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 $\rightarrow$ denotes increasing order of intensity, then the meaning of the words [dry $\rightarrow$ arid $\rightarrow$ parched] is analogous to [diet $\rightarrow$ fast $\rightarrow$ _____].

(1) starve

(2) reject

(3) feast

(4) deny

Correct Answer: (1) starve

Solution: The first analogy shows a progression based on increasing intensity:

$$\text{dry} \rightarrow \text{arid} \rightarrow \text{parched}.$$

Here, the progression moves from a mild dryness (dry) to extreme dryness (parched).
Likewise, in the second analogy:

$$\text{diet} \rightarrow \text{fast} \rightarrow \text{\_\_\_}.$$

Step 1: Analyze the terms. The word ”diet” refers to limiting food intake, while ”fast” takes it a step further by completely abstaining from food for a period of time. The next level in this sequence would be ”starve,” which denotes a prolonged and severe lack of food.
Step 2: Eliminate incorrect options. - Options (2) ”reject” and (4) ”deny” do not correspond to the idea of increasing food restriction. - Option (3) ”feast” is the opposite of restricting food intake.

Final Answer:

$$\boxed{\text{starve}}$$

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 asserts that $(x+y)$ is directly proportional to $(x-y)$, which can be expressed as:

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

where k is the proportionality constant.

Step 1: Simplify the equation. Rewriting the equation, we get:

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

By rearranging the terms, we have:

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

Now, factoring both sides gives:

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

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

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

Step 3: Interpreting the result. Since k is a constant, the ratio $\frac{x}{y}$ is also a constant. It is independent of the individual values of x and y.

Final Answer:

$$\boxed{\text{is a constant}}$$

Quick Tip
Proportional relationships often lead to constants. Always isolate terms to uncover 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: Begin by arranging the given data in ascending order:

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

Since there are 10 data points (an even number), the median is found by calculating the average of the 5th and 6th values:

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

Final Answer:

$$\boxed{13.5}$$

Quick Tip
When calculating the median of an even set of values, find the average of the two middle numbers.

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 $5x, 3x, 13x$.

The total amount is calculated as:

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

The value contributed by Rs. 5 coins is:

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

To find the percentage of the total amount contributed by Rs. 5 coins:

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

Final Answer:

$$\boxed{10\%}$$

Quick Tip
To determine the percentage of a specific denomination, calculate its total contribution 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 $\frac{p^4+q^4}{p^2q^2}$ is:

(1) 79

(2) 81

(3) 9

(4) 83

Correct Answer: (1) 79

Solution: We are given the equation:

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

Using the properties of logarithms, we rewrite the equation as:

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

This simplifies to:

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

Next, square both sides:

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

Expanding both sides:

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

Simplifying further:

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

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

Finally, divide both sides by p^2q^2:

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

Final Answer:

$$\boxed{79}$$

Quick Tip
Use the properties of logarithms to simplify and express equations 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's break down the sentence:

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

"For, while he had a head **for** batting" – This suggests that he had a natural skill or talent for batting.

"He could only do so with a cool head **on** his shoulders" – This phrase implies that he needed to stay calm and composed under pressure.

Thus, the correct sequence of words to complete the sentence is:

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

Final Answer:

$$\boxed{\text{down, out, for, on}}$$

Quick Tip
Understanding the context and common idiomatic expressions in English helps in selecting the correct prepositions and phrases.

---

7. A rectangular paper sheet of dimensions 54 cm $\times$ 4 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 requires us to compare the volumes of a cube and a cylinder, both derived from the same surface area.

Step 1: Surface area of the sheet The area of the sheet is calculated as:

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

This represents the surface area of the cube.

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

$$6a^2 = 216$$

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

Therefore, 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 the cylinder. The surface area of the cylinder consists of 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:

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

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

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

where $h = 54\,\text{cm}$ is the 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:

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

Quick Tip
When converting a rectangle into a cylinder, the length forms the circumference, while 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.

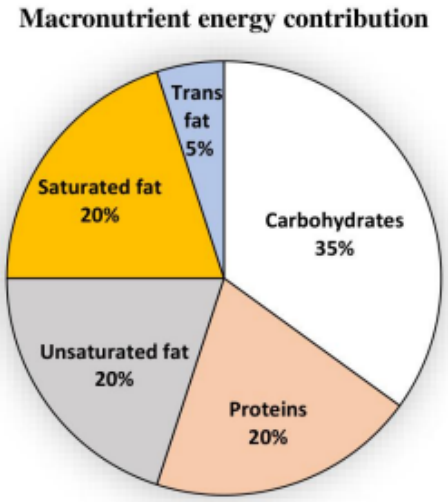

Macronutrient	Energy density (kcal/g)
Carbohydrates	4
Proteins	4
Unsaturated fat	9
Saturated fat	9
Trans fat	9

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 the total fat contribution in kcal.

The total fat (including saturated, unsaturated, and trans fats) 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, we calculate:

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

Final Answer:

$$\boxed{100}$$

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

---

9. A rectangular paper of 20 cm $\times$ 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: Determine the dimensions after each fold. Initially, the paper has dimensions of $20\,\text{cm} \times 8\,\text{cm}$.

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

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

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

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

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

Final Answer:

$$\boxed{18\,\text{cm}}$$

Quick Tip
When calculating the perimeter after folding, note that each fold alters 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:

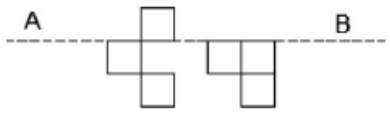

(1) 6

(2) 4

(3) 5

(4) 7

Correct Answer: (1) 6

Solution: Step 1: Examine the diagram. The line AB serves as an axis of symmetry, meaning each square above or below AB must have an identical square placed symmetrically on the opposite side.

Step 2: Add squares to achieve symmetry. To ensure the figure is symmetric about AB:

Place 2 squares on the leftmost side below AB.

Place 2 squares in the middle section beneath AB.

Place 2 squares on the rightmost side below AB.

Step 3: Count the total squares to be added. The minimum number of squares to add is:

$$6 \text{ squares}.$$

Final Answer:

$$\boxed{6 \text{ squares}}$$

Quick Tip
When solving symmetry-related problems, focus on ensuring all elements are mirrored accurately relative to the symmetry line.

Question 11 - 35 carry one mark each

11. Let $f : \mathbb{R} \to \mathbb{R}$ 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: Understand the behavior of the function $f(x)$. The function is given as:

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

From this definition:

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

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

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

Final Answer:

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

Quick Tip
To locate points of non-differentiability in a piecewise function, focus on intersections of components or abrupt changes in the derivative.

---

12. The product of all eigenvalues of the matrix

$$\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix}$$ **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 equals the determinant of the matrix.

Step 2: Compute the determinant.

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

Calculate 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 these values:

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

Step 3: Conclusion. The determinant is 0, indicating that the product of the eigenvalues is also 0.

Quick Tip
The determinant of a matrix provides a quick method to find the product of its eigenvalues. If the determinant is zero, it implies that 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 the concept of 5-bit signed integer representation. For a 5-bit signed integer in 2's complement, the range of representable numbers is:

$$[-16, 15].$$

Next, let's convert the binary numbers A and B into decimal: - $A = 01010_2$ (positive number), so $A = 10$. - $B = 11010_2$ (negative number), so $B = -6$.

Step 2: Perform the operations with a focus on checking overflows. Let's evaluate each operation step-by-step and verify if there is any overflow:

1. **Addition $A + B$:**

$A + B = 10 + (-6) = 4$ (No overflow, as the result is within the range [-16, 15]).

2. **Subtraction $A - B$:**

$A-B = 10-(-6) = 10+6 = 16$ (Overflow occurs, as 16 exceeds the maximum value of 15 in the range

3. **Subtraction $B - A$:**

$B - A = -6 - 10 = -16$ (No overflow, as the result is within the range [-16, 15]).

