

[SOL] Practice Questions for COMP285

Final

The following are questions meant to help you practice, and cannot be submitted for a grade.

Important Notes

- *It is meant to give you a chance to do some practice questions after having reviewed the slides, quizzes, in-class exercises, homeworks, etc.*
- *It should give a rough sense of some ways questions might be posed, though there's no guarantee that the actual final will have the exact same format (at a minimum, one difference is that the actual final will show the point values associated with the questions).*
- *It should give a rough idea of the level of mastery expected generally, though more/less mastery may be expected for any given topic.*

Thanks for reading the notes above - the big picture thing is that I want to be sure you use this resource appropriately, while at the same time **do not neglect the many other more comprehensive resources!**

Measuring Performance

1. What is the worst-case time complexity and space complexity of the below? Remember to provide a tight upper bound with Big-O, and **justify your answer**.

```
algorithm findSumOfUniqueOccurrences
input: vector<int> nums
output: sum of each element in nums that appears exactly once

numOccurrences = new unordered_map of int to int
for each element n in nums
    if numOccurrences.contains(n)
        numOccurrences[n] += 1
    else
        numOccurrences[n] = 1

answer = 0
for each key-value pair k, v in numOccurrences
    if v == 1
        answer += k
return answer
```

Time complexity: n = the size of nums. The first for-loop is $O(n)$ with $O(1)$ work happening inside the loop (the if-statements and key checks / key updates are constant time). The second for-loop goes through all key-value pairs, which, in the worst case, is $O(n)$ as well with $O(1)$ work happening inside. So this is $O(n) + O(n)$ work which simplifies to $O(n)$.

Space Complexity: n = the size of nums. We are creating a new data structure which, in the worst case, will have to store $O(n)$ keys and values. So, our space complexity is $O(n)$.

2. What is the worst case **time complexity** and **space complexity** of the below? Remember to provide a tighter upper bound with Big-O, and **justify your answer**.

```
void doSomething(std::vector<int>& nums) {
    for (int i = 0; i < nums.size(); i++) {
        for (int j = 0; j < 10; j++) {
            std::cout << "hi" << std::endl;
        }
    }

    for (int i = 0; i < 10; i++) {
        for (int j = 0; j < nums.size(); j++) {
            for (int k = j; k < nums.size(); k++) {
                std::cout << "hello" << std::endl;
            }
        }
    }
}
```

}

Time complexity: n = the size of `nums`. For the first double-nested for loop, the innermost $O(1)$ statement runs $O(n * 10)$ times, which simplifies to $O(n)$. The triple-nested for-loop ultimately simplifies to $O(n^2)$. This is because the two innermost loops will run $1 + 2 + 3 + \dots + n - 1 = O(n^2)$ times, and the outerloop multiplies that by 10, but we drop the multiplicative constant. The amount of work done on the inside is constant. So we have $O(n) + O(n^2)$, which simplifies to $O(n^2)$.

Space complexity: n = the size of `nums`. We are not creating any additional space as there are no new data structures created and no recursive stack frames that build up. So, our space complexity is $O(1)$.

3. Bob has analyzed the best-case time complexity for binary search as $O(1)$ because there could be an array with one element in it containing exactly the desired target. If Bob is correct, explain. If not, what is the correct best-case time complexity and what input produces it?

The best-case time complexity is indeed $O(1)$, but it is not shown by the example Bob provided. We need to find the best-case time complexity when run on a large input. So, because binary search starts by looking at the middle element, if the middle element was the target we were searching for (even with a very large array) then we would only do $O(1)$ work. An example of such an input would be $vec = \{1, 2, 3, 4, 5, \dots, 99, 100\}$ with a target = 50.

Trees


4. Order the following data structures from most general to most specific:
Trees, Red-Black Trees, Binary Trees, Graphs, Binary Search Trees

Graphs, Trees, Binary Trees, Binary Search Trees, Red-Black Trees

For questions 5 - 8, refer to the information below.

