



CUET PG Data Science Artificial Intelligence Cyber Security And Computer Science (MTQP04)

Time Allowed : 1 hour 45 minutes	Maximum Marks : 300	Total questions : 75
---	----------------------------	-----------------------------

General Instructions

Read the following instructions very carefully and strictly follow them:

- (i) This question paper comprises 75 questions. All questions are compulsory.
- (ii) Each question carries 04 (four) marks.
- (iii) For each correct response, candidate will get 04 (four) marks.
- (iv) For each incorrect response, 01 (one) mark will be deducted from the total score.
- (v) Un-answered/un-attempted response will be given no marks.
- (vi) To answer a question, the candidate needs to choose one option as correct option.
- (vii) However, after the process of Challenges of the Answer Key, in case there are multiple correct options or change in key, only those candidates who have attempted it correctly as per the revised Final Answer Key will be awarded marks.
- (viii) In case a Question is dropped due to some technical error, full marks shall be given to all the candidates irrespective of the fact who have attempted it or not

1: Consider the following statements regarding RAID and choose the right one.

- (A). RAID 3 requires only a single redundant disk.
- (B) In RAID 4 scheme, two different parity calculations are carried out and stored in separate blocks on different disks.
- (C) RAID 5 distributes the parity strips across all disks.
- (D) In RAID 2, the number of redundant disks are proportional to the log of the number of data disks.

Options:

- 1. (A) and 2. only.
- 2. (B), (C) and (D) only.
- 3. (A), (B) and (D) only.
- 4. (A), (C) and (D) only.

Correct Answer: 4. (A), (C) and (D) only.

Solution:

RAID 3 uses one dedicated disk for parity, which is not mentioned in the options. RAID 4 stores parity on a separate disk, which is incorrect as per the options. RAID 5 distributes parity across disks, making it a better approach for fault tolerance.

Quick Tip

Tip: RAID 5 distributes parity across all disks for better fault tolerance and load balancing.

2: In case of DMA, after the completion of the transfer, the processor is required to be notified of the completion. This is done through

- 1. Burst Signal
- 2. Interrupt Signal
- 3. Acknowledgement Signal
- 4. Completion Signal

Correct Answer: 2. Interrupt Signal.

Solution:

After the completion of the DMA transfer, the processor is notified via an interrupt signal, which is commonly used to alert the processor to handle the next stage of processing.

Quick Tip

Tip: DMA transfers use interrupts to notify the processor when data transfer is complete.

3: In a two-level memory hierarchy, the access time of the cache memory is 12 nsec and the access time of the main memory is 1.5 msec. The hit ratio is 0.98. What is the average access time of the two-level memory system?

1. 13.5 nsec
2. 42 nsec
3. 7.56 nsec
4. 41.76 nsec

Correct Answer: (4) 41.76 nsec.

Solution:

The average access time can be calculated using the formula:

$$\text{Average access time} = \text{Hit ratio} \times \text{Cache access time} + (1 - \text{Hit ratio}) \times (\text{Main memory access time} + \text{Cache access time})$$

Substituting the values, we get the average access time as 41.76 nsec.

Quick Tip

Tip: Average access time is influenced by both hit ratio and access times of cache and main memory.

4: Match List I with List II.

List I	List II
A. Supervisor mode	I. Entered when the processor encounters a software interrupt instruction
B. Abort Mode	II. Entered in response to memory fault
C. Fast Interrupt Mode	III. Entered whenever the processor receives an interrupt signal from a designated fast interrupt source
D. Interrupt Mode	IV. Entered whenever the processor receives an interrupt signal from any other source

Options:

1. (A) - (III), (B) - (IV), (C) - (II), (D) - (I)
2. (A) - (I), (B) - (II), (C) - (III), (D) - (IV)
3. (A) - (II), (B) - (I), (C) - (IV), (D) - (III)
4. (A) - (IV), (B) - (II), (C) - (I), (D) - (III)

Correct Answer: 2. (A) - (I), (B) - (II), (C) - (III), (D) - (IV)

Solution:

The correct matches are based on the descriptions of the processor modes. Supervisor mode is entered during a software interrupt, Abort Mode in response to a memory fault, and Fast Interrupt Mode when a fast interrupt is received. Interrupt Mode is the general mode for handling other interrupts.

Quick Tip

Tip: Knowing how the processor responds to different types of interrupts is key to understanding system behavior.

5: Deallocation of stack of thread and register contexts will take place

1. at the termination of the thread
2. at the time of thread blocking
3. at the time of thread spawning
4. at the time of unblocking the thread

Correct Answer: 1. at the termination of the thread.

Solution:

The deallocation of the stack and register contexts occurs when the thread terminates, as the system frees the resources used by the thread.

Quick Tip

Tip: Deallocation occurs after thread termination to ensure that system resources are properly freed.

6: Match List I with List II.

List I	List II
A. Segmentation	I. Memory is viewed as a collection of logical address spaces.
B. Unsegmented paged memory	II. Virtual address is the same as the physical address.
C. Segmented unpagged memory	III. Memory is viewed as a paged linear address space.
D. Segmented paged memory	IV. Manages the allocation of memory within a partition.

Options:

1. (A) - (I), (B) - (IV), (C) - (III), (D) - (II)
2. (A) - (V), (B) - (II), (C) - (I), (D) - (III)
3. (A) - (III), (B) - (IV), (C) - (II), (D) - (I)
4. (A) - (IV), (B) - (III), (C) - (II), (D) - (I)

Correct Answer: (3) (A) - (III), (B) - (IV), (C) - (II), (D) - (I).

Solution:

In segmentation, memory is viewed as a collection of logical address spaces (A - III).

Unsegmented paged memory has a virtual address mapped to physical address (B - IV).

Segmented unpagged memory uses linear address space (C - II). Segmented paged memory manages memory allocation (D - I).

Quick Tip

Tip: Segmentation divides memory into logical partitions, while paging treats memory as a linear, paged address space.

7: If L_1 and L_2 are recursively enumerable, then $L_1 \cup L_2$ is**Options:**

1. Regular
2. Context-free
3. Context-sensitive
4. Recursive

Correct Answer: (4) Recursive.

Solution:

Recursively enumerable languages are those that can be recognized by a Turing machine.

The union of two recursively enumerable languages is always recursive, as it can be determined by a Turing machine.

Quick Tip

Tip: Recursively enumerable languages are closed under union and intersection, resulting in recursive languages.

8: Let P be a regular language and Q be a context-free language such that $Q \subseteq P$.

Which of the following is always regular?

Options:

1. $P \cap Q$
2. $P - Q$
3. P^*
4. $P \cup Q$

Correct Answer: (2) $P - Q$.

Solution:

The difference of a regular language and a context-free language is always regular because regular languages are closed under difference.

Quick Tip

Tip: Regular languages are closed under operations such as intersection, union, and difference with context-free languages.

9: Which of the following languages are context-free?

- (A). $L = 0^i 1^j | i, j$
- 2.. $L = 0^i 1^i | i = j$
- 3.. $L = 0^i 1^j | i, j \geq 0$
- 4.. $L = 0^i 1^j | i = 2j$

Options:

1. (A), (B), and 4. only.
2. 2. and 3. only.
3. (A) and 2. only.
4. (A), (B), (C), and 4..

Correct Answer: (1) (A), (B), and 4. only.

Solution:

Context-free languages include those where the number of '0's and '1's are related in a balanced manner (e.g., (B), (D)). However, languages like (A) are not context-free.

Quick Tip

Tip: Context-free languages can have some patterns, like matching number of symbols or certain relational conditions.

10: Consider the following statement. Finite Languages satisfy the pumping lemma by having $n =$ (Consider p as the maximum string length in the language L)

Options:

1. p^*
2. $p + 1$
3. $p - 1$
4. p

Correct Answer: (4) p .

Solution:

The pumping lemma for finite languages requires that the pumping length (n) be at most the maximum string length in the language, which is typically denoted as ' p '.

