


2nd Edition



Frequently Asked Technical Interview Questions and Answers

Contents

Frequently Asked Interview Questions and Answers	6
Technical Questions and Answers	7
1. General	9
2. Programming Fundamentals	11
3. Analytical.....	16
4. Algorithms and Coding.....	28
5. C/C++ & Java	42
C/C++.....	42
Java	47
6. Databases.....	57
7. Internet Technology	59
8. Wireless Technology	62
General	62
J2ME	63
Brew	65
9. Out of the Box Questions	68



2. Programming Fundamentals

How can computer technology be integrated into an elevator system for a hundred story office building? How do you optimize for availability? How would variation of traffic over a typical work week, floor, or time of day affect this? How would you test this system?

Implement a model similar to thread scheduling models found in most multitasking OS.

Some points to consider:

- Some floor may have higher priorities, like the executive's floor for instance. Calls from those floors would always be serviced first.
- The elevator should only stop at a floor for x number of seconds, to avoid monopolizing the elevator. This time would be decreased during busy hours and increased during slow hours.
- If, for some reason, a floor has an elevator call outstanding for some long period of time, that floor's priority gets boosted so that it's serviced immediately.
- Based on weight (or # of people or some other similar metric) the elevator car could know when it's full and ignore calls until there is room for at least 2 people. Stopping for just one person isn't worth it for a full elevator.
- When the elevator is idle, it can go to the floor that will need the elevator the most or that which has the highest priority. This will save on the wait time at that floor (e.g: The ground floor).
- Since it is a 100 story bldg, it is highly likely that many elevators are servicing it. So there needs to be coordination between the elevators to service properly. The elevators could be programmed so that they only serve certain floors or floor ranges. This, divide and conquer strategy can be optimized by reducing the number of floors an elevator serves based on floor location (e.g.: The higher the floor, the less number of floors an elevator serves).
- In front of the entrance door of each floor, there could be placed a BIG poster/caricature depicting the health benefits of climbing the stairs!! This could reduce the elevator traffic.

List some of the problems unique to distributed databases?

Among other things, the following points must be considered while designing a distributed database:

- Storing data (fragmentation/replication, horizontal/vertical table partitioning)
- Catalog management (naming, data independence)
- Query processing (cost-based optimization, semi-join)
- Updating data (synchronous/asynchronous)

3. Analytical

Why is a manhole cover round?

There are many possible answers to this question. The interviewer is trying to test your creative thinking and analytical skills, so be inventive. Some answers are:

- The Man hole cover is round so it will not fall into the hole. If it were square it would be possible for the cover to fall into the hole if it was tilted at the proper angle. Being round there is no way, without breaking the cover, for it to fall into the hole. Round covers are more easily transported by one person because they can be rolled on their edge.
- The first person who designed it chose it to be round. Since he was the "first", it became a standard. People are more likely to follow standards rather than question them.
- The circular shape prevents the stress concentration at the edges, as circle does not have any edges.
- Round manhole will expand (heat & cold) evenly and not crack the surrounding concrete.
- Round holes are easier to bore.

If given a rectangular cake with a rectangular piece removed (any size or orientation), how would you cut the remainder of the cake into two equal halves with one straight cut of a knife?

Two possible answers are:

- Find the centers of both the original cake and the removed piece. Cut the cake along the line connecting these two centers. As this line cuts the original cake and the removed piece in half, the remainder is two equal halves.

Cut the entire cake in half horizontally (i.e. parallel to the table). This will get you two even halves.

There are four people who need to cross a bridge at night. The bridge is only wide enough for two people to cross at once. There is only one flashlight for the entire group. When two people cross, they must cross at the slower member's speed. All four people must cross the bridge in 17 minutes, since the bridge will collapse in exactly that amount of time. Here are the times each member takes to cross the bridge:

- Person A: 1 minute
- Person B: 2 minutes
- Person C: 5 minutes
- Person D: 10 minutes

If Person A and C crossed the bridge initially, 5 minutes would elapse, because Person C takes 5 minutes to cross. Then Person A would have to come back to the other side of the bridge, taking another minute, or six minutes total. Now, if Person A and D crossed the bridge next, it would take them 10 minutes, totaling to 16

4. Algorithms and Coding

You are developing a web browser (something like e.g. Netscape, etc.) and need to display all visited links on a page. The visited links need to use a different color than that used to display scheme than the unvisited links. Now, given a history of links you have visited before, how would you go about writing the piece of code that makes the determination if you have seen this link before? Answer or not? The answer could be a simple string comparison, but then think about the time it will take for the client to render any HTML page. Alternatively, so, given a history of URLs, come-up with an elegant way (algorithm, data structures, etc.) to make the determination if a given link already exists in the history list?

Using a Hash Table is probably the most efficient way to do this. You can use several hashing algorithms. For example, checksum of a link can be used as a key for hashing. This will ensure $O(1)$ order provided a good checksum algorithm is used. Whenever a page loads, we can parse all URL's from the page, take their checksum, and compare them with the hashtable. Whichever links match are then displayed in a different color.

Since web pages can have multiple URLs pointing to them, as a web browser developer how can you make sure you have never seen the same content before?

Make a list (or a binary tree) of hashes using MD5, SHA1 or a similar hash/digest algorithm of the pages you've visited. Then compare the digest of the current page to the hashes in the tree. A hashtable is good here too! A hashtable of other longer hashes is a quick, easy, and efficient solution.

Write a function to print all the possible permutations of a string. Now, modify the algorithm to discard duplicates.

