

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 [dry → arid → parched] is analogous to [diet → fast → _____].

Which one of the given options is appropriate to fill the blank?

- (A) starve
- (B) reject
- (C) feast
- (D) deny

Q.2 If two distinct non-zero real variables x and y are such that $(x + y)$ is proportional to $(x - y)$ then the value of $\frac{x}{y}$

- (A) depends on xy
- (B) depends only on x and not on y
- (C) depends only on y and not on x
- (D) is a constant

Q.3 Consider the following sample of numbers:

9, 18, 11, 14, 15, 17, 10, 69, 11, 13

The median of the sample is

- (A) 13.5
- (B) 14
- (C) 11
- (D) 18.7

Q.4 The number of coins of ₹1, ₹5, and ₹10 denominations that a person has are in the ratio 5:3:13. Of the total amount, the percentage of money in ₹5 coins is

- (A) 21%
- (B) $14\frac{2}{7}\%$
- (C) 10%
- (D) 30%

Q.5 For positive non-zero real variables p and q , if

$$\log (p^2 + q^2) = \log p + \log q + 2 \log 3,$$

then, the value of $\frac{p^4 + q^4}{p^2 q^2}$ is

- (A) 79
- (B) 81
- (C) 9
- (D) 83

GATE 2024

Q.6 – Q.10 Carry TWO marks Each

Q.6 In the given text, the blanks are numbered (i)–(iv). Select the best match for all the blanks.

Steve was advised to keep his head _____ (i) _____ before heading _____ (ii) _____ to bat; for, while he had a head _____ (iii) _____ batting, he could only do so with a cool head _____ (iv) _____ his shoulders.


- (A) (i) down (ii) down (iii) on (iv) for
- (B) (i) on (ii) down (iii) for (iv) on
- (C) (i) down (ii) out (iii) for (iv) on
- (D) (i) on (ii) out (iii) on (iv) for

Q.7 A rectangular paper sheet of dimensions $54 \text{ cm} \times 4 \text{ cm}$ is taken. The two longer edges of the sheet are joined together to create a cylindrical tube. A cube whose surface area is equal to the area of the sheet is also taken.

Then, the ratio of the volume of the cylindrical tube to the volume of the cube is

- (A) $1/\pi$
- (B) $2/\pi$
- (C) $3/\pi$
- (D) $4/\pi$

Q.8 The pie chart presents the percentage contribution of different macronutrients to a typical 2,000 kcal diet of a person.



The typical energy density (kcal/g) of these macronutrients is given in the table.

Macronutrient	Energy density (kcal/g)
Carbohydrates	4
Proteins	4
Unsaturated fat	9
Saturated fat	9
Trans fat	9

The total fat (all three types), in grams, this person consumes is

- (A) 44.4
- (B) 77.8
- (C) 100
- (D) 3,600

Q.11 – Q.35 Carry ONE mark Each

Q.11 Let $f: \mathbb{R} \rightarrow \mathbb{R}$ be a function such that $f(x) = \max\{x, x^3\}$, $x \in \mathbb{R}$, where \mathbb{R} is the set of all real numbers. The set of all points where $f(x)$ is NOT differentiable is

- (A) $\{-1, 1, 2\}$
- (B) $\{-2, -1, 1\}$
- (C) $\{0, 1\}$
- (D) $\{-1, 0, 1\}$

Q.12 The product of all eigenvalues of the matrix $\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix}$ is

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

Q.13 Consider a system that uses 5 bits for representing signed integers in 2's complement format. In this system, two integers A and B are represented as $A=01010$ and $B=11010$. Which one of the following operations will result in either an arithmetic overflow or an arithmetic underflow?

(A) $A + B$

(B) $A - B$

(C) $B - A$

(D) $2 * B$

Q.14 Consider a permutation sampled uniformly at random from the set of all permutations of $\{1, 2, 3, \dots, n\}$ for some $n \geq 4$. Let X be the event that 1 occurs before 2 in the permutation, and Y the event that 3 occurs before 4. Which one of the following statements is TRUE?

(A) The events X and Y are mutually exclusive

(B) The events X and Y are independent

(C) Either event X or Y must occur

(D) Event X is more likely than event Y

Q.15 Which one of the following statements is FALSE?

- (A) In the cycle stealing mode of DMA, one word of data is transferred between an I/O device and main memory in a stolen cycle
- (B) For bulk data transfer, the burst mode of DMA has a higher throughput than the cycle stealing mode
- (C) Programmed I/O mechanism has a better CPU utilization than the interrupt driven I/O mechanism
- (D) The CPU can start executing an interrupt service routine faster with vectored interrupts than with non-vectored interrupts

Q.16 A user starts browsing a webpage hosted at a remote server. The browser opens a single TCP connection to fetch the entire webpage from the server. The webpage consists of a top-level index page with multiple embedded image objects. Assume that all caches (e.g., DNS cache, browser cache) are all initially empty. The following packets leave the user's computer in some order.

- (i) HTTP GET request for the index page
- (ii) DNS request to resolve the web server's name to its IP address
- (iii) HTTP GET request for an image object
- (iv) TCP SYN to open a connection to the web server

Which one of the following is the CORRECT chronological order (earliest in time to latest) of the packets leaving the computer ?

- (A) (iv), (ii), (iii), (i)
- (B) (ii), (iv), (iii), (i)
- (C) (ii), (iv), (i), (iii)
- (D) (iv), (ii), (i), (iii)

Q.17 Given an integer array of size N , we want to check if the array is sorted (in either ascending or descending order). An algorithm solves this problem by making a single pass through the array and comparing each element of the array only with its adjacent elements. The worst-case time complexity of this algorithm is

- (A) both $O(N)$ and $\Omega(N)$
- (B) $O(N)$ but not $\Omega(N)$
- (C) $\Omega(N)$ but not $O(N)$
- (D) neither $O(N)$ nor $\Omega(N)$

Q.18 Consider the following C program:

```
#include <stdio.h>

int main() {
    int a = 6;
    int b = 0;
    while(a < 10) {
        a = a / 12 + 1;
        a += b;}
    printf("%d", a);
    return 0;}
```

Which one of the following statements is CORRECT?

- (A) The program prints 9 as output
- (B) The program prints 10 as output
- (C) The program gets stuck in an infinite loop
- (D) The program prints 6 as output

Q.19 Consider the following C program:





```
#include <stdio.h>
void fX();
int main(){
    fX();
    return 0;}

void fX(){
    char a;
    if((a=getchar()) != '\n')
        fX();
    if(a != '\n')
        putchar(a);}
```

Assume that the input to the program from the command line is 1234 followed by a newline character. Which one of the following statements is CORRECT?

- (A) The program will not terminate
- (B) The program will terminate with no output
- (C) The program will terminate with 4321 as output
- (D) The program will terminate with 1234 as output

Q.20 Let S be the specification: "Instructors teach courses. Students register for courses. Courses are allocated classrooms. Instructors guide students." Which one of the following ER diagrams CORRECTLY represents S?



- (A) (i)
- (B) (ii)
- (C) (iii)
- (D) (iv)

Q.21 In a B+ tree, the requirement of at least half-full (50%) node occupancy is relaxed for which one of the following cases?

- (A) Only the root node
- (B) All leaf nodes
- (C) All internal nodes
- (D) Only the leftmost leaf node

Q.22 Which of the following statements about a relation **R** in first normal form (1NF) is/are TRUE ?

- (A) **R** can have a multi-attribute key
- (B) **R** cannot have a foreign key
- (C) **R** cannot have a composite attribute
- (D) **R** cannot have more than one candidate key

Q.23 Let L_1, L_2 be two regular languages and L_3 a language which is not regular. Which of the following statements is/are *always* TRUE?

- (A) $L_1 = L_2$ if and only if $L_1 \cap \overline{L_2} = \phi$
- (B) $L_1 \cup L_3$ is not regular
- (C) $\overline{L_3}$ is not regular
- (D) $\overline{L_1} \cup \overline{L_2}$ is regular

Q.24 Which of the following statements about threads is/are TRUE?

- (A) Threads can only be implemented in kernel space
- (B) Each thread has its own file descriptor table for open files
- (C) All the threads belonging to a process share a common stack
- (D) Threads belonging to a process are by default not protected from each other

Q.25 Which of the following process state transitions is/are NOT possible?

- (A) Running to Ready
- (B) Waiting to Running
- (C) Ready to Waiting
- (D) Running to Terminated


Q.26 Which of the following is/are Bottom-Up Parser(s)?

- (A) Shift-reduce Parser
- (B) Predictive Parser
- (C) LL(1) Parser
- (D) LR Parser

Q.27 Let A and B be two events in a probability space with $P(A) = 0.3$, $P(B) = 0.5$, and $P(A \cap B) = 0.1$. Which of the following statements is/are TRUE?

- (A) The two events A and B are independent
- (B) $P(A \cup B) = 0.7$
- (C) $P(A \cap B^c) = 0.2$, where B^c is the complement of the event B
- (D) $P(A^c \cap B^c) = 0.4$, where A^c and B^c are the complements of the events A and B , respectively

Q.28 Consider the circuit shown below where the gates may have propagation delays. Assume that all signal transitions occur instantaneously and that wires have no delays. Which of the following statements about the circuit is/are CORRECT?



- (A) With no propagation delays, the output Y is always logic Zero
- (B) With no propagation delays, the output Y is always logic One
- (C) With propagation delays, the output Y can have a transient logic One after X transitions from logic Zero to logic One
- (D) With propagation delays, the output Y can have a transient logic Zero after X transitions from logic One to logic Zero

Q.29 TCP client P successfully establishes a connection to TCP server Q. Let N_P denote the sequence number in the SYN sent from P to Q. Let N_Q denote the acknowledgement number in the SYN ACK from Q to P. Which of the following statements is/are CORRECT?

- (A) The sequence number N_P is chosen randomly by P
- (B) The sequence number N_P is always 0 for a new connection
- (C) The acknowledgement number N_Q is equal to N_P
- (D) The acknowledgement number N_Q is equal to $N_P + 1$

Q.30 Consider a 5-stage pipelined processor with Instruction Fetch (IF), Instruction Decode (ID), Execute (EX), Memory Access (MEM), and Register Writeback (WB) stages. Which of the following statements about *forwarding* is/are CORRECT?

- (A) In a pipelined execution, forwarding means the result from a source stage of an earlier instruction is passed on to the destination stage of a later instruction
- (B) In forwarding, data from the output of the MEM stage can be passed on to the input of the EX stage of the next instruction
- (C) Forwarding cannot prevent all pipeline stalls
- (D) Forwarding does not require any extra hardware to retrieve the data from the pipeline stages

Q.31 Which of the following fields is/are modified in the IP header of a packet going out of a network address translation (NAT) device from an internal network to an external network?

- (A) Source IP
- (B) Destination IP
- (C) Header Checksum
- (D) Total Length

Q.32 Let A and B be non-empty finite sets such that there exist one-to-one and onto functions (i) from A to B and (ii) from $A \times A$ to $A \cup B$. The number of possible values of $|A|$ is _____

Q.33 Consider the operator precedence and associativity rules for the *integer* arithmetic operators given in the table below.

Operator	Precedence	Associativity
+	Highest	Left
-	High	Right
*	Medium	Right
/	Low	Right

The value of the expression $3 + 1 + 5 * 2 / 7 + 2 - 4 - 7 - 6 / 2$ as per the above rules is _____

Q.34 The number of spanning trees in a *complete* graph of 4 vertices labelled A, B, C, and D is _____

Q.35 Consider the following two relations, $R(A,B)$ and $S(A,C)$:

R	
A	B
10	20
20	30
30	40
30	50
50	95

S	
A	C
10	90
30	45
40	80

The total number of tuples obtained by evaluating the following expression

$$\sigma_{B < C}(R \bowtie_{R.A=S.A} S)$$

is _____

Q.36 – Q.65 Carry TWO marks Each

Q.36 Consider a network path P—Q—R between nodes P and R via router Q. Node P sends a file of size 10^6 bytes to R via this path by splitting the file into chunks of 10^3 bytes each. Node P sends these chunks one after the other without any wait time between the successive chunk transmissions. Assume that the size of extra headers added to these chunks is negligible, and that the chunk size is less than the MTU.

Each of the links P—Q and Q—R has a bandwidth of 10^6 bits/sec, and negligible propagation latency. Router Q immediately transmits every packet it receives from P to R, with negligible processing and queueing delays. Router Q can simultaneously receive on link P—Q and transmit on link Q—R.

Assume P starts transmitting the chunks at time $t = 0$. Which one of the following options gives the time (*in seconds, rounded off to 3 decimal places*) at which R receives all the chunks of the file?

- (A) 8.000
- (B) 8.008
- (C) 15.992
- (D) 16.000

Q.37 Consider the following syntax-directed definition (SDD).

$S \rightarrow DHTU$	$\{ S.val = D.val + H.val + T.val + U.val; \}$
$D \rightarrow "M"D_1$	$\{ D.val = 5 + D_1.val; \}$
$D \rightarrow \epsilon$	$\{ D.val = -5; \}$
$H \rightarrow "L"H_1$	$\{ H.val = 5 * 10 + H_1.val; \}$
$H \rightarrow \epsilon$	$\{ H.val = -10; \}$
$T \rightarrow "C"T_1$	$\{ T.val = 5 * 100 + T_1.val; \}$
$T \rightarrow \epsilon$	$\{ T.val = -5; \}$
$U \rightarrow "K"$	$\{ U.val = 5; \}$

Given "MMLK" as the input, which one of the following options is the CORRECT value computed by the SDD (in the attribute $S.val$)?

- (A) 45
- (B) 50
- (C) 55
- (D) 65

- Q.38 Consider the following grammar G , with S as the start symbol. The grammar G has three incomplete productions denoted by (1), (2), and (3).

$$\begin{aligned} S &\rightarrow daT \mid \underline{\hspace{1cm}} \text{ (1)} \\ T &\rightarrow aS \mid bT \mid \underline{\hspace{1cm}} \text{ (2)} \\ R &\rightarrow \underline{\hspace{1cm}} \text{ (3)} \mid \epsilon \end{aligned}$$

The set of terminals is $\{a, b, c, d, f\}$. The FIRST and FOLLOW sets of the different non-terminals are as follows.

$$\text{FIRST}(S) = \{c, d, f\}, \text{FIRST}(T) = \{a, b, \epsilon\}, \text{FIRST}(R) = \{c, \epsilon\}$$

$$\text{FOLLOW}(S) = \text{FOLLOW}(T) = \{c, f, \$\}, \text{FOLLOW}(R) = \{f\}$$

Which one of the following options CORRECTLY fills in the incomplete productions?