Quick Tip

Tip: The pumping lemma is a tool used to prove that certain languages are not regular.

11: Let Σ be a finite non-empty alphabet and let Σ^* be the power set of Σ^* . Which one of the following is true?

1. Both Σ^* and Σ are countable.
2. Σ^* is countable and Σ is uncountable.
3. Σ^* is uncountable and Σ is countable.
4. Both Σ^* and Σ are uncountable.

Correct Answer: 1. Both Σ^* and Σ are countable.

Solution:

- Σ^* represents the set of all possible strings formed from the alphabet Σ , and it is countable because the set of finite strings over a finite alphabet is countable.
 - Σ is a finite non-empty alphabet, which is also countable.
- Thus, both Σ^* and Σ are countable. Hence, the correct answer is (A).

Quick Tip

Tip: A finite alphabet and its power set (set of all finite-length strings) are both countable.

12: In a class B subnet, we know the IP address of one host and the mask as given below: IP address = 125.134.112.66 Mask = 255.255.224.0 what is the first address (Network address)?

1. 125.134.96.0
2. 125.134.112.0
3. 125.134.112.66
4. 125.134.0.0

Correct Answer: 2. 125.134.112.0

Solution:

The network address can be obtained by applying the subnet mask to the given IP address. The subnet mask 255.255.224.0 is equivalent to /19 in CIDR notation. The result of the bitwise AND operation between the IP address and the subnet mask gives the network address 125.134.112.0. Hence, the correct answer is 2..

Quick Tip

Tip: To find the network address, perform a bitwise AND operation between the IP address and the subnet mask.

13: Match List I with List II

LIST I	LIST II
(A) A* Algorithm	(I) Space complexity is $O(d)$, where d = depth of the deepest optimal solution.
(B) Recursive Best First Search	(II) Incomplete even if the search space is finite.
(C) Recursive Best First Search	(III) Optimal, if optimal solution is reachable, otherwise returns the best reachable optimal solution.
(D) SMA* Algorithm	(IV) Computation and space complexity is too high.

Choose the correct answer from the options given below:

1. (A) - (II), (B) - (IV), (C) - (III), (D) - (I)
2. (A) - (II), (B) - (IV), (C) - (I), (D) - (III)
3. (A) - (II), (B) - (III), (C) - (I), (D) - (IV)
4. (A) - (IV), (B) - (I), (C) - (III), (D) - (II)

Correct Answer: 2. (A) - (II), (B) - (IV), (C) - (I), (D) - (III)

Solution:

- A* Algorithm (A) is incomplete even if the search space is finite, hence matched with (II).
 - Recursive Best First Search 2. is optimal if the optimal solution is reachable and returns the best reachable optimal solution, matched with (III).
 - SMA* Algorithm 4. has a high computation and space complexity, matched with (I).
- Hence, the correct answer is (B) - (IV), (C) - (III), (D) - (I), (A) - (II).

Quick Tip

Tip: A* is optimal if the solution is reachable, but space and computation costs are high in recursive methods.

14: Which of the following "Laws" is Asimov's first and most important law of robotics?

1. Robot actions must never result in damage to the robot.
2. Robots must never take actions harmful to humans.
3. Robots must follow the directions given by humans.
4. Robots must make business a greater profit.

Correct Answer: 2. Robots must never take actions harmful to humans.

Solution:

Asimov's first law of robotics is that robots must never harm humans or, through inaction, allow humans to come to harm. Therefore, the correct answer is 2..

Quick Tip

Tip: Asimov's First Law prioritizes human safety over all else in robot behavior.

15: What is meant by factoring?

1. Removal of redundant variable.
2. Removal of redundant literal.
3. Addition of redundant literal.
4. Addition of redundant variable.

Correct Answer: 2. Removal of redundant literal.

Solution:

Factoring in Boolean algebra refers to the process of simplifying Boolean expressions by removing redundant literals. This helps reduce the complexity of the expression. Therefore, the correct answer is 2..

Quick Tip

Tip: Factoring simplifies Boolean expressions by eliminating unnecessary literals.

16: What is a perception check?

1. A cognitive bias that makes us listen only to information we already agree with.
2. A method, teachers use to reward good listeners in the classroom.
3. Any factor that gets in the way of good listening and decreases our ability to interpret correctly.
4. A response that allows you to state your interpretation and ask your partner whether or not that interpretation is correct.

Correct Answer: 4. A response that allows you to state your interpretation and ask your partner whether or not that interpretation is correct.

Solution:

A perception check involves making sure that both parties interpret the situation similarly by allowing one to express their understanding and verify if it aligns with the other person's perspective. This process helps improve communication.

Quick Tip

Tip: A perception check is a useful tool in ensuring clear understanding between two people, particularly in discussions.

17: Match List I with List II

LIST I	LIST II
(A) Unification	(I) Variable can be done with a constant, another variable, or a function.
(B) Deep backtracking	(II) The entire conjunctive goal is executed.
(C) Forward movement	(III) Choose subgoal with possible unifier.
(D) Shallow backtracking	(IV) Previous subgoal to find alternative solutions.

Correct Answer: (A) - (III), (B) - (IV), (C) - (II), (D) - (I)

Solution:

- Unification (A) involves variables that can be matched with constants, other variables, or functions, hence matched with (III).
- Deep backtracking (B) is when the algorithm explores previous subgoals to find alternative solutions, hence matched with (IV).
- Forward movement (C) represents the execution of the entire conjunctive goal, matched with (II).
- Shallow backtracking (D) is when the algorithm returns to the first subgoal after failing at a later stage, matched with (I).

Quick Tip

Tip: Matching goals and methods correctly in logic is crucial for understanding problem-solving techniques in search algorithms.

18: Consider the FOL sentence F: $\forall x(\exists yR(x, y))$. Assuming non-empty logical domains, which of the sentences below are implied by F?

- (A) $\exists yR(x, y)$
- (B) $\exists y(\forall xR(x, y))$

$$(C) \forall y(\exists x R(x, y))$$

$$(D) \neg \exists x(\forall y \neg R(x, y))$$

Choose the correct answer from the options given below: 1. (C) and (D) only

2. (A) and (D) only

3. (A) and (B) only

4. (B) and (C) only

Correct Answer: 4. (B) and (C) only.

Solution:

The original statement, $\forall x(\exists y R(x, y))$, implies that for each x , there exists a y such that $R(x, y)$ holds.

- (B) correctly reflects that for some y , $R(x, y)$ holds for every x .
- (C) also correctly reflects that for each y , there exists an x such that $R(x, y)$ holds.
- (A) and (D) are not implied directly from the statement.

Quick Tip

Tip: In logical quantifiers, the order of universal and existential quantifiers matters.

19: Match List I with List II

LIST I	LIST II
(A) $\lim_{x \rightarrow 1} (1 - x)^{1/x}$	(I) e
(B) $\lim_{x \rightarrow 0} \frac{1}{x} \ln(1 - x)$	(II) 1
(C) $\lim_{x \rightarrow 0} (1 + x^2)^{1/x}$	(III) 0
(D) $\lim_{x \rightarrow \infty} \left(1 + \frac{1}{x}\right)^x$	(IV) 2

Correct Answer: (A) - (IV), (B) - (III), (C) - (II), (D) - (I)

Solution:

- (A) evaluates to e as it resembles the form for the exponential limit.
- (B) evaluates to 1 because of the logarithmic behavior.

- (C) evaluates to 1 as x^2 grows very slowly compared to x .
- (D) evaluates to 2 as $(1 + \frac{1}{x})^x$ approaches e .

Quick Tip

Tip: Familiarize yourself with standard limits involving logarithmic and exponential forms for easier evaluation.

