

CS 270 (Spring 2022) — Assignment 4

Due: Friday, September 30, by 23:00

[Important Note: To allow us to return sample solutions two days before the midterm, this assignment cannot be submitted more than three (3) days late, i.e., no later than Monday, October 3, by 23:00.]

General Instructions. The following assignment is meant to be challenging, and we anticipate that it will take most of you at least 10–15 hours to complete, so please allow yourself plenty of time to work on it. (This assignment may be a bit more substantial than the third; in particular, Question 3 is not easy.) We highly recommend reading the entire assignment right away — you never know when inspiration will strike. Please provide a formal mathematical proof for all your claims, and present runtime guarantees for your algorithms using asymptotic (big- $O/\Omega/\Theta$) notation, unless stated otherwise. You may assume that all basic arithmetic operations (multiplication, subtraction, division, comparison, etc.) take constant time.

Collaboration. Please carefully check the collaboration policy on the course website. When in doubt, ask an instructor.

Consulting outside sources. Please carefully check the policy regarding outside sources on the course website. Again, when in doubt, ask an instructor.

Submission. Homework submission will be through the Gradescope system. Instructions and links have been provided through the course website and Piazza. The only accepted format is PDF. Starting with this homework, only typed solutions will be accepted, and we highly recommend producing your solutions using \LaTeX (the text markup language we are also using for this assignment).

Recommended practice problems (do not hand in): KT Problems 5.1, 5.3, 5.5, 5.6.

(1) [1+1+1+2+2=7 points]

Solve the following recurrences. In all cases, assume $T(1)$ is some positive constant. You should be able to find a matching upper and lower bound on $T(n)$. In other words, your answer should be of the form $T(n) = \Theta(\dots)$. You do not need to provide formal proofs, just a brief justification for your answer. As we often do in class, you may sacrifice formality by ignoring issues with rounding integers; in all these recurrences, those issues will not matter.

(a) $T(n) = 9T(n/3) + c_1 \cdot n^2$

(b) $T(n) = 10T(n/5) + c_2 \cdot n^{1.5}$

(c) $T(n) = 3T(n/4) + c_3 \cdot \sqrt{n}$

(d) $T(n) = T(n/3) + T(\frac{2}{3}n) + c_4 \cdot n$

(e) $T(n) = T(n/4) + T(n/2) + c_5 \cdot n$

(2) [3+7=10 points]

Consider an array of integers $A = [a_1, \dots, a_n]$ with $a_1 < a_n$.

(a) Show that there exists an index $i \in \{1, \dots, n-1\}$ such that $a_i < a_{i+1}$. In other words, there exist two consecutive entries of the array of which the first is smaller than the second.

(b) Give and analyze an algorithm which finds such a pair of entries in $O(\log n)$ time. Your algorithm should output $i \in \{1, \dots, n-1\}$ such that $a_i < a_{i+1}$. You may assume that the array A has already been loaded into memory. You may also assume that comparing two entries of the array is a primitive operation (i.e., takes constant time).

[Hint: The first part may be useful here.]

(3) [10 points]

Consider the following situation. At the end of the semester, David has two arrays a, b of weighted average scores of the students in the two sections of CSCI 270. Let's say that both sections have the same size n , which you may assume is a power of 2 if it simplifies things for you.

He has already sorted the students by score, so $a[1] \leq a[2] \leq \dots \leq a[n]$, and $b[1] \leq b[2] \leq \dots \leq b[n]$. However, he has not yet done any comparison of students between the two sections: it is possible that all numbers in the a array are larger, or all numbers in the b array are larger, or they are mixed in any possible way.

David now wants to know how the "most middle" student performed. That is, he wants to find a student i (out of the $2n$ total) such that at least n students had scores no higher than i , and at least n students had scores no lower than i . There may be multiple such students; for instance, if all students achieved the same score, then any student would do. Give an algorithm that always runs in time $O(\log n)$ and finds one such student, by outputting the array (a or b) and the index i where the student occurs in that array.

(4) [0 points]

Chocolate Problem: 1 chocolate bar

Reminder: If you solve a chocolate problem (which you can do in groups of size up to 3), please e-mail David with the solution — do not submit it on Gradescope. Also, feel free to list preferences or dietary restrictions for/against particular types of chocolate.

In discussion section, you saw how to generalize the idea of Binary Search to trees: each tree has a node v such that when you query that node and the answer points to one of the subtrees from that node, there are at most $n/2$ nodes in that subtree. Thus, repeating this querying at most $\log_2(n)$ times, you can find any node in a tree from the answers to such queries.

Here, we want to generalize the idea further, from trees to arbitrary undirected graphs. In non-tree graphs, it does not make sense to talk about "the subgraph containing the node". For instance, if your graph is a cycle, then no matter which node you query, you can go around the cycle in either direction to get to any other node.

So we will clarify the answer as follows: when a node v is queried, and the correct answer is t , the answer will reveal an edge out of v that lies on a *shortest* path from v to t . If there are multiple shortest paths from v to t (with different edges out of v), then any of them could be returned.

You now play the following game: both you and your opponent know the undirected graph G . Your opponent picks a node $t \in G$, without telling you what it is. In each round, you get to point to a vertex v . If $v = t$, then the game is over. Otherwise, your opponent reveals an edge e incident on v that lies on a shortest path from v to t . The game repeats until you've found t . Your goal is to get there with few queries. Prove the following:

Lemma 1. *For every graph $G = (V, E)$, there exists a vertex v you can query such that never mind which edge e incident on v your opponent reveals, the set of remaining vertices that are consistent with this answer shrinks by at least a factor 2. That is, if S is the set of all nodes t such that e is on a shortest path from v to t , then $|S| \leq |V|/2$.*

To round out the answer, show how to use the lemma to guarantee that you can find the node t in at most $\log_2 n$ queries — this part is basically trivial once you have proved the lemma.