We want to write `findSumOfTreeLeaves`, which finds the sum of the values of all leaf nodes given the root of a binary tree.

5. What should `findSumOfBinaryTreeLeaves` return when called on the tree below?



$$5 + 5 + 8 = 18$$

6. What would an in-order traversal print out for the tree above?

5, 3, -3, 5, 1, 8

7. Complete `findSumOfBinaryTreeLeaves` below **with recursion**. (Do not worry about the distinction between `->` and `.` for accessing variables / functions in the pseudocode. Also, you can access fields as follows):
- `root.isLeaf` - returns a boolean for whether or not the node `root` is a leaf
 - `root.value` - returns the value of node `root`.
 - `root.left` - returns the left child.
 - `root.right` - returns the right child.

```
algorithm findSumOfBinaryTreeLeaves
```

```
  input: TreeNode root which represents the root of the tree
```

```
  output: the sum of the value of all tree leaves
```

```
if root == nullptr
    return 0
```

```
if root.isLeaf
    return root.value
leftSum = findSumOfBinaryTreeLeaves(root.left)
rightSum = findSumOfBinaryTreeLeaves(root.right)
return leftSum + rightSum
```

8. Assume the function above is run on a balanced **Red-Black tree**. What is the **time complexity** and **space complexity** of this recursive algorithm? Provide a tight upper-bound with Big-O in terms of n (number of nodes in the tree) and justify your answer.

Time Complexity: there will be $O(n)$ recursive calls made (one for each node in the tree) and the work done within each recursive call is $O(1)$ (simple condition checks and addition). The fact that the tree is balanced does not change the fact that we have to visit each node. So, the total time complexity would be $O(n)$.

Space complexity: following the recursion tree, the tallest that the stack frames will build up before beginning to be popped off is proportional to the height of a tree. The height of a balanced binary tree is $O(\log(n))$. Each of the recursive stack frames stores $O(1)$ data. So, the total space complexity would be $O(\log(n))$.

9. Explain whether it is more time efficient to be running "search" on a balanced BST or unbalanced BST. Provide a tight upper-bound with Big-O in terms of n (number of nodes in the tree) for each case.

Search is more efficient on a balanced BST, as in the worst-case, we will have to go all the way down to a leaf, which would be $O(\text{height of tree})$. The height of the tree in the balanced BST case is $\log(n)$ while in a completely unbalanced tree it would be n . So it is more efficient to run search on a balanced BST $O(\log(n))$ rather than not ($O(n)$).

10. Provide a tight upper-bound with Big-O in terms of n (number of nodes in the tree) for the time complexity of the **best-case** for search on a BST. Describe an input for which this happens.

$O(1)$. This can occur on a large BST when the target we are searching for is equivalent to the root.

Graphs

11. Which of the below are true? Select all that apply.

- Bellman-Ford can be used even if there are negative-weight edges.
- Dijkstra's can be used even if there are negative-weight edges.
- When both can apply, we should prefer Dijkstra's over Bellman-Ford for time efficiency.
- Dijkstra's is an example of a Greedy Algorithm.

For questions 12 - 15, refer to the below.

There are n cities, some of which are connected by roads. If there is a road from a to b , and from b to c , then we say that a and c are connected indirectly. A province is a group of directly or indirectly connected cities, and no other cities outside the group are reachable.

12. In order to solve this problem, we can represent this as a graph. What are the nodes and edges?

The nodes are each of the cities and there is an undirected edge from a to b if there is a road connecting cities a and b .

13. We want to find the total number of provinces. How would you solve this problem leveraging graph algorithms we've covered? Explain which algorithm(s) you would use in words AND how you would use it.

We could find the number of connected components using BFS / DFS. We will choose BFS arbitrarily. We increment a global `totalProvinces` counter then call BFS on an unvisited node, marking all nodes reachable from the selected node as visited. We repeat the previous step, continually incrementing `totalProvinces` each time we have to launch a BFS that will span an entire connected component. At the end, we return `totalProvinces` as our answer.