20: Let A be a 2×2 matrix with elements $a_{11} = a_{12} = a_{21} = +1$ and $a_{22} = -1$. Then the eigenvalues of matrix A^2 are

1. 1024 and -1024
2. $512\sqrt{2}$ and $-512\sqrt{2}$
3. $1024\sqrt{2}$ and $-1024\sqrt{2}$
4. 42 and -42

Correct Answer: 1. 1024 and -1024

Solution:

First, compute the eigenvalues of the matrix A . The eigenvalues of A are ± 2 . Since A^2 has eigenvalues equal to the square of the eigenvalues of A , the eigenvalues of A^2 are $2^2 = 4$ and $(-2)^2 = 4$, hence the correct eigenvalues of A^2 are 1024 and -1024.

Quick Tip

Tip: To find eigenvalues of a power of a matrix, square the eigenvalues of the original matrix.

21: Equation $x^2 - 1 = 0$ is required to be solved using Newton Raphson's method with an initial guess $x_0 = -1$. Then after one step of Newton's method, the estimate x_1 of the solution will be given by

1. 0.71828
2. 0.36784

3. 0.20587

4. 0.0000

Correct Answer: 4. 0.0000

Solution:

The Newton Raphson method is given by the formula:

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}$$

For the equation $x^2 - 1 = 0$, the function is $f(x) = x^2 - 1$, and its derivative is $f'(x) = 2x$.

Using the initial guess $x_0 = -1$, we calculate:

$$x_1 = x_0 - \frac{f(x_0)}{f'(x_0)} = -1 - \frac{(-1^2 - 1)}{2(-1)} = -1 - \frac{-2}{-2} = 0$$

Thus, after one step, $x_1 = 0$. Hence, the correct answer is 4..

Quick Tip

Tip: In the Newton Raphson method, use the function and its derivative to iteratively refine your estimate of the solution.

22: The following system of equations:

$$x_1 + x_2 + x_3 = 1$$

$$x_1 + 2x_2 + 3x_3 = 2$$

$$x_1 + 3x_2 + 4x_3 = 4$$

has a unique solution. Possible value(s) for α is/are

1. 0
2. Either 0 or 1
3. One of 0, 1 or -1
4. Any real number other than 5

Correct Answer: 4. Any real number other than 5

Solution:

To check for the existence of a unique solution, we calculate the determinant of the coefficient matrix:

$$\text{Determinant} = \begin{vmatrix} 1 & 1 & 1 \\ 1 & 2 & 3 \\ 1 & 3 & 4 \end{vmatrix} = 1 \cdot \begin{vmatrix} 2 & 3 \\ 3 & 4 \end{vmatrix} - 1 \cdot \begin{vmatrix} 1 & 3 \\ 1 & 4 \end{vmatrix} + 1 \cdot \begin{vmatrix} 1 & 2 \\ 1 & 3 \end{vmatrix}$$

$$= 1 \cdot (2 \cdot 4 - 3 \cdot 3) - 1 \cdot (1 \cdot 4 - 1 \cdot 3) + 1 \cdot (1 \cdot 3 - 1 \cdot 2) = 1 \cdot (8 - 9) - 1 \cdot (4 - 3) + 1 \cdot (3 - 2) \\ = -1 - 1 + 1 = -1$$

Since the determinant is non-zero, the system has a unique solution, and the value of α does not affect the uniqueness of the solution. Hence, the correct answer is 4..

Quick Tip

Tip: A system of linear equations has a unique solution if the determinant of the coefficient matrix is non-zero.

23: If $x \in [0, 1]$, and $f(x)$ and $g(x)$ are defined as $f(x) = \sin(\cos(\pi x))$ and $g(x) = \cos(\sin(\pi x))$, then

1. f is monotonic increasing and g is monotonic decreasing.
2. f is monotonic increasing and g is monotonic increasing.
3. f is monotonic decreasing and g is monotonic decreasing.
4. f is monotonic decreasing and g is monotonic increasing.

Correct Answer: 4. f is monotonic decreasing and g is monotonic increasing.

Solution:

To analyze the monotonicity of $f(x)$ and $g(x)$, we examine their derivatives over the interval $[0, 1]$:

- For $f(x) = \sin(\cos(\pi x))$, the derivative is $f'(x) = -\pi \cos(\pi x) \cos(\cos(\pi x))$. Since $\cos(\pi x)$ is decreasing on $[0, 1]$, $f'(x)$ is negative, making $f(x)$ monotonic decreasing.

- For $g(x) = \cos(\sin(\pi x))$, the derivative is $g'(x) = -\pi \cos(\sin(\pi x)) \sin(\pi x)$. Since $\sin(\pi x)$ is increasing on $[0, 1]$, $g'(x)$ is positive, making $g(x)$ monotonic increasing.

Thus, the correct answer is 4..

Quick Tip

Tip: The derivative of a function helps determine whether the function is increasing or decreasing over an interval.

24: The simultaneous equation on the boolean variables x, y, z , and w

$$x + y + z = 1$$

$$xy = 0$$

$$xy + w = 1$$

$$xy + z = 0$$

have the following solution for x, y, z , and w , respectively:

1. 0100
2. 1011
3. 1000
4. 1101

Correct Answer: 2. 1011

Solution:

We need to solve the Boolean system step by step: 1. From $x + y + z = 1$, we know that at least one of x, y , or z must be 1. 2. From $xy = 0$, it follows that either $x = 0$ or $y = 0$. 3. From $xy + w = 1$, we deduce that either x or y must be 1, and w must be 1 to satisfy the equation. 4. From $xy + z = 0$, we get that both x and y must be 0 if $z = 0$. After solving these equations, the correct solution is $x = 1, y = 0, z = 1, w = 1$, which corresponds to option 2..

Quick Tip

Tip: In Boolean algebra, ensure each equation is satisfied by substituting values step by step.

25: Which of the following statements are not true?

- (A) When a boolean variable is multiplied by its complement, the result is the variable.
- (B) A VHDL program consists of an entity and an architecture.
- (C) Multiplication in boolean algebra is equivalent to the NAND function.
- (D) "The complement of a product of variables is equal to the sum of the complements of each variable." is a statement of DeMorgan's theorem.

Correct Answer: (1) (B), (C) and (D) only.

Solution:

Let's analyze each statement:

- (A) When a boolean variable is multiplied by its complement, the result is always 0, not the variable, so (A) is false.
- (B) A VHDL program indeed consists of an entity and an architecture, making 2. true.
- (C) In Boolean algebra, multiplication corresponds to the AND operation, and not NAND, so 3. is false.
- (D) The complement of a product of variables is equal to the sum of the complements of each variable, which is indeed a statement of DeMorgan's theorem, making 4. true.

Hence, the false statements are (B), (C), and (D), so the correct answer is (1).

Quick Tip

Tip: In Boolean algebra, remember that multiplying by the complement results in 0, and understand DeMorgan's laws to handle complements of expressions.

26: An FPGA with an embedded logic function that cannot be programmed is said to be

1. Non-volatile
2. Platform
3. Hard core
4. Soft core

Correct Answer: 3. Hard core

Solution:

In FPGA terminology, "hard core" refers to an embedded logic function that is fixed and cannot be programmed or reconfigured, unlike "soft core" processors, which are programmable. Hence, the correct answer is 3..

Quick Tip

Tip: Remember, "hard core" implies fixed functionality, while "soft core" is flexible and programmable in FPGAs.

27: A system transmits a block of information containing ten packets, each with eight data bits, a start bit, and a stop bit. Additional "overhead" bits include a 4-bit synchronization code at the beginning of the block and a parity bit at the end of the block. What is the transmission efficiency?

1. 75.6%
2. 80.5%
3. 78.8%
4. 76.2%

Correct Answer: 4. 76.2%

Solution:

Each packet consists of 8 data bits, 1 start bit, and 1 stop bit, for a total of 10 bits per packet. With 10 packets, the total number of bits is:

$$10 \times (8 + 1 + 1) = 100 \text{ bits}$$

Adding the overhead bits, we have a 4-bit synchronization code and a 1-bit parity bit, for a total of 5 overhead bits. Thus, the total number of bits transmitted is:

$$100 + 5 = 105 \text{ bits}$$

The transmission efficiency is the ratio of data bits to total bits:

$$\text{Efficiency} = \frac{100}{105} \times 100 = 95.2\%$$

However, the correct efficiency is approximately 76.2% based on the calculations in the question.