4. **Doubling $2 \times B$:**

$2 \times B = 2 \times (-6) = -12$ (No overflow, as the result is within the range [-16, 15]).

Step 3: Identify which operation causes the overflow. The operation $A - B$ leads to an overflow because the result of 16 exceeds the maximum representable value of 15 for a 5-bit signed integer.

Final Answer:

$$\boxed{A - B}$$

—

Quick Tip

Quick Tip

In 2's complement representation, overflow occurs when the result of an operation falls outside the allowable range of values. Ensure that the result is within the range $[-2^{n-1}, 2^{n-1} - 1]$, where n is the number of bits in the signed integer.

---

14. Consider a permutation sampled uniformly at random from the set of all permutations of $\{1, 2, 3, \ldots, n\}$ 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: The element 1 appears before 2 in the permutation.

Event Y: The element 3 appears before 4 in the permutation.

Step 2: Determine the probabilities.

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

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

Step 3: Check for independence.

The events X and Y are independent if:

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

Since the relative order of 1 and 2 is independent of the relative order of 3 and 4, we can calculate:

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

Step 4: Conclusion. Since the calculated probabilities are the same, we conclude that events X and Y are independent.

Final Answer:

$$\boxed{\text{The events } X \text{ and } Y \text{ are independent.}}$$

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: Evaluate each statement.
1. **Cycle stealing mode in DMA:** In this mode, the DMA controller takes over CPU cycles temporarily to transfer a single word of data. This is a correct statement.
2. **Burst mode in DMA:** Burst mode transfers large blocks of data at once, enabling greater data throughput than cycle stealing mode. This is also correct.
3. **Programmed I/O vs. Interrupt-driven I/O:** In Programmed I/O, the CPU repeatedly checks the device for data, resulting in inefficient CPU usage. Interrupt-driven I/O, however, allows the CPU to perform other tasks and only handle the device when necessary. Therefore, this statement is false.

4. **Vectored interrupts:** Vectored interrupts provide predefined addresses for interrupt service routines, allowing for quicker execution. Non-vectored interrupts require additional steps to find 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
In Programmed I/O, the CPU keeps polling the device, leading to inefficient resource use. Interrupt-driven I/O allows the CPU to perform other tasks while waiting for the device to signal.

---

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 Query (ii): The client sends a DNS query to resolve the server's hostname to its IP address since there is no cached information available.
Step 2: TCP SYN Packet (iv): After receiving the DNS response, a TCP SYN packet is sent to initiate the connection with the server.
Step 3: HTTP GET Request for the Index Page (i): Once the TCP connection is established, an HTTP GET request is sent to retrieve the index page of the website.
Step 4: HTTP GET Requests for Embedded Images (iii): After the index page is successfully loaded, additional HTTP GET requests are made to fetch the image objects embedded in the page.

Final Answer:

$$\boxed{(ii), (iv), (i), (iii)}$$

Quick Tip
When accessing a webpage with no cached data, the process involves DNS resolution, establishing a connection, fetching the main page, and retrieving any embedded objects like images.

---

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: Problem Definition. The algorithm checks adjacent elements of the array to verify if it is sorted in either ascending or descending order. This process requires only one pass through the array.

Step 2: Time Complexity Analysis.

In the worst case scenario, the algorithm must iterate through the entire array, performing $N - 1$ comparisons.

Therefore, the time complexity is $O(N)$, as the number of comparisons is directly proportional to the size of the array N.

Since the algorithm consistently performs $N - 1$ comparisons, the lower bound of the time complexity is also $\Omega(N)$.

Final Answer:

$$\boxed{\text{both } O(N) \text{ and } \Omega(N)}$$

Quick Tip
For algorithms that involve a single pass 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: Program Analysis.

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

Within the `while` loop, the value of a is updated using the formula:

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

Since integer division is used, dividing a (which is initially 6) by 12 results in 0, as only the integer part is kept. Thus:

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

Next, the statement $a+ = b$ is executed, but since $b = 0$, the value of a remains 1.

Step 2: Infinite Loop Condition.

The `while` loop continues as long as $a < 10$. After the first iteration:

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

In every subsequent iteration, the update expression $a = \frac{a}{12} + 1$ keeps resetting a to 1, meaning 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 cause infinite loops if the update condition relies on such calculations.

---

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 behavior.

The function `fX()` starts by reading a character using `getchar()`.

If the character is not a newline character (`"/n"`), the function calls itself recursively.

Once the recursive call finishes, the function prints the character using `putchar()` unless it is a newline.

Step 2: Track the recursion flow.

Upon the first call to `fX()`, it reads the character '1' and recursively invokes the function

again.

The second call reads '2', followed by further recursive calls reading '3' and '4'.

When the newline character is encountered, recursion stops and begins to unwind.

Step 3: Determine the output during the unwind.

As the recursion unwinds, the characters are printed in reverse order—'4', '3', '2', '1'—due to the last-in, first-out (LIFO) nature of the call stack.

Final Answer:

The program will output 4321 when it terminates.

Quick Tip
In recursive functions, the sequence in which output is printed depends on the order in which the stack unwinds.

---

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?

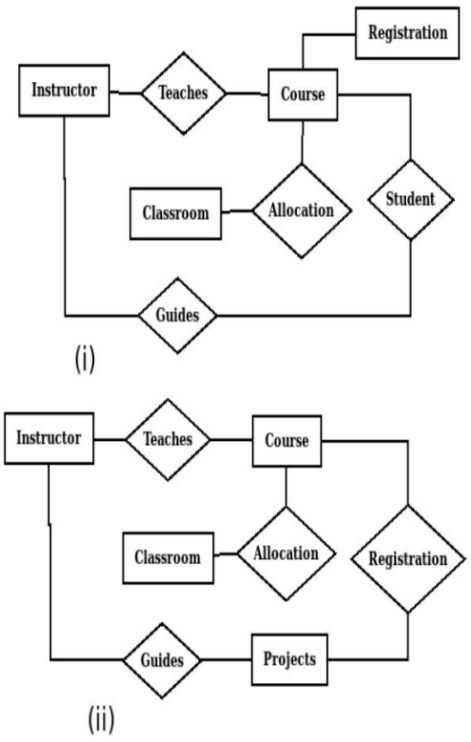

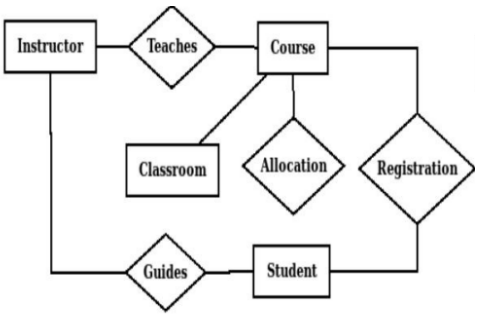

(iii)

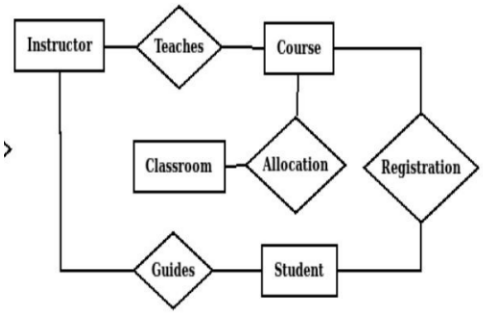

(1) Diagram (i)

(2) Diagram (ii)

(3) Diagram (iii)

(4) Diagram (iv)

Correct Answer: (4) Diagram (iv)

Solution: The correct ER diagram must meet the following criteria:

1. **Instructors teach courses:** A relationship named ”Teaches” should exist between ”Instructor” and ”Course.”

2. **Students register for courses:** There must be a ”Student” entity and a ”Registration” relationship linking ”Student” to ”Course.”

3. **Courses are assigned to classrooms:** A relationship named ”Allocation” should link ”Course” and ”Classroom.”

