

1 Tech interview questions: nodes and graphs

1. Write a method which prints a binary tree in level-order (i.e. it first prints the root, then the children of the root at level 1, etc.)
2. Write an iterator for a binary tree which returns only the leaves in order (i.e. from left to right).
3. You have two very large binary trees containing many nodes. The first tree A contains millions of nodes; the second tree B contains hundreds to thousands. Create an algorithm to decide if B is a subtree of A .

Note: the tree B would be a subtree of A if there exists some node n in A where n and all its children are exactly equivalent to B .

4. What if we have many different trees and want to check each one to see if they're a subtree of A ? How could we do this efficiently? Assume each of these trees contain only a few hundred to thousand elements.
5. Suppose we have a graph containing both directed and undirected edges. If we ignore the undirected edges, we know for certain that the directed edges taken together do *not* introduce acycle.

Your job is to implement an algorithm that assigns each of the undirected edges a direction such the graph remains acycle once there are no more undirected edges.

6. Implement an algorithm that checks whether a graph is a tree.

2 Tech interview questions: general

1. Implement a method named `fizzBuzz` that prints out the numbers from 1 to 100. If the number is a multiple of 3, print out "Fizz" instead of the number. If the number is a multiple of 5, print out "Buzz" instead. If the number is a multiple of 15, print out "Fizzbuzz" instead.

2. Imagine you have a chess knight on a old-school phone dialpad. Given some starting position and a length, how many phone numbers can the knight dial?

For example, if the knight starts on 1 and dials numbers of length 3, it can dial '161', '167', '160', '181', and '183' for a total of 5 numbers. Describe the time and space complexity of your solution.

3. Given some stream of data (let's say ints), we would like to return a (true) random subset of size K of that stream. Assume that the stream is very large – we are not able to store every single int we receive in memory.

4. Suppose we have some steadily incoming stream of data. We want to be able to find the median of all the data we've seen so far as quickly as possible.

How would you do this? Can you figure out a solution to do so in $O(1)$ time?

5. Suppose we want a map that keeps track of not only key-value pairs, but also keeps track of an *expiry time* per each pair. If the expiry time for a pair is past, that pair should no longer be considered a part of the map.
6. Suppose we want to implement an algorithm that accepts an integer and returns the equivalent Roman numerals. For example, the number '597' would be DXCVII in Roman numeral.

For a review of how Roman numerals work, see <https://tinyurl.com/yb736cb5>