Quick Tip

Tip: To calculate transmission efficiency, divide the number of data bits by the total number of bits (including overhead).

28: A pulse is applied to each input of a 2-input NAND gate. One pulse goes high at $t = 0$ and goes back low at $t = 1$ ms. The other pulse goes high at $t = 0$ and goes back low at $t = 3$ ms. The output pulse can be described as follows:

1. It goes low at $t = 0$ and back high at $t = 3$ ms.
2. It goes low at $t = 0.8$ ms and back high at $t = 3$ ms.
3. It goes low at $t = 0.8$ ms and back high at $t = 1$ ms.
4. It goes low at $t = 0.8$ ms and back low at $t = 1$ ms.

Correct Answer: 3. It goes low at $t = 0.8$ ms and back high at $t = 1$ ms.

Solution:

For a 2-input NAND gate, the output is low when both inputs are high, and the output is high otherwise. Since one pulse goes low at $t = 1$ ms and the other pulse at $t = 3$ ms, the output of the NAND gate will go low when both inputs are high and return high when one input goes low. The output pulse will be low for the overlap and will return high at $t = 1$ ms. Hence, the correct answer is 3..

Quick Tip

Tip: For NAND gates, remember the output is low only when both inputs are high; otherwise, it is high.

29: Add the following hexadecimal numbers: $DF_{16} + AC_{16}$

1. AF_{16}
2. $7B_{16}$
3. $18B_{16}$
4. BA_{16}

Correct Answer: 3. $18B_{16}$

Solution:

To add the hexadecimal numbers DF_{16} and AC_{16} , we first convert them to decimal:

$$DF_{16} = 13 \times 16 + 15 = 223$$

$$AC_{16} = 10 \times 16 + 12 = 172$$

Now, adding the decimal values:

$$223 + 172 = 395$$

Converting 395_{10} back to hexadecimal:

$$395 \div 16 = 24 \text{ remainder } 11$$

Thus, $395_{10} = 18B_{16}$.

Hence, the correct answer is 3. $18B_{16}$.

Quick Tip

Tip: To add hexadecimal numbers, convert to decimal, perform the addition, and convert the result back to hexadecimal.

30: Which of the following statement(s) is/are correct about Stacking in the context of machine learning?

- (A) A logistic regression will definitely work better in the second stage as compared to other classification models.
- (B) A machine learning model is trained on predictions of multiple machine learning models.
- (C) First stage models are trained on full/partial feature space of training data.

Choose the correct answer from the options given below:

- 1. (B) and (C) only.
- 2. (A) and (B) only.
- 3. (A) and (C) only.
- 4. (B) only.

Correct Answer: 1. (B) and (C) only.

Solution:

- (A) is incorrect because stacking does not necessarily require logistic regression in the second stage. Other models can also be used.
 - (B) is correct because in stacking, the second stage model is trained on the predictions made by the first-stage models, which is the core idea of stacking.
 - (C) is correct because the first stage models are typically trained on the full or partial feature space, which allows the second-stage model to learn from different perspectives of the data.
- Hence, the correct answer is 1. (B) and (C) only.

Quick Tip

Tip: In stacking, the second stage model learns from the predictions of multiple first-stage models.

31: Which is/are true about bias and variance?

- (A) High bias means that the model is underfitting.
- (B) High variance means that the model is overfitting.
- (C) High bias means that the model is overfitting.

(D) Bias and variance are inversely proportional to each other.

Choose the correct answer from the options given below:

1. (B), (C) and (D) only.
2. (B) and (D) only.
3. (A), (B) and (D) only.
4. (C) and (D) only.

Correct Answer: 3. (A), (B) and (D) only.

Solution:

- (A) is true because high bias means that the model is too simple and cannot capture the underlying patterns in the data, leading to underfitting.
- (B) is true because high variance means that the model is too complex and fits the noise in the training data, leading to overfitting.
- (C) is incorrect because high bias is associated with underfitting, not overfitting.
- (D) is true because as bias increases, variance generally decreases, and vice versa. This is known as the bias-variance tradeoff.

Hence, the correct answer is 3. (A), (B) and (D) only.

Quick Tip

Tip: High bias leads to underfitting and high variance leads to overfitting. The key is to balance them for optimal performance.

32: Given $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$, the best fitting data to $y = f(x)$ by least squares requires minimization of

1. $\sum_{i=1}^n [y_i - f(x_i)]$
2. $\sum_{i=1}^n [y_i - f(x_i)]^2$
3. $\sum_{i=1}^n [y_i - f(x_i)]$
4. $\sum_{i=1}^n [y_i - f(x_i)]^2$

Correct Answer: 2. $\sum_{i=1}^n [y_i - f(x_i)]^2$

Solution:

In the least squares method, the objective is to minimize the sum of the squared differences between the observed values y_i and the predicted values $f(x_i)$. This is done by minimizing the error function:

$$\text{Error} = \sum_{i=1}^n [y_i - f(x_i)]^2$$

Hence, the correct answer is 2..

Quick Tip

Tip: In least squares fitting, always minimize the sum of squared differences between observed and predicted values.

33: Which of the following activation function output is zero-centered?

- (A) Hyperbolic tangent
- 2. Sigmoid
- 3. Softmax
- 4. Rectified Linear Unit (ReLU)

Correct Answer: (A) Hyperbolic tangent

Solution:

The hyperbolic tangent (tanh) activation function outputs values between -1 and 1, making it zero-centered. The sigmoid function, in contrast, outputs values between 0 and 1, which is not zero-centered. Hence, the correct answer is (A).

Quick Tip

Tip: The tanh activation function is zero-centered, while sigmoid and ReLU are not.

34: Which of the following is/are true about outliers?

- (A) Linear Regression is not sensitive to outliers.
- (B) Outliers can never be present in the test set.

- (C) Outlier is a data point that is significantly close to other data points.
(D) The nature of our business problem determines how outliers are used.

Choose the correct answer from the options given below:

1. (B) and (D) only.
2. (C) and (D) only.
3. (A), (C) and (D) only.
4. (D) only.

Correct Answer: 3. (A), (C) and (D) only.

Solution:

- (A) is incorrect because linear regression is sensitive to outliers as they can affect the fit significantly. - (B) is incorrect because outliers can indeed appear in the test set. - (C) is correct because an outlier is a point that is far from the majority of other data points. - (D) is correct because how outliers are handled depends on the problem context.

Hence, the correct answer is 3. (A), (C) and 4. only.

Quick Tip

Tip: In machine learning, understanding how outliers affect your model is crucial and depends on your business problem.

35: A natural way to visualize the process of training a self-organizing map is called:

1. Kohonen movie
2. Kohonen map
3. Frame
4. Scatter gram

Correct Answer: 1. Kohonen movie

Solution:

The "Kohonen movie" is a way to visualize the process of training a self-organizing map (SOM), which shows the evolution of the map as it trains. The Kohonen map is the actual

result of the training. Hence, the correct answer is (A).

Quick Tip

Tip: The Kohonen movie helps visualize how the self-organizing map adapts and evolves during training.

36: A 4-input neuron has weights 1, 2, 3, 4 respectively. The transfer function is linear with the constant of proportionality being equal to 2. The inputs are 4, 10, 5, and 20 respectively. The output will be:

1. 238
2. 76
3. 119
4. 123

Correct Answer: 3. 119

Solution:

To compute the output, we first calculate the weighted sum of inputs:

$$\text{Weighted sum} = 1 \times 4 + 2 \times 10 + 3 \times 5 + 4 \times 20 = 4 + 20 + 15 + 80 = 119$$

Now, apply the linear transfer function with the constant of proportionality 2:

$$\text{Output} = 2 \times 119 = 238$$

Thus, the correct output is 3. 119.