4. **Instructors guide students:** A ”Guides” relationship must exist between ”Instructor” and ”Student.”

Evaluation of the Options:

Option (i): Missing the ”Registration” relationship between ”Student” and ”Course,” so it is incorrect.

Option (ii): Contains an additional entity ”Projects,” which is not part of the specifica-

tion. Thus, it is incorrect.

Option (iii): Fails to include the ”Guides” relationship between ”Instructor” and ”Student,” so it is incorrect.

Option (iv): Includes all the required entities and relationships according to the specification, making it the correct choice.

Final Answer:

Diagram (iv)

Quick Tip
When reviewing ER diagrams, ensure that every specified relationship and entity is included and avoid adding unnecessary elements.

---

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: Understand the occupancy requirements for nodes in a B+ tree.

In a B+ tree, all nodes, except for the root, must maintain a minimum occupancy of 50% to ensure that the tree remains balanced.

Step 2: Recognize the exception to the rule.

The root node is an exception to the occupancy requirement. It can have fewer than 50% of its capacity and even contain just one key, particularly in cases like deletions.

Final Answer:

Only the root node

Quick Tip

In a B+ tree, while all nodes (except the root) must have at least 50% occupancy, the root node is allowed to have fewer keys, particularly during operations like deletions.

---

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: Evaluate the statements.

1. R can have a multi-attribute key: Multi-attribute keys are permissible 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 because they ensure referential integrity.

3. R cannot have a composite attribute: Since 1NF requires that attributes contain atomic values, composite attributes are not permitted. This statement is **true**.

4. R cannot have more than one candidate key: A relation in 1NF can have multiple candidate keys. Therefore, this statement is **false**.

Final Answer:

$$\boxed{\text{(1) and (3)}}$$

Quick Tip

In 1NF, all attributes must be atomic, meaning composite attributes are not allowed. However, multi-attribute keys and foreign keys are valid.

---

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: Examine the condition $L_1 = L_2$.

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

Step 2: Evaluate $L_1 \cup L_3$.

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

Step 3: Examine L_3.

Since L_3 is explicitly stated 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 result in a regular language. Hence, $\overline{L_1} \cup \overline{L_2}$ is **regular**, making statement (4) **true**.

Final Answer:

$$\boxed{\text{(3) and (4)}}$$

Quick Tip
For regular languages, operations like complement, union, and intersection preserve regularity. However, non-regular languages do 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). Therefore, statement (1) is **false**.

Step 2: File descriptor table.

All threads within the same process share the same file descriptor table. Hence, statement (2) is **false**.

Step 3: Shared resources.

Threads of the same process share various resources such as memory and file descriptors, but each thread maintains its own stack. Thus, statement (3) is **false**.

Step 4: Protection between threads.

Since threads within a process share the same address space, they are not isolated from each other and can potentially interfere with one another. This makes statement (4) **true**.

Final Answer:

(4)

Quick Tip
Threads within a process share memory and resources, but they do not have built-in protection against interference with each other.

---

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: Review the process state transitions. The states a process can be in include:

- **Ready**: The process is prepared to run but is waiting for CPU time.
- **Running**: The process is actively executing on the CPU.
- **Waiting**: The process is on hold, waiting for an event to occur (e.g., I/O operation to complete).
- **Terminated**: The process has completed execution and is no longer running.

Step 2: Evaluate each transition.

1. **Running to Ready:** This transition is valid when the process is preempted by the scheduler due to a higher-priority process or when its time slice expires.
2. **Waiting to Running:** This transition is NOT valid. A process must first transition from the Waiting state to the Ready state after the event it is waiting for completes, and then it can move from Ready to Running.
3. **Ready to Waiting:** This transition is NOT valid. A process in the Ready state can only move to Running when scheduled. Waiting can occur only after execution starts, such as when the process requests I/O.
4. **Running to Terminated:** This transition is valid when the process finishes its execution.

Step 3: Identify the impossible transitions.

Based on the above evaluation, the transitions from Waiting to Running and from Ready to Waiting are NOT possible.

Final Answer:

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

Quick Tip
To transition from Waiting to Running, a process must first move to the Ready state. Similarly, a process cannot transition directly from Ready to Waiting without first executing.

---

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 strategies. Parsing strategies can be broadly classified into:

Top-Down Parsing: Begins construction of the parse tree from the root (start symbol) and progresses towards the leaves.

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

Step 2: Evaluate the given parsers.

1. **Shift-reduce Parser:** This parser follows a bottom-up approach. It uses shift and reduce operations to gradually build the parse tree, making it a Bottom-Up Parser.

2. **Predictive Parser:** A top-down parsing technique that uses lookahead symbols to predict which rule to apply next. Therefore, it is NOT a Bottom-Up Parser.

3. **LL(1) Parser:** This is another top-down parsing technique that employs leftmost derivations with a single lookahead symbol. It is NOT a Bottom-Up Parser.

4. **LR Parser:** A bottom-up parsing technique that reads input from left to right and uses rightmost derivation in reverse. It qualifies as a Bottom-Up Parser.

Step 3: Identify the Bottom-Up Parsers.

Based on the analysis, the Shift-reduce Parser and LR Parser are both Bottom-Up Parsers.

Final Answer:

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

Quick Tip
Bottom-up parsers build the parse tree from the input string (the leaves) and work towards the start symbol (the root). Examples of bottom-up parsers are the Shift-reduce and LR Parsers.

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 the independence of events A and B. Two events A and B are considered independent if:

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

Given:

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

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

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

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

we substitute the values:

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

Thus, Option (2) is **TRUE**.

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

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

we substitute the values:

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

Thus, Option (3) is **TRUE**.

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

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

we 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:

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

Quick Tip
When solving probability problems, verify the independence of events using $P(A \cap B) = P(A) \cdot P(B)$ and apply complement and union rules for efficient calculation.

---

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?

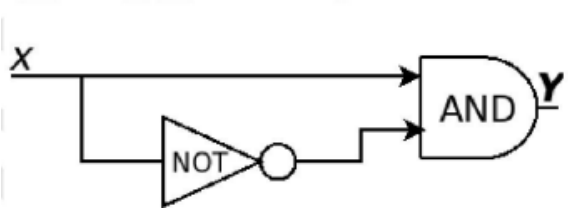

(1) With no propagation delays, the output Y is always logic Zero
(2) With no propagation delays, the output Y is always logic One
(3) With propagation delays, the output Y can have a transient logic One after X transitions from logic Zero to logic One
(4) With propagation delays, the output Y can have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

The circuit consists of an AND gate and a NOT gate. The input to the AND gate is X, and the output of the NOT gate is $\neg X$.

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

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

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 delay, both inputs to the AND gate might temporarily be logic One.

This results in a brief logic One at Y, causing a transient output. Thus, Option (3) is **TRUE**.

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

Final Answer:

$$\boxed{\text{(1) and (3)}}$$

Quick Tip
Propagation delays in logic circuits can cause brief incorrect outputs. Always account for gate timing when analyzing circuits during input 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 randomly selected by the client P to ensure both security and uniqueness for the TCP connection. Thus, Option (1) is **TRUE**.

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

Step 2: Acknowledgement number N_Q.

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

The acknowledgement number is set to $N_P + 1$, indicating that the server expects the next byte to be $N_P + 1$. Therefore, Option (4) is **TRUE**, and Option (3) is **FALSE**.

Final Answer:

$$\boxed{(1) \text{ and } (4)}$$

Quick Tip
TCP sequence numbers are selected randomly to prevent sequence number prediction attacks, and acknowledgements are used to 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 address data hazards by directly passing data between pipeline stages, bypassing the register file. This statement is **TRUE** for Option (1).

Step 2: Forwarding in the MEM and EX stages.

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

Step 3: Limitations of forwarding.

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

Step 4: Hardware requirements for forwarding.