14. Now suppose the roads have associated distances, and we want to find the cost of the shortest path between a city m and n , if it exists. Which algorithm(s) that we've covered would be most appropriate and efficient to solve this problem? Explain which algorithm(s) you would use in words AND how you would use it.

We could run Dijkstra's starting at m and directly use the result corresponding to the cost to reach n at the end of Dijkstra's as our answer.

15. One province leader is interested in building an underground tunnel between the 2 closest cities. Which algorithm(s) that we've covered would be most appropriate and efficient to solve this problem? Explain which algorithm(s) you would use in words AND how you would use it.

We could use an all-pairs shortest path algorithm like Floyd-Warshall to find the pairwise shortest distances between each of the cities. We would then look at the pair of cities with the minimum distance and return this pair of cities as our answer.

MSTs



16. Given the graph above, which of the following are valid *Spanning Trees*? Note that we are not looking for *Minimum* Spanning Trees. Select all that apply.

- A-D, A-F, B-C, C-F
- A-B, A-D, A-F, C-F, E-F
- A-C, B-D, C-F, D-F, E-F
- A-D, A-F, C-F, D-F, E-F
- A-C, A-F, B-D, C-F, D-F, E-F

17. Given the above graph, what's the weight of the **minimum** spanning tree?

2

For the following questions, assume we run the following version of Prim's on the graph above.

algorithm Prim's

Input: Weighted, Undirected, connected Graph $G=(V,E)$ with edge weights w_e

Output: A Tree $T=(V,E')$, with $E' \subseteq E$ that minimizes the edge weight sum

for all $u \in V$:

$\text{cost}(u) = \infty$

$\text{prev}(u) = \text{nil}$

Pick any initial node u_0

$\text{cost}(u_0) = 0$

$\text{unvisited} = \text{makequeue}(V)$ (priority queue, using cost-values as keys)

while unvisited is not empty:

$v = \text{extractmin}(\text{unvisited})$

 for each $\{v, z\} \in E$:

 if $\text{cost}(z) > w(v, z)$ and $\text{unvisited.contains}(z)$:

$\text{cost}(z) = w(v, z)$

$\text{prev}(z) = v$

$\text{decreasekey}(\text{unvisited}, z)$

18. Which of the following will be the correct cost array after two iterations of the while loop?

Assume we start at node D. Select one.

- $\text{cost}(A) = -2, \text{cost}(B) = 1, \text{cost}(C) = \text{inf}, \text{cost}(D) = 0, \text{cost}(E) = \text{inf}, \text{cost}(F) = 9$
- $\text{cost}(A) = -2, \text{cost}(B) = 5, \text{cost}(C) = 3, \text{cost}(D) = 0, \text{cost}(E) = \text{inf}, \text{cost}(F) = 9$
- $\text{cost}(A) = -2, \text{cost}(B) = 1, \text{cost}(C) = 3, \text{cost}(D) = 0, \text{cost}(E) = \text{inf}, \text{cost}(F) = 8$
- $\text{cost}(A) = -2, \text{cost}(B) = 1, \text{cost}(C) = 3, \text{cost}(D) = 0, \text{cost}(E) = \text{inf}, \text{cost}(F) = 9$

19. If you run Prim's algorithm **starting at node A**, list the sequence of nodes you'll remove from the priority queue in order.

A, D, B, C, F, E

20. What's the order of edges visited by Kruskal's? Include edges that are skipped (put "skipped" next to the edge if applicable), and assume we stop after picking $|V|-1$ edges.


E-F, A-D, B-D, A-C, C-F

21. In general, which of the following is true about Kruskal's and Prim's? Select all that apply.


- Kruskal's incorrectly outputs on graphs with negative edges
- Kruskal's correctly outputs on graphs with negative cycles
- Prim's will produce a correct result on graphs with negative cycles
- Prim's and Kruskal's will output different MSTs on the same graph even if they follow the same tie-breaking convention (smallest labeled nodes/edges first)

Network Flow

Consider the flow network below.