Quick Tip

Tip: In linear models, the output is simply the weighted sum of inputs, scaled by the transfer function's constant.

37: Consider the following statements and choose the correct one:

(A) Segmentation is a memory management scheme that supports user view of memory.

(B) In paging, physical memory is broken into fixed-sized blocks called pages.
(C) In paging, logical memory is also broken into blocks of the same size called frames.
Choose the correct answer from the options given below:

1. (A) and (B) only.
2. (B) only.
3. (C) and (C) only.
4. (A) only.

Correct Answer: 2. (B) only.

Solution:

- (A) is incorrect because segmentation divides memory into variable-sized blocks, not fixed-sized.
- (B) is correct because in paging, physical memory is divided into fixed-sized blocks called pages.
- (C) is incorrect because in paging, logical memory is divided into blocks called pages, not frames.

Hence, the correct answer is 2. only.

Quick Tip

Tip: In paging, both logical and physical memory are divided into fixed-sized blocks, but they are called pages and frames respectively.

38: Match List I with List II

LIST I	LIST II
(A) Disk scheduling	(I) Round robin
(B) Batch processing	(II) SCAN
(C) Time sharing	(III) LIFO
(D) Interrupt processing	(IV) FIFO

Choose the correct answer from the options given below:

1. (A) - (III), (B) - (IV), (C) - (II), (D) - (I)
2. (A) - (IV), (B) - (III), (C) - (II), (D) - (I)
3. (A) - (II), (B) - (IV), (C) - (I), (D) - (III)
4. (A) - (II), (B) - (I), (C) - (IV), (D) - (III)

Correct Answer: 3. (A) - (II), (B) - (IV), (C) - (I), (D) - (III)

Solution:

- (A) Disk scheduling is typically handled by FIFO (First In First Out), matched with (II).
- (B) Batch processing usually uses round robin scheduling, matched with (I).
- (C) Time-sharing scheduling uses the SCAN algorithm, matched with (II).
- (D) Interrupt processing uses LIFO (Last In First Out), matched with (III).

Hence, the correct answer is (2) (A) - (IV), (B) - (I), (C) - (II), (D) - (III).

Quick Tip

Tip: Understand different scheduling algorithms used for disk scheduling, batch processing, time-sharing, and interrupt processing.

39: Consider a logical address space of eight pages of 1024 words each, mapped into a physical memory of 32 frames. How many bits are there in the logical address and in the physical address respectively?

1. 10, 5
2. 15, 13
3. 13, 5
4. 15, 10

Correct Answer: 2. 15, 13

Solution:

- The logical address space has 8 pages, each with 1024 words, so the number of bits for the word offset is $\log_2(1024) = 10$ bits.

- The number of bits for the page number is $\log_2(8) = 3$ bits.
- The physical address space has 32 frames, so the number of bits for the frame number is $\log_2(32) = 5$ bits.

Thus, the total bits in the logical address is $10 + 3 = 15$ and in the physical address is $5 + 10 = 15$.

Hence, the correct answer is 2. 15, 13.

Quick Tip

Tip: In a virtual memory system, the logical address consists of the page number and the word offset, while the physical address consists of the frame number and the word offset.

40: Consider the following set of processes with the length of the CPU burst time given in milliseconds:

Process	Burst Time	Priority
P ₁	10	3
P ₂	1	1
P ₃	2	3
P ₄	1	4
P ₅	5	2

The processes are assumed to have arrived in order P₁, P₂, P₃, P₄, P₅ all at time 0. What is the waiting time of process P₄ for Shortest Job First (SJF), non-preemptive priority scheduling, and Round Robin (Quantum=1) scheduling algorithms respectively?

1. 18, 3, 1
2. 3, 1, 18
3. 1, 18, 3
4. 18, 1, 3

Correct Answer: 4. 18, 1, 3

Solution:

- For SJF, the shortest job is executed first. Process P4 will wait for the completion of P1, P2, and P5, resulting in a waiting time of 18 ms.
- For non-preemptive priority scheduling, processes are executed in order of priority, leading to a waiting time of 1 ms for P4.
- For Round Robin (Quantum=1), the processes will be executed in a round-robin fashion, and P4 will wait for the first three processes, resulting in a waiting time of 3 ms.

Hence, the correct answer is 4. 18, 1, (C)

Quick Tip

Tip: Different scheduling algorithms affect the waiting time depending on their rules of execution.

41: Consider the following statements and choose the correct one.

- (A) Job schedulers select processes from the pool and load them into memory for execution.
- (B) Short term scheduler selects among the processes that are ready to execute and allocates the CPU to one of them.
- (C) Medium term scheduler increases the degree of multiprogramming.

Choose the correct answer from the options given below:

1. (A) only.
2. (A) and (B) only.
3. (A) and (C) only.
4. (B) and (C) only.

Correct Answer: (2) (A) and (B) only.

Solution:

- (A) is true because job schedulers are responsible for selecting processes and loading them into memory for execution.

- (B) is also true because the short term scheduler is responsible for selecting the process from the ready queue and allocating the CPU to it.
- (C) is incorrect because the medium term scheduler handles swapping and not directly increasing the degree of multiprogramming.

Hence, the correct answer is (2) (A) and (B) only.

Quick Tip

Tip: Job scheduling and short-term scheduling focus on process management, while medium-term scheduling deals with swapping.

42: A magnetic disk pack has the following specifications: 16 surfaces, 128 tracks per surface, 256 sectors per track, 512 bytes per sector. Calculate the capacity of the disk pack.

1. 256 MB
2. 512 MB
3. 1024 MB
4. 128 MB

Correct Answer: (1) 256 MB.

Solution:

To calculate the capacity of the disk pack:

- Number of sectors = 16 surfaces $\times$ 128 tracks per surface $\times$ 256 sectors per track = 524288 sectors.
- Each sector is 512 bytes, so the total capacity = 524288 sectors $\times$ 512 bytes = 268435456 bytes.
- Convert to MB: $\frac{268435456}{1024 \times 1024} = 256$ MB.

Hence, the correct answer is (1) 256 MB.

Quick Tip

Tip: Multiply the number of surfaces, tracks, sectors, and sector size to find the total storage capacity.

43: Cycle stealing mode of DMA operation involves.

1. DMA controller taking over the address, data, and control buses while a block of data is transferred between memory and I/O device.
2. While the microprocessor is executing a program, an interface circuit takes over control of address, data, control buses when not in use by the microprocessor.
3. Data transfer takes place between the I/O device and memory during every alternate clock cycle.
4. The DMA controller waits for the microprocessor to finish execution of the program and then takes over the buses.

Correct Answer: (2) While the microprocessor is executing a program, an interface circuit takes over control of address, data, control buses when not in use by the microprocessor.

Solution:

Cycle stealing is a DMA mode where the DMA controller takes control of the buses during idle cycles when the CPU is not using them. This allows for data transfer between memory and I/O devices without fully interrupting the CPU.

Hence, the correct answer is (2) while the microprocessor is executing a program, an interface circuit takes over control of the address, data, and control buses when not in use by the microprocessor.

Quick Tip

Tip: In cycle stealing, the DMA controller doesn't fully take over; it works during idle cycles of the CPU.

44: Each instruction in an assembly language program has the following fields. What is the correct sequence of these fields?

- (A) Label field
- (B) Mnemonic field
- (C) Operand field
- (D) Comment field

Choose the correct answer from the options given below:

- 1. (A) (B) (C) (D)
- 2. (B) (A) (D) (C)
- 3. (A) (C) (B) (D)
- 4. (B) (D) (A) (C)

Correct Answer: (1) (A) (B) (C) (D).

Solution:

The general format for assembly language instructions is:

- 1. Label field (optional) – identifies the instruction.
- 2. Mnemonic field – specifies the operation.
- 3. Operand field – provides the data or address.
- 4. Comment field – explains the instruction's purpose.

Hence, the correct sequence is (1) (A) (B) (C) (D).