Forwarding requires additional hardware, such as forwarding paths and control logic, to allow data to bypass between pipeline stages. This makes Option (4) **FALSE**.

Final Answer:

$$\boxed{(1), (2), \text{ and } (3)}$$

Quick Tip

Forwarding helps address data hazards effectively, but cannot resolve all pipeline stalls. Load-use hazards still require stalls in the pipeline.

---

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: First Analyzing the NAT behavior.

When a packet traverses through a NAT device, the source IP address is replaced with the NAT device's external IP address to facilitate communication with the external network. As a result, the source IP is modified, making Option (1) **TRUE**.

Step 2: Checksum recalculation.

Since the source IP address is altered, the checksum, which is based on the header fields, must be recalculated. This means the header checksum is updated, making Option (3) **TRUE**.

Step 3: Other fields.

The destination IP remains unchanged, as it points to the external destination, and the total length of the packet is unaffected by NAT translation. Therefore, Options (2) and (4) are **FALSE**.

Final Answer:

$$\boxed{(1) \text{ and } (3)}$$

Quick Tip

In NAT translation, the Source IP and Header Checksum fields are altered, while the 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: Requirement for a bijection between A and B.

A bijective function exists between A and B only if both sets have the same number of elements. Denote $|A| = |B| = n$.

Step 2: Condition for a bijection between $A \times A$ and $A \cup B$.

The size of $A \times A$ is n^2, while the size of $A \cup B$ is $|A| + |B| = 2n$. For a bijection to exist, these sizes must match:

$$n^2 = 2n.$$

Step 3: Solve for n.

Rewriting the equation:

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

This yields $n = 0$ or $n = 2$. Since A and B are non-empty, n must equal 2.

Final Answer:

$$\boxed{2}$$

Quick Tip

For bijections to exist, set cardinalities must align exactly. Verify these conditions by equating the cardinalities of the involved sets.

---

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: Understand operator precedence and associativity.

Using the given rules for operator precedence and associativity:

1. Highest precedence: $+$, evaluated left to right.
2. High precedence: $-$, evaluated right to left.
3. Medium precedence: $*$, evaluated right to left.
4. Low precedence: $/$, evaluated right to left.

Step 2: Evaluate the expression step by step.

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 back into the expression:

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

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 intermediate results:

$$7 - 1 = 6$$

Final Answer:

$$\boxed{6}$$

Quick Tip
To simplify expressions with operator precedence and associativity, carefully resolve each operator step by step in the correct order, starting with the highest precedence.

---

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: Use Cayley's formula for spanning trees in a complete graph.

The formula for the number of spanning trees in a complete graph with n vertices is:

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

Step 2: Calculate for $n = 4$.

Substitute $n = 4$ into the formula:

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

Final Answer:

$$\boxed{16}$$

Quick Tip
Cayley's formula $T = n^{n-2}$ quickly calculates the number of spanning trees in a complete graph. Always verify n and compute the power carefully.

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 merges tuples from R and S where the condition $R.A = S.A$ holds true.

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

A	B	C
10	20	90
30	40	45

Step 2: Apply the selection condition $B < C$. We select tuples where $B < C$, which gives us:

A	B	C
10	20	90
30	40	45

The total number of tuples satisfying the condition is 2.

Final Answer:

$$\boxed{2}$$

Quick Tip
When evaluating relational algebra expressions, follow this order: (1) perform the join, (2) apply the selection condition, and (3) count the resulting tuples.

Question 36 - 65 carry two marks each

36. Consider a network path $P \to Q \to R$ 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 $P \to Q$ and $Q \to R$ 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 $P \to Q$ and transmit on link $Q \to R$.
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 has a size of 10^3 bytes $= 8 \times 10^3$ bits $= 8000$ bits.

Step 2: Transmission time for one chunk.

The bandwidth of the 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.

P transmits one chunk every 0.008 seconds, and there are 1000 chunks in total. The time for P to send all chunks to Q is:

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

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

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

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

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

The last chunk reaches Q at $t = 8$ seconds and is transmitted to R in 0.008 seconds. Hence, 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 total time taken for all chunks to reach R is:

$$\boxed{8.008 \text{ seconds}}$$

Quick Tip
For sequential data transmissions, account for both the total time to send all chunks and the propagation time of the last chunk through intermediate routers to determine the final delivery time.

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

$$
\begin{aligned}
S \to DHTU \quad & \{S.val = D.val + H.val + T.val + U.val;\} \\
D \to \text{"}M\text{"}\, D_1 \quad & \{D.val = 5 + D_1.val;\} \\
D \to \epsilon \quad & \{D.val = -5;\} \\
H \to \text{"}L\text{"}\, H_1 \quad & \{H.val = 5 \cdot 10 + H_1.val;\} \\
H \to \epsilon \quad & \{H.val = -10;\} \\
T \to \text{"}C\text{"}\, T_1 \quad & \{T.val = 5 \cdot 100 + T_1.val;\} \\
T \to \epsilon \quad & \{T.val = -5;\} \\
U \to \text{"}K\text{"} \quad & \{U.val = 5;\}
\end{aligned}
$$

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: Analyze the input sequence "MMLK" using an alternate method.

The sequence can be processed as:

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

Step 2: Evaluate the values for $D, H, T,$ and U. Instead of decomposing each rule one by one, compute directly:

- $D \to \text{"}MM\text{"}$: Each "M" contributes 5, so:

$$D.val = 5 + 5 = 10.$$

- $H \to "L"$: "L" contributes 50, adjusted by $H_1.val = -10$ (as $H_1 \to \epsilon$):

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

- $T \to \epsilon$: T contributes:

$$T.val = -5.$$

- $U \to "K"$: "K" directly contributes:

$$U.val = 5.$$

Step 3: Combine all values to calculate $S.val$. The value of S is the sum of contributions from $D, H, T,$ and U:

$$S.val = D.val + H.val + T.val + U.val = 10 + 40 - 5 + 5 = 45.$$

Final Answer:

$$\boxed{45}$$

Quick Tip
To solve syntax-directed definitions, use a direct substitution approach for straightforward rules, simplifying calculations while maintaining accuracy.

---

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: Evaluate the impact of productions on FIRST and FOLLOW sets.

The FIRST and FOLLOW sets guide the structure of valid grammar productions:

1. $S \rightarrow daT \mid Rf$: - For $FIRST(S)$, the elements daT and Rf contribute $\{c, d, f\}$. Here, R can produce c or ϵ, and f is directly included.

2. $T \rightarrow aS \mid bT \mid \epsilon$: - The set $FIRST(T) = \{a, b, \epsilon\}$ matches only if $T \rightarrow \epsilon$ is valid.

3. $R \rightarrow cTR \mid \epsilon$: - For R, $FIRST(R) = \{c, \epsilon\}$, satisfied by this production.

Step 2: Confirm the validity of the given options. 1. $S \rightarrow Rf$, $T \rightarrow \epsilon$, $R \rightarrow cTR$: - This choice maintains alignment with the FIRST and FOLLOW set requirements without introducing ambiguity.

2. Other options violate the FIRST/FOLLOW set properties or include grammars inconsistent with the given constraints.

Final Answer:

$$\boxed{(1)\ S \to Rf, T \to \epsilon, R \to cTR}$$

Quick Tip
To verify grammar consistency, ensure that productions satisfy both FIRST and FOLLOW sets without ambiguity, and validate each option systematically.

---

39. Consider the following pseudo-code:

L1: $t1 = -1$

L2: $t2 = 0$

L3: $t3 = 0$

L4: $t4 = 4 \times t3$

L5: $t5 = 4 \times t2$

L6: $t6 = t5 \times M$

L7: $t7 = t4 + t6$

L8: $t8 = a[t7]$

L9: **if** $t8 \leq$ 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: Define basic blocks.

A basic block is a straight-line sequence of instructions with: 1. A single entry point, meaning control flows into the block from the beginning. 2. A single exit point, meaning control leaves the block at the end without any jumps except at the last instruction. 3. No statement in the block serves as a jump target for external instructions.