- (A) (1) $S \rightarrow Rf$ (2) $T \rightarrow \epsilon$ (3) $R \rightarrow cTR$
- (B) (1) $S \rightarrow fR$ (2) $T \rightarrow \epsilon$ (3) $R \rightarrow cTR$
- (C) (1) $S \rightarrow fR$ (2) $T \rightarrow cT$ (3) $R \rightarrow cR$
- (D) (1) $S \rightarrow Rf$ (2) $T \rightarrow cT$ (3) $R \rightarrow cR$

Q.39 Consider the following pseudo-code.

```
L1:      t1 = -1
L2:      t2 = 0
L3:      t3 = 0
L4:      t4 = 4 * t3
L5:      t5 = 4 * t2
L6:      t6 = t5 * M
L7:      t7 = t4 + t6
L8:      t8 = a[t7]
L9:      if t8 <= max goto L11
L10:     t1 = t8
L11:     t3 = t3 + 1
L12:     if t3 < M goto L4
L13:     t2 = t2 + 1
L14:     if t2 < N goto L3
L15:     max = t1
```

Which one of the following options CORRECTLY specifies the number of basic blocks and the number of instructions in the largest basic block, respectively ?

- (A) 6 and 6
- (B) 6 and 7
- (C) 7 and 7
- (D) 7 and 6

- Q.40 Consider the following two threads T1 and T2 that update two shared variables a and b. Assume that initially $a = b = 1$. Though context switching between threads can happen at any time, each statement of T1 or T2 is executed atomically without interruption.

T1	T2
$a = a + 1;$	$b = 2 * b;$
$b = b + 1;$	$a = 2 * a;$

Which one of the following options lists all the possible combinations of values of a and b after both T1 and T2 finish execution?

- (A) $(a = 4, b = 4); (a = 3, b = 3); (a = 4, b = 3)$
- (B) $(a = 3, b = 4); (a = 4, b = 3); (a = 3, b = 3)$
- (C) $(a = 4, b = 4); (a = 4, b = 3); (a = 3, b = 4)$
- (D) $(a = 2, b = 2); (a = 2, b = 3); (a = 3, b = 4)$

- Q.41 An array [82, 101, 90, 11, 111, 75, 33, 131, 44, 93] is heapified. Which one of the following options represents the first three elements in the heapified array?

- (A) 82, 90, 101
- (B) 82, 11, 93
- (C) 131, 11, 93
- (D) 131, 111, 90

Q.42 Consider the following recurrence relation:

$$T(n) = \begin{cases} \sqrt{n}T(\sqrt{n}) + n & \text{for } n \geq 1, \\ 1 & \text{for } n = 1. \end{cases}$$

Which one of the following options is CORRECT?

- (A) $T(n) = \Theta(n \log \log n)$
- (B) $T(n) = \Theta(n \log n)$
- (C) $T(n) = \Theta(n^2 \log n)$
- (D) $T(n) = \Theta(n^2 \log \log n)$

Q.43 Consider a binary min-heap containing 105 distinct elements. Let k be the index (in the underlying array) of the maximum element stored in the heap. The number of possible values of k is

- (A) 53
- (B) 52
- (C) 27
- (D) 1

Q.44 The symbol \rightarrow indicates functional dependency in the context of a relational database. Which of the following options is/are TRUE?

- (A) $(X, Y) \rightarrow (Z, W)$ implies $X \rightarrow (Z, W)$
- (B) $(X, Y) \rightarrow (Z, W)$ implies $(X, Y) \rightarrow Z$
- (C) $((X, Y) \rightarrow Z$ and $W \rightarrow Y)$ implies $(X, W) \rightarrow Z$
- (D) $(X \rightarrow Y$ and $Y \rightarrow Z)$ implies $X \rightarrow Z$

Q.45 Let G be a directed graph and T a depth first search (DFS) spanning tree in G that is rooted at a vertex v . Suppose T is also a breadth first search (BFS) tree in G , rooted at v . Which of the following statements is/are TRUE for *every* such graph G and tree T ?

- (A) There are no back-edges in G with respect to the tree T
- (B) There are no cross-edges in G with respect to the tree T
- (C) There are no forward-edges in G with respect to the tree T
- (D) The only edges in G are the edges in T

Q.46 Consider the following read-write schedule S over three transactions T_1 , T_2 , and T_3 , where the subscripts in the schedule indicate transaction IDs:

$$S: r_1(z); w_1(z); r_2(x); r_3(y); w_3(y); r_2(y); w_2(x); w_2(y);$$

Which of the following transaction schedules is/are conflict equivalent to S ?

(A) $T_1T_2T_3$

(B) $T_1T_3T_2$

(C) $T_3T_2T_1$

(D) $T_3T_1T_2$

Q.47 Consider a Boolean expression given by $F(X, Y, Z) = \sum(3,5,6,7)$.

Which of the following statements is/are CORRECT?

(A) $F(X, Y, Z) = \Pi(0,1,2,4)$

(B) $F(X, Y, Z) = XY + YZ + XZ$

(C) $F(X, Y, Z)$ is independent of input Y

(D) $F(X, Y, Z)$ is independent of input X

Q.48 Consider the following C function definition.

```
int f(int x, int y) {  
    for (int i=0; i<y; i++) {  
        x=x+x+y;  
    }  
    return x;  
}
```


Which of the following statements is/are TRUE about the above function?

- (A) If the inputs are $x=20$, $y=10$, then the return value is greater than 2^{20}
- (B) If the inputs are $x=20$, $y=20$, then the return value is greater than 2^{20}
- (C) If the inputs are $x=20$, $y=10$, then the return value is less than 2^{10}
- (D) If the inputs are $x=10$, $y=20$, then the return value is greater than 2^{20}

Q.49 Let A be any $n \times m$ matrix, where $m > n$. Which of the following statements is/are TRUE about the system of linear equations $Ax = \mathbf{0}$?

- (A) There exist at least $m - n$ linearly independent solutions to this system
- (B) There exist $m - n$ linearly independent vectors such that every solution is a linear combination of these vectors
- (C) There exists a non-zero solution in which at least $m - n$ variables are 0
- (D) There exists a solution in which at least n variables are non-zero

- Q.50 Consider the 5-state DFA M accepting the language $L(M) \subset (0 + 1)^*$ shown below. For any string $w \in (0 + 1)^*$ let $n_0(w)$ be the number of 0's in w and $n_1(w)$ be the number of 1's in w .



Which of the following statements is/are FALSE?

- (A) States 2 and 4 are distinguishable in M
- (B) States 3 and 4 are distinguishable in M
- (C) States 2 and 5 are distinguishable in M
- (D) Any string w with $n_0(w) = n_1(w)$ is in $L(M)$

Q.51 The chromatic number of a graph is the minimum number of colours used in a *proper* colouring of the graph. Let G be any graph with n vertices and chromatic number k . Which of the following statements is/are always TRUE?

- (A) G contains a complete subgraph with k vertices
- (B) G contains an independent set of size at least n/k
- (C) G contains at least $k(k - 1)/2$ edges
- (D) G contains a vertex of degree at least k

Q.52 Consider the operators \diamond and \square defined by $a \diamond b = a + 2b$, $a \square b = ab$, for positive integers. Which of the following statements is/are TRUE?

- (A) Operator \diamond obeys the associative law
- (B) Operator \square obeys the associative law
- (C) Operator \diamond over the operator \square obeys the distributive law
- (D) Operator \square over the operator \diamond obeys the distributive law

Q.53 Consider two set-associative cache memory architectures: **WBC**, which uses the write back policy, and **WTC**, which uses the write through policy. Both of them use the LRU (*Least Recently Used*) block replacement policy. The cache memory is connected to the main memory. Which of the following statements is/are TRUE?

- (A) A read miss in **WBC** never evicts a dirty block
- (B) A read miss in **WTC** never triggers a write back operation of a cache block to main memory
- (C) A write hit in **WBC** can modify the value of the dirty bit of a cache block
- (D) A write miss in **WTC** always writes the victim cache block to main memory before loading the missed block to the cache

Q.54 Consider a 512 GB hard disk with 32 storage surfaces. There are 4096 sectors per track and each sector holds 1024 bytes of data. The number of cylinders in the hard disk is _____

Q.55 The baseline execution time of a program on a 2 GHz single core machine is 100 nanoseconds (*ns*). The code corresponding to 90% of the execution time can be fully parallelized. The overhead for using an additional core is 10 *ns* when running on a multicore system. Assume that all cores in the multicore system run their share of the parallelized code for an equal amount of time.

The number of cores that minimize the execution time of the program is _____

Q.56 A given program has 25% load/store instructions. Suppose the ideal CPI (cycles per instruction) without any memory stalls is 2. The program exhibits 2% miss rate on instruction cache and 8% miss rate on data cache. The miss penalty is 100 cycles. The speedup (*rounded off to two decimal places*) achieved with a perfect cache (i.e., with NO data or instruction cache misses) is _____

Q.57 Consider the following code snippet using the `fork()` and `wait()` system calls. Assume that the code compiles and runs correctly, and that the system calls run successfully without any errors.

```
int x = 3;
while(x > 0) {
    fork();
    printf("hello");
    wait(NULL);
    x--;
}
```

The total number of times the `printf` statement is executed is _____

Q.58 Consider the entries shown below in the forwarding table of an IP router. Each entry consists of an IP prefix and the corresponding next hop router for packets whose destination IP address matches the prefix. The notation “/N” in a prefix indicates a subnet mask with the most significant N bits set to 1.

Prefix	Next hop router
10.1.1.0/24	R1
10.1.1.128/25	R2
10.1.1.64/26	R3
10.1.1.192/26	R4

This router forwards 20 packets each to 5 hosts. The IP addresses of the hosts are 10.1.1.16, 10.1.1.72, 10.1.1.132, 10.1.1.191, and 10.1.1.205 . The number of packets forwarded via the next hop router R2 is _____

Q.59 Let $G = (V, \Sigma, S, P)$ be a context-free grammar in Chomsky Normal Form with $\Sigma = \{a, b, c\}$ and V containing 10 variable symbols including the start symbol S . The string $w = a^{30}b^{30}c^{30}$ is derivable from S . The number of steps (application of rules) in the derivation $S \rightarrow^* w$ is _____

Q.60 The number of edges present in the forest generated by the DFS traversal of an undirected graph G with 100 vertices is 40. The number of connected components in G is _____

Q.61 Consider the following two regular expressions over the alphabet $\{0,1\}$:

$$r = 0^* + 1^*$$


$$s = 01^* + 10^*$$

The total number of strings of length less than or equal to 5, which are neither in r nor in s , is _____

Q.62 Consider a memory management system that uses a page size of 2 KB. Assume that both the physical and virtual addresses start from 0. Assume that the pages 0, 1, 2, and 3 are stored in the page frames 1, 3, 2, and 0, respectively. The physical address (*in decimal format*) corresponding to the virtual address 2500 (*in decimal format*) is _____

Q.63 A bag contains 10 red balls and 15 blue balls. Two balls are drawn randomly without replacement. Given that the first ball drawn is red, the probability (*rounded off to 3 decimal places*) that both balls drawn are red is _____

Q.64 Consider a digital logic circuit consisting of three 2-to-1 multiplexers M1, M2, and M3 as shown below. X1 and X2 are inputs of M1. X3 and X4 are inputs of M2. A, B, and C are select lines of M1, M2, and M3, respectively.



For an instance of inputs $X1=1$, $X2=1$, $X3=0$, and $X4=0$, the number of combinations of A, B, C that give the output $Y=1$ is _____

Q.65 Consider sending an IP datagram of size 1420 bytes (including 20 bytes of IP header) from a sender to a receiver over a path of two links with a router between them. The first link (sender to router) has an MTU (Maximum Transmission Unit) size of 542 bytes, while the second link (router to receiver) has an MTU size of 360 bytes. The number of fragments that would be delivered at the receiver is _____

GATE 2024 CS and IT Set 1 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 [dry → arid → parched] is analogous to [diet → fast → -----].

- (1) starve
- (2) reject
- (3) feast
- (4) deny

Correct Answer: (1) starve

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

dry → arid → parched.

This progression moves from a low level of dryness (dry) to extreme dryness (parched).

Similarly, in the second analogy:

diet → fast → ----.

Step 1: Analyze the terms. "Diet" implies restricting certain foods. "Fast" intensifies this by avoiding all food for a specific time. The next level of intensity would be "starve," which represents an extreme and prolonged lack of food.

Step 2: Eliminate incorrect options. - (2) "reject" and (4) "deny" do not align with the context of increasing food restriction. - (3) "feast" is the opposite of restriction.

Final Answer:

starve

Quick Tip

When solving analogies, focus on the relationship described in the first set and apply it consistently to the second set.

2. If two distinct non-zero real variables x and y are such that $(x + y)$ is proportional to $(x - y)$, then the value of $\frac{x}{y}$ is:

- (1) depends on xy
- (2) depends only on x and not on y
- (3) depends only on y and not on x
- (4) is a constant

Correct Answer: (4) is a constant

Solution: The given condition states that $(x + y)$ is proportional to $(x - y)$, which means:

$$x + y = k(x - y),$$

where k is a proportionality constant.

Step 1: Simplify the equation. Rewriting the equation:

$$x + y = kx - ky.$$

Rearranging terms:

$$x - kx = -ky - y.$$

Factoring:

$$x(1 - k) = -y(1 + k).$$

Step 2: Solve for $\frac{x}{y}$. Divide both sides by $y(1 - k)$ (assuming $1 - k \neq 0$):

$$\frac{x}{y} = -\frac{1 + k}{1 - k}.$$

Step 3: Interpret the result. Since k is a constant, $\frac{x}{y}$ is also a constant. It does not depend on x or y individually.

Final Answer:

is a constant

Quick Tip

Proportional relationships often simplify to constants. Always isolate terms to identify constants or dependencies.

3. Consider the following sample of numbers:

9, 18, 11, 14, 15, 17, 10, 69, 11, 13

The median of the sample is:

- (1) 13.5
- (2) 14
- (3) 11
- (4) 18.7

Correct Answer: (1) 13.5

Solution: First, sort the given data in ascending order:

9, 10, 11, 11, 13, 14, 15, 17, 18, 69

Since there are 10 numbers (even count), the median is the average of the 5th and 6th values:

$$\frac{13 + 14}{2} = 13.5.$$

Final Answer:

13.5

Quick Tip

To find the median of an even number of values, average the two middle values.

4. The number of coins of Rs.1, Rs.5, and Rs.10 denominations that a person has are in the ratio 5:3:13. Of the total amount, the percentage of money in Rs.5 coins is:

- (1) 21%
- (2) $14\frac{2}{7}\%$
- (3) 10%
- (4) 30%

Correct Answer: (3) 10%

Solution: Let the number of coins be proportional to $5x, 3x, 13x$.

