

The Tech Interview Playbook: From DSA to System Design

© Chinmoy Mukherjee 2025-2045 no part of this document can be used without explicit written permission from the author.

The Tech Interview Playbook: From DSA to System Design

Introduction: Navigating the Tech Interview Landscape

Chapter 1: The Bedrock - Data Structures & Algorithms (DSA)

1.1 Linked Lists: The Building Blocks

Question: How do you reverse a linked list?

1.2 Trees: Hierarchical Data

Question: How do you find the k-th smallest element in a Binary Search Tree?

1.3 Arrays and Strings: Foundational Manipulation

Question: Given a string, find the first non-repeated character in it.

Question: You have an array of size n containing integers in the range 1 to n. Each number is present at least once, except for two missing numbers. How do you find them?

Chapter 2: Language & Platform Deep Dives

2.1 Core OOP Principles in a Modern Context

Question: Describe the principles of OOPS.

2.2 Java Collections and Concurrency

Question: What is the difference between ArrayList and Vector?

Question: What is the difference between HashMap and Hashtable?

2.3 Code Optimization Principles

Question: How would you optimize the following code? `for(int i=0; i<10000; i++) { x = y * 2; }`

Question: How would you optimize this code block? `int x=1, a[5]; for(int i=0; i<5; i++) { x = x*a[i]; }`

Chapter 3: Databases in the Age of Big Data

3.1 Mastering SQL: Beyond the Basics

Question: What is the difference between the WHERE and HAVING clauses?

Question: What is the difference between IN and EXISTS?

3.2 Indexing and Performance

Question: What is a clustered index, and how many can a table have?

Question: When should you not use an index?

3.3 The Great Debate: SQL vs. NoSQL

Chapter 4: The Main Event - System Design

4.1 Designing a Distributed Word Counter

4.2 Designing a URL Shortener (e.g., TinyURL)

Chapter 5: Writing Production-Ready Code

Chapter 6: Architecting for the Cloud: micro-services & Migration

Understanding micro-services

Benefits of micro-services

Decomposing an Application into Services

Handling Distributed Transactions: The Saga Pattern

Common micro-service Patterns

Cloud Migration Strategies: The 6 Rs

Chapter 7: Platform & Technology Deep Dive

- Java Concurrency & Features
- Distributed Systems & Messaging
- Chapter 8: Brain Teasers - The Classics
 - Data Platforms
 - The 25 Horses Problem
 - The Island of Liars and Truth-tellers
- Chapter 9: Architecting for the Cloud: micro-services & Migration
 - Common micro-service Patterns
 - Spring Framework & Spring Boot
 - Features of Spring Cloud
 - Distributed Systems & Messaging
 - Containers & Orchestration
 - Data Platforms
- Chapter 10: The Behavioral Deep Dive
 - Q: Tell me about a time you had a conflict with a coworker. How did you resolve it?
 - Q: How do you stay up-to-date with new technologies?
- Conclusion: The Continuous Learner

Introduction: Navigating the Tech Interview Landscape

The landscape for interview has evolved significantly; where rote memorization of algorithms might have once sufficed, today's interviews are structured as a collaborative conversation. They are designed to assess not just your technical knowledge, but more importantly, how you think. Companies today are looking for engineers who can analyze complex trade-offs, design resilient systems, and articulate their thought process with exceptional clarity. The focus has shifted from finding a single, perfect answer to exploring the solution space like a true engineer.

Essential modern software concepts such as distributed system design, cloud-native architecture, advanced concurrency, and API design. This book will not merely provide you with answers; it will equip you with the mental frameworks and contextual understanding necessary to demonstrate the engineering mindset that top-tier companies are actively searching for.

Chapter 1: The Bedrock - Data Structures & Algorithms (DSA)

Data Structures and Algorithms continue to be the absolute foundation of software engineering and, by extension, the technical interview. A robust understanding of DSA serves as a proxy for your capacity to write efficient, optimized, and scalable code. However, in today's landscape, simply arriving at a correct solution is no longer sufficient. You are expected to perform a thorough analysis of its performance, clearly articulate the time and space complexity using Big O notation, and confidently discuss the trade-offs of alternative approaches you considered and why you discarded them.

1.1 Linked Lists: The Building Blocks

Linked lists are a fundamental data structure, often used to assess your core understanding of pointers, memory allocation, and iteration.

Question: How do you reverse a linked list?

This is arguably the most iconic linked list question. You should be prepared to solve it both iteratively and recursively, understanding the trade-offs of each method.

Iterative Solution: This is the most common and space-efficient method, utilizing three pointers: previous, current, and next.

Initialize

previous to NULL and current to the head of the list.

Iterate as long as

current is not NULL.

Inside the loop, first store the next node to prevent losing the rest of the list:

next = current->link.

Reverse the pointer of the current node to point backward:

current->link = previous.

Move the pointers one step forward for the next iteration:

previous = current, current = next.

When the loop terminates,

previous will be pointing to the new head of the reversed list.

Analysis:

Time Complexity: $O(n)$, as every node must be visited exactly once.

Space Complexity: $O(1)$, because only a constant number of pointers are used, regardless of the list's size.

Recursive Solution: This approach is often considered more elegant but is less space-efficient due to the overhead of the call stack.

Base Case: If the head is NULL or the list has only one node (`head->link == NULL`), it is already reversed, so return the head.

Recursive Step: Call the reverse function on the rest of the list: `newHead = reverse(head->link)`. This will reverse the remainder of the list and return its new head.

Reverse Pointer: After the recursive call returns, `head->link` still points to the next node. That next node's link should now point back to head. So, you set

```
head->link->link = head.
```

Break the original forward link by setting

```
head->link = NULL.
```

Return the

`newHead` that was returned from the initial recursive call.

Analysis:

Time Complexity: $O(n)$, as each node is still visited once.

Space Complexity: $O(n)$, due to the recursion depth. Each function call adds a frame to the call stack, which could lead to a stack overflow for very long lists.

1.2 Trees: Hierarchical Data

Trees, particularly Binary Search Trees (BSTs), are used to test your grasp of recursion and hierarchical data management.

Question: How do you find the k-th smallest element in a Binary Search Tree?

The solution hinges on the key property of a BST: an in-order traversal visits its nodes in ascending sorted order.

Solution: In-Order Traversal

Perform a recursive in-order traversal (Left, Root, Right).

Maintain a counter to track the number of nodes visited.

When the counter reaches the value of

k , the current node's value is the k -th smallest element.

Analysis:

Time Complexity: $O(n)$ in the worst case (a skewed tree) and $O(k)$ on average for a balanced tree, as you can stop after finding the element.

Space Complexity: $O(n)$ for the recursion stack in a worst-case skewed tree, or $O(\log n)$ for a balanced tree.

Modern Follow-up Question: If this operation is performed frequently, how can you optimize it?

Optimized Solution: Augment the `TreeNode` structure to include a `leftSubtreeSize` field, which stores the count of nodes in its left subtree. This count must be updated during every insertion and deletion.

To find the k -th smallest element, you can then traverse the tree in

$O(\log n)$ time (on a balanced tree) by using the counts to decide whether to go left or right, avoiding a full traversal.