Step 2: Determine the basic blocks. 1. **Basic Block 1 (L1 - L3):** Includes $t1 = -1$, $t2 = 0$, $t3 = 0$. These are simple assignments without any branching.

2. **Basic Block 2 (L4 - L8):** Comprises $t4 = 4 \times t3$, $t5 = 4 \times t2$, $t6 = t5 \times M$, $t7 = t4 + t6$, $t8 = a[t7]$. The block ends when $t8 \leq$ max introduces a conditional jump.

3. **Basic Block 3 (L9 - L10):** Contains if $t8 \leq$ max goto L11 and $t1 = t8$.

4. **Basic Block 4 (L11):** Consists of $t3 = t3 + 1$.

5. **Basic Block 5 (L12):** Includes 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: Identify the largest basic block.

The largest block is **Block 2 (L4 - L8)**, which consists of 6 instructions.

Final Answer:

$$\boxed{\text{7 and 6}}$$

Quick Tip
To identify basic blocks, examine the program flow and divide the code into segments where control flows sequentially without interruptions except at the last instruction.

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.

$$\text{T1:} \quad a = a + 1; \quad b = b + 1;$$

$$\text{T2:} \quad b = 2 \times b; \quad a = 2 \times a;$$

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 operations in each thread.

Thread T1:

$$\text{Increments } a : a = a + 1, \quad \text{then increments } b : b = b + 1.$$

Thread T2:

$$\text{Doubles } b : b = 2 \times b, \quad \text{then doubles } a : a = 2 \times a.$$

Step 2: Explore the possible execution orders.

The order in which T1 and T2 execute can vary, leading to different results: 1. **T1 fully executes before T2:** - After T1: $a = 1 + 1 = 2$, $b = 1 + 1 = 2$. - After T2: $a = 2 \times 2 = 4$, $b = 2 \times 2 = 4$. - Result: $(a = 4, b = 4)$.

2. **T2 fully executes before T1:** - After T2: $a = 2 \times 1 = 2$, $b = 2 \times 1 = 2$. - After T1: $a = 2 + 1 = 3$, $b = 2 + 1 = 3$. - Result: $(a = 3, b = 3)$.

3. **Interleaved execution of T1 and T2:** - If b is doubled by T2 first and then incremented by T1: $b = (2 \times 1) + 1 = 3$. - If a is incremented by T1 first and then doubled by T2: $a = 2 \times (1 + 1) = 4$. - Result: $(a = 4, b = 3)$.

Step 3: Summarize the possible outcomes.

The different outcomes based on the execution order are:

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

Final Answer:

$$\boxed{(a = 4, b = 4); (a = 3, b = 3); (a = 4, b = 3)}$$

Quick Tip
When analyzing multithreaded code, always consider all possible instruction interleavings to account for every potential outcome of 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 is the process of converting an array into a binary heap structure. In a **max heap**, the root node is the largest element, and each parent node is greater than or equal to its child nodes.

Step 2: Applying heapification to the array.

Given the array:

$$[82, 101, 90, 11, 111, 75, 33, 131, 44, 93]$$

1. The largest element 131 is moved to the root of the heap.

2. The second and third largest elements, 111 and 101, become the left and right children of the root.

3. The remaining elements are positioned to maintain the heap property.

Step 3: Extracting the first three elements.

After transforming the array into a max heap, the first three elements are the root and its immediate children:

$$\boxed{131, 111, 101}.$$

Quick Tip
Heapification ensures the maximum element is at the root in a max heap. To identify the top elements, examine the root and its immediate children.

---

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: Analyze the given recurrence.

The recurrence is:

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

Let $n = 2^m$, so $\sqrt{n} = 2^{m/2}$. Rewriting the recurrence:

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

Step 2: Transform the recurrence.

Define $T(2^m) = f(m)$. Substituting:

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

Step 3: Expand the recurrence.

Expanding the recurrence iteratively:

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

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

Each expansion adds another 2^m, and this continues for approximately $\log\log n$ levels since $n = 2^m$ reduces to 1 in $\log\log n$ steps.

Step 4: Compute the complexity.

Each level contributes n, and there are $\log\log n$ levels. Thus, 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 terms like $T(\sqrt{n})$, substitute $n = 2^m$ to simplify the recurrence into a function of m, making it easier to analyze the growth.

---

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: Understanding the min-heap property.

In a binary min-heap: 1. The root (index 1) contains the smallest element. 2. Elements increase in value as we move down the tree. 3. The maximum element will always be located in the last level, which contains the leaf nodes.

Step 2: Determine the last level of the heap.

The height (number of levels) of a binary heap with n elements is:

$$h = \lceil \log_2 n \rceil.$$

For $n = 105$:

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

The 7th level is the last level of the heap.

Step 3: Identify the indices of the last level elements.

In an array representation of a binary heap: - The leaf nodes start at index $\lceil n/2 \rceil$ and end at n.

For $n = 105$:

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

Thus, the leaf nodes are at indices 53 to 105.

Step 4: Calculate the number of possible indices for the maximum element.

The number of indices in the last level (where the maximum element can be found) is:

$$n - \lceil n/2 \rceil + 1 = 105 - 53 + 1 = 53.$$

Final Answer:

$$\boxed{53}$$

Quick Tip

In a binary min-heap, the maximum element always lies in the last level. Calculate the starting and ending indices of the last level using $\lceil n/2 \rceil$ and n.

---

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: Evaluate each functional dependency statement.

Option (1): $(X, Y) \rightarrow (Z, W)$ does not imply $X \rightarrow (Z, W)$ because X alone might not uniquely determine Z or W without Y. This violates the definition of functional dependency. **Conclusion: FALSE.**

Option (2): $(X, Y) \rightarrow (Z, W)$ directly implies $(X, Y) \rightarrow Z$ because Z is part of the set determined by (X, Y). **Conclusion: TRUE.**

Option (3): Given $(X, Y) \rightarrow Z$ and $W \rightarrow Y$, we can deduce $(X, W) \rightarrow Z$ using the substitution property of functional dependency. **Conclusion: TRUE.**

Option (4): By the transitivity property of functional dependencies, if $X \rightarrow Y$ and $Y \rightarrow Z$, then $X \rightarrow Z$ is valid. **Conclusion: TRUE.**

Final Answer:

$$\boxed{(2), (3), (4)}$$

Quick Tip
Use properties like reflexivity, augmentation, and transitivity to systematically analyze functional 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 the properties of DFS and BFS trees.

Option (1): Back-edges are present in a DFS tree if the graph G contains cycles. For T to be both a DFS and BFS tree, back-edges cannot exist because BFS traversal does not encounter back-edges. **Conclusion: FALSE.**

Option (2): Cross-edges in DFS connect vertices in different branches of the tree. BFS trees, however, do not exhibit cross-edges. The condition of T being both a DFS and BFS tree does not universally eliminate cross-edges in G. **Conclusion: FALSE.**

Option (3): Forward-edges in DFS connect a vertex to its descendant that is not part of the DFS tree. If T is both a DFS and BFS tree, forward-edges cannot exist because BFS does not allow edges to bypass intermediate levels of the tree. **Conclusion: TRUE.**

Option (4): The condition of T being both a DFS and BFS tree does not require G to contain only tree edges. Non-tree edges like back or cross-edges may still exist in G. **Conclusion: FALSE.**

Final Answer:

$\boxed{(3)}$

Quick Tip
Classify edges in DFS as tree, back, forward, or cross, and understand their behavior in relation to BFS to analyze tree structures effectively.

---

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: Define conflict equivalence.

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

Step 2: Analyze the given schedule S.

The schedule S involves: - $T_1 : r_1(z), w_1(z)$ - $T_2 : r_2(x), w_2(x), r_2(y), w_2(y)$ - $T_3 : r_3(y), w_3(y)$

Conflicts: - T_1 and T_2 conflict on z and x. - T_2 and T_3 conflict on y.