Total amount:

$$5x \cdot 1 + 3x \cdot 5 + 13x \cdot 10 = 5x + 15x + 130x = 150x.$$

Value contributed by Rs.5 coins:

$$3x \cdot 5 = 15x.$$

Percentage of total amount:

$$\frac{15x}{150x} \times 100 = 10\%.$$

Final Answer:

10%

Quick Tip

To find the percentage of a specific denomination, calculate its contribution to the total value and multiply by 100.

5. For positive non-zero real variables p and q , if

$$\log(p^2 + q^2) = \log p + \log q + 2 \log 3,$$

then, the value of $\frac{p^4 + q^4}{p^2 q^2}$ is:

- (1) 79
- (2) 81
- (3) 9
- (4) 83

Correct Answer: (1) 79

Solution: Given the equation:

$$\log(p^2 + q^2) = \log p + \log q + 2 \log 3$$

Rewrite using logarithm properties:

$$\log(p^2 + q^2) = \log(pq) + \log 9$$

This simplifies to:

$$p^2 + q^2 = 9pq$$

Squaring both sides:

$$(p^2 + q^2)^2 = (9pq)^2$$

$$p^4 + q^4 + 2p^2q^2 = 81p^2q^2$$

$$p^4 + q^4 = 81p^2q^2 - 2p^2q^2$$

$$p^4 + q^4 = 79p^2q^2$$

Dividing by p^2q^2 :

$$\frac{p^4 + q^4}{p^2q^2} = 79$$

Final Answer:

79

Quick Tip

Use logarithmic properties to simplify equations and express them in terms of multiplication or addition.

Question 6 - 10 carry two mark each

6. In the given text, the blanks are numbered (i)–(iv). Select the best match for all the blanks.

Steve was advised to keep his head ___ (i) ___ before heading ___ (ii) ___ to bat; for, while he had a head ___ (iii) ___ batting, he could only do so with a cool head ___ (iv) ___ his shoulders.

- (1) down, down, for, for
- (2) on, down, for, on
- (3) down, out, for, on
- (4) on, out, on, for

Correct Answer: (3) down, out, for, on

Solution: Let us analyze the sentence:

Steve was advised to keep his head **down** before heading **out** to bat – This means staying focused and cautious.

For, while he had a head **for** batting – This indicates he had a natural aptitude or talent for batting.

He could only do so with a cool head **on** his shoulders – This phrase means remaining calm and composed under pressure.

Thus, the correct sequence of words to fill the blanks is:

- (i) down, (ii) out, (iii) for, (iv) on

Final Answer:

down, out, for, on

Quick Tip

Understanding the context and common idiomatic expressions in English helps to choose the correct prepositions and phrases.

7. A rectangular paper sheet of dimensions $54\text{ cm} \times 4\text{ cm}$ is taken. The two longer edges of the sheet are joined together to create a cylindrical tube. A cube whose surface area is equal to the area of the sheet is also taken.

Then, the ratio of the volume of the cylindrical tube to the volume of the cube is:

- (1) $\frac{1}{\pi}$
- (2) $\frac{2}{\pi}$
- (3) $\frac{3}{\pi}$

$$(4) \frac{4}{\pi}$$

Correct Answer: (1) $\frac{1}{\pi}$

Solution: The problem involves comparing the volumes of a cube and a cylinder derived from the same surface area.

Step 1: Surface area of the sheet

$$\text{Area of sheet} = 54 \times 4 = 216 \text{ cm}^2$$

This area is given as the surface area of the cube.

Step 2: Side of the cube Let the side of the cube be a . The surface area of a cube is:

$$6a^2 = 216$$
$$a^2 = \frac{216}{6} = 36 \quad \implies \quad a = 6 \text{ cm.}$$

Hence, the **volume of the cube** is:

$$\text{Volume of cube} = a^3 = 6^3 = 216 \text{ cm}^3.$$

Step 3: Radius of the cylinder The same surface area is used to form a cylinder. The surface area of the cylinder includes the lateral surface area and the areas of the two circular bases:

$$\text{Circumference of top and bottom circles} = 4 \quad (\text{from the lateral surface}).$$

The circumference is given by:

$$2\pi R_c = 4 \quad \implies \quad R_c = \frac{2}{\pi}.$$

Step 4: Volume of the cylinder The volume of the cylinder is:

$$\text{Volume of cylinder} = \pi R_c^2 h,$$

where $h = 54$ cm (height of the cylinder). Substituting $R_c = \frac{2}{\pi}$:

$$\text{Volume of cylinder} = \pi \left(\frac{2}{\pi}\right)^2 \times 54 = \pi \cdot \frac{4}{\pi^2} \cdot 54 = \frac{4 \cdot 54}{\pi} = \frac{216}{\pi} \text{ cm}^3.$$

Step 5: Ratio of volumes The ratio of the volumes of the cylinder to the cube is:

$$\text{Ratio} = \frac{\text{Volume of cylinder}}{\text{Volume of cube}} = \frac{\frac{216}{\pi}}{216} = \frac{1}{\pi}$$

Final Answer:


$$\frac{1}{\pi}$$

Quick Tip

When converting a rectangle to a cylinder, the length forms the circumference, and the height remains constant.

8. The pie chart presents the percentage contribution of different macronutrients to a typical 2,000 kcal diet of a person. The typical energy density (kcal/g) of these macronutrients is given in the table.

Macronutrient energy contribution



Macronutrient	Energy density (kcal/g)
Carbohydrates	4
Proteins	4
Unsaturated fat	9
Saturated fat	9
Trans fat	9

The total fat (all three types), in grams, this person consumes is:

- (1) 44.4
- (2) 77.8
- (3) 100
- (4) 3,600

Correct Answer: (3) 100

Solution: Step 1: Calculate total fat contribution in kcal. The total fat (saturated, unsaturated, and trans fat) contributes:

$$20\% + 20\% + 5\% = 45\%.$$

$$\text{Total fat contribution} = 0.45 \times 2000 = 900 \text{ kcal.}$$

Step 2: Convert kcal to grams. Since the energy density of fat (all types) is 9 kcal/g:

$$\text{Total fat consumed} = \frac{900}{9} = 100 \text{ g.}$$

Final Answer:

100

Quick Tip

To calculate total macronutrient intake in grams, divide the total kcal contribution by the energy density of the nutrient.

9. A rectangular paper of 20 cm × 8 cm is folded 3 times. Each fold is made along the line of symmetry, which is perpendicular to its long edge. The perimeter of the final folded sheet (in cm) is:

- (1) 18
- (2) 24
- (3) 20
- (4) 21

Correct Answer: (1) 18

Solution: Step 1: Dimensions after each fold. Initially, the dimensions of the paper are $20\text{ cm} \times 8\text{ cm}$.

First fold: The paper is folded along the 20 cm side, resulting in dimensions $10\text{ cm} \times 8\text{ cm}$.

Second fold: The paper is folded along the new 10 cm side, resulting in dimensions $5\text{ cm} \times 8\text{ cm}$.

Third fold: The paper is folded along the new 8 cm side, resulting in dimensions $5\text{ cm} \times 4\text{ cm}$.

Step 2: Calculate the perimeter. The perimeter of the final folded sheet is:

$$P = 2 \times (\text{length} + \text{breadth}) = 2 \times (5 + 4) = 18\text{ cm}.$$


Final Answer:

18 cm

Quick Tip

When calculating the perimeter after folding, remember that folding changes only one dimension at a time.

10. The least number of squares to be added in the figure to make AB a line of symmetry is:



- (1) 6
- (2) 4
- (3) 5
- (4) 7

Correct Answer: (1) 6

Solution: Step 1: Analyze the figure. The line AB must be a line of symmetry, meaning that for every square above or below AB, there must be a corresponding square on the opposite side, symmetrically placed.

Step 2: Add squares. To make the figure symmetric about AB:

Add 2 squares to the leftmost side below AB.

Add 2 squares to the middle section below AB.

Add 2 squares to the rightmost side below AB.

Step 3: Total squares to be added. The least number of squares required is:

6 squares.

Final Answer:

6 squares

Quick Tip

For symmetry problems, focus on mirroring all elements relative to the axis of symmetry.

Question 11 - 35 carry one mark each

11. Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a function such that $f(x) = \max\{x, x^3\}$, $x \in \mathbb{R}$, where \mathbb{R} is the set of all real numbers. The set of all points where $f(x)$ is NOT differentiable is:

- (1) $\{-1, 1, 2\}$
- (2) $\{-2, -1, 1\}$
- (3) $\{0, 1\}$
- (4) $\{-1, 0, 1\}$

Correct Answer: (4) $\{-1, 0, 1\}$

Solution: Step 1: Analyze the function $f(x)$. The function is defined as:

$$f(x) = \max\{x, x^3\}.$$

This implies:

For $x \geq 1$, $f(x) = x$, as $x > x^3$.

For $-1 \leq x \leq 1$, $f(x) = x^3$, as $x^3 > x$.

For $x \leq -1$, $f(x) = x$, as $x > x^3$.

Step 2: Points of non-differentiability. The function $f(x)$ is not differentiable at the points where the two components x and x^3 intersect. Solve $x = x^3$:

$$x^3 - x = 0 \implies x(x^2 - 1) = 0 \implies x = 0, x = 1, x = -1.$$

Final Answer:

$$\{-1, 0, 1\}$$

Quick Tip

To find points of non-differentiability for a piecewise-defined function, analyze where the two components intersect or where the derivative changes abruptly.

12. The product of all eigenvalues of the matrix

$$\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix} \text{ is:}$$

- (1) -1
- (2) 0
- (3) 1
- (4) 2

Correct Answer: (2) 0

Solution: Step 1: Determine the product of eigenvalues. The product of all eigenvalues of a square matrix is equal to the determinant of the matrix.

Step 2: Calculate the determinant.

$$\det \left(\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix} \right) = 1 \begin{vmatrix} 5 & 6 \\ 8 & 9 \end{vmatrix} - 2 \begin{vmatrix} 4 & 6 \\ 7 & 9 \end{vmatrix} + 3 \begin{vmatrix} 4 & 5 \\ 7 & 8 \end{vmatrix}.$$

Evaluate each minor determinant:

$$\begin{vmatrix} 5 & 6 \\ 8 & 9 \end{vmatrix} = (5)(9) - (6)(8) = 45 - 48 = -3,$$

$$\begin{vmatrix} 4 & 6 \\ 7 & 9 \end{vmatrix} = (4)(9) - (6)(7) = 36 - 42 = -6,$$

$$\begin{vmatrix} 4 & 5 \\ 7 & 8 \end{vmatrix} = (4)(8) - (5)(7) = 32 - 35 = -3.$$

Substitute back:

$$\det = 1(-3) - 2(-6) + 3(-3) = -3 + 12 - 9 = 0.$$

Step 3: Conclusion. The determinant is 0, so the product of eigenvalues is 0.

Quick Tip

The determinant of a matrix is a quick way to find the product of all eigenvalues. If the determinant is zero, at least one eigenvalue is zero.

13. Consider a system that uses 5 bits for representing signed integers in 2's complement format. In this system, two integers A and B are represented as $A = 01010$ and $B = 11010$. Which one of the following operations will result in either an arithmetic overflow or an arithmetic underflow?

- (1) $A + B$
- (2) $A - B$
- (3) $B - A$
- (4) $2 * B$

Correct Answer: (2) $A - B$

Solution: Step 1: Understand 5-bit signed integer representation. The range of a 5-bit signed integer in 2's complement is:

$$[-16, 15].$$

Convert A and B to decimal:

$A = 01010$ (positive): $A = 10$.

$B = 11010$ (negative): $B = -6$.

Step 2: Perform the operations. 1. $A + B$:

$$10 + (-6) = 4 \quad (\text{within range, no overflow/underflow}).$$

2. $A - B$:

$$10 - (-6) = 10 + 6 = 16 \quad (\text{exceeds the range, causes overflow}).$$

3. $B - A$:

$$-6 - 10 = -16 \quad (\text{within range, no overflow/underflow}).$$

4. $2 * B$:

$$2 \times (-6) = -12 \quad (\text{within range, no overflow/underflow}).$$

Step 3: Identify the problematic operation. The operation $A - B$ causes an overflow as the result exceeds the maximum representable value 15.

Final Answer:

$$\boxed{A - B}$$

Quick Tip

In a 2's complement system, arithmetic overflow occurs if the result exceeds the range $[-2^{n-1}, 2^{n-1} - 1]$, where n is the number of bits.

14. Consider a permutation sampled uniformly at random from the set of all permutations of $\{1, 2, 3, \dots, n\}$ for some $n \geq 4$. Let X be the event that 1 occurs before 2 in the permutation, and Y the event that 3 occurs before 4. Which one of the following statements is TRUE?

- (1) The events X and Y are mutually exclusive.
- (2) The events X and Y are independent.
- (3) Either event X or Y must occur.
- (4) Event X is more likely than event Y .

Correct Answer: (2) The events X and Y are independent.

Solution: Step 1: Define the events X and Y .

Event X : 1 occurs before 2 in the permutation.

Event Y : 3 occurs before 4 in the permutation.

Step 2: Determine probabilities.

For any two distinct elements in a permutation, the probability of one occurring before the other is:

$$P(X) = P(Y) = \frac{1}{2}.$$

Step 3: Check independence.

Events X and Y are independent if:

$$P(X \cap Y) = P(X) \cdot P(Y).$$

Since the order of 1 and 2 is independent of the order of 3 and 4, we have:

$$P(X \cap Y) = \frac{1}{4}, \quad P(X) \cdot P(Y) = \frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}.$$

Step 4: Conclusion. Events X and Y are independent as the probabilities match.

Final Answer:

The events X and Y are independent.

Quick Tip

Independence of events means the outcome of one event does not influence the probability of the other.

15. Which one of the following statements is FALSE?

(1) In the cycle stealing mode of DMA, one word of data is transferred between an I/O device and main memory in a stolen cycle.

(2) For bulk data transfer, the burst mode of DMA has a higher throughput than the cycle stealing mode.

(3) Programmed I/O mechanism has a better CPU utilization than the interrupt-driven I/O mechanism.

(4) The CPU can start executing an interrupt service routine faster with vectored interrupts than with non-vectored interrupts.

Correct Answer: (3) Programmed I/O mechanism has a better CPU utilization than the interrupt-driven I/O mechanism.

Solution: Step 1: Analyze the statements.

1. **Cycle stealing mode in DMA:** In this mode, the DMA controller temporarily "steals" cycles from the CPU to transfer one word of data. This is a valid statement.

2. **Burst mode in DMA:** Burst mode transfers large blocks of data at once, resulting in higher throughput compared to cycle stealing mode. This is also a valid statement.

3. **Programmed I/O vs. Interrupt-driven I/O:** Programmed I/O requires the CPU to continuously poll the device, leading to inefficient CPU utilization. Interrupt-driven I/O, on the other hand, allows the CPU to perform other tasks until the device signals it. Hence, this statement is FALSE.