22. Draw the edges and weights of the residual graph for the flow network below.





23. Identify an augmenting path in the residual graph if there is one. Note: we are looking for a sequence of vertex labels here that starts with s and ends at t.

s, v₁, v₃, t

24. Update the flows on the original network accordingly with a drawing below.



25. Suppose we are optimizing a toy manufacturing facility. For a toy to be successfully created, it must go through an A-machine which creates the toy and then a B-machine which packages it. We also have conveyor belts that can move toys from A-machines to B-machines. Each conveyor belt has a max speed given in toy(s) per minute (tpm). Each machine also has a max tpm that they produce. Draw a transformed flow network for the toy factory below for which the max flow will tell us how many toys we can produce when running optimally. You do not need to label the vertices, but there must be a source s and a sink t , and all edge capacities must be specified.




26. What does the max flow on the transformed graph represent?


It represents the maximum number of toys that can be produced per minute.

Greedy, Recursion

A kiwi bird is trying to cross a river by hopping across rocks in the stream. The kiwi can hop at most k feet at a time, and rocks are placed at various parts of the stream. For example, if our kiwi can hop at most 3 feet, then it'll take 4 hops to make it across this river on the following set of rocks:



If the rocks were positioned differently, our kiwi wouldn't be able to make it across since it wouldn't be able to jump far enough to safely make it to the next rock (assuming $k = 3$).



We want to write an algorithm named `findMinBirdHops`, which takes in as input the k and vector `<int> rocks` that describe the rock positions and the distance to the far shore, and computes the minimum number of hops needed to cross the river. It should return -1 if it isn't possible for our bird to safely make it to the other side of the river. Here are some examples:

```
// should output 4 (start, rock 3, rock 6, rock 8, end)
```

```
findMinBirdHops(3, {0, 1, 3, 4, 5, 6, 8, 10})
```

```
// should output 3 (start, rock 4, rock 8, end)
```

```
findMinBirdHops(4, {0, 1, 3, 4, 5, 6, 8, 10})
```

```
// should output -1 (after rock 3, no rocks near enough to continue)
```

```
findMinBirdHops(3, {0, 1, 3, 7, 8, 10})
```

Note: The kiwi can always hop less than k feet. The first element in the array is always 0 (the start). The last element in the array should always be the distance to the far side of the river (the end).

27. What should findMinBirdHops output in this case? 6
findMinBirdHops(3, {0, 1, 3, 4, 7, 8, 11, 12})

28. What should birdHops output in this case? 1
findMinBirdHops(5, {0, 4})

29. What should birdHops output in this case? 0
findMinBirdHops(3, {0, 4})

30. Describe a greedy algorithm that solves this problem. You do not need to write code, but we will expect you to be detailed enough so that we could design an algorithm from your response. Please include how you would update and return the minimum # of hops. Keep track of minHops as a variable. At each step, hop as far forward as much as possible, ensuring that we can (if we cannot, return -1). Every time we hop forward, increment minHops. Afterward, we return minHops as our final answer.

31. Now let's try to write a brute-force recursive version. Given the following code, fill in the blanks so that it will correctly compute the minimum hops.

```
algorithm findBirdHops
  input: int k, vector<int> of N rocks
  output: optimal # of bird hops
  return birdHopsHelper(_____, _____, _____) // k, rocks, 0
```

```
algorithm findBirdHopsHelper
  input: int k, vector<int> of N rocks, int index
  output: <left unspecified for this question>
  if _____ // N - 1 == index
    return 0
  minHops = MAX_INT // predefined constant
  for i = index + 1, index + 2, ... N - 1
    if rocks[index] + k >= rocks[i]
      result = _____ // findBirdHopsHelper(k, rocks, i)
      if result != -1
        minHops = min(minHops, 1 + result)
  if minHops == MAX_INT
    return _____ // -1
  return minHops
```


Greedy, Dynamic Programming

Consider the following problem:

Imagine you had a vending machine that dispensed n different types of snacks. Snack i has $\text{calories}[i]$ calories and costs $\text{price}[i]$ dollars. Assume the vending machine will never run out of any snack (i.e. you can buy as many of one type of snack as you would like). Given X dollars, what's the most amount of calories you can buy from the machine?

