate Value/Address} = 32 \text{ bits.}$$

Step 2: Determine the number of bits for the registers. The processor has 50 architectural registers. The number of bits required to uniquely identify a register is:

$$\text{Register bits} = \lceil \log_2(50) \rceil = 6 \text{ bits.}$$

Thus, each register field (DST Register, SRC Register1, SRC Register2) uses 6 bits.

Step 3: Determine the number of bits for the OPCODE field. The ISA contains 150 distinct instructions, equally divided into R-type and I-type. Since 1 bit in the OPCODE field is used to distinguish between instruction types, the remaining bits encode the operations:

$$\text{OPCODE bits} = \lceil \log_2(150) \rceil = 8 \text{ bits.}$$

Step 4: Calculate the number of bits for the UNUSED field (R-type) and Immediate Value/Address field (I-type).

- **R-type:** From the total 32 bits, the UNUSED field is:

$$X = 32 - (\text{OPCODE} + 3 \times \text{Register bits}) = 32 - (8 + 3 \times 6) = 6 \text{ bits.}$$

- **I-type:** The Immediate Value/Address field uses:

$$Z = 32 - (\text{OPCODE} + 2 \times \text{Register bits}) = 32 - (8 + 2 \times 6) = 12 \text{ bits.}$$

Step 5: Compute $X + 2Y + Z$. Substituting the values of X , Y , and Z :

$$X + 2Y + Z = 6 + 2 \times 8 + 12 = 34.$$

Final Answer:

34

Quick Tip

The OPCODE size is determined by the total number of instructions. Ensure register and immediate field sizes fit within the 32-bit instruction format.

62. Let L_1 be the language represented by the regular expression $b^*ab^*(ab^*ab^*)^*$ and $L_2 = \{w \in (a+b)^* \mid |w| \leq 4\}$, where $|w|$ denotes the length of string w . The number of strings in L_2 which are also in L_1 is

Correct Answer: 15

Solution:

Step 1: Define the set of language symbols.

Let the set of language symbols be:

$$\Sigma = \{a, b\} \implies |\Sigma| = 2.$$

Step 2: Define the regular language L_1 .

The regular language L_1 is defined as the set of strings where the number of a 's is odd.

Step 3: Analyze the strings with an odd number of a 's.

For all strings with the alphabet set $\{a, b\}$ of length n , half of them will have an odd number of a 's. This can be expressed as:

$$\frac{2^n}{2},$$

where the numerator 2^n represents the total number of strings of length n (based on the number of input symbols).

Step 4: Calculate the total number of strings in L_1 for all lengths up to 4.

We sum up half of the total strings of all lengths (up to 4) in L_1 :

$$\frac{2^0}{2} + \frac{2^1}{2} + \frac{2^2}{2} + \frac{2^3}{2} + \frac{2^4}{2}.$$

Simplify the terms:

$$0 + 1 + 2 + 4 + 8 = 15.$$

Final Answer: The total number of strings in L_1 (up to length 4) is:

$$\boxed{15}.$$

Quick Tip

When intersecting a regular expression and a finite language, enumerate all strings in the finite language and check which satisfy the regular expression.

63. Let Z_n be the group of integers $\{0, 1, 2, \dots, n-1\}$ with addition modulo n as the group operation. The number of elements in the group $Z_2 \times Z_3 \times Z_4$ that are their own inverses is

Correct Answer: 4

Solution: Step 1: Define the property of an element being its own inverse. An element x in a group is its own inverse if:

$$x + x \equiv 0 \pmod{n}.$$

This simplifies to:

$$2x \equiv 0 \pmod{n}.$$

Step 2: Analyze each group.

- Z_2 : The elements are $\{0, 1\}$.

$$2x \equiv 0 \pmod{2} \implies x = 0, 1.$$

- Z_3 : The elements are $\{0, 1, 2\}$.

$$2x \equiv 0 \pmod{3} \implies x = 0.$$

- Z_4 : The elements are $\{0, 1, 2, 3\}$.

$$2x \equiv 0 \pmod{4} \implies x = 0, 2.$$

Step 3: Combine the results for $Z_2 \times Z_3 \times Z_4$. The total number of elements in $Z_2 \times Z_3 \times Z_4$ is:

$$|Z_2 \times Z_3 \times Z_4| = 2 \times 3 \times 4 = 24.$$

For an element $(x_2, x_3, x_4) \in Z_2 \times Z_3 \times Z_4$ to be its own inverse:

- $x_2 \in \{0, 1\}$ (2 choices).
- $x_3 \in \{0\}$ (1 choice).
- $x_4 \in \{0, 2\}$ (2 choices).

The total number of such elements is:

$$2 \times 1 \times 2 = 4.$$

Final Answer:

4

Quick Tip

To find elements that are their own inverses in a group, solve $2x \equiv 0$ modulo the group order for each component.