4. **Vectored interrupts:** Vectored interrupts have predefined addresses for the interrupt service routines, allowing faster execution compared to non-vectored interrupts, which require additional steps to determine the address. This statement is TRUE.

Final Answer:

Programmed I/O mechanism has a better CPU utilization than the interrupt-driven I/O mechanism.

Quick Tip

Programmed I/O involves continuous polling by the CPU, leading to inefficient use of resources. Interrupt-driven I/O allows better CPU utilization by enabling multitasking.

16. A user starts browsing a webpage hosted at a remote server. The browser opens a single TCP connection to fetch the entire webpage from the server. The webpage consists of a top-level index page with multiple embedded image objects. Assume that all caches

(e.g., DNS cache, browser cache) are all initially empty. The following packets leave the user's computer in some order:

1. HTTP GET request for the index page
2. DNS request to resolve the web server's name to its IP address
3. HTTP GET request for an image object
4. TCP SYN to open a connection to the web server

Which one of the following is the **CORRECT** chronological order (earliest in time to latest) of the packets leaving the computer?

- (1) (iv), (ii), (iii), (i)
- (2) (ii), (iv), (iii), (i)
- (3) (ii), (iv), (i), (iii)
- (4) (iv), (ii), (i), (iii)

Correct Answer: (3) (ii), (iv), (i), (iii)

Solution:Step 1: DNS request (ii): The client first sends a DNS request to resolve the server's hostname to its IP address, as no cache is available.

Step 2: TCP SYN (iv): After the DNS response, a TCP SYN packet is sent to establish a connection with the server.

Step 3: HTTP GET request for the index page (i): Once the connection is established, an HTTP GET request for the index page is sent.

Step 4: HTTP GET request for image objects (iii): After receiving the index page, HTTP GET requests for the embedded image objects are sent.

Final Answer:

(ii), (iv), (i), (iii)

Quick Tip

The sequence of operations while accessing a webpage without cached data involves DNS resolution, connection establishment, fetching the main content, and then retrieving embedded objects.

17. Given an integer array of size N , we want to check if the array is sorted (in either ascending or descending order). An algorithm solves this problem by making a single pass through the array and comparing each element of the array only with its adjacent elements. The worst-case time complexity of this algorithm is:

- (1) both $O(N)$ and $\Omega(N)$
- (2) $O(N)$ but not $\Omega(N)$
- (3) $\Omega(N)$ but not $O(N)$
- (4) neither $O(N)$ nor $\Omega(N)$

Correct Answer: (1) both $O(N)$ and $\Omega(N)$

Solution:Step 1: Define the problem. The algorithm compares adjacent elements of the array to check if it is sorted in either ascending or descending order. This requires a single pass through the array.

Step 2: Time complexity analysis.

In the worst case, the algorithm must traverse the entire array of size N , making $N - 1$ comparisons.

Thus, the time complexity is $O(N)$, as the number of operations is proportional to N .

Since the algorithm always requires $N - 1$ comparisons, the lower bound is also $\Omega(N)$.

Final Answer:

both $O(N)$ and $\Omega(N)$

Quick Tip

For algorithms requiring a single traversal through an array, the time complexity is both $O(N)$ (upper bound) and $\Omega(N)$ (lower bound).

18. Consider the following C program:

```
#include <stdio.h>
int main() {
    int a = 6;
```

```

int b = 0;
while(a < 10) {
    a = a / 12 + 1;
    a += b;
}
printf("%d", a);
return 0;
}

```

Which one of the following statements is CORRECT?

- (1) The program prints 9 as output.
- (2) The program prints 10 as output.
- (3) The program gets stuck in an infinite loop.
- (4) The program prints 6 as output.

Correct Answer: (3) The program gets stuck in an infinite loop.

Solution: Step 1: Analyze the program.

Initially, $a = 6$ and $b = 0$.

Inside the `while` loop, the value of a is updated as:

$$a = \frac{a}{12} + 1.$$

In integer arithmetic, dividing a (initially 6) by 12 results in 0 since only the integer part is retained. Hence:

$$a = 0 + 1 = 1.$$

The next statement is:

$$a += b.$$

Since $b = 0$, a remains 1.

Step 2: Infinite loop condition.

The condition for the `while` loop is $a < 10$. After the first iteration:

$$a = 1 \quad (\text{remains unchanged as } b = 0).$$

In subsequent iterations, a is repeatedly updated to 1 (because $a/12 + 1 = 1$), and the loop never terminates.

Final Answer:

The program gets stuck in an infinite loop.

Quick Tip

In C, integer division truncates the decimal part, which can lead to unintended infinite loops if the update condition depends on such operations.

19. Consider the following C program:

```
#include <stdio.h>
void fX();
int main() {
    fX();
    return 0;
}
```

```
void fX() {
    char a;
    if((a = getchar()) != '\n')
        fX();
    if(a != '\n')
        putchar(a);
}
```

Assume that the input to the program from the command line is 1234 followed by a newline character. Which one of the following statements is CORRECT?

(1) The program will not terminate.

- (2) The program will terminate with no output.
- (3) The program will terminate with 4321 as output.
- (4) The program will terminate with 1234 as output.

Correct Answer: (3) The program will terminate with 4321 as output.

Solution: Step 1: Understand the recursive function.

The function `fX()` reads a character using `getchar()`.

If the character is not a newline ("`\n`"), `fX()` is called recursively.

After returning from the recursive call, the character is printed using `putchar()` if it is not a newline.

Step 2: Analyze the flow of recursion.

When `fX()` is first called, it reads the first character, '1', and calls itself recursively.

The second call reads '2', and again calls itself recursively.

This continues for '3' and '4'. When `getchar()` reads the newline ("`\n`"), the recursion stops and begins to unwind.

Step 3: Output generation during unwinding.

As the recursion unwinds, each character ('4', '3', '2', '1') is printed in reverse order of their input due to the stack behavior of recursion.





Final Answer:

The program will terminate with 4321 as output.

Quick Tip

In recursive functions with delayed output, the stack unwinding order determines the sequence of printed characters.

20. Let S be the specification: "Instructors teach courses. Students register for courses. Courses are allocated classrooms. Instructors guide students." Which one of the following ER diagrams CORRECTLY represents S ?



- (1) Diagram (i)
- (2) Diagram (ii)
- (3) Diagram (iii)
- (4) Diagram (iv)

Correct Answer: (4) Diagram (iv)

Solution: The correct ER diagram must fulfill the following requirements:

1. Instructors teach courses: The diagram should include a relationship named "Teaches" between "Instructor" and "Course."
2. Students register for courses: There must be an entity for "Student" and a relationship named "Registration" linking "Student" and "Course."
3. Courses are allocated classrooms: There should be a relationship named "Allocation" between "Course" and "Classroom."
4. Instructors guide students: The diagram must include a relationship named "Guides" between "Instructor" and "Student."

Analysis of the Options:

Option (i): Does not include the "Registration" relationship between "Student" and "Course."
Incorrect.

Option (ii): Contains an additional entity "Projects," which is not part of the specification.
Incorrect.

Option (iii): Does not include the relationship "Guides" between "Instructor" and "Student."
Incorrect.

Option (iv): Includes all the required entities and relationships as per the specification. Correct.

Final Answer:

Diagram (iv)

Quick Tip

When analyzing ER diagrams, ensure all relationships and entities in the specification are represented without introducing extraneous components.

21. In a B+ tree, the requirement of at least half-full (50%) node occupancy is relaxed for which one of the following cases?

- (1) Only the root node
- (2) All leaf nodes
- (3) All internal nodes

(4) Only the leftmost leaf node

Correct Answer: (1) Only the root node

Solution:

Step 1: Define the conditions for node occupancy in a B+ tree.

In a B+ tree, all nodes except the root must maintain at least 50% occupancy to ensure the balanced structure of the tree.

Step 2: Identify the exception.

The root node is an exception to the occupancy rule. It is allowed to have fewer keys, including only one key, during operations such as deletions.

Final Answer:

Only the root node

Quick Tip

In a B+ tree, the root node is exempt from the minimum occupancy rule, but all other nodes must maintain at least 50% occupancy.

22. Which of the following statements about a relation R in first normal form (1NF) is/are TRUE?

- (1) R can have a multi-attribute key
- (2) R cannot have a foreign key
- (3) R cannot have a composite attribute
- (4) R cannot have more than one candidate key

Correct Answer: (1) and (3)

Solution:Step 1: Analyze the statements.

1. R can have a multi-attribute key: Multi-attribute keys are allowed in 1NF as long as the attributes are atomic. This statement is **true**.
2. R cannot have a foreign key: This statement is **false**. Foreign keys are allowed in 1NF as they maintain referential integrity.

3. R cannot have a composite attribute: In 1NF, attributes must store atomic values, so composite attributes are not allowed. This statement is **true**.
4. R cannot have more than one candidate key: A relation in 1NF can have multiple candidate keys. This statement is **false**.

Final Answer:

(1) and (3)

Quick Tip

In 1NF, atomicity of attributes is essential. Composite attributes are not allowed, but multi-attribute keys and foreign keys are permitted.

23. Let L_1, L_2 be two regular languages and L_3 a language which is not regular. Which of the following statements is/are always TRUE?

- (1) $L_1 = L_2$ if and only if $L_1 \cap \overline{L_2} = \phi$
- (2) $L_1 \cup L_3$ is not regular
- (3) $\overline{L_3}$ is not regular
- (4) $\overline{L_1} \cup \overline{L_2}$ is regular

Correct Answer: (3) and (4)

Solution: Step 1: Analyze $L_1 = L_2$ condition.

For $L_1 = L_2$, the condition $L_1 \cap \overline{L_2} = \phi$ must hold. However, this is not a sufficient condition to prove equality, so statement (1) is **false**.

Step 2: Analyze $L_1 \cup L_3$.

If L_3 is non-regular, $L_1 \cup L_3$ may or may not be non-regular, depending on L_3 . Hence, statement (2) is **false**.

Step 3: Analyze L_3 .

Since L_3 is given as non-regular, statement (3) is **true**.

Step 4: Analyze $\overline{L_1} \cup \overline{L_2}$.

The complement and union of two regular languages are regular. Thus, $\overline{L_1} \cup \overline{L_2}$ is **regular**, making statement (4) **true**.

Final Answer:

(3) and (4)

Quick Tip

For regular languages, operations such as complement, union, and intersection preserve regularity, while non-regular languages may not guarantee these properties.

24. Which of the following statements about threads is/are TRUE?

- (1) Threads can only be implemented in kernel space
- (2) Each thread has its own file descriptor table for open files
- (3) All the threads belonging to a process share a common stack
- (4) Threads belonging to a process are by default not protected from each other

Correct Answer: (4)

Solution: Step 1: Threads in kernel/user space.

Threads can be implemented either in kernel space (kernel threads) or in user space (user threads). Statement (1) is **false**.

Step 2: File descriptor table.

Threads of the same process share the same file descriptor table, making statement (2) **false**.

Step 3: Shared resources.

Threads belonging to the same process share memory, file descriptors, and other resources. However, each thread has its own stack. Thus, statement (3) is **false**.

Step 4: Protection between threads.

Threads of a process share the same address space and are not protected from each other. Hence, statement (4) is **true**.

Final Answer:

(4)

Quick Tip

Threads share memory and resources within the same process but are not protected from interfering with each other by default.

25. Which of the following process state transitions is/are NOT possible?

- (1) Running to Ready
- (2) Waiting to Running
- (3) Ready to Waiting
- (4) Running to Terminated

Correct Answer: (2) Waiting to Running, (3) Ready to Waiting

Solution:

Step 1: Understand process state transitions. Process states include:

- **Ready:** The process is ready to run but is waiting for CPU time.
- **Running:** The process is currently executing on the CPU.
- **Waiting:** The process is waiting for an event (e.g., I/O completion).
- **Terminated:** The process has finished execution.

Step 2: Analyze each transition.

1. **Running to Ready:** This is possible when the process is preempted by the scheduler due to a higher-priority process or time slice expiration.
2. **Waiting to Running:** This is NOT possible. A process must first transition from the Waiting state to the Ready state when the event it is waiting for completes. It can then move from Ready to Running.
3. **Ready to Waiting:** This is NOT possible. A process in the Ready state can only transition to the Running state when scheduled. Waiting occurs after execution begins (e.g., when it requests I/O).
4. **Running to Terminated:** This is possible when the process completes its execution.

Step 3: Conclude the impossible transitions.

From the above analysis, the transitions Waiting to Running and Ready to Waiting are NOT possible.

Final Answer:

(2) Waiting to Running, (3) Ready to Waiting

Quick Tip

A process must pass through the Ready state to transition from Waiting to Running. Similarly, a process cannot directly transition from Ready to Waiting.

26. Which of the following is/are Bottom-Up Parser(s)?

- (1) Shift-reduce Parser
- (2) Predictive Parser
- (3) LL(1) Parser
- (4) LR Parser

Correct Answer: (1) Shift-reduce Parser, (4) LR Parser

Solution:

Step 1: Understand parsing techniques. Parsing techniques are broadly categorized as:

Top-Down Parsing: Constructs the parse tree starting from the root (start symbol) and moves toward the leaves.

Bottom-Up Parsing: Constructs the parse tree starting from the leaves (input string) and moves toward the root.

Step 2: Analyze the given parsers.

1. **Shift-reduce Parser:** This is a bottom-up parsing technique. It uses shift and reduce operations to construct the parse tree. Thus, it is a Bottom-Up Parser.
2. **Predictive Parser:** This is a top-down parsing technique. It uses lookahead symbols to predict which production to apply. It is NOT a Bottom-Up Parser.
3. **LL(1) Parser:** This is a top-down parsing technique based on leftmost derivations with one symbol of lookahead. It is NOT a Bottom-Up Parser.

4. **LR Parser:** This is a bottom-up parsing technique that uses left-to-right scanning and rightmost derivation in reverse. It is a Bottom-Up Parser.

Step 3: Conclude the correct parsers.

From the analysis, Shift-reduce Parser and LR Parser are Bottom-Up Parsers.

Final Answer:

(1) Shift-reduce Parser, (4) LR Parser

Quick Tip

Bottom-up parsers construct the parse tree starting from the input string (leaves) and work their way toward the start symbol (root). Examples include Shift-reduce Parser and LR Parser.

27. Let A and B be two events in a probability space with $P(A) = 0.3$, $P(B) = 0.5$, and $P(A \cap B) = 0.1$. Which of the following statements is/are TRUE?