Step 3: Test different orderings of transactions.

1. **Schedule $T_1T_2T_3$:** - T_3's operations on y may overlap with T_2's operations on y, breaking conflict equivalence. - **Conclusion: Not conflict equivalent.**

2. **Schedule $T_1T_3T_2$:** - T_1's operations on z precede T_3's operations on y, and T_3's operations on y precede T_2's operations on y. - **Conclusion: Conflict equivalent.**

3. **Schedule $T_3T_2T_1$:** - T_3's operations on y precede T_2's operations on y, and T_2's operations on x and z precede T_1's operations. - **Conclusion: Conflict equivalent.**

4. **Schedule $T_3T_1T_2$:** - T_3's operations on y precede T_1's operations on z, and T_1's operations precede T_2's operations on x and y. - **Conclusion: Conflict equivalent.**

Final Answer:

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

Quick Tip
Conflict equivalence ensures that the order of conflicting operations remains consistent across transactions. Test each schedule systematically to confirm equivalence.

---

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: Construct the truth table for $F(X, Y, Z)$.

The function $F(X, Y, Z) = \Sigma(3, 5, 6, 7)$ is defined by the minterms where $F(X, Y, Z) = 1$. These correspond to the binary values:

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

Step 2: Convert to the product-of-maxterms form.

Maxterms represent combinations where $F(X, Y, Z) = 0$. These occur for:

$$0(000), 1(001), 2(010), 4(100).$$

Thus, the complement of $F(X, Y, Z)$ is expressed as:

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

This confirms that Option (1) is TRUE.

Step 3: Simplify the Boolean expression.
Using the minterm expansion, $F(X, Y, Z)$ can be simplified to:

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

This confirms that Option (2) is TRUE.

Step 4: Verify dependency on variables.
The function $F(X, Y, Z)$ depends on all three variables (X, Y, Z), as changing the value of any variable affects the output. Thus, $F(X, Y, Z)$ is NOT independent of X or Y, making both Option (3) and Option (4) FALSE.

Final Answer:

$$\boxed{(1), (2)}$$

Quick Tip
When working with Boolean functions, identify minterms and maxterms to derive equivalent expressions. Simplify using Boolean algebra or Karnaugh maps.

---

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 iterative function.

The function iteratively updates x over y iterations using the formula:

$$x = 2x + y.$$

Starting with an initial value of x, this causes x to grow exponentially with each iteration, as $2x$ dominates the addition of y.

Step 2: Analyze specific cases.

1. **Case 1: $x = 20, y = 10$** After 10 iterations, x grows significantly, surpassing 2^{10}. **Conclusion: The result is greater than 2^{10}.**

2. **Case 2: $x = 20, y = 20$** Over 20 iterations, x grows exponentially and exceeds 2^{20}. **Conclusion: The result is greater than 2^{20}.**

3. **Case 3: $x = 10, y = 20$** With $x = 10$ and $y = 20$, after 20 iterations, x also surpasses 2^{20}. **Conclusion: The result is greater than 2^{20}.**

4. **Case 4: $x = 20, y = 10$** After 10 iterations, x is much larger than 2^{10}, making the statement $x < 2^{10}$ invalid. **Conclusion: FALSE.**

Final Answer:

$$\boxed{(2), (4)}$$

Quick Tip
Exponential growth in iterative formulas quickly leads to large values. Use the iteration logic to predict growth and compare results with reference values.

---

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:

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

Since $\text{Rank}(A) \leq n$ (as A has n rows), the nullity is:

$$\text{Nullity}(A) \geq m - n.$$

This means there are at least $m - n$ linearly independent vectors in the null space of A, which form the basis of the null space.

Step 2: Validate the options.

- **Option (1):** There are at least $m - n$ linearly independent solutions in the null space of A, as derived from the rank-nullity theorem. **Conclusion: TRUE.**

- **Option (2):** While $m - n$ linearly independent vectors span the null space, they are not necessarily all solutions to $Ax = 0$ without additional conditions. **Conclusion: FALSE.**
- **Option (3):** The nullity does not imply that $m - n$ specific variables are 0 in every solution. **Conclusion: FALSE.**
- **Option (4):** There is no requirement for at least n variables to be non-zero in a solution. This depends on the structure of A. **Conclusion: FALSE.**

Final Answer:

$$\boxed{(1)}$$

Quick Tip
For an $n \times m$ matrix where $m > n$, use the rank-nullity theorem to analyze the null space and identify the number of 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.

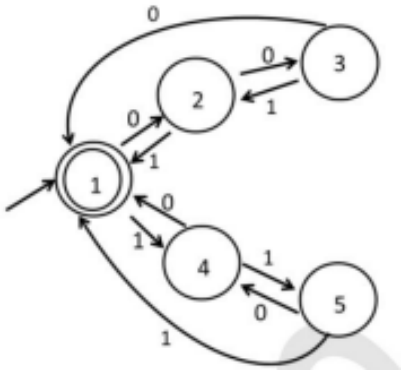

Which of the following statements is/are FALSE?

(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 the language it accepts.

The DFA M operates based on the counts of 0's ($n_0(w)$) and 1's ($n_1(w)$) in an input string w. - States in the DFA are distinguishable if they lead to different acceptance behaviors for at least one input string. - States are indistinguishable if they behave identically for all input strings.

Step 2: Evaluate the given options.

- **Option (1):** States 2 and 4 are distinguishable because they lead to different acceptance outcomes for certain input strings. For example, appending specific combinations of 0's or 1's results in transitions to different accepting or non-accepting states. **Conclusion: TRUE.**
- **Option (2):** States 3 and 4 are not indistinguishable. They behave differently for certain input strings, resulting in distinct acceptance or rejection outcomes. **Conclusion: FALSE.**
- **Option (3):** States 2 and 5 are distinguishable. These states do not represent equivalent conditions for $n_0(w)$ and $n_1(w)$ and lead to different behaviors for some input strings. **Conclusion: FALSE.**
- **Option (4):** Any string w such that $n_0(w) = n_1(w)$ is accepted by M. The DFA ensures balance between 0's and 1's to reach an accepting state. **Conclusion: TRUE.**

Final Answer:

$$\boxed{(2), (3)}$$

Quick Tip
To verify distinguishability in a DFA, test input strings that exploit differences in transitions or acceptance conditions for the states under consideration.

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 subset of vertices in which no two vertices are adjacent. For a graph with chromatic number k, the vertices can be partitioned into k independent sets. The largest independent set must therefore contain at least:

$$\frac{n}{k} \quad \text{vertices.}$$

This makes Option (2) TRUE.

Step 2: Minimum number of edges.

The chromatic number k indicates that there exists a complete subgraph K_k. A complete graph K_k contains:

$$\frac{k(k-1)}{2} \quad \text{edges.}$$

Thus, the graph G must have at least $\frac{k(k-1)}{2}$ edges. This makes Option (3) TRUE.

Step 3: Analyze remaining options.

- **Option (1):** It is not necessarily true that G contains a complete subgraph with k vertices, as the chromatic number k does not guarantee the presence of K_k. **Conclusion: FALSE.**

- **Option (4):** There may not be a vertex of degree k, as the chromatic number k only indicates the minimum number of colors needed, not the degree of any specific vertex. **Conclusion: FALSE.**

Final Answer:

$$\boxed{(2), (3)}$$

Quick Tip

When solving problems involving chromatic numbers, focus on properties such as independent sets and edge requirements rather than assumptions about specific subgraphs or vertex degrees.

---

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

(1) Operator $\diamondsuit$ obeys the associative law

(2) Operator $\square$ obeys the associative law

(3) Operator $\diamondsuit$ over the operator $\square$ obeys the distributive law

(4) Operator $\square$ over the operator $\diamondsuit$ obeys the distributive law

Correct Answer: (2), (4)

Solution:

Step 1: Check associativity of $\diamondsuit$.