Below is a sample implementation that prints all possible permutations of a string:

```
#include <iostream.h>
#include <string.h>

const int MAX_STR = 20;

void CopyStr(char *s2, char *s1, int i)
{
    for (int j = 0, k = 0, len = strlen(s2); j < len; j++) {
        if (i != j) {
            s1[k++] = s2[j];
        }
    }
}
```

5. C/C++ & Java

C/C++

Explain the mechanism of virtual functions and virtual function tables.

Whenever a class member function is declared as virtual, the compiler creates a virtual table in memory which contains all function pointers that are declared as virtual in that class. This enables run time polymorphism (i.e. finding out the desired function at run time). Virtual function tables also have an additional pointer in the object to the vtable. As this additional pointer and the vtable increases the size of the object, a class designer needs to be judicious about declaring functions virtual. The sequence of events upon calling a method on the base object pointer is:

- Get vtable pointer (this vtable pointer points to the beginning of the vtable).
- Get the function pointers in the vtable using offset.

Invoke the function indirectly through the vtable pointer.

Given a singly linked list and a pointer to a certain node in the list, how would you delete that node in constant time?

There are two things that we require to delete a node: the previous node's address and the next node. In this case we only know the current nodes address. So to delete it without breaking the list:

Assume that the node comprises of

Data = The data

Next = Pointer to the next node

Also, assume that our linked list looks like this and we only know the current node address.

```
PREVIOUS NODE -> CURRENT NODE -> NEXT NODE
```

```
//Now Copy the contents of the next node to the current node...that's it
```

```
nextnode=currentnode->next
```

```
currentnode->data = nextnode->data
```

```
currentnode->next = nextnode->next
```

This does not work if the current node is the last node in the list.

What are recursive functions? What are the advantages and disadvantages of recursive algorithms?

A function that calls itself repeatedly, satisfying some condition, is called a Recursive Function. We can't actually declare that recursive or non-recursive algorithms are good or bad. Some problems inherently are better suited for recursion.

6. Databases

What is the difference between a "where" clause and a "having" clause?

"Where" is a restriction statement. You use where clause to restrict data being accessed from the database. Where clause is used before result is retrieved. But having clause is used after retrieving the data. Having clause is a kind of filtering command. You should always use a WHERE clause in preference to a HAVING clause, if possible.

What is a "transaction"? Why are they necessary?

A transaction is a series of data manipulation operations mostly triggered from the client end that must be committed into the database as one whole operation. This means that a transaction is complete only when all the series of data manipulations involved have been successfully committed into the database. All transactions must follow the ACID rules.

- **Atomic** - All operations execute or No operations execute.
- **Consistent** - After the operations execute, database is in a consistent state as it was before execution.
- **Isolated** - Appropriate locks are placed on shared data in order to isolate the operation on the data.
- **Durability** - Data must be persisted on hard disk, so that if there is a crash the data is durable.

What is "normalization"? "Denormalization"? Why do you sometimes want to denormalize?

Normalizing data means eliminating redundant information from a table and organizing the data so that future changes to the table are easier. Denormalization means allowing redundancy in a table. The main benefit of denormalization is improved performance with simplified data retrieval and manipulation. This is done by reduction in the number of joins needed for data processing.

What types of join algorithms can you have?

There are three kinds of JOINS in SQL Server, Nested table, Merge and Hash joins. All of the above depend upon the resources available, MERGE and HASH consume lots of memory; so, a system that is low on memory would sparingly resort to the above two join methodologies.

Nested table joins are based on the fact that indexes exist on nested tables in a SELECT query, for each row in the outer table, a scan is done for the rows available on the inner tables' rows. If a match is found, that row is retrieved.

7. Internet Technology

You have two machines, both connected to the internet. One you are told has a web server running on it, the other doesn't have a web browser of any kind. Using the machine without the browser, how can you tell that the web server is running on the other machine?

The common mistake people make here is they answer "ping the machine". That doesn't check anything other than to see if the machine is visible on the network. It has nothing to do with the web server itself. To find out if the web server is running, just telnet to port 80. If the port 80 is open, it doesn't mean that it is the web server listening to that port yet. To make sure it is the web server, you can type a primitive HTTP request like "GET /" and press "enter" twice. If it is the web server, then you should get the index page HTML code on your screen.

Explain how ping works.

The Packet Internet Groper (ping) utility on the host machine sends datagrams, or packets, addressed to the Internet Protocol (IP) address of the target host using the Internet Control Message Protocol (ICMP) echo request command. If the specified host is operational, it will respond using the ICMP echo reply. The ping utility will then display statistics about the number of lost packets and the amount of time it takes for the response to return from the target host. This utility only verifies that the bottom two layers, Network Interface Layer and Internet Layer, of the TCP/IP stack on the target machine are operational.

How does Ethernet work? How big is an Ethernet address?

Ethernet is normally a shared media LAN. All stations on the segment share the total bandwidth, which is either 10 Mbps (Ethernet), 100 Mbps (Fast Ethernet) or 1000 Mbps (Gigabit Ethernet). With switched Ethernet, each sender and receiver pair has the full bandwidth.

Ethernet uses the CSMA/CD technology to broadcast each frame onto the physical medium (wire, fiber etc.). All stations attached to the Ethernet are "listening," and the station with the matching destination address accepts the frame and checks for errors. It is a data link protocol (MAC layer protocol) and functions at layers 1 and 2 of the OSI model. A unique number is assigned to each Ethernet network adapter. It is a 48-bit number maintained by the IEEE.

What is a bridge? A router? A gateway?

A bridge is an intermediate system used to connect two LANs that use similar LAN protocols. The bridge acts as an address filter, picking-up packets from one LAN that are intended for a destination on another LAN and passing those packets on. It operates at layer 2 of the OSI model.

A router used to connect two networks that may or may not be similar. Router employs an internet protocol present in each router and each end system of the network. It operates at layer 3 of the OSI model.