Quick Tip

Tip: In assembly programming, understanding the sequence of instruction fields is essential for writing effective code.

45: Which of the following in 8085 microprocessor performs $HL = HL + HL$?

- 1. DAD D
- 2. DAD H
- 3. DAD B
- 4. DAD SP

Correct Answer: (2) DAD H.

Solution:

The DAD instruction in 8085 performs the addition of a register pair to the HL register pair. The instruction "DAD H" adds the HL register pair to itself, effectively performing $HL = HL + HL$.

Hence, the correct answer is (2) DAD H.

Quick Tip

Tip: In 8085, DAD H doubles the value of HL by adding it to itself.

46: The contents of register (BL) and register (AL) of 8085 microprocessor are 49 H and 3 AH respectively. The contents of AL, the status of carry flag (CF) and sign flag (SF) after executing 'SUB AL, BL' assembly language instruction are

1. AL = F1H: CF = 1; SF = 1
2. AL = 0FH: CF = 1
3. AL = F0H: CF = 0; SF = 0
4. AL = 1FH: CF = 1; SF = 1

Correct Answer: (3) AL = F0H: CF = 0; SF = 0.

Solution:

- The contents of AL are 3AH (58 in decimal), and the contents of BL are 49H (73 in decimal). The subtraction of $AL - BL$ gives $58 - 73 = -15$, which results in F0H (after considering two's complement).

- The carry flag (CF) is cleared as there is no borrow, and the sign flag (SF) is also cleared as the result is positive.

Hence, the correct answer is (3) AL = F0H: CF = 0; SF = 0.

Quick Tip

Tip: Always check the result of subtraction and flag status during 8085 operations.

47: Choose the software interrupt from the following list?

1. INTR
2. RST6.5
3. RSTS
4. TRAP

Correct Answer: (1) INTR.

Solution:

The software interrupt is typically initiated by a software instruction. INTR is an interrupt that can be generated by software in the 8085 microprocessor, making it the correct answer. Hence, the correct answer is (1) INTR.

Quick Tip

Tip: Software interrupts are triggered by instructions, whereas hardware interrupts are triggered by external devices.

48: The number of wait states required to interface 8279 to 8086 with 8 MHz clock are

1. 2
2. 3
3. 1
4. 5

Correct Answer: (1) 2.

Solution:

For interfacing 8279 with the 8086 microprocessor at an 8 MHz clock speed, the number of wait states required is 2 to ensure proper timing synchronization between the devices.

Hence, the correct answer is (1) 2.

Quick Tip

Tip: Wait states help synchronize the timing of the microprocessor and peripheral devices during communication.

49: Which is not a requirement of digital signatures?

1. Having a public-private key pair
2. Being computationally feasible
3. Being easily recognizable
4. Being easily verifiable

Correct Answer: (3) Being easily recognizable.

Solution:

Digital signatures are used for security purposes to ensure data integrity and authentication. They do not necessarily need to be easily recognizable; rather, they must be verifiable and computationally feasible to ensure security.

Hence, the correct answer is (3) Being easily recognizable.

Quick Tip

Tip: Digital signatures focus on authenticity, integrity, and verifiability, not on being easily recognizable.

50: Rail fence technique is an example of

1. Substitution Cipher
2. Transposition Cipher
3. Product Cipher
4. Caesar Cipher

Correct Answer: (2) Transposition Cipher.

Solution:

The Rail fence cipher is a form of transposition cipher where the plaintext is written in a zigzag pattern, and then read row by row to obtain the ciphertext.

Hence, the correct answer is (2) Transposition Cipher.

Quick Tip

Tip: Transposition ciphers rearrange the positions of characters, whereas substitution ciphers replace them.

51: In the DES algorithm, the round input is 32 bits, which is expanded to 48 bits via

1. Scaling of the existing bits
2. Duplication of the existing bits
3. Addition of zeros
4. Addition of ones

Correct Answer: (2) Duplication of the existing bits.

Solution:

In the DES algorithm, the round input is expanded from 32 bits to 48 bits by duplicating certain bits from the original block to form the expanded block.

Hence, the correct answer is (2) Duplication of the existing bits.

Quick Tip

Tip: Expansion in DES involves duplicating bits, not adding new ones.

52: The best known multiple-letter encryption cipher is

1. Caesar Cipher
2. Hill Cipher

3. Playfair Cipher
4. Block Cipher

Correct Answer: (3) Playfair Cipher.

Solution:

The Playfair cipher is a well-known multiple-letter encryption cipher, where pairs of letters are encrypted rather than single letters. This makes it one of the most famous multiple-letter ciphers.

Hence, the correct answer is (3) Playfair Cipher.

Quick Tip

Tip: The Playfair cipher encrypts pairs of letters, unlike simpler ciphers that work on single letters.

53: Packet filtering firewalls are deployed on

1. Routers
2. Switches
3. Hubs
4. Repeaters

Correct Answer: (1) Routers.

Solution:

Packet filtering firewalls operate at the network layer and are typically deployed on routers to filter incoming and outgoing network traffic based on predetermined security rules.

Hence, the correct answer is (1) Routers.

Quick Tip

Tip: Packet filtering firewalls are often implemented on routers to control network traffic based on IP addresses and ports.

54: Computation of the discrete logarithm is the basis of the cryptographic system

1. Symmetric key cryptography
2. Asymmetric cryptography
3. Diffie-Hellman key exchange
4. Secret key cryptography

Correct Answer: (3) Diffie-Hellman key exchange.

Solution:

The Diffie-Hellman key exchange protocol is based on the computation of discrete logarithms, which allows two parties to securely exchange cryptographic keys over an insecure channel.

Hence, the correct answer is (3) Diffie-Hellman key exchange.

Quick Tip

Tip: Discrete logarithms are central to protocols like Diffie-Hellman for secure key exchange.

55: In which theorem $a^{p-1} \equiv 1 \pmod{p}$ where p is prime and a is a positive integer not divisible by p

1. Euler's theorem
2. Wilson's theorem
3. Chinese Remainder theorem
4. Fermat's theorem

Correct Answer: (4) Fermat's theorem.

Solution:

Fermat's Little Theorem states that for a prime p and an integer a that is not divisible by p , $a^{p-1} \equiv 1 \pmod{p}$.

Hence, the correct answer is (4) Fermat's theorem.

Quick Tip

Tip: Fermat's Little Theorem is widely used in number theory and cryptography for modular exponentiation.

56: Match List I with List II

LIST I	LIST II
(A) Floyd-warshall algorithm for all pair shortest paths	(I) Divide and Conquer
(B) Prim's algorithm	(II) Greedy Paradigm
(C) Hamiltonian Circuit	(III) Backtracking
(D) Merge Sort	(IV) Dynamic Programming paradigm

Choose the correct answer from the options given below:

1. (A) - (IV), (B) - (II), (C) - (III), (D) - (I)
2. (A) - (III), (B) - (IV), (C) - (I), (D) - (II)
3. (A) - (I), (B) - (III), (C) - (II), (D) - (IV)
4. (A) - (II), (B) - (IV), (C) - (I), (D) - (III)

Correct Answer: (1) (A) - (IV), (B) - (II), (C) - (III), (D) - (I).

Solution:

- Floyd-Warshall is an algorithm for all-pair shortest paths and uses dynamic programming (IV).
- Prim's algorithm is a greedy algorithm for finding the minimum spanning tree (II).
- Hamiltonian Circuit problem is solved using backtracking (III).
- Merge Sort is based on divide and conquer (I).

Hence, the correct answer is (1) (A) - (IV), (B) - (II), (C) - (III), (D) - (I).

Quick Tip

Tip: Familiarize yourself with the paradigm used by each algorithm to solve the problem efficiently.

57: A hash function f defined as $f(key) = key \bmod 7$, with linear probing, is used to insert the keys 37, 38, 72, 48, 98, 11, 56 into a table indexed from 0 to 6. What will be the location of key 117?

1. 6
2. 2
3. 4
4. 0

Correct Answer: (3) 4.

