IAE ACADEMY INSTITUTE FOR ENGINEERS

CORE JOBS · INTERVIEW QUESTIONS

IAE Academy free study material · Updated 30 October 2026

Data Structures and Algorithms Interview Questions

INSIDE

- 24 questions with short answers
- 6 numericals with the working
- Grouped by topic

Companies do not publish their interview questions; these are ours.

Online, with links to the notes: **iaeacademy.in/study/data-structures-algorithms/interview-questions**

FREE COUNSELLING

Preparing for GATE or a core job? Talk to IAE Academy.

Or call **+91 84204 36524**

Every day, 10:30 am to 7:30 pm

Chat on WhatsApp

Arrays, stacks, queues and linked lists

1. How does a stack convert an infix expression to postfix?

Scan the expression left to right. Operands go straight to the output. An operator first pops operators from the top of the stack while they have higher precedence (or equal precedence, if it is left-associative), stopping at an opening bracket, then is pushed. An opening bracket is pushed, a closing bracket pops until its opening bracket, and at the end the stack is emptied to the output.

2. A 2-D integer array has 20 rows and 30 columns, stored in row-major order from base address 1000, with 4 bytes per element and indices starting at 0. Find the address of A[5][7].

The address of A[5][7] is 1628.

> **WORKING**
>
> 1. Elements before row 5: $5 \times 30 = 150$
> 2. Offset within row 5 is 7, so the element index is $150 + 7 = 157$
> 3. Address $= 1000 + 157 \times 4$
> 4. $= 1000 + 628 = 1628$

3. What is the difference between an array and a singly linked list in access and insertion cost?

An array gives $O(1)$ access by index but inserting in the middle costs $O(n)$ because elements shift. A singly linked list needs $O(n)$ to reach the k-th element but inserts $O(1)$ once the position is known. The list also spends memory on pointers and has poorer cache behaviour.

Trees, binary search trees and binary heaps

4. In a binary tree, how is the number of leaves related to the number of nodes with two children, and why?

The number of leaves is one more than the number of nodes with two children: $n_0 = n_2 + 1$. A tree with n nodes has $n - 1$ edges, which equals $n_1 + 2n_2$. Substituting $n = n_0 + n_1 + n_2$ gives the result.

5. A binary tree has 9 nodes with two children and 6 nodes with exactly one child. How many nodes does it have in total?

The tree has 25 nodes.

> **WORKING**
>
> 1. Leaves: $n_0 = n_2 + 1 = 9 + 1 = 10$
> 2. Total $n = n_0 + n_1 + n_2$
> 3. $n = 10 + 6 + 9 = 25$

6. Why does building a binary heap from an array of n elements take linear time rather than n log n?

Build-heap runs sift-down on each internal node from the bottom up, and the total is $\Theta(n)$. A node at height h costs $O(h)$, and about $n/2^{h+1}$ nodes sit at that height. The sum of $h \cdot n/2^{h+1}$ over all heights is $O(n)$, because most nodes are near the leaves.

7. What does an AVL tree guarantee, and how does it keep that guarantee?

An AVL tree keeps the heights of every node's two subtrees within 1 of each other, so its height is $O(\log n)$ and search, insert and delete all cost $O(\log n)$. After an insert or delete, rotations restore the balance. A plain binary search tree can degrade to height $n - 1$ on sorted input.

Asymptotic complexity and recurrences

8. What is the difference between Big O, Omega and Theta notation?

Big O gives an upper bound on growth, Omega gives a lower bound, and Theta gives both at once, so the function grows at exactly that rate up to constants. Saying an algorithm is $O(n^2)$ allows it to be much faster, while $\Theta(n^2)$ says its cost really is quadratic.

9. Consider `for (i = 1; i < n; i = i * 2) for (j = 0; j < n; j++) count++;` with n = 1024. What is the final value of count, starting from 0?

The statement runs 10240 times.

> **WORKING**
>
> 1. Outer loop: $i = 1, 2, 4, \ldots, 512$, stopping when $i \geq 1024$
> 2. Outer iterations $= \log_2 1024 = 10$
> 3. The inner loop runs $n = 1024$ times per outer iteration
> 4. Total $= 10 \times 1024 = 10240$

10. Solve $T(n) = 2T(n/2) + n$ using the Master theorem.

Here $a = 2$, $b = 2$ and $f(n) = n$. Since $n^{\log_b a} = n^{\log_2 2} = n$ matches $f(n)$, the second case applies and $T(n) = \Theta(n \log n)$. This is the recurrence of merge sort: two half-size subproblems and a linear combine step.

11. Why does asymptotic analysis ignore constants and lower-order terms?

It describes how cost grows as the input becomes large, and for large n the dominant term decides the growth. Constants depend on the machine and compiler, so they say little about the algorithm itself. For small inputs they still matter: a quadratic method with a small constant can beat an $n \log n$ one.

Searching, sorting and hashing

12. What is the difference between merge sort and quicksort?

Merge sort is $O(n \log n)$ in the worst case, is stable, and needs $O(n)$ extra space for arrays. Quicksort is $O(n^2)$ in the worst case but $O(n \log n)$ on average, sorts in place apart from its recursion stack, and is usually not stable. A bad pivot choice causes its worst case.

13. Why can no comparison-based sort beat $\Omega(n \log n)$ comparisons in the worst case?

A comparison sort can be drawn as a binary decision tree in which every possible ordering of the input reaches its own leaf. With n distinct keys there are at least $n!$ leaves, so the tree height is at least $\log_2 n! = \Theta(n \log n)$. The longest path is the worst-case number of comparisons.