- (1) The two events A and B are independent
- (2) $P(A \cup B) = 0.7$
- (3) $P(A \cap B^c) = 0.2$, where B^c is the complement of the event B
- (4) $P(A^c \cap B^c) = 0.4$, where A^c and B^c are the complements of the events A and B , respectively

Correct Answer: (2) $P(A \cup B) = 0.7$, (3) $P(A \cap B^c) = 0.2$

Solution:

Step 1: Check independence of events A and B . Two events A and B are independent if:

$$P(A \cap B) = P(A) \cdot P(B).$$

Here:

$$P(A \cap B) = 0.1, \quad P(A) \cdot P(B) = 0.3 \cdot 0.5 = 0.15.$$

Since $P(A \cap B) \neq P(A) \cdot P(B)$, the events are NOT independent. Option (1) is FALSE.

Step 2: Compute $P(A \cup B)$. Using the formula:

$$P(A \cup B) = P(A) + P(B) - P(A \cap B),$$

substitute the values:

$$P(A \cup B) = 0.3 + 0.5 - 0.1 = 0.7.$$

Thus, Option (2) is TRUE.

Step 3: Compute $P(A \cap B^c)$. Using the complement rule:

$$P(A \cap B^c) = P(A) - P(A \cap B).$$

Substitute the values:

$$P(A \cap B^c) = 0.3 - 0.1 = 0.2.$$

Thus, Option (3) is TRUE.

Step 4: Compute $P(A^c \cap B^c)$. Using the complement rule:

$$P(A^c \cap B^c) = 1 - P(A \cup B).$$

Substitute the value of $P(A \cup B)$:

$$P(A^c \cap B^c) = 1 - 0.7 = 0.3.$$

Since the option states $P(A^c \cap B^c) = 0.4$, Option (4) is FALSE.


Final Answer:

$$(2) P(A \cup B) = 0.7, (3) P(A \cap B^c) = 0.2$$

Quick Tip

When solving probability problems, always verify independence using $P(A \cap B) = P(A) \cdot P(B)$ and use complement and union rules to simplify calculations.

28. Consider the circuit shown below where the gates may have propagation delays. Assume that all signal transitions occur instantaneously and that wires have no delays. Which of the following statements about the circuit is/are CORRECT?



- (1) With no propagation delays, the output Y is always logic Zero
- (2) With no propagation delays, the output Y is always logic One
- (3) With propagation delays, the output Y can have a transient logic One after X transitions from logic Zero to logic One
- (4) With propagation delays, the output Y can have a transient logic Zero after X transitions from logic One to logic Zero

Correct Answer: (1) With no propagation delays, the output Y is always logic Zero, (3) With propagation delays, the output Y can have a transient logic One after X transitions from logic Zero to logic One

Solution:

Step 1: Analyze the circuit without propagation delays.

The circuit includes an AND gate and a NOT gate. The input to the AND gate is X and the output of the NOT gate (i.e., $\neg X$).

The AND gate has two inputs: X and $\neg X$. For any value of X :

$$Y = X \wedge \neg X.$$

Since $X \wedge \neg X = 0$ for all values of X , the output Y is always logic Zero without propagation delays. Thus, Option (1) is TRUE, and Option (2) is FALSE.

Step 2: Analyze the circuit with propagation delays. When X transitions from logic Zero to logic One:

- 1. The NOT gate takes some time to invert the signal from X to $\neg X$.
- 2. During this propagation delay, both inputs of the AND gate could temporarily be logic One.

This causes a transient logic One at Y . Thus, Option (3) is TRUE.

When X transitions from logic One to logic Zero, the NOT gate's propagation delay does not lead to a transient logic Zero at Y , because the AND gate will have at least one input at logic Zero. Thus, Option (4) is FALSE.

Final Answer:

(1) and (3)

Quick Tip

Propagation delays in logic gates can lead to transient outputs. Always consider the gate timing when analyzing circuit behavior during transitions.

29. TCP client P successfully establishes a connection to TCP server Q . Let N_P denote the sequence number in the SYN sent from P to Q . Let N_Q denote the acknowledgement number in the SYN ACK from Q to P . Which of the following statements is/are CORRECT?

- (1) The sequence number N_P is chosen randomly by P
- (2) The sequence number N_P is always 0 for a new connection
- (3) The acknowledgement number N_Q is equal to N_P
- (4) The acknowledgement number N_Q is equal to $N_P + 1$

Correct Answer: (1) The sequence number N_P is chosen randomly by P , (4) The acknowledgement number N_Q is equal to $N_P + 1$

Solution:

Step 1: Sequence number N_P .

The sequence number in the SYN segment is chosen randomly by the client P to ensure security and uniqueness for the TCP connection. Hence, Option (1) is TRUE.

Option (2) is FALSE because N_P is not always 0; it is chosen randomly for each new connection.

Step 2: Acknowledgement number N_Q .

In the SYN-ACK packet sent by the server Q , the acknowledgement number N_Q acknowledges the receipt of N_P .

The acknowledgement number $N_Q = N_P + 1$, indicating that the next byte expected by the server is $N_P + 1$. Hence, Option (4) is TRUE, and Option (3) is FALSE.

Final Answer:

(1) and (4)

Quick Tip

TCP sequence numbers are chosen randomly to prevent sequence number prediction attacks, and acknowledgements indicate the next expected byte in the sequence.

30. Consider a 5-stage pipelined processor with Instruction Fetch (IF), Instruction Decode (ID), Execute (EX), Memory Access (MEM), and Register Writeback (WB) stages. Which of the following statements about forwarding is/are CORRECT?

- (1) In a pipelined execution, forwarding means the result from a source stage of an earlier instruction is passed on to the destination stage of a later instruction
- (2) In forwarding, data from the output of the MEM stage can be passed on to the input of the EX stage of the next instruction
- (3) Forwarding cannot prevent all pipeline stalls
- (4) Forwarding does not require any extra hardware to retrieve the data from the pipeline stages

Correct Answer: (1) In a pipelined execution, forwarding means the result from a source stage of an earlier instruction is passed on to the destination stage of a later instruction, (2) In forwarding, data from the output of the MEM stage can be passed on to the input of the EX stage of the next instruction, (3) Forwarding cannot prevent all pipeline stalls

Solution:

Step 1: Define forwarding.

Forwarding is a technique used in pipelined processors to resolve data hazards by directly passing data from one pipeline stage to another, bypassing the register file. This is TRUE for Option (1).

Step 2: Forwarding in MEM and EX stages.

Forwarding can occur between the output of the MEM stage of an earlier instruction and the input of the EX stage of a later instruction if the result is required before it is written back. This is TRUE for Option (2).

Step 3: Limitations of forwarding.

Forwarding cannot resolve hazards arising from load instructions where the data is not available until the MEM stage. This limitation leads to some pipeline stalls. Thus, Option (3) is TRUE.

Step 4: Hardware requirements for forwarding.

Forwarding requires additional hardware, such as forwarding paths and control logic, to enable bypassing data between pipeline stages. Hence, Option (4) is FALSE.

Final Answer:

(1), (2), and (3)

Quick Tip

Forwarding is an effective mechanism to handle data hazards but cannot resolve all stalls. Load-use hazards often require pipeline stalls.

31. Which of the following fields is/are modified in the IP header of a packet going out of a network address translation (NAT) device from an internal network to an external network?

- (1) Source IP
- (2) Destination IP
- (3) Header Checksum
- (4) Total Length

Correct Answer: (1) Source IP, (3) Header Checksum

Solution:

Step 1: Analyze NAT behavior.

When a packet passes through a NAT device, the source IP address is replaced with the NAT device's external IP address to allow for communication with the external network. Hence, the Source IP is modified, making Option (1) TRUE.

Step 2: Checksum recalculation.

Since the source IP in the header is modified, the checksum, which is computed based on the

header fields, must also be recalculated. Thus, the Header Checksum is modified, making Option (3) TRUE.

Step 3: Other fields.

The Destination IP remains unchanged, as it points to the intended external destination. Similarly, the Total Length of the packet is not affected by NAT translation. Thus, Options (2) and (4) are FALSE.

Final Answer:

(1) and (3)

Quick Tip

In NAT translation, the Source IP and Header Checksum fields are typically modified, while fields like Destination IP and Total Length remain unchanged.

32. Let A and B be non-empty finite sets such that there exist one-to-one and onto functions (i) from A to B and (ii) from $A \times A$ to $A \cup B$. The number of possible values of $|A|$ is _____.

Correct Answer: 2

Solution:

Step 1: Condition for bijection from A to B .

For a one-to-one and onto function to exist from A to B , the cardinalities of A and B must be equal. Let $|A| = |B| = n$.

Step 2: Condition for bijection from $A \times A$ to $A \cup B$.

The cardinality of $A \times A$ is n^2 , and the cardinality of $A \cup B$ is $|A| + |B| = 2n$. For a one-to-one and onto function to exist, these cardinalities must be equal:

$$n^2 = 2n.$$

Step 3: Solve the equation.

Simplify:

$$n(n - 2) = 0.$$

Thus, $n = 0$ or $n = 2$. Since A and B are non-empty, $n = 2$.

Final Answer:

2

Quick Tip

The existence of bijections imposes strict equality conditions on the cardinalities of sets. Solve carefully using properties of set operations.

33. Consider the operator precedence and associativity rules for the integer arithmetic operators given in the table below.

Operator	Precedence	Associativity
+	Highest	Left
-	High	Right
*	Medium	Right
/	Low	Right

The value of the expression $3 + 1 + 5 * 2/7 + 2 - 4 - 7 - 6/2$ as per the above rules is

Correct Answer: 6

Solution:

Step 1: Apply operator precedence and associativity.

Using the given precedence and associativity:

1. Highest precedence: +, Left associativity.
2. High precedence: -, Right associativity.
3. Medium precedence: *, Right associativity.
4. Low precedence: /, Right associativity.

Step 2: Evaluate the expression.

The expression is:

$$3 + 1 + 5 * 2/7 + 2 - 4 - 7 - 6/2$$

Begin with * and / (medium and low precedence, evaluated right to left):

$$5 * 2 = 10, \quad 10/7 = 1, \quad 6/2 = 3$$

Substitute these results:

$$3 + 1 + 1 + 2 - 4 - 7 - 3$$

Evaluate + (highest precedence, left to right):

$$3 + 1 = 4, \quad 4 + 1 = 5, \quad 5 + 2 = 7$$

Evaluate - (high precedence, right to left):

$$7 - 4 = 3, \quad 3 - 7 = -4, \quad -4 - 3 = -7$$

Add all results:

$$7 - 1 = 6$$

Final Answer:

6

Quick Tip

When solving arithmetic expressions with operator precedence and associativity, always evaluate operators step by step based on their precedence, starting from the highest.

34. The number of spanning trees in a complete graph of 4 vertices labelled A, B, C, and D is

Correct Answer: 16

Solution: Step 1: Recall the formula for spanning trees in a complete graph.

The number of spanning trees in a complete graph with n vertices is given by Cayley's formula:

$$T = n^{n-2}.$$

Step 2: Substitute the value of n .

For $n = 4$:

$$T = 4^{4-2} = 4^2 = 16.$$

Final Answer:

16

Quick Tip

Cayley's formula $T = n^{n-2}$ is applicable for calculating the number of spanning trees in a complete graph. For $n = 4$, the result is 16.

35. Consider the following two relations, $R(A, B)$ and $S(A, C)$:

Relation $R(A, B)$:

A	B
10	20
20	30
30	40
30	50
50	95

Relation $S(A, C)$:

A	C
10	90
30	45
40	80

The total number of tuples obtained by evaluating the following expression:

$$\sigma_{B < C}(R \bowtie_{R.A=S.A} S)$$

is

Correct Answer: 2

Solution: Step 1: Perform the natural join $R \bowtie_{R.A=S.A} S$. The natural join combines tuples from R and S where $R.A = S.A$.

$R \bowtie_{R.A=S.A} S$:

A	B	C
10	20	90
30	40	45

Step 2: Apply the selection condition $B < C$. Select tuples where $B < C$:

A	B	C
10	20	90
30	40	45

The total number of tuples satisfying the condition is 2.

Final Answer:

2

Quick Tip

For evaluating relational algebra expressions, follow the order: (1) perform joins, (2) apply selections, and (3) count the resulting tuples.

Question 36 - 65 carry two marks each

36. Consider a network path $P \rightarrow Q \rightarrow R$ between nodes P and R via router Q . Node P sends a file of size 10^6 bytes to R via this path by splitting the file into chunks of 10^3 bytes each. Node P sends these chunks one after the other without any wait time between the successive chunk transmissions. Assume that the size of extra headers added to these chunks is negligible, and that the chunk size is less than the MTU.

Each of the links $P \rightarrow Q$ and $Q \rightarrow R$ has a bandwidth of 10^6 bits/sec, and negligible propagation latency. Router Q immediately transmits every packet it receives from P to R , with negligible processing and queueing delays. Router Q can simultaneously receive on link $P \rightarrow Q$ and transmit on link $Q \rightarrow R$.

Assume P starts transmitting the chunks at time $t = 0$. Which one of the following options gives the time (in seconds, rounded off to 3 decimal places) at which R receives all the chunks of the file?

- (1) 8.000
- (2) 8.008
- (3) 15.992
- (4) 16.000

Correct Answer: (2) 8.008

Solution:

Step 1: File size and chunk breakdown.

The total file size is 10^6 bytes, and each chunk is 10^3 bytes. The number of chunks is:

$$\text{Number of chunks} = \frac{\text{Total file size}}{\text{Chunk size}} = \frac{10^6}{10^3} = 1000.$$

Each chunk is 10^3 bytes = 8×10^3 bits = 8000 bits.

Step 2: Transmission time for one chunk.

The bandwidth of each link is 10^6 bits/sec. The time to transmit one chunk is:

$$t_{\text{chunk}} = \frac{\text{Chunk size (in bits)}}{\text{Bandwidth (in bits/sec)}} = \frac{8000}{10^6} = 0.008 \text{ seconds.}$$

Step 3: Time for P to send all chunks to Q .

Since P transmits one chunk every 0.008 seconds and there are 1000 chunks, the total time for P to finish sending all chunks to Q is:

$$t_P = 1000 \times t_{\text{chunk}} = 1000 \times 0.008 = 8 \text{ seconds.}$$

The last chunk is sent by P to Q at $t = 8$ seconds.

Step 4: Time for Q to send the last chunk to R .

Router Q immediately forwards chunks to R as it receives them. The transmission time for one chunk from Q to R is also:

$$t_{\text{chunk}} = 0.008 \text{ seconds.}$$

The last chunk reaches Q at $t = 8$ seconds. Router Q transmits it to R in 0.008 seconds. Thus, the last chunk reaches R at:

$$t_{\text{total}} = t_P + t_{\text{chunk}} = 8 + 0.008 = 8.008 \text{ seconds.}$$

Step 5: Final result. The time at which R receives all the chunks is 8.008 seconds.

Quick Tip

For problems involving sequential chunk transmissions, always compute the total transmission time by considering the time to send all chunks and the time for the last chunk to propagate through the network.