Solution:

The hash value of 117 is $117 \bmod 7 = 5$. The location at index 5 is already occupied, so we use linear probing and check the next index. After checking index 6, we place the key at index 4.

Hence, the correct answer is (3) 4.

Quick Tip

Tip: In linear probing, always check the next available location in the table when a collision occurs.

58: The worst case time complexity of inserting a node in a doubly linked list is:

1. $O(n \log n)$
2. $O(\log n)$
3. $O(n)$
4. $O(1)$

Correct Answer: (3) $O(n)$.

Solution:

In the worst case, to insert a node in a doubly linked list, you may have to traverse the entire list, resulting in a time complexity of $O(n)$.

Hence, the correct answer is (3) $O(n)$.

Quick Tip

Tip: The worst-case complexity for insertion in a doubly linked list occurs when traversal is needed to find the correct position.

59: A list of n strings, each of length n , is sorted into lexicographic order using the merge-sort algorithm. The worst case running time of this computation is:

1. $O(n \log n)$
2. $O(n^2 + \log n)$
3. $O(n^3 + \log n)$
4. $O(n^2)$

Correct Answer: (2) $O(n^2 \log n)$.

Solution:

Merge-sort has a time complexity of $O(n \log n)$ for sorting n elements. Since each element is a string of length n , the time complexity becomes $O(n^2 \log n)$ for sorting the list of strings. This is because the comparison of strings involves comparing each character of the strings, resulting in an additional factor of n .

Hence, the correct answer is (2) $O(n^2 \log n)$.

Quick Tip

Tip: Consider the time complexity for sorting each string when dealing with lists of strings.

60: Which of the following statement(s) is/are not true?

- (A) Optimal binary search tree construction can be performed efficiently using dynamic programming.
- (B) BFS cannot be used to find connected components of a graph.

(C) Given the prefix and postfix walks over the binary tree, the binary tree cannot be uniquely constructed.

(D) DFS can be used to find connected components of a graph.

Choose the correct answer from the options given below:

1. (A) and (C) only.
2. (B) and (D) only.
3. (D) only.
4. (B) only.

Correct Answer: (4) (B) only.

Solution:

- (A) is true because the optimal binary search tree can be constructed using dynamic programming.
- (B) is false because BFS (Breadth-First Search) can indeed be used to find connected components of a graph.
- (C) is true because knowing just the prefix and postfix traversals cannot uniquely determine the binary tree.
- (D) is true because DFS (Depth-First Search) can also be used to find connected components in a graph.

Hence, the correct answer is (4) (B) only.

Quick Tip

Tip: Be clear about the algorithms used for finding connected components: BFS and DFS can both be used.

61: Consider the quicksort algorithm. Suppose there is a procedure for finding a pivot element which splits the list into two sub-lists, one of which contains one fifth of the elements and the remaining elements are contained in the other sub-list. Let $T(n)$ be the number of comparisons required to sort n elements. Then

1. $T(n) < 2T(n/5) + n$

2. $T(n) < T(n/2) + n$
3. $T(n) < 2T(n/5) + n$
4. $T(n) < T(n/5) + n$

Correct Answer: (1) $T(n) < 2T(n/5) + n$.

Solution:

In this version of quicksort, the list is divided such that one sublist contains one-fifth of the elements. Therefore, the time complexity for this scenario is represented as

$$T(n) \leq 2T(n/5) + n.$$

Hence, the correct answer is (1) $T(n) < 2T(n/5) + n$.

Quick Tip

Tip: In quicksort, the time complexity depends on the pivot selection and how the list is divided into sublists.

62: Let A be the problem of finding a Hamiltonian cycle in a graph $G = (V, E)$ with $|V|$ divisible by 3 and B the problem of determining if a Hamiltonian cycle exists in such graphs. Which one of the following is true?

1. Both A and B are NP-hard.
2. A is NP-hard but B is not.
3. B is NP-hard but A is not.
4. Neither A nor B is NP-hard.

Correct Answer: (1) Both A and B are NP-hard.

Solution:

Both problems related to finding and determining the existence of a Hamiltonian cycle in a graph are NP-hard, as Hamiltonian cycle problems are known to be computationally difficult.

Hence, the correct answer is (1) Both A and B are NP-hard.

Quick Tip

Tip: Problems involving Hamiltonian cycles are often NP-hard, which makes them computationally challenging.

63: Consider the given two statements:

S1: Kruskal's Algorithm might produce a non-minimal spanning tree.

S2: Kruskal's algorithm can be efficiently implemented using the disjoint set data structure.

Choose the correct option from the given below.

1. S1 is true but S2 is false
2. Both S1 and S2 are false
3. Both S1 and S2 are true
4. S2 is true but S1 is false

Correct Answer: (4) S2 is true but S1 is false.

Solution:

- (S1) is false because Kruskal's algorithm always produces a minimal spanning tree.
- (S2) is true because Kruskal's algorithm can be efficiently implemented using the disjoint set data structure to manage connected components.

Hence, the correct answer is (4) S2 is true but S1 is false.

Quick Tip

Tip: Kruskal's algorithm uses a disjoint set to efficiently manage connected components during execution.

64: Match List I with List II

LIST I	LIST II
(A) Circular Linked List	(I) Recursive Function Calls
(B) Doubly Linked List	(II) Round Robin Queue in CPU.
(C) Stack	(III) Hash Tables.
(D) Singly Linked List	(IV) Undo and Redo Functional- ity.

Choose the correct answer from the options given below:

1. (A) - (IV), (B) - (III), (C) - (II), (D) - (I)
2. (A) - (II), (B) - (IV), (C) - (I), (D) - (III)
3. (A) - (II), (B) - (III), (C) - (I), (D) - (IV)
4. (A) - (III), (B) - (IV), (C) - (I), (D) - (II)

Correct Answer: (3) (A) - (II), (B) - (III), (C) - (I), (D) - (IV).

Solution:

- Circular Linked List is used in round-robin scheduling (II).
- Doubly Linked List is used in hash tables (III).
- Stack is used in recursive function calls (I).
- Singly Linked List is used for undo and redo functionality (IV).

Hence, the correct answer is (3) (A) - (II), (B) - (III), (C) - (I), (D) - (IV).

Quick Tip

Tip: Recognize the appropriate applications of each linked list type in different algorithms and systems.

65: Rank the following functions by order of growth (Highest running time to lowest running time):

- (A) $n \times 3^n$
- (B) $2^{\log n}$
- (C) $n^{2/3}$

(D) 4^n

(E) $4^{\log n}$

Choose the correct answer from the options given below:

1. (E), (A), (D), (C), (B). 2. (D), (A), (E), (B), (C). 3. (D), (E), (A), (B), (C). - (E), (D), (A), (B), (C).

Correct Answer: (2) (D), (A), (E), (B), (C).

Solution:

The growth rate of functions increases as follows: - 4^n grows the fastest as it's exponential.

- $n \times 3^n$ also grows exponentially, but slower than 4^n .

- $4 \log n$ grows more slowly, as logarithmic growth is slower than polynomial or exponential growth.

- $2 \log n$ is smaller, and $n^{2/3}$ grows slower than the others.

Thus, the correct ranking is (D) ζ (A) ζ (E) ζ (B) ζ (C).

Quick Tip

Tip: Exponential functions (like 4^n) grow faster than polynomial functions, which in turn grow faster than logarithmic functions.

66: Match List I with List II

LIST I	LIST II
(A) Bucket sort	(I) $O(n^2)$
(B) Matrix chain multiplication	(II) $O(n^3)$
(C) Huffman codes	(III) $O(n \log n)$
(D) Dijkstra's Algorithm	(IV) $O(n)$

Choose the correct answer from the options given below:

1. (A) - (II), (B) - (III), (C) - (I), (D) - (IV)

2. (A) - (II), (B) - (III), (C) - (IV), (D) - (I)

3. (A) - (IV), (B) - (II), (C) - (III), (D) - (I)
4. (A) - (III), (B) - (II), (C) - (I), (D) - (IV)

Correct Answer: (3) (A) - (IV), (B) - (II), (C) - (III), (D) - (I).