For example, given the following array `calories` and array `price`, if $X = 10$, then the total calories we could buy is 48 (buy the $i=0$ snack once and the $i=3$ snack twice).

```
calories = {30, 14, 16, 9}
price =    {6, 3, 4, 2}
```

32. Consider the following greedy algorithm:

Pick the snack that has the best calorie to price ratio and buy as many of them as possible until you no longer have enough money to do so. Then buy the snack with the next best ratio as much as possible. Each time, add the total calories gained from the purchases to a `totalCalories` variable. Repeat this process until you don't have enough money to buy any snack. Return `totalCalories`. The calorie to price ratio of snack i is $\text{calories}[i]/\text{price}[i]$.

Given `calorie` and `price` arrays above, what would this algorithm output when $X = 10$?

44. We get here by picking snack 0 first since it has 5 calories/\$ and then snack 1 with 4.66 calories/\$.

33. Now consider another greedy algorithm:

Pick the snack that has the most calories and buy as many of them as possible until you no longer have enough money to do so. Then buy the snack with the next most calories as much as possible. Each time, add the total calories gained from the purchases to a `totalCalories` variable. Repeat this process until you don't have enough money to buy any snack. Return `totalCalories`.

Given `calorie` and `price` arrays above, what would this algorithm output when $X = 10$?

46. We get here by picking snack 0 first then snack 3.

34. Using the answers from above, describe why both greedy approaches won't work in this case. (Proof by counterexample)

Since neither greedy algorithm produces the optimum solution based on the above, we cannot say that either greedy approach will work.

Now let's examine a Dynamic Program algorithm to solve this problem.

35. Let's say we initialize an array of zeros called DP of length X+1. How should we define DP[j]? We are looking for a sentence which describes what DP[j] represents in words.

- DP[j] is the most amount of calories you can get by buying only snack j
- DP[j] is the most amount of calories you can get with j dollars
- DP[j] is the calories earned if you add the calories of all snacks from 0 to j
- DP[j] is the total price of buying one of each snack from 0 to j
- None of the above

36. What should our base case(s) be?

DP[0] = 0. We should also initialize everything else to 0 in our table before we solve the problem bottom-up.

37. Describe the recursive subproblem for this problem in the blank below.

DP[j] = max(DP[j], calories[i] + DP[j-price[i]])
for all i from 0 to n-1

(Recall that n is the # of types of snacks, aka price.size())

38. What special conditions, if any, do we need to watch out for? What order should we solve the subproblems in, and what will we return as our final answer?

If j - price[i] is out-of-range for a given i, we do not want to include that operation. This can be done with an if-statement in the code.

We should solve the subproblems from 0, 1, ... X. We return DP[X] as our final answer.

39. What is the time complexity of this DP approach?


The number of subproblems is O(x) where x is the value of X. The time per sub-problem is O(n) where n is the size of price and calories. So, our total runtime is O(n * x).

40. We can also solve DP questions top-down with memoization. In your own words, explain what memoization is and when it is useful.

Memoization is the process by which we save results from previously computed subproblems for re-use later. It is only useful if we have overlapping subproblems in our recursive tree.

Otherwise, there are no savings to be gained from saving results that will only be used once.

Dynamic Programming



Consider the following problem:

A robot is located at the top-left corner of a $m \times n$ grid (marked 'Start' in the diagram below). The robot can only move either down or right at any point in time. The robot is trying to reach the bottom-right corner of the grid (marked 'Finish' in the diagram below). But, there are some squares which are magnet traps (gray circles) that the robot cannot step on. For inputs, we are given the dimensions m and n , as well as a vector $\langle \text{pair} \langle \text{int}, \text{int} \rangle \rangle$ representing the coordinates of all magnet traps. How many possible unique paths are there?