37. Consider the following syntax-directed definition (SDD).

$$S \rightarrow DHTU \quad \{S.val = D.val + H.val + T.val + U.val;\}$$

$$D \rightarrow "M" D_1 \quad \{D.val = 5 + D_1.val;\}$$

$$D \rightarrow \epsilon \quad \{D.val = -5;\}$$

$$H \rightarrow "L" H_1 \quad \{H.val = 5 \cdot 10 + H_1.val;\}$$

$$H \rightarrow \epsilon \quad \{H.val = -10;\}$$

$$T \rightarrow "C" T_1 \quad \{T.val = 5 \cdot 100 + T_1.val;\}$$

$$T \rightarrow \epsilon \quad \{T.val = -5;\}$$

$$U \rightarrow "K" \quad \{U.val = 5;\}$$

Given the input "MMLK", which one of the following options is the CORRECT value computed by the SDD (in the attribute $S.val$)?

- (1) 45
- (2) 50
- (3) 55
- (4) 65

Correct Answer: (1) 45

Solution:

Step 1: Break down the input sequence "MMLK".

The input corresponds to:

$$S \rightarrow DHTU \quad \text{where } D = "MM", H = "L", T = \epsilon, U = "K".$$

Step 2: Calculate the values for D , H , T , and U .

- $D = "MM"$: Using the production rule $D \rightarrow "M" D_1$:

$$D.val = 5 + D_1.val.$$

The first "M" contributes 5, and the second "M" contributes 5, resulting in:

$$D.val = 5 + 0 = 5.$$

- $H = "L"$: Using the production rule $H \rightarrow "L" H_1$:

$$H.val = 5 \cdot 10 + H_1.val.$$

Here, $H_1 \rightarrow \epsilon$ contributes -10 , so:

$$H.val = 50 - 10 = 40.$$

- $T = \epsilon$: Using the production rule $T \rightarrow \epsilon$:

$$T.val = -5.$$

- $U = "K"$: Using the production rule $U \rightarrow "K"$:

$$U.val = 5.$$

Step 3: Compute $S.val$.

Using the rule $S.val = D.val + H.val + T.val + U.val$, substitute the calculated values:

$$S.val = 5 + 40 - 5 + 5 = 45.$$

Final Answer:

45

Quick Tip

When solving syntax-directed definition problems, evaluate each production rule systematically and substitute values step by step to avoid errors.

38. Consider the following grammar G , with S as the start symbol. The grammar G has three incomplete productions denoted by (1), (2), and (3).

$$S \rightarrow daT \mid (1)$$

$$T \rightarrow aS \mid bT \mid (2)$$

$$R \rightarrow (3) \mid \epsilon$$

The set of terminals is $\{a, b, c, d, f\}$. The FIRST and FOLLOW sets of the different non-terminals are as follows:

$$\text{FIRST}(S) = \{c, d, f\}, \text{FIRST}(T) = \{a, b, \epsilon\}, \text{FIRST}(R) = \{c, \epsilon\},$$

$$\text{FOLLOW}(S) = \text{FOLLOW}(T) = \{c, f, \$\}, \text{FOLLOW}(R) = \{f\}.$$

Which one of the following options CORRECTLY fills in the incomplete productions?

$$(1) S \rightarrow Rf, T \rightarrow \epsilon, R \rightarrow cTR$$

$$(2) S \rightarrow fR, T \rightarrow \epsilon, R \rightarrow cTR$$

$$(3) S \rightarrow fR, T \rightarrow cT, R \rightarrow cR$$

$$(4) S \rightarrow Rf, T \rightarrow cT, R \rightarrow cR$$

Correct Answer: (1) $S \rightarrow Rf, T \rightarrow \epsilon, R \rightarrow cTR$

Solution:

Step 1: Analyze the FIRST and FOLLOW sets.

The FIRST and FOLLOW sets are crucial for constructing a grammar that satisfies the given properties.

1. $S \rightarrow daT \mid Rf$: - The $\text{FIRST}(S)$ must include Rf and daT . This matches the FIRST set $\{c, d, f\}$, as R can produce c and ϵ , and f is present.

2. $T \rightarrow aS \mid bT \mid \epsilon$: - The $\text{FIRST}(T)$ is $\{a, b, \epsilon\}$, which matches the given options only if $T \rightarrow \epsilon$ is included.

3. $R \rightarrow cTR \mid \epsilon$: - The $\text{FIRST}(R)$ includes c and ϵ , which is achieved by this production.

Step 2: Validate the correctness of options.

1. $S \rightarrow Rf, T \rightarrow \epsilon, R \rightarrow cTR$: - This option satisfies all FIRST and FOLLOW set requirements.

2. Other options either include productions that violate the FIRST or FOLLOW set constraints or introduce inconsistencies in the grammar structure.

Final Answer:

$$(1) S \rightarrow Rf, T \rightarrow \epsilon, R \rightarrow cTR$$

Quick Tip

When dealing with grammar and FIRST/FOLLOW sets, ensure that all productions align with the constraints imposed by the sets to maintain consistency and avoid ambiguity.

39. Consider the following pseudo-code:

```
L1:  t1 = -1
L2:  t2 = 0
L3:  t3 = 0
L4:  t4 = 4 × t3
L5:  t5 = 4 × t2
L6:  t6 = t5 × M
L7:  t7 = t4 + t6
L8:  t8 = a[t7]
L9:  if t8 ≤ max goto L11
L10: t1 = t8
L11: t3 = t3 + 1
L12: if t3 < M goto L4
L13: t2 = t2 + 1
L14: if t2 < N goto L3
L15: max = t1
```

Which one of the following options CORRECTLY specifies the number of basic blocks and

the number of instructions in the largest basic block, respectively?

- (1) 6 and 6
- (2) 6 and 7
- (3) 7 and 7
- (4) 7 and 6

Correct Answer: (4) 7 and 6

Solution:

Step 1: Understanding basic blocks.

A basic block is a sequence of consecutive statements in the code where:

- 1. Control enters at the beginning and exits at the end without any possibility of branching except at the end.
- 2. No statements in a basic block are the target of any jump.

Step 2: Identifying the basic blocks.

- 1. Basic Block 1 (L1 - L3): This includes instructions $t1 = -1$, $t2 = 0$, $t3 = 0$. These are straightforward assignments with no branching.
- 2. Basic Block 2 (L4 - L8): Instructions $t4 = 4 \times t3$, $t5 = 4 \times t2$, $t6 = t5 \times M$, $t7 = t4 + t6$, and $t8 = a[t7]$ form a block. The condition $t8 \leq \max$ introduces a jump, ending the block.
- 3. Basic Block 3 (L9 - L10): Contains the conditional jump if $t8 \leq \max$ goto L11 and the assignment $t1 = t8$.
- 4. Basic Block 4 (L11): Contains $t3 = t3 + 1$.
- 5. Basic Block 5 (L12): Contains if $t3 < M$ goto L4.
- 6. Basic Block 6 (L13): Contains $t2 = t2 + 1$.
- 7. Basic Block 7 (L14 - L15): Includes if $t2 < N$ goto L3 and $\max = t1$.

Step 3: Count the instructions in the largest basic block.

The largest basic block is Block 2, which contains 6 instructions (L4 to L8).

Final Answer:

7 and 6

Quick Tip

When analyzing basic blocks, focus on identifying the start and endpoints of blocks based on conditional jumps and labels.

40. Consider the following two threads T1 and T2 that update two shared variables a and b . Assume that initially $a = 1$ and $b = 1$. Though context switching between threads can happen at any time, each statement of T1 or T2 is executed atomically without interruption.

T1: $a = a + 1;$ $b = b + 1;$

T2: $b = 2 \times b;$ $a = 2 \times a;$

Which one of the following options lists all the possible combinations of values of a and b after both T1 and T2 finish execution?

(1) $(a = 4, b = 4); (a = 3, b = 3); (a = 4, b = 3)$

(2) $(a = 3, b = 4); (a = 4, b = 3); (a = 3, b = 3)$

(3) $(a = 4, b = 4); (a = 4, b = 3); (a = 3, b = 4)$

(4) $(a = 2, b = 2); (a = 2, b = 3); (a = 3, b = 4)$

Correct Answer: (1) $(a = 4, b = 4); (a = 3, b = 3); (a = 4, b = 3)$

Solution:

Step 1: Analyze the threads.

Thread T1:

Increases a by 1, i.e., $a = a + 1$.

Increases b by 1, i.e., $b = b + 1$.

Thread T2:

Doubles b , i.e., $b = 2 \times b$.

Doubles a , i.e., $a = 2 \times a$.

Step 2: Possible execution orders.

The statements in T1 and T2 can interleave in various ways, leading to the following combi-

nations of a and b :

1. T1 executes completely, then T2 executes completely: $a = 2 \times (1+1) = 4$, $b = 2 \times (1+1) = 4$.

Result: $(a = 4, b = 4)$.

2. T2 executes completely, then T1 executes completely: $a = (2 \times 1) + 1 = 3$, $b = (2 \times 1) + 1 = 3$.

Result: $(a = 3, b = 3)$.

3. T1 and T2 interleave: For example, if b is doubled first and then incremented: $b = (2 \times 1) + 1 = 3$, $a = 2 \times (1 + 1) = 4$. Result: $(a = 4, b = 3)$.

Step 3: Final combinations.

The possible combinations are:

$$(a = 4, b = 4), (a = 3, b = 3), (a = 4, b = 3).$$

Final Answer:

$$(a = 4, b = 4); (a = 3, b = 3); (a = 4, b = 3)$$

Quick Tip

In multithreading problems, consider all possible interleavings of instructions to determine the range of outcomes for shared variables.

41. An array [82, 101, 90, 11, 111, 75, 33, 131, 44, 93] is heapified. Which one of the following options represents the first three elements in the heapified array?

(1) 82, 90, 101

(2) 82, 11, 93

(3) 131, 11, 93

(4) 131, 111, 90

Correct Answer: (4) 131, 111, 90

Solution:

Step 1: Understanding heapification.

Heapification transforms an array into a binary heap. In a **max heap**, the root node is the largest element, and every parent node is larger than its child nodes.

Step 2: Applying heapification to the array.

Starting with [82, 101, 90, 11, 111, 75, 33, 131, 44, 93]:

1. The largest element 131 becomes the root.
2. The next largest elements, 111 and 90, are placed as the left and right children of the root, respectively.

Step 3: First three elements.

After heapification, the first three elements in the max heap are:

131, 111, 90

Quick Tip

In a max heap, the root is always the largest element, and the heap property ensures that each subtree also satisfies this rule.

42. Consider the following recurrence relation:

$$T(n) = \begin{cases} \sqrt{n}T(\sqrt{n}) + n & \text{for } n \geq 1, \\ 1 & \text{for } n = 1. \end{cases}$$

Which one of the following options is CORRECT?

- (1) $T(n) = \Theta(n \log \log n)$
- (2) $T(n) = \Theta(n \log n)$
- (3) $T(n) = \Theta(n^2 \log n)$
- (4) $T(n) = \Theta(n^2 \log \log n)$

Correct Answer: (1) $T(n) = \Theta(n \log \log n)$

Solution:

Step 1: Recurrence Analysis.

The given recurrence is:

$$T(n) = \sqrt{n}T(\sqrt{n}) + n.$$

Let $n = 2^m$. Then, $\sqrt{n} = 2^{m/2}$, and the recurrence becomes:

$$T(2^m) = 2^{m/2}T(2^{m/2}) + 2^m.$$

Step 2: Substitution.

Substitute $T(2^m) = f(m)$. The recurrence now becomes:

$$f(m) = 2^{m/2} f(m/2) + 2^m.$$

Step 3: Expanding the Recurrence.

Expand the recurrence:

$$f(m) = 2^{m/2} \left(2^{m/4} f(m/4) + 2^{m/2} \right) + 2^m.$$

$$f(m) = 2^{m/2} \cdot 2^{m/4} f(m/4) + 2^m + 2^m.$$

This continues for k levels, with each level contributing 2^m . The number of levels is approximately $\log \log n$.

Step 4: Complexity.

The total complexity is:

$$T(n) = \Theta(n \log \log n).$$

Final Answer:

$$\boxed{\Theta(n \log \log n)}$$

Quick Tip

For recurrence relations involving nested square roots, use transformations like $n = 2^m$ to simplify and determine the number of levels.

43. Consider a binary min-heap containing 105 distinct elements. Let k be the index (in the underlying array) of the maximum element stored in the heap. The number of possible values of k is:

- (1) 53
- (2) 52
- (3) 27
- (4) 1

Correct Answer: (1) 53

Solution:**Step 1:** Min-Heap Property.

In a binary min-heap, the root (index 1) contains the minimum element, and the elements increase in value as we move down the tree. The maximum element will always be in the last level of the heap.

Step 2: Calculating the Last Level.

The number of levels in a binary heap with n elements is $\lceil \log_2 n \rceil$. For $n = 105$:

$$\lceil \log_2 105 \rceil = 7.$$

The 7th level is the last level, and it contains the remaining elements.

Step 3: Indices of Last Level Elements.

In the array representation of a binary heap, the indices of the elements in the last level start from 64 and end at 105. The maximum element will always be one of these leaf nodes.

Step 4: Leaf Nodes in the Array.

The leaf nodes in a binary heap correspond to the indices from $\lceil n/2 \rceil$ to n . For $n = 105$:

$$\lceil 105/2 \rceil = 53.$$

Thus, the possible indices of the maximum element range from 53 to 105, which gives:

$$105 - 53 + 1 = 53 \text{ possible values for } k.$$

Final Answer:

53

Quick Tip

In a binary heap, the maximum element always resides in the last level. Use the total elements and levels to find its possible indices.

44. The symbol \rightarrow indicates functional dependency in the context of a relational database. Which of the following options is/are TRUE?

- (1) $(X, Y) \rightarrow (Z, W)$ implies $X \rightarrow (Z, W)$
- (2) $(X, Y) \rightarrow (Z, W)$ implies $(X, Y) \rightarrow Z$
- (3) $((X, Y) \rightarrow Z$ and $W \rightarrow Y)$ implies $(X, W) \rightarrow Z$
- (4) $(X \rightarrow Y$ and $Y \rightarrow Z)$ implies $X \rightarrow Z$

Correct Answer: (2), (3), (4)

Solution:

Step 1: Analyze each functional dependency statement.

- **Option (1):** $(X, Y) \rightarrow (Z, W)$ does not necessarily imply $X \rightarrow (Z, W)$. For example, if X does not uniquely determine Z or W without Y , this fails. Hence, this is FALSE.
- **Option (2):** $(X, Y) \rightarrow (Z, W)$ implies $(X, Y) \rightarrow Z$, as Z is part of the attributes determined by (X, Y) . This is TRUE.
- **Option (3):** From the given dependencies $(X, Y) \rightarrow Z$ and $W \rightarrow Y$, we can infer $(X, W) \rightarrow Z$ by substitution. This is TRUE.
- **Option (4):** The transitivity property of functional dependency ensures $X \rightarrow Y$ and $Y \rightarrow Z$ implies $X \rightarrow Z$. This is TRUE.