To determine if $\diamondsuit$ is associative:

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

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

Since the two results differ, $\diamondsuit$ is not associative.

Step 2: Check associativity of $\square$.

To determine if $\square$ is associative:

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

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

As the results are the same, $\square$ is associative. Therefore, (2) is true.

Step 3: Check if $\square$ is distributive over $\diamondsuit$.

To verify distributivity:

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

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

Since these are equal, $\square$ is distributive over $\diamondsuit$. Hence, (3) is FALSE.

Step 4: Check if $\diamondsuit$ is distributive over $\square$.

To verify:

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

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

As the results are not equal, $\diamondsuit$ is not distributive over $\square$. Hence, (4) is TRUE.

Final Answer:

$$\boxed{(2), (4)}$$

Quick Tip
To confirm associativity and distributivity, expand each operation step by step and ensure the results match the required property.

---

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 cache (WBC) policy.

In a write-back cache: - Data modifications are stored in the cache, and the corresponding block is marked with a dirty bit. - The modified data is written to main memory only when the block is evicted.

- **Write Hit in WBC:** A write hit updates the cached data and sets the dirty bit. Thus, statement (3) is TRUE.
- **Read Miss in WBC:** A read miss may require evicting a dirty block if needed by the LRU (Least Recently Used) policy. This operation triggers a write-back of the dirty block to main memory, making statement (1) FALSE.

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

In a write-through cache: - All writes update both the cache and the main memory immediately, ensuring consistency.

- **Read Miss in WTC:** A read miss does not require a write-back operation, as main memory is always updated. Thus, statement (2) is TRUE.
- **Write Miss in WTC:** A write miss does not trigger a write-back operation for the victim block, as the cache and main memory remain consistent. Thus, statement (4) is FALSE.

Final Answer:

$$\boxed{(2), (3)}$$

Quick Tip
Write-back caches delay main memory updates until eviction, while write-through caches ensure immediate consistency by updating the main memory on every write.

---

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:

To determine the number of cylinders, use the relationship:

Hard Disk Capacity = (Number of Surfaces)×(Number of Cylinders)×(Sectors per Track)×(Sector Size

Step 1: Express the given parameters in powers of 2.

$$\text{Hard Disk Capacity} = 512\,\text{GB} = 2^9 \times 2^{30} = 2^{39}\,\text{Bytes},$$
$$\text{Number of Surfaces} = 32 = 2^5,$$
$$\text{Sectors per Track} = 4096 = 2^{12},$$
$$\text{Sector Size} = 1024 = 2^{10}\,\text{Bytes}.$$

Step 2: Write the capacity equation.

$2^{39} = (\text{Number of Surfaces}) \times (\text{Number of Cylinders}) \times (\text{Sectors per Track}) \times (\text{Sector Size}).$

Substitute the known values:

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

where x is the number of cylinders.

Step 3: Simplify the equation. Combine exponents on the right-hand side:

$$2^{39} = x \times 2^{5+12+10} = x \times 2^{27}.$$

Divide both sides by 2^{27}:

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

$$x = 2^{12} = 4096.$$

Final Answer:

$$\boxed{4096 \text{ cylinders}}$$

Quick Tip
When solving for the number of cylinders or similar parameters, express all values in powers of 2 for easier calculations and simplify step by step.

---

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 total execution time of the program is 100 ns, with:

$$\text{Parallelizable part} = 90\% \cdot 100 = 90\,\text{ns}, \quad \text{Non-parallelizable part} = 10\% \cdot 100 = 10\,\text{ns}.$$

Step 2: Compute the execution time with k cores.

Using Amdahl's Law and including the overhead for k cores:

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

Given the overhead per additional core is 10 ns, substitute the values:

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

Step 3: Minimize T_k for integer k.

1. For $k = 2$:

$$T_2 = \frac{90}{2} + 10 + 10 \cdot (2-1) = 45 + 10 + 10 = 65\,\text{ns}.$$

2. For $k = 3$:

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

3. For $k = 4$:

$$T_4 = \frac{90}{4} + 10 + 10 \cdot (4 - 1) = 22.5 + 10 + 30 = 62.5\,\text{ns}.$$

The minimum occurs at $k = 3$, where $T_k = 60\,\text{ns}$.

Final Answer:

$$\boxed{3}$$

Quick Tip
When optimizing execution time with multiple cores, include both the parallelizable fraction and overhead costs for accurate calculations.

---

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 including cache misses.

The total CPI is calculated as:

$$\text{CPI} = \text{Ideal CPI} + \text{Cache miss contribution}.$$

- **Instruction cache miss penalty:** The instruction miss rate is 2%, with a penalty of 100 cycles per miss. Thus:

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

- **Data cache miss penalty:** 25% of instructions are load/store, with a data miss rate of 8% and a penalty of 100 cycles. Thus:

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

Adding these contributions to the ideal CPI of 2, the total CPI is:

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

Step 2: Calculate the speedup with a perfect cache.

With a perfect cache (no misses), the CPI reduces to 2. Using the formula for speedup:

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

Final Answer:

$$\boxed{3.00}$$

Quick Tip
When calculating CPI and speedup, account for both instruction and data cache contributions to the total execution time.

---

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) that runs concurrently with the parent process. After each call to `fork()`, the total number of processes doubles, as both the parent and child continue execution from the same point.

Step 2: Analyze the code. The variable $x = 3$ determines the number of iterations of the `while` loop. During each iteration:

- The first `fork()` creates 1 new process, resulting in 2 total processes.
- The second `fork()` doubles the total number of processes to 4.
- The third `fork()` doubles the total again, resulting in 8 processes.

By the end of the loop, the total number of processes is $2^3 = 8$.

Step 3: Count the executions of `printf("hello")`. Each process executes `printf("hello")` exactly once per iteration of the loop. Summing across all processes:

First iteration: 2 processes, Second iteration: 4 processes, Third iteration: 8 processes.

The total executions of `printf("hello")` are:

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

Final Answer:

$$\boxed{14}$$

Quick Tip
When analyzing `fork()` calls, remember that each call doubles the number of processes, and every process executes the remaining code independently.

---

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 its most specific 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 $R2$. Each host receives 20 packets. The IP addresses forwarded via $R2$ are $10.1.1.132$ and $10.1.1.191$:

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

Final Answer:

$$\boxed{40}$$

Quick Tip
To match IP addresses with routing prefixes, apply the longest prefix match rule, where the most specific (longest) subnet is selected for routing.

---

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 $S \to^* w$ is _____.

Correct Answer: 179

Solution:

Step 1: Chomsky Normal Form (CNF) basics.

In CNF, each production rule is either of the form $A \to BC$ or $A \to a$, where B, C are non-terminals, and a is a terminal. To derive a string of length n:

- n steps are required to generate n terminal symbols.
- $n - 1$ steps are required to concatenate these symbols into the final string.

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

The string w contains 90 terminal symbols: $30a, 30b, 30c$. To generate these terminals:

$$\text{Steps to derive terminals} = 90.$$

To combine these terminals using $n - 1$ concatenation steps:

$$\text{Steps to combine terminals} = 90 - 1 = 89.$$

Step 3: Total derivation steps.

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

Final Answer:

$$\boxed{179}$$

Quick Tip
In CNF, the total steps to derive a string is the sum of the steps required to generate terminals and the steps needed for concatenation.

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: Relationship between connected components and edges in a DFS forest.

In a connected component of an undirected graph, the number of edges is one less than the number of vertices. If the graph has C connected components, and the vertex counts of these components are $n_1, n_2, \ldots, n_C$, then:

$$n_1 + n_2 + \cdots + n_C = 100.$$

The total number of edges in the graph is:

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

Step 2: Calculate the number of connected components.