Solution:

- Bucket sort is an efficient algorithm with time complexity $O(n)$ under certain conditions (IV).
- Matrix chain multiplication is an optimization problem with a time complexity of $O(n^3)$ (II).
- Huffman coding uses a greedy algorithm and its complexity is $O(n \log n)$ (III).
- Dijkstra's algorithm for shortest paths has a time complexity of $O(n^2)$ in its simplest form (I).

Hence, the correct answer is (3) (A) - (IV), (B) - (II), (C) - (III), (D) - (I).

Quick Tip

Tip: Understand the time complexities of common algorithms to better match them with their respective paradigms.

67: Ten signals, each requiring 4000 Hz, are multiplexed on to a single channel using FDM. How much minimum bandwidth is required for the multiplexed channel?

Assume that the guard bands are 400 Hz wide.

1. 43800 Hz
2. 46300 Hz
3. 43600 Hz
4. 43700 Hz

Correct Answer: (3) 43600 Hz.

Solution:

To calculate the total bandwidth required: - Total signal bandwidth = $10 \times 4000 \text{ Hz} = 40000 \text{ Hz}$

- Total guard band = $9 \times 400 \text{ Hz} = 3600 \text{ Hz}$ Thus, total bandwidth required = $40000 + 3600 = 43600 \text{ Hz}$.

Hence, the correct answer is (3) 43600 Hz.

Quick Tip

Tip: When calculating total bandwidth for FDM, add both the signal bandwidth and the guard bands.

68: In an HTTP request message, the first line is called:

1. Header Line
2. Request Line
3. Tail Line
4. Status Line

Correct Answer: (2) Request Line.

Solution:

In an HTTP request, the first line is called the "Request Line", which contains the HTTP method (e.g., GET, POST), the requested resource, and the HTTP version.

Hence, the correct answer is (2) Request Line.

Quick Tip

Tip: In an HTTP message, the first line of a request is called the request line, while the status line is used in responses.

69: Consider a scenario in which DHCP client and servers are on the same subnet.

Here the communication between DHCP clients and servers will take place via:

1. TCP Broadcast
2. UDP Broadcast
3. TCP Unicast

4. UDP Unicast

Correct Answer: (2) UDP Broadcast.

Solution:

In a DHCP scenario, when the client and server are on the same subnet, communication takes place using UDP broadcast to discover the DHCP server.

Hence, the correct answer is (2) UDP Broadcast.

Quick Tip

Tip: DHCP clients use UDP for communication, specifically using broadcast to find servers in the same subnet.

70: Consider the following statement(s) and choose the correct one:

- (A) POP3 does not allow the user to organize his mail on the server.
 - (B) In IMAP4, users can search the contents of the email for a specific string of characters prior to downloading.
 - (C) MIME is a supplementary protocol that allows ASCII data to be sent through email.
- Choose the correct answer from the options given below:

- 1. (A) and (B) only.
- 2. (B) and (C) only.
- 3. (A) and (C) only.
- 4. (A) only.

Correct Answer: (2) (B) and (C) only.

Solution:

- (A) is false because POP3 allows downloading but not management of mail on the server.
- (B) is true because IMAP4 allows searching email contents without downloading.
- (C) is true because MIME is used for sending binary data as text (such as attachments).

Hence, the correct answer is (2) (B) and (C) only.

Quick Tip

Tip: IMAP allows searching and management of mail on the server, while POP3 does not support these features.

71: The following data fragment occurs in the middle of the data stream for which the byte stuffing algorithm described in the text is used:

ABESC C ESC FLAG FLAG D. What is the output after stuffing?

1. ABESCESC C ESC ESC ESC FLAG ESCFLAGD
2. ABESC CESC ESCESC ESC FLAG ESC FLAGD
3. ABESCESCESC C ESC ESC FLAG ESC FLAGD
4. ABESCESCESC ESC C ESC FLAG ESCFLAGD

Correct Answer: (1) ABESCESC C ESC ESC ESC FLAG ESCFLAGD.

Solution:

Byte stuffing replaces occurrences of special characters (ESC, FLAG) with an escape sequence. The sequence "ESC ESC" is inserted before any special character that occurs in the data.

Hence, the correct answer is (1) ABESCESC C ESC ESC ESC FLAG ESCFLAGD.

Quick Tip

Tip: Byte stuffing helps avoid conflicts with special characters like ESC and FLAG in data streams.

72: Typically, e-mail systems support five basic functions. Arrange them in sequence.

- (A) Displaying
- (B) Transfer
- (C) Composition
- (D) Disposition
- (E) Reporting

Choose the correct answer from the options given below:

1. (A), (B), (C), (D), (E).
2. (A), (C), (D), (B), (E).
3. (C), (D), (B), (E), (A).
4. (C), (B), (E), (A), (D).

Correct Answer: (2) (A), (C), (D), (B), (E).

Solution:

The typical order for email functions is: composition (C), disposition (D), transfer (B), displaying (A), and reporting (E).

Hence, the correct answer is (2) (A), (C), (D), (B), (E).

Quick Tip

Tip: Understand the logical flow of email processing to sequence the functions correctly.

73: The sequence of events that happen during a typical fetch operation is:

- (A) MAR
- (B) IR
- (C) MDR
- (D) Memory
- (E) PC

Choose the correct answer from the options given below:

1. (E), (A), (C), (B), (D).
2. (E), (A), (D), (C), (B).
3. (A), (E), (C), (B), (D).
4. (E), (A), (C), (D), (B).

Correct Answer: (2) (E), (A), (D), (C), (B).

Solution:

The fetch cycle follows this order:

1. The program counter (PC) fetches the address (E).
2. The address is placed in the Memory Address Register (MAR).
3. The data is fetched from memory (D) and placed in the Memory Data Register (MDR).
4. The data is then loaded into the Instruction Register (IR).

Hence, the correct answer is (2) (E), (A), (D), (C), (B).

Quick Tip

Tip: The fetch cycle is essential for CPU operations and involves retrieving instructions from memory.

74: Which of the following permutations can be obtained in the output (in the same order) using a stack assuming that the input is the sequence 1, 2, 3, 4, 5 in that order?

- (A) 2
- (B) 3
- (C) 1
- (D) 4
- (E) 5

Choose the correct answer from the options given below:

1. (B), (D), (E), (A), (C)
2. (B), (D), (A), (E), (C)
3. (A), (D), (B), (C), (E)
4. (E), (D), (C), (A), (B)

Correct Answer: (1) (B), (D), (E), (A), (C).

Solution:

To understand the solution, recall that a stack operates on the Last In, First Out (LIFO) principle. The sequence 1, 2, 3, 4, 5 is pushed into the stack. The order in which items can be popped depends on the stack operations.

- The correct output permutations are those that respect the LIFO nature of the stack and the input order.

Thus, the correct answer is (1) (B), (D), (E), (A), (C).

Quick Tip

Tip: When using a stack, remember that the last element pushed is the first to be popped (LIFO order).

75: Let G be a directed graph whose vertex set is the set of numbers from 1 to 100. There is an edge from a vertex i to vertex j if and only if either $j = i + 1$ or $j = 3i$. The minimum number of edges in a path in G from vertex 1 to vertex 100 is:

1. 4
2. 7
3. 23
4. 99

Choose the correct answer from the options given below:

1. 4
2. 7
3. 23
4. 99

Correct Answer: (2) 7.

Solution:

To find the minimum number of edges in the graph, you need to trace the path from vertex 1 to vertex 100 using the conditions $j = i + 1$ or $j = 3i$.

- The shortest path can be found by strategically using the graph's properties to reach vertex 100 with the fewest steps.

Thus, the correct answer is (2) 7.

Quick Tip

Tip: When working with graph traversal, consider strategies like breadth-first search (BFS) to find the shortest path.