Final Answer:

(2), (3), (4)

Quick Tip

Functional dependency follows reflexivity, augmentation, and transitivity. Carefully apply these properties to analyze dependencies.

45. Let G be a directed graph and T a depth first search (DFS) spanning tree in G that is rooted at a vertex v . Suppose T is also a breadth first search (BFS) tree in G , rooted at v . Which of the following statements is/are TRUE for every such graph G and tree T ?

- (1) There are no back-edges in G with respect to the tree T
- (2) There are no cross-edges in G with respect to the tree T
- (3) There are no forward-edges in G with respect to the tree T
- (4) The only edges in G are the edges in T

Correct Answer: (3)

Solution:

Step 1: Analyze properties of DFS and BFS trees.

Option (1): Back-edges are present in a DFS tree if cycles exist in G . However, for T to be both a DFS and BFS tree, back-edges cannot exist. Hence, this is FALSE as it depends on the structure of G .

Option (2): Cross-edges appear in DFS when visiting previously visited vertices in another branch. In a BFS, no cross-edges exist by definition. However, the dual nature of T does not guarantee this universally. Hence, this is FALSE.

Option (3): Forward-edges are edges from a vertex to its descendant in DFS but not part of the DFS tree. Since T is both a DFS and BFS tree, forward-edges cannot exist. This is TRUE.

Option (4): The existence of T as both DFS and BFS does not limit G to contain only tree edges. Non-tree edges (e.g., cross or back-edges) may still exist. Hence, this is FALSE.

Final Answer:

(3)

Quick Tip

Understanding edge classifications in DFS (tree, back, forward, cross) is crucial to analyzing tree structures in directed graphs.

46. Consider the following read-write schedule S over three transactions $T_1, T_2,$ and $T_3,$ where the subscripts in the schedule indicate transaction IDs:

$$S : r_1(z); w_1(z); r_2(x); r_3(y); w_3(y); r_2(y); w_2(x); w_2(y);$$

Which of the following transaction schedules is/are conflict equivalent to S ?

- (1) $T_1T_2T_3$
- (2) $T_1T_3T_2$
- (3) $T_3T_2T_1$
- (4) $T_3T_1T_2$

Correct Answer: (2), (3), (4)

Solution:

Step 1: Conflict equivalence definition.

Two schedules are conflict equivalent if they preserve the same order of conflicting operations (read/write or write/write) across all transactions.

Step 2: Analyze the given schedule S .

Transaction T_1 performs $r_1(z); w_1(z)$. Transaction T_2 performs $r_2(x); w_2(x); r_2(y); w_2(y)$. Transaction T_3 performs $r_3(y); w_3(y)$. Conflicts occur between T_1 and T_2 on z and x , and between T_3 and T_2 on y .

Step 3: Test each schedule.

For $T_1T_2T_3$, conflicts are not preserved, as T_3 's operations on y may overlap incorrectly with T_2 . For $T_1T_3T_2$, conflicts are preserved, maintaining the correct order of operations between T_1, T_3 , and T_2 . For $T_3T_2T_1$, conflicts are preserved, maintaining correct order of operations between T_3, T_2 , and T_1 . For $T_3T_1T_2$, conflicts are preserved, maintaining correct order of operations between T_3, T_1 , and T_2 .

Final Answer:

$$\boxed{T_1T_3T_2, T_3T_2T_1, T_3T_1T_2}$$

Quick Tip

Conflict equivalence ensures consistency in read/write operations across transactions. Always analyze dependencies and conflicting operations carefully.

47. Consider a Boolean expression given by $F(X, Y, Z) = \Sigma(3, 5, 6, 7)$. Which of the following statements is/are CORRECT?

- (1) $F(X, Y, Z) = \Pi(0, 1, 2, 4)$

(2) $F(X, Y, Z) = XY + YZ + XZ$

(3) $F(X, Y, Z)$ is independent of input Y

(4) $F(X, Y, Z)$ is independent of input X

Correct Answer: (1), (2)

Solution:

Step 1: Derive the truth table for $F(X, Y, Z)$.

The given $F(X, Y, Z) = \Sigma(3, 5, 6, 7)$ corresponds to minterms where the function evaluates to 1. These minterms are:

$$3(011), 5(101), 6(110), 7(111).$$

Step 2: Convert to product-of-maxterms.

Maxterms where $F(X, Y, Z) = 0$ are 0(000), 1(001), 2(010), 4(100). Thus, the complement of $F(X, Y, Z)$ is given by $\Pi(0, 1, 2, 4)$. This makes Option (1) TRUE.

Step 3: Simplify the Boolean expression.

Using minterms, the Boolean expression can be simplified as:

$$F(X, Y, Z) = XY + YZ + XZ.$$

This makes Option (2) TRUE.

Step 4: Check independence of X or Y .

The function depends on all three variables (X, Y, Z), as flipping any variable changes the output. Thus, $F(X, Y, Z)$ is NOT independent of X or Y , making Options (3) and (4) FALSE.

Final Answer:

(1), (2)

Quick Tip

For Boolean expressions, convert between minterm and maxterm forms to verify equivalences. Use truth tables or Karnaugh maps for simplification.

48. Consider the following C function definition:

```
int f(int x, int y) {
    for (int i = 0; i < y; i++) {
        x = x + x + y;
    }
    return x;
}
```

Which of the following statements is/are TRUE about the above function?

- (1) If the inputs are $x = 20, y = 10$, then the return value is greater than 2^{20}
- (2) If the inputs are $x = 20, y = 20$, then the return value is greater than 2^{20}
- (3) If the inputs are $x = 20, y = 10$, then the return value is less than 2^{10}
- (4) If the inputs are $x = 10, y = 20$, then the return value is greater than 2^{20}

Correct Answer: (2), (4)

Solution:

Step 1: Understand the function.

The function iterates y times. In each iteration, the value of x is updated as $x = x + x + y$, which simplifies to $x = 2x + y$. Starting with the initial value of x , the function exponentially increases x over y iterations.

Step 2: Analyze the cases.

For $x = 20, y = 10$: After 10 iterations, the value of x becomes significantly large, and it exceeds $2^{10} = 1024$. Thus, the return value is greater than 2^{10} .

For $x = 20, y = 20$: After 20 iterations, x grows even more exponentially and far exceeds 2^{20} . Thus, the return value is greater than 2^{20} .

For $x = 10, y = 20$:

Similarly, with $x = 10$ and $y = 20$, the function results in x exceeding 2^{20} .

$x = 20, y = 10$: The return value is never less than 2^{10} , so this case is invalid.

Final Answer:

(2), (4)

Quick Tip

Exponential growth in iterative processes leads to extremely large values quickly. Always analyze the iteration logic carefully for such functions.

49. Let A be any $n \times m$ matrix, where $m > n$. Which of the following statements is/are TRUE about the system of linear equations $Ax = 0$?

- (1) There exist at least $m - n$ linearly independent solutions to this system
- (2) There exist $m - n$ linearly independent vectors such that every solution is a linear combination of these vectors
- (3) There exists a non-zero solution in which at least $m - n$ variables are 0
- (4) There exists a solution in which at least n variables are non-zero

Correct Answer: (1)

Solution:

Step 1: Analyze the null space of A .

For an $n \times m$ matrix A , the rank-nullity theorem states:

$$\text{Nullity}(A) = m - \text{Rank}(A).$$

Given $m > n$, the rank of A is at most n , so the nullity is at least $m - n$.

Step 2: Validate the options.

Option (1): There are at least $m - n$ linearly independent solutions in the null space of A , which is correct.

Option (2): While $m - n$ vectors span the null space, this does not imply they are solutions to $Ax = 0$. Incorrect.

Option (3): There is no guarantee that $m - n$ variables in the solution are 0. Incorrect.

Option (4): There is no guarantee that at least n variables are non-zero in the solution. Incorrect.


Final Answer:

(1)

Quick Tip

For underdetermined systems ($m > n$), the null space dimension is key to determining the number of linearly independent solutions.

50. Consider the 5-state DFA M accepting the language $L(M) \subseteq (0 + 1)^*$ shown below. For any string $w \in (0 + 1)^*$, let $n_0(w)$ be the number of 0's in w and $n_1(w)$ be the number of 1's in w .



Which of the following statements is/are FALSE?

- (1) States 2 and 4 are distinguishable in M
- (2) States 3 and 4 are distinguishable in M
- (3) States 2 and 5 are distinguishable in M
- (4) Any string w with $n_0(w) = n_1(w)$ is in $L(M)$

Correct Answer: (2), (3)

Solution:

Step 1: Analyze the DFA and its language.

The DFA M accepts strings based on a condition related to the counts of 0's and 1's. Transitions lead to different states depending on the parity (even or odd) of $n_0(w)$ and $n_1(w)$. Distinguishability of states is determined by whether they lead to different acceptance behavior for any input string.

Step 2: Verify each statement.

Option (1): States 2 and 4 are distinguishable because they behave differently for certain

strings. For example, adding more 1's or 0's leads to different final states. This is true.

Option (2): States 3 and 4 are indistinguishable because they behave identically for all input strings. This is false.

Option (3): States 2 and 5 are indistinguishable because they represent equivalent conditions for $n_0(w)$ and $n_1(w)$. This is false.

Option (4): Any string w with $n_0(w) = n_1(w)$ is accepted by M because the DFA ensures balance between the counts of 0's and 1's. This is true.

Final Answer:

(2), (3)

Quick Tip

To check distinguishability in a DFA, consider input strings that lead to different acceptance behaviors. Indistinguishable states behave identically for all input strings.

51. The chromatic number of a graph is the minimum number of colours used in a proper colouring of the graph. Let G be any graph with n vertices and chromatic number k . Which of the following statements is/are always TRUE?

- (1) G contains a complete subgraph with k vertices
- (2) G contains an independent set of size at least n/k
- (3) G contains at least $k(k-1)/2$ edges
- (4) G contains a vertex of degree at least k

Correct Answer: (2), (3)

Solution:

Step 1: Independent set size.

An independent set is a set of vertices such that no two vertices are adjacent. For a graph with chromatic number k , the graph can be partitioned into k independent sets. The largest independent set must have at least n/k vertices. Hence, (2) is true.

Step 2: Number of edges.

The chromatic number k implies there is at least a complete subgraph K_k . A complete sub-

graph with k vertices has $k(k - 1)/2$ edges. Therefore, G must contain at least $k(k - 1)/2$ edges. Hence, (3) is true.

Step 3: Remaining options.

Option (1) is not necessarily true because the graph may not contain a complete subgraph with k vertices.

Option (4) is also not necessarily true because there may not be a vertex of degree k .

Final Answer:

(2), (3)

Quick Tip

For chromatic number problems, analyze independent sets and edge requirements based on graph properties.

52. Consider the operators \diamond and \square defined by $a \diamond b = a + 2b$ and $a \square b = ab$, for positive integers. Which of the following statements is/are TRUE?

- (1) Operator \diamond obeys the associative law
- (2) Operator \square obeys the associative law
- (3) Operator \diamond over the operator \square obeys the distributive law
- (4) Operator \square over the operator \diamond obeys the distributive law

Correct Answer: (2), (4)

Solution:

Step 1: Associativity of \diamond .

To check if \diamond is associative:

$$(a \diamond b) \diamond c = (a + 2b) + 2c = a + 2b + 2c,$$

$$a \diamond (b \diamond c) = a + 2(b + 2c) = a + 2b + 4c.$$

Since the results are not equal, \diamond is not associative.

Step 2: Associativity of \square .

To check if \square is associative:

$$(a \square b) \square c = (ab) \square c = (ab)c = abc,$$

$$a \square (b \square c) = a \square (bc) = a(bc) = abc.$$

Since the results are equal, \square is associative. Hence, (2) is true.

Step 3: Distributivity of \square over \diamond .

To check if \square is distributive over \diamond :

$$a \square (b \diamond c) = a \cdot (b + 2c) = ab + 2ac,$$

$$(a \square b) \diamond (a \square c) = (ab) \diamond (ac) = ab + 2ac.$$

Since the results are equal, \square is distributive over \diamond . Hence, (3) is false.

Step 4: Distributivity of \diamond over \square .

To check if \diamond is distributive over \square :

$$a \diamond (b \square c) = a + 2(bc),$$

$$(a \diamond b) \square (a \diamond c) = (a + 2b)(a + 2c) \neq a + 2(bc).$$

Since the results are not equal, \diamond is not distributive over \square . Hence, (4) is true.

Final Answer:

(2), (4)

Quick Tip

Always check associativity and distributivity systematically by verifying the properties for all combinations of the operations.

53. Consider two set-associative cache memory architectures: *WBC*, which uses the write-back policy, and *WTC*, which uses the write-through policy. Both of them use the LRU (Least Recently Used) block replacement policy. The cache memory is connected to the main memory. Which of the following statements is/are TRUE?

- (1) A read miss in *WBC* never evicts a dirty block
- (2) A read miss in *WTC* never triggers a write-back operation of a cache block to main memory
- (3) A write hit in *WBC* can modify the value of the dirty bit of a cache block
- (4) A write miss in *WTC* always writes the victim cache block to main memory before loading the missed block to the cache

Correct Answer: (2), (3)

Solution:

Step 1: Analyze the write-back policy (WBC).

In a write-back cache, the dirty bit is set whenever a write modifies the cached data. The modified data is not written back to the main memory until the block is evicted. Therefore:

A write hit in *WBC* modifies the value of the dirty bit, making statement (3) true.

A read miss in *WBC* may evict a dirty block if required by the LRU policy, making statement (1) false.

Step 2: Analyze the write-through policy (WTC).

In a write-through cache, every write operation immediately updates both the cache and the main memory. Therefore:

A read miss in *WTC* does not trigger a write-back operation, as all writes have already updated the main memory. Hence, statement (2) is true.

A write miss in *WTC* does not necessarily write the victim cache block to the main memory, as the cache contents are always consistent with the main memory. Hence, statement (4) is false.

Final Answer:

(2), (3)

Quick Tip

Understand the behavior of write-back and write-through policies. Write-back caches delay updates to main memory, while write-through caches ensure consistency with immediate updates.

54. Consider a 512 GB hard disk with 32 storage surfaces. There are 4096 sectors per track and each sector holds 1024 bytes of data. The number of cylinders in the hard disk is _____ .