14. A hash table has 7 slots, numbered 0 to 6, with hash function h(k) = k mod 7 and linear probing with step 1. Keys 10, 17 and 24 are inserted in that order into an empty table. In which slot does 24 land?

Key 24 lands in slot 5.

WORKING

1. $10 \bmod 7 = 3$, so 10 goes to slot 3
2. $17 \bmod 7 = 3$ is taken, so probe slot 4: 17 goes there
3. $24 \bmod 7 = 3$ is taken, slot 4 is taken, so probe slot 5
4. Slot 5 is free

15. Binary search runs on a sorted array of 1000 elements, comparing the key with one middle element per step. At most how many comparisons are needed?

At most 10 comparisons are needed.

WORKING

1. Each comparison halves the range, so the worst case is $\lfloor \log_2 n \rfloor + 1$
2. $\log_2 1000 \approx 9.97$, so the floor is 9
3. $9 + 1 = 10$

Greedy, dynamic programming and divide-and-conquer

16. What is the difference between a greedy algorithm and dynamic programming?

A greedy algorithm makes one locally best choice at each step and never revisits it, so it is fast but correct only for problems where that choice can be proved safe. Dynamic programming solves overlapping subproblems once, stores the answers and combines them, so it finds the optimum for problems with optimal substructure, including many where greedy fails.

17. Why does choosing items by highest value-to-weight ratio fail for 0/1 knapsack, although it works for fractional knapsack?

Items cannot be split in 0/1 knapsack, so the best-ratio choices can leave unusable capacity. With capacity 50 and (weight, value) items (10, 60), (20, 100), (30, 120), the ratio rule takes the first two for 160, but items 2 and 3 give 220. Fractional knapsack can fill the gap.

18. How does the dynamic programming solution for the longest common subsequence work?

Let $L[i][j]$ be the LCS length of the first i characters of one string and the first j of the other. If the characters match, $L[i][j] = L[i-1][j-1] + 1$; otherwise it is the larger of $L[i-1][j]$ and $L[i][j-1]$. Filling the table takes $O(mn)$ time for lengths m and n.

Graphs

19. What is the difference between BFS and DFS?

BFS uses a queue and visits vertices level by level, so it finds shortest paths by edge count in an unweighted graph. DFS uses a stack or recursion and goes deep first, giving discovery and finish times used for edge

classification, topological sort and strongly connected components. Both take $O(V + E)$ with adjacency lists.

20. Why does Dijkstra's algorithm fail on graphs with negative edge weights?

Dijkstra finalises a vertex when it is removed from the queue, assuming no later path can be shorter. A negative edge can make a later path cheaper, so a finalised distance may be wrong. Bellman-Ford handles negative edges in $O(VE)$ and also detects negative cycles.

21. What is the difference between Kruskal's and Prim's algorithms for minimum spanning trees?

Kruskal sorts all edges and adds the cheapest one that does not form a cycle, using union-find, in $O(E \log E)$. Prim grows a single tree from a start vertex, always adding the cheapest edge leaving it, in $O(E \log V)$ with a heap. Both are correct because of the cut property.

22. The Floyd-Warshall algorithm runs on a graph with 100 vertices. How many relaxation steps does it perform, counting one per triple of vertices?

It performs 1000000 relaxation steps.

> **WORKING**
>
> 1. Floyd-Warshall has three nested loops over k, i and j
> 2. Each loop runs over all V vertices
> 3. Steps $= V^3 = 100^3 = 1000000$

Programming in Python

23. Why does a mutable default argument such as `def f(x, L=[])` behave unexpectedly in Python?

The default value is created once, when the function is defined, not on each call. So every call that omits the argument uses the same list, and appends build up across calls. The usual fix is to default to `None` and create a new list inside the function.

24. What does `-7 // 2` give in Python 3, and why?

It gives -4. The operator `//` is floor division, which rounds toward negative infinity, not toward zero. The remainder follows the same rule, so `-7 % 2` is 1, and the identity `(a // b) * b + a % b == a` always holds.

FROM IAE ACADEMY

Preparing for GATE or a core job?

IAE Academy teaches GATE in seven papers in Ballygunge, Kolkata, and online. Counselling is free: ask us about batches, fees or your preparation.

FREE COUNSELLING

Chat on WhatsApp

Call +91 84204 36524

Open every day, 10:30 am to 7:30 pm.

Scan to chat on WhatsApp

GATE COURSES

Classroom batches in Ballygunge and online batches. Fees and details are on the course pages.

Scan: courses

See the courses

FREE ON IAEACADEMY.IN

Notes and solved previous-year questions for Data Structures and Algorithms, topic by topic.

Scan: the notes

Open the notes

Found this useful? Send it to a friend preparing for core-job interviews.

Send to a friend on WhatsApp

IAE Academy free study material. Companies do not publish their interview questions; these are ours. Prepared and published by IAE Academy, Kolkata. Free to read and free to share, unaltered and with attribution.

IAE ACADEMY INSTITUTE FOR ENGINEERS

GATE coaching in Kolkata since 2012 · EC, EE, IN, CS and IT, DA, BT and CE · in the classroom and online

227A Rash Behari Ave, Ballygunge, Kolkata 700019

+91 84204 36524 · iaeacademy.in · Updated 30 October 2026

Get directions