

Due on February 1, 2017

60 points total

General Directions: If you are asked to provide an algorithm, you should clearly define each step of the procedure, establish its correctness, and then analyze its overall running time (for this assignment, this means arguing why your algorithm achieves the target running time specified by the question). There is no need to write pseudo-code; an unambiguous description of your algorithm in plain text will suffice.

Please make sure that you **START EARLY**. Some of these questions might take you several hours to solve. Since this is a mathematical/theoretical class, the ratio between how much time you spend thinking and how long it takes you to write your solution will be high (at least higher than what you might be used to in coding assignments).

All the answers must be typed, preferably using LaTeX. If you are unfamiliar with LaTeX, you are strongly encouraged to learn it. However, answers typed in other text processing software and properly converted to a pdf file will also be accepted. Before submitting the pdf file, please make sure that it can be opened using any standard pdf reader (such as Acrobat Reader) and your entire answer is readable. **Handwritten answers or pdf files that cannot be opened will not be graded and will not receive any credit.**

Finally, please read the detailed collaboration policy given on the course website. You are **not** allowed to discuss homework problems in groups of more than 3 students. **Failure to adhere to these guidelines will be promptly reported to the relevant authority without exception.**

Problem 1 (25 points)

In class, we learned a linear time algorithm that finds the median of an array of numbers. We now wish to solve a more general problem using this algorithm. The input is an (unsorted) array of numbers x_1, x_2, \dots, x_n , each with an associated positive weight (we denote the weight of x_i by w_i) such that the sum of weights is 1 (i.e., $\sum_{i=1}^n w_i = 1$). Suppose ϕ defines an ordering of indices such that $x_{\phi(1)}, x_{\phi(2)}, \dots, x_{\phi(n)}$ is sorted. The *weighted median* of the numbers is the value $x_{\phi(m)}$ such that $\sum_{i=1}^{m-1} w_{\phi(i)} < 0.5$ and $\sum_{i=1}^m w_{\phi(i)} \geq 0.5$.

- (a) (10 points) Give an algorithm that finds the weighted median of an array in $O(n)$ time. You may assume that you have an algorithm to find the median of an array (without weights) in $O(n)$ time (e.g., the algorithm you learned in class). This algorithm can be used as a “black box” (like a function call) and need not be described in your solution. Remember that the array is unsorted and sorting the array at any point during the algorithm will take $\Omega(n \log n)$ time.
- (b) (5 points) Alice Algorithmix lives on a long, narrow island somewhere off the coast of North Carolina where all the cities on the island are on a single straight line. The island government has decided to build an airport that will serve all the cities and has asked Alice to come up with a location that would minimize the average travel time to the airport for the residents of the island. (The travel time of a resident is proportional to her distance from the airport.) Can you help Alice solve the problem in $O(n)$ time if there are n cities? You can assume that Alice gives you the locations of the cities and their respective populations.
- (c) (10 points) What if instead of Alice Algorithmix, you were asked this question by Bob Bitfiddler, who is a student at Duke and wants to find a location for an airport that will serve all residents of Durham county? You are given the coordinates of all neighborhoods in the county and their respective populations. As earlier, the goal is to minimize the average travel time of residents to the airport, where the travel time is now proportional to the “Manhattan distance” defined as $|x_1 - x_2| + |y_1 - y_2|$ for two points (x_1, y_1) and (x_2, y_2) . Your algorithm should run in $O(n)$ time if there are n neighborhoods.

Hint: The last two parts are related to the first part. You can use an algorithm you designed for the first part as a black box in the last two parts.

Problem 2 (35 points)

Since you managed to solve her previous problem so successfully, your old friend Alice Algorithmix asks for your help once again.

- (a) (10 points) This time, she has been contracted by the department of geological survey who want to find the highest altitude in the island. Since this is essentially a volcanic island and a narrow one at that, the altitude data that Alice has is just an array of n numbers that increase till some index and decrease thereafter. Formally, if x_1, x_2, \dots, x_n are the altitudes, there

exists some index i such that $x_j > x_{j-1}$ for all $j \leq i$ and $x_j < x_{j-1}$ for all $j > i$. Clearly, x_i is the altitude that Alice needs to find. Can you help her find this index i in $O(\log n)$ time?

- (b) (25 points) If Alice comes, can Ben be far behind? Now, our beloved Duke student Bob Bitfiddler has a $n \times n$ 2-D array on his hands which gives the altitudes of all points in Durham county (let's assume Durham is square). But, Durham is not quite a volcano! So, the altitudes form an arbitrary matrix of numbers. Bob would have ideally liked to find the highest point in Durham, but since he does not have the nice property that Alice had, he settles for an easier target: find a *locally* highest point, i.e., an entry in the $n \times n$ 2-D array that is as large as both its horizontal neighbors *and* both its vertical neighbors. (Bob hopes that this will at least be a hillock if not a mountain!) Formally, M is an $n \times n$ array, then $M[i, j]$ is a locally highest point if $M[i, j] \geq \max(M[i-1, j], M[i+1, j], M[i, j-1], M[i, j+1])$. Note that an entry on the boarder of the grid (which will only have 2 or 3 valid vertical/horizontal neighbors) just needs to be as large as its valid neighbors in order to be considered a locally highest point. Can you help Bob find such an entry in $O(n)$ time? (An $O(n \log n)$ solution will get partial credit.)