41. Let's say we initialize a 2D table of size m by n . How should we define $DP[i][j]$ in words?

$DP[i][j]$ is the number of paths we can take to reach coordinate i, j on the grid.

42. What should our base case(s) be?

$DP[i][0] = 0$ for $i = 0 \dots m$

$DP[0][j] = 0$ for $j = 0 \dots n$

$DP[i][j] = 0$ for all i, j in our magnet traps list (or this can be taken care of in later answers).

43. Describe the recursive subproblem for this problem.

$DP[i][j] = DP[i-1][j] + DP[i][j-1]$

44. What special conditions, if any, do we need to watch out for? What order should we solve the subproblems in and what should we return as our final answer?

We should not update magnet squares with the recurrence above, as there should be no way to reach them (see base case).


We should solve the subproblems row-by-row or column-by-column.

We should return $DP[m-1][n-1]$ as our final answer.

45. What is the time complexity of this DP approach?

The number of subproblems is $O(m * n)$ and each subproblem is just an addition which takes $O(1)$ time. So our total runtime is $O(m * n)$.

Brute Force



Before smartphones, people would text with a keypad. There was a predictive technology called “T9” which would allow you to press a key once for any letter that might be on that number, and T9 would show you some word options. For example, pressing 2, 2, 8 might show “bat”, “cat”, or “act”.

Given a vector<int> of digits (2 - 9), we want to return all possible letter combinations that the number could represent (including non-valid dictionary words).

Example:

{2, 3} → {"ad", "ae", "af", "bd", "be", "bf", "cd", "ce", "cf"}

46. Complete the below (pseudocode is fine).

```
unordered_map<int, vector<string>> DIGIT_TO_LETTERS = {
    {2, {"a", "b", "c"}},
    {3, {"d", "e", "f"}},
    {4, {"g", "h", "i"}},
    {5, {"j", "k", "l"}},
    {6, {"m", "n", "o"}},
    {7, {"p", "q", "r", "s"}},
    {8, {"t", "u", "v"}},
    {9, {"w", "x", "y", "z"}},
};

void letterCombinationsHelper(vector<int> digits,
                             int i, const string currAnswer,
                             vector<string>& answers) {
    if (digits.size() == i) {
        answers.push_back(currAnswer);
        return;
    }
}
```

```
for each letter in DIGIT_TO_LETTERS[digits[i]]
  // choose / unchoose not needed because we make a copy of the string
  letterCombinationsHelper(digits, i + 1, currAnswer + letter, answers) // explore
```

```
}
```

```
vector<string> letterCombinations(vector<int> digits) {
  vector<string> answers = {};
  letterCombinationsHelper( digits, 0, "", answers )
  return answers;
}
```

47. Suppose we have a function `prefixExists(const string currAnswer)` which checks an `unordered_set` in $O(1)$ time to see whether `currAnswer` is the prefix of a valid English word. Would you do this check after the base case of `letterCombinationsHelper` or filter the results at the end of `letterCombinations` before returning the answer? Justify your answer.

We should check after the base case, so that we can prune the search space as we are exploring to actually improve the runtime in practice. If we filter afterwards, we do not save any time.

P, NP, and More

48. Given a problem that does not have a polynomial time verifier, which complexity class must this problem belong to? Select all that apply.

- P
- NP
- NP-Complete
- None of the Above

49. If a problem is in NP, which of the following COULD be true? Select all that apply.

- The solution to this problem could be verifiable in polynomial time
- The problem could be solved in polynomial time
- The problem can only be verified in exponential time

50. If $P = NP$, which of the following statements are true? Select all that apply.

- The best algorithms for NP problems will be exponential.
- All problems in NP will be solvable in polynomial time.
- RSA Encryption would be broken