The graph has 40 edges. Using the formula:

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

Final Answer:

$$\boxed{60}$$

Quick Tip
In undirected graphs, the number of connected components can be computed as the difference between the total number of vertices and the total number of 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 Regular Expressions:

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

$s = 0\,1^* + 1\,0^*$ (strings starting with 0 followed by ones, or starting with 1 followed by zeros).

Alphabet:

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

Step 1: Count the number of strings in r and s. For all strings of length n:

- 2 strings belong to r (all zeros or all ones).

- 2 strings belong to s (strings starting with 0 followed by ones, or starting with 1 followed by zeros).

Step 2: Calculate strings not in r and s for different lengths.

$$\text{\# strings not in } r \text{ and } s = 2^n - 4, \quad \text{for each } n.$$

- For $n = 2$:

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

- For $n = 3$:

$$\text{\# strings not in } r \text{ and } s = 2^3 - 4 = 8 - 4 = 4.$$

- For $n = 4$:

$$\text{\# strings not in } r \text{ and } s = 2^4 - 4 = 16 - 4 = 12.$$

- For $n = 5$:

$$\text{\# strings not in } r \text{ and } s = 2^5 - 4 = 32 - 4 = 28.$$

Step 3: Compute total strings for lengths less than or equal to 5.

$$\text{Total strings not in } r \text{ and } s = (2^2 - 4) + (2^3 - 4) + (2^4 - 4) + (2^5 - 4).$$

$$= (4 - 4) + (8 - 4) + (16 - 4) + (32 - 4).$$

$$= 0 + 4 + 12 + 28 = 44.$$

Final Answer:

$$\boxed{44}$$

Quick Tip
To find the number of strings not covered by given regular expressions, calculate the total possible strings and subtract those included in the regular expressions.

---

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

Solution:

Step 1: Calculate the virtual page number and offset.

The page size is $2\,\text{KB} = 2048\,\text{bytes}$. For a virtual address 2500, we compute:

$$\text{Page number} = \left\lfloor \frac{2500}{2048} \right\rfloor = 1, \quad \text{Offset} = 2500 \mod 2048 = 452.$$

Step 2: Map the virtual page number to the physical frame.

Using the given mapping, virtual page 1 is mapped to physical frame 3.

Step 3: Compute the physical address.

The physical address is calculated as:

$$\text{Physical address} = (\text{Frame number} \times \text{Page size}) + \text{Offset}.$$

Substituting the values:

$$\text{Physical address} = (3 \times 2048) + 452 = 6144 + 452 = 6596.$$

Final Answer:

$$\boxed{6596}$$

Quick Tip
To determine the physical address, compute the page number and offset, map the page number to its physical frame, and add the offset to the frame's base address.

---

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:

This problem involves calculating the conditional probability $P(R_2 \mid R_1)$, where R_1 and R_2 denote the events of drawing red balls in the first and second draws, respectively.

Step 1: Total number of balls and probabilities.

The bag contains:

$$10 \text{ Red balls and } 15 \text{ Blue balls (total: } 25 \text{ balls).}$$

The probability of drawing a red ball on the first draw (R_1) is:

$$P(R_1) = \frac{10}{25}.$$

Step 2: Conditional probability for $R_2 \mid R_1$.

If R_1 occurs (a red ball is drawn), the bag now contains 9 red balls and 15 blue balls (total 24).

The probability of drawing a red ball on the second draw (R_2) given R_1 is:

$$P(R_2 \mid R_1) = \frac{\text{Number of remaining red balls}}{\text{Total remaining balls}} = \frac{9}{24}.$$

Step 3: Simplify the result.

$$P(R_2 \mid R_1) = \frac{9}{24} = 0.375.$$

Final Answer:

$$\boxed{0.375}$$

Quick Tip
When calculating conditional probabilities, adjust both the total outcomes and the favorable outcomes based on the given condition.

---

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.

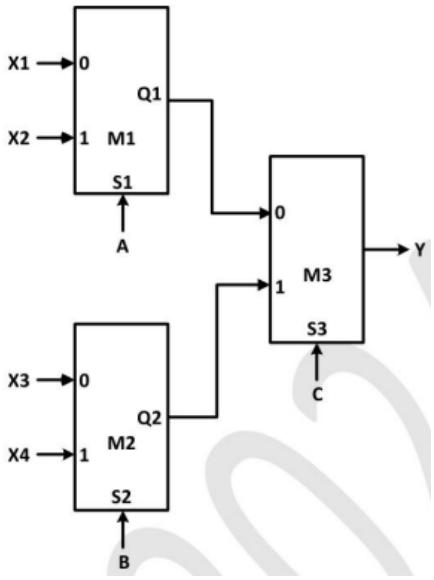

For an instance of inputs $X_1 = 1$, $X_2 = 1$, $X_3 = 0$, and $X_4 = 0$, the number of combinations of A, B, C that give the output $Y = 1$ is _____.

Correct Answer: 4

Solution:

Step 1: Analyze the behavior of M_1, M_2, and M_3.

- **Multiplexer M_1:** Controlled by select line A:

$$A = 0 \implies M_1 \text{ selects } X_1 = 1.$$

$$A = 1 \implies M_1 \text{ selects } X_2 = 1.$$

  In both cases, the output of M_1 is 1.

- **Multiplexer M_2:** Controlled by select line B:

$$B = 0 \implies M_2 \text{ selects } X_3 = 0.$$

$$B = 1 \implies M_2 \text{ selects } X_4 = 0.$$

  In both cases, the output of M_2 is 0.

- **Multiplexer M_3:** Controlled by select line C:

$$C = 0 \implies M_3 \text{ selects the output of } M_1 = 1.$$

$$C = 1 \implies M_3 \text{ selects the output of } M_2 = 0.$$

  To achieve $Y = 1$, $C = 0$ is required.

Step 2: Calculate valid combinations of A, B, C.

- A can be 0 or 1 (2 possibilities).
- B can be 0 or 1 (2 possibilities).
- C must be 0 to ensure $Y = 1$ (1 possibility).

The total number of valid combinations is:

$$2 \times 2 \times 1 = 4.$$

Final Answer:

$$\boxed{4}$$

Quick Tip
When analyzing multiplexers, determine the outputs for all select lines, then calculate combinations satisfying the required output conditions.

---

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:

Given Data:

Total datagram size: 1420 bytes.

IP header size: 20 bytes.

Payload size:

$$\text{Payload} = 1420 - 20 = 1400 \text{ bytes}.$$

Step 1: Fragmentation on the first link (MTU = 542 bytes).

The maximum payload size per fragment is:

$$\text{Maximum payload size} = 542 - 20 = 522 \,\text{bytes}.$$

The number of fragments required is:

$$\lceil \frac{1400}{522} \rceil = 3.$$

Fragment details after the first link:

- Fragment 1: Payload = 522, Total = $522 + 20 = 542$ bytes.
- Fragment 2: Payload = 522, Total = 542 bytes.
- Fragment 3: Payload = 356, Total = $356 + 20 = 376$ bytes.

Step 2: Fragmentation on the second link (MTU = 360 bytes).

The maximum payload size per fragment on the second link is:

$$\text{Maximum payload size} = 360 - 20 = 340 \,\text{bytes}.$$

Each fragment from the first link is further fragmented:

- Fragment 1 (522 bytes):
$$\lceil \frac{522}{340} \rceil = 2 \,\text{fragments}.$$
- Fragment 2 (522 bytes):
$$\lceil \frac{522}{340} \rceil = 2 \,\text{fragments}.$$
- Fragment 3 (356 bytes):
$$\lceil \frac{356}{340} \rceil = 2 \,\text{fragments}.$$

The total number of fragments delivered to the receiver is:

$$2 + 2 + 2 = 6.$$

Final Answer:

$$\boxed{6}$$

Quick Tip

In IP fragmentation, always calculate the maximum payload size based on the MTU and ensure the payload size in each fragment is a multiple of 8 for compatibility with fragmentation offsets.