64. Consider a 32-bit system with a 4 KB page size and page table entries of size 4 bytes each. Assume $1 \text{ KB} = 2^{10}$ bytes. The OS uses a 2-level page table for memory management, with the page table containing an outer page directory and an inner page table. The OS allocates a page for the outer page directory upon process creation. The OS uses demand paging when allocating memory for the inner page table, i.e., a page of the inner page table is allocated only if it contains at least one valid page table entry. An active process in this system accesses 2000 unique pages during its execution, and none of the pages are swapped out to disk. After it completes the page accesses, let X denote the minimum and Y denote the maximum number of pages across the two levels of the page table of the process. The value of $X + Y$ is ____.

Correct Answer: 1028

Solution:

Given:

- System address space (VA) = 32 bits
- Page size (PS) = 4 KB
- Total pages used by a process (N) = 2000 pages
- Page table entry size (PTE) = 4 B

It is given in the question that the OS allocates a **page** (means single page) for the **outer page directory** upon process creation.

Step 1: Calculate the total number of bits needed for the outer page table.

The total number of bits needed to represent entries in the outer page table is:

$$\log \left(\frac{PS}{PTE} \right) = \log \left(\frac{2^{12}}{2^2} \right) = \log(2^{10}) = 10 \text{ bits.}$$

Step 2: Calculate the total number of bits needed for the inner page table.

The total number of bits needed to represent entries in the inner page table is:

$$32 \text{ (total VA bits)} - 10 \text{ (outer PT bits)} - 12 \text{ (page offset bits)} = 10 \text{ bits.}$$

Step 3: Calculate X .

In each inner page table, we can store 2^{10} entries. The minimum number of page tables (PT) needed to store 2000 pages will be:

$$\lceil \frac{2000}{2^{10}} \rceil = 2 \text{ } PT.$$

Thus:

$$X = 2 + 1 \text{ (for the outer PT)} = 3.$$

Step 4: Calculate Y .

For occupying the maximum pages in a 2-level page table, we need to place at least one PTE from 2000 pages in every page of the inner page table. This equates to 2^{10} PT . Thus:

$$Y = 2^{10} + 1 \text{ (for the outer PT)} = 1025.$$

Step 5: Calculate $X + Y$.

The value of $X + Y$ is:

$$X + Y = 3 + 1025 = 1028.$$

Final Answer:

$$\boxed{1028}.$$

Quick Tip

Always calculate the number of page table entries per inner page and analyze the allocation policy for minimum and maximum scenarios.

65. Consider the following augmented grammar, which is to be parsed with an SLR parser. The set of terminals is $\{a, b, c, d, \#, @\}$.

$$S' \rightarrow S \quad S \rightarrow SS \mid Aa \mid bAc \mid BC \mid bBa \quad A \rightarrow d\# \quad B \rightarrow @$$

Let $I_0 = \text{CLOSURE}(\{S' \rightarrow \cdot S\})$. The number of items in the set $\text{GOTO}(I_0, S)$ is _____.

Correct Answer: 9

Solution:

Step 1: Analyze the given DFA diagram.

The given DFA represents the items in the closure of I_0 and the GOTO function for the symbol S .

Step 2: Identify the items in the closure of I_0 .

The closure of I_0 contains the following items:

- $S \rightarrow \cdot S$
- $S \rightarrow \cdot SS$
- $S \rightarrow \cdot Aa$
- $S \rightarrow \cdot bAc$
- $S \rightarrow \cdot Bc$
- $S \rightarrow \cdot bAb$
- $A \rightarrow \cdot d\#$
- $B \rightarrow \cdot @$

Step 3: Apply the GOTO function on the symbol S .

When the GOTO function is applied to S , the following transitions occur:

- $S \rightarrow S \cdot$
- $S \rightarrow S \cdot S$
- $S \rightarrow SS \cdot$
- $S \rightarrow Aa \cdot$
- $S \rightarrow bAc \cdot$
- $S \rightarrow Bc \cdot$
- $S \rightarrow bAb \cdot$

- $A \rightarrow d\#.$
- $B \rightarrow @.$

Step 4: Count the total number of items.

From the diagram, the GOTO function $GOTO(\text{closure}(I_0), S)$ results in 9 items:

$S \rightarrow S., S \rightarrow S \cdot S, S \rightarrow SS., S \rightarrow Aa., S \rightarrow bAc., S \rightarrow Bc., S \rightarrow bAb., A \rightarrow d\#., B \rightarrow @. .$

Final Answer: From the above DFA, we can conclude that:

$GOTO(\text{closure}(I_0), S)$ contains 9 items.

Quick Tip

To compute the number of items in a GOTO set, always analyze the closure operation and include all applicable productions for the given grammar symbol.