Correct Answer: 4096

Solution: Given that:

$$\text{Hard disk capacity} = 512 \text{ GB} = 2^9 \times 2^{30} = 2^{39} \text{ Bytes,}$$

$$\text{Number of surfaces} = 32 = 2^5,$$

$$\text{Sectors per track} = 4096 = 2^{12},$$

$$\text{Sector size (capacity)} = 1024 = 2^{10} \text{ Bytes,}$$

$$\text{Number of cylinders} = x.$$

Substituting the values:

$$2^{39} = 2^5 \times x \times 2^{12} \times 2^{10}.$$

Simplifying:

$$2^{39} = x \times 2^{12+10+5}.$$

$$2^{39} = x \times 2^{27}.$$

$$x = \frac{2^{39}}{2^{27}} = 2^{12} = 4096.$$

Final Answer:

4096 cylinders

Quick Tip

To solve such problems, ensure the total capacity is divided by the per-cylinder capacity, considering all surfaces and tracks. Break down calculations step by step for accuracy.

55. The baseline execution time of a program on a 2 GHz single core machine is 100 ns. The code corresponding to 90% of the execution time can be fully parallelized. The overhead for using an additional core is 10 ns when running on a multicore system. Assume that all cores in the multicore system run their share of the parallelized code for an equal amount of time. The number of cores that minimize the execution time of the program is

Correct Answer: 3

Solution:

Step 1: Divide the execution time into parallel and serial components. The program's execution time is 100 ns. The parallelizable part is 90%, and the non-parallelizable part is 10%:

$$\text{Parallelizable time} = 90 \text{ ns}, \quad \text{Non-parallelizable time} = 10 \text{ ns}.$$

Step 2: Calculate the execution time for k cores. Using Amdahl's Law and considering overhead, the execution time is given by:

$$T_k = \frac{\text{Parallelizable time}}{k} + \text{Non-parallelizable time} + (k - 1) \times \text{Overhead}.$$

Substitute values:

$$T_k = \frac{90}{k} + 10 + 10 \times (k - 1).$$

Step 3: Minimize T_k . For $k = 3$:

$$T_3 = \frac{90}{3} + 10 + 10 \times (3 - 1) = 30 + 10 + 20 = 60 \text{ ns}.$$

Checking $k = 2$ and $k = 4$, T_3 is minimum.

Final Answer:

3

Quick Tip

Amdahl's Law helps calculate speedup and execution time in parallelized programs. Always include overhead costs when minimizing execution time.

56. A given program has 25% load/store instructions. Suppose the ideal CPI (cycles per instruction) without any memory stalls is 2. The program exhibits 2% miss rate on instruction cache and 8% miss rate on data cache. The miss penalty is 100 cycles. The speedup (rounded off to two decimal places) achieved with a perfect cache (i.e., with NO data or instruction cache misses) is

Correct Answer: 3.00

Solution:

Step 1: Calculate the effective CPI with cache misses. The total CPI is given by:

$$\text{CPI} = \text{Ideal CPI} + \text{Miss contribution.}$$

For instruction cache:

$$\text{Instruction miss penalty} = 0.02 \times 100 = 2 \text{ cycles.}$$

For data cache (25% load/store):

$$\text{Data miss penalty} = 0.25 \times 0.08 \times 100 = 2 \text{ cycles.}$$

Total CPI:

$$\text{CPI} = 2 + 2 + 2 = 6.$$

Step 2: Compute speedup with a perfect cache. With no cache misses, the CPI reduces to 2.

The speedup is:

$$\text{Speedup} = \frac{\text{CPI with cache misses}}{\text{CPI without cache misses}} = \frac{6}{2} = 3.00.$$

Final Answer:

3.00

Quick Tip

Understand Amdahl's law for parallelization and analyze cache performance using CPI contributions from misses.

57. Consider the following code snippet using the `fork()` and `wait()` system calls. Assume that the code compiles and runs correctly, and that the system calls run successfully without any errors.

```
int x = 3;
while (x > 0) {
    fork();
    printf("hello");
    wait(NULL);
    x--;
}
```

The total number of times the `printf` statement is executed is

Correct Answer: 14

Solution:

Step 1: Understand the behavior of the `fork()` system call. The `fork()` system call creates a new process (child) from the parent process. The total number of processes created depends on how many times `fork()` is executed.

Step 2: Analyze the code. Initially, $x = 3$. In the first iteration of the `while` loop, one new process is created, leading to 2 processes. In the second iteration, each of these 2 processes creates a new process, resulting in $2 + 2 = 4$ processes. In the third iteration, each of the 4 processes creates another process, resulting in $4 + 4 = 8$ processes.

Step 3: Count the total executions of `printf("hello")`.

Each process executes `printf("hello")` exactly once for each iteration of the loop. Summing across all iterations:

$$2 + 4 + 8 = 14.$$

Final Answer:

14

Quick Tip

Each `fork()` doubles the number of processes, and each process executes subsequent code. Always account for all processes when analyzing output.

58. Consider the entries shown below in the forwarding table of an IP router. Each entry consists of an IP prefix and the corresponding next hop router for packets whose destination IP address matches the prefix. The notation “/N” in a prefix indicates a subnet mask with the most significant *N* bits set to 1.

Prefix	Next hop router
10.1.1.0/24	R1
10.1.1.128/25	R2
10.1.1.64/26	R3
10.1.1.192/26	R4

This router forwards 20 packets each to 5 hosts. The IP addresses of the hosts are 10.1.1.16, 10.1.1.72, 10.1.1.132, 10.1.1.191, and 10.1.1.205. The number of packets forwarded via the next hop router *R2* is

Correct Answer: 40

Solution:

Step 1: Match each IP address to a prefix.

- 10.1.1.16: Matches 10.1.1.0/24 (R1).
- 10.1.1.72: Matches 10.1.1.64/26 (R3).
- 10.1.1.132: Matches 10.1.1.128/25 (R2).
- 10.1.1.191: Matches 10.1.1.128/25 (R2).
- 10.1.1.205: Matches 10.1.1.192/26 (R4).

Step 2: Count packets forwarded via *R2*. 20 packets are forwarded to each host. The hosts 10.1.1.132 and 10.1.1.191 are forwarded via *R2*:

$$20 \text{ (to 10.1.1.132)} + 20 \text{ (to 10.1.1.191)} = 40.$$

Final Answer:

40

Quick Tip

When matching IP addresses to routing prefixes, always choose the most specific prefix (longest match).

59. Let $G = (V, \Sigma, S, P)$ be a context-free grammar in Chomsky Normal Form with $\Sigma = \{a, b, c\}$ and V containing 10 variable symbols including the start symbol S . The string $w = a^{30}b^{30}c^{30}$ is derivable from S . The number of steps (application of rules) in the derivation $S \rightarrow^* w$ is

Correct Answer: 179

Solution:

Step 1: Understand Chomsky Normal Form (CNF).

In CNF, every production is of the form $A \rightarrow BC$ or $A \rightarrow a$. To derive a string of length n , we need $n - 1$ steps for concatenation and n steps to derive each terminal symbol.

Step 2: Derive $w = a^{30}b^{30}c^{30}$.

To generate 90 terminals ($30a + 30b + 30c$), 90 steps are needed.

To combine these terminals, $90 - 1 = 89$ steps are needed.

Step 3: Total steps.

$$\text{Total steps} = 90 + 89 = 179.$$

Final Answer:

179

Quick Tip

In CNF, the number of steps to derive a string is equal to the number of terminals plus the number of concatenations required.

60. The number of edges present in the forest generated by the DFS traversal of an undirected graph G with 100 vertices is 40. The number of connected components in G is

Correct Answer: 60

Solution:

Step 1: Understand the relationship between connected components and edges in a DFS forest. In a connected component, the number of edges is one less than the number of vertices. For C connected components, the number of vertices in each component can be written as:

$$n_1, n_2, \dots, n_C,$$

where $n_1 + n_2 + \dots + n_C = 100$. The total number of edges in the graph is:

$$\text{Total edges} = (n_1 - 1) + (n_2 - 1) + \dots + (n_C - 1) = 100 - C.$$

Step 2: Compute the number of connected components. Given that the number of edges is 40:

$$100 - C = 40 \implies C = 60.$$

Final Answer:

60

Quick Tip

For an undirected graph, the number of connected components can be calculated as the difference between the number of vertices and edges in the DFS forest.

61. Consider the following two regular expressions over the alphabet $\{0, 1\}$:

$$r = 0^* + 1^* \quad \text{and} \quad s = 01^* + 10^*.$$

The total number of strings of length less than or equal to 5, which are neither in r nor in s , is

Correct Answer: 44

Solution:

Given:

$$r = 0^* + 1^* \quad (\text{either all zeros or all ones}),$$

$$s = 01^* + 10^* \quad (\text{start with zero, then all ones, or start with one, then all zeros}).$$

Alphabet:

$$\Sigma = \{0, 1\}.$$

For all strings of length n , we have:

- 2 strings from r (all zeros or all ones),
- 2 strings from s (strings starting with 0 followed by ones, or starting with 1 followed by zeros).

For Σ^2 :

$$\# \text{ strings not in } r \text{ and } s = 2^2 - 4 = 4 - 4 = 0.$$

For Σ^3 :

$$\# \text{ strings not in } r \text{ and } s = 2^3 - 4 = 8 - 4 = 4.$$

For Σ^4 :

$$\# \text{ strings not in } r \text{ and } s = 2^4 - 4 = 16 - 4 = 12.$$

For Σ^5 :

$$\# \text{ strings not in } r \text{ and } s = 2^5 - 4 = 32 - 4 = 28.$$

Total number of strings not in r and not in s for string lengths less than or equal to 5:

$$= (2^2 - 4) + (2^3 - 4) + (2^4 - 4) + (2^5 - 4).$$

$$= (4 - 4) + (8 - 4) + (16 - 4) + (32 - 4).$$

$$= 0 + 4 + 12 + 28 = 44.$$

Final Answer:

44

Quick Tip

For regular expressions, calculate the total strings and subtract the covered strings to find the complement.

62. Consider a memory management system that uses a page size of 2 KB. Assume that both the physical and virtual addresses start from 0. Assume that the pages 0, 1, 2, and 3 are stored in the page frames 1, 3, 2, and 0, respectively. The physical address (in decimal format) corresponding to the virtual address 2500 (in decimal format) is

Correct Answer: 6596

Step 1: Calculate the virtual page number and offset. The page size is 2 KB = 2048 bytes. The virtual address 2500 can be divided into:

$$\text{Page number} = \left\lfloor \frac{2500}{2048} \right\rfloor = 1, \quad \text{Offset} = 2500 \bmod 2048 = 452.$$

Step 2: Map the virtual page to the physical frame. From the given mapping:

$$\text{Page 1} \rightarrow \text{Frame 3.}$$

Step 3: Compute the physical address. The physical address is:

$$\text{Physical address} = \text{Frame number} \times \text{Page size} + \text{Offset} = 3 \times 2048 + 452 = 6596.$$

Final Answer:

6596

Quick Tip

To find the physical address, calculate the page number, map it to the frame, and add the offset.

63. A bag contains 10 red balls and 15 blue balls. Two balls are drawn randomly without replacement. Given that the first ball drawn is red, the probability (rounded off to 3 decimal places) that both balls drawn are red is _____.

Correct Answer: 0.150

Solution: The problem involves calculating the probability $P(R_2 | R_1)$, where R_1 and R_2 represent the drawing of red balls in the first and second draws, respectively.

Step 1: Analyze the outcomes

Total number of balls: 25 (10 Red + 15 Blue).

Probability of drawing R_1 (a red ball in the first draw):

$$P(R_1) = \frac{10}{25}.$$

Step 2: Conditional probabilities

If R_1 is drawn, there are now 9 red balls and 15 blue balls left in the bag (total: 24).

The probability of drawing R_2 (a red ball in the second draw given R_1):

$$P(R_2 | R_1) = \frac{9}{24}.$$

Step 3: Final calculation

$$P(R_2 | R_1) = 0.375.$$


Final Answer

0.375

Quick Tip

For conditional probabilities, adjust the total outcomes and remaining events based on the condition provided.

64. Consider a digital logic circuit consisting of three 2-to-1 multiplexers M_1 , M_2 , and M_3 as shown below. X_1 and X_2 are inputs of M_1 . X_3 and X_4 are inputs of M_2 . A , B , and C are select lines of M_1 , M_2 , and M_3 , respectively.



For an instance of inputs $X_1 = 1$, $X_2 = 1$, $X_3 = 0$, and $X_4 = 0$, the number of combinations of A , B , C that give the output $Y = 1$ is

Correct Answer: 4

Solution:

M_1 is controlled by select line A :

If $A = 0$, M_1 selects $X_1 = 1$.

If $A = 1$, M_1 selects $X_2 = 1$.

In both cases, the output of M_1 is 1.

M_2 is controlled by select line B :

If $B = 0$, M_2 selects $X_3 = 0$.

If $B = 1$, M_2 selects $X_4 = 0$.

In both cases, the output of M_2 is 0.

M_3 is controlled by select line C :

If $C = 0$, M_3 selects the output of $M_1 = 1$.

If $C = 1$, M_3 selects the output of $M_2 = 0$.

For $Y = 1$, $C = 0$ is required.

The number of valid combinations of A , B , C is determined as follows:

A can take either 0 or 1 (2 possibilities).

B can take either 0 or 1 (2 possibilities).

C must be 0 (1 possibility).

Total combinations = $2 \times 2 \times 1 = 4$.

Final Answer:

4

Quick Tip

When dealing with IP fragmentation: Subtract the IP header size from the MTU to get the payload size.

65. Consider sending an IP datagram of size 1420 bytes (including 20 bytes of IP header) from a sender to a receiver over a path of two links with a router between them. The first link (sender to router) has an MTU of 542 bytes, while the second link (router to receiver) has an MTU of 360 bytes. The number of fragments that would be delivered at the receiver is

Correct Answer: 4

Solution:

Total size of the datagram: 1420 bytes.

IP header size: 20 bytes.

Payload size: $1420 - 20 = 1400$ bytes.

Step 1: Fragmentation on the first link (MTU = 542 bytes)

Maximum payload size per fragment = $542 - 20 = 522$ bytes.

Number of fragments = $\lceil \frac{1400}{522} \rceil = 3$.

Sizes of fragments after the first link:

Fragment 1: $522 + 20 = 542$ bytes.

Fragment 2: $522 + 20 = 542$ bytes.

Fragment 3: $356 + 20 = 376$ bytes.

Step 2: Fragmentation on the second link (MTU = 360 bytes)

Maximum payload size per fragment = $360 - 20 = 340$ bytes.

Each of the three fragments is further fragmented:

Fragment 1: $\lceil \frac{522}{340} \rceil = 2$ fragments.

Fragment 2: $\lceil \frac{522}{340} \rceil = 2$ fragments.

Fragment 3: $\lceil \frac{356}{340} \rceil = 2$ fragments.

Total fragments delivered to the receiver = $2 + 2 + 2 = 6$.

Final Answer:

6

Quick Tip

When dealing with IP fragmentation: Ensure that the payload size is a multiple of 8 for compatibility with fragmentation offsets.
