

TUTORIAL SHEET 6

1. **[KT-Chapter6]** Suppose we want to replicate a file over a collection of n servers, labeled S_1, S_2, \dots, S_n . To place a copy of the file at server S_i results in a placement cost of c_i , for an integer $c_i > 0$. Now, if a user requests the file from server S_i , and no copy of the file is present at S_i , then the servers $S_{i+1}, S_{i+2}, S_{i+3}, \dots$ are searched in order until a copy of the file is finally found, say at server S_j , where $j > i$. This results in an access cost of $j - i$. (Note that the lower-indexed servers S_{i-1}, S_{i-2}, \dots are not consulted in this search.) The access cost is 0 if S_i holds a copy of the file. We will require that a copy of the file be placed at server S_n , so that all such searches will terminate, at the latest, at S_n . We would like to place copies of the files at the servers so as to minimize the sum of placement and access costs. Formally, we say that a configuration is a choice, for each server S_i with $i = 1, 2, \dots, n - 1$, of whether to place a copy of the file at S_i or not. (Recall that a copy is always placed at S_n .) The total cost of a configuration is the sum of all placement costs for servers with a copy of the file, plus the sum of all access costs associated with all n servers. Give a polynomial-time algorithm to find a configuration of minimum total cost.
2. **[KT-Chapter6]** Suppose we are given a directed graph $G = (V, E)$, with costs on edges – the costs may be positive or negative, but every cycle in the graph has positive cost. We are also given two nodes v and w in the graph G . Give an efficient algorithm to compute the number of shortest $v - w$ paths in G (the algorithm should NOT list the paths; it should just output the number of such paths).
3. **[KT-Chapter6]** Consider the following inventory problem. You are running a store that sells some large product (let us assume you sell trucks), and predictions tell you the quantity of sales to expect over the next n months. Let d_i denote the number of sales you expect in month i . We will assume that all sales happen at the beginning of the month, and trucks that are not sold are stored until the beginning of the next month. You can store at most S trucks, and it costs C to store a single truck for a month. You receive shipments of trucks by placing orders for them, and there is a fixed ordering fee of K each time you place an order (regardless of the number of trucks you order). You start out with no trucks. The problem is to design an algorithm that decides how to place orders so that you satisfy all the demands $\{d_i\}$, and minimize the costs. In summary:
 - There are two parts to the cost. First, storage: it costs C for every truck on hand that is not needed that month. Second, ordering fees: it costs K for every order placed.
 - In each month you need enough trucks to satisfy the demand d_i , but the amount left over after satisfying the demand for the month should not exceed the inventory limit S .

Give an algorithm that solves this problem in time that is polynomial in n and S .

4. **[Dasgupta, Papadimitriou, Vazirani -Chapter6]** You are given a string of n characters $s[1..n]$, which you believe to be a corrupted text document in which all punctuation has vanished (so that it looks something like “itwasthebestoftimes...”). You wish to reconstruct the document using a dictionary, which is available in the form of a Boolean function $\text{dict}()$: for any string w , $\text{dict}(w)$ outputs true if w is a valid word false otherwise. Give a dynamic programming algorithm that determines whether the string $s[]$ can be reconstituted as a sequence of valid words. The running time should be at most $O(n^2)$, assuming each call to $\text{dict}()$ takes unit time.
5. **[Dasgupta, Papadimitriou, Vazirani -Chapter6]** We are given a checkerboard which has 4 rows and n columns, and has an integer written in each square. We are also given a set of $2n$ pebbles, and we want to place some or all of these on the checkerboard (each pebble can be placed on exactly one square) so as to maximize the sum of the integers in the squares that are covered by pebbles. There is one constraint: for a placement of pebbles to be legal, no two of them can be on horizontally or vertically adjacent squares (diagonal adjacency is fine). Give an $O(n)$ time algorithm to find an optimal placement of the pebbles.
6. **[KT-Chapter6]** Suppose you’re consulting for a small computation-intensive investment company, and they have the following type of problem that they want to solve over and over. A typical instance of the problem is: they are doing a simulation in which they look at n consecutive days of a given stock, at some point in the past. Let us number the days $i = 1, 2, \dots, n$. For each day i , they have a price $p(i)$ per share for the stock on that day. (We will assume for simplicity that the price was fixed during each day.) Suppose during this time period, they wanted to buy 1000 shares on some day, and sell all these shares on some (later) day. They want to know: when should they have bought and when should they have sold in order to have made as much money as possible? (If there was no way to make money during the n days, you should report this instead.)

Example: Suppose $n = 3, p(1) = 9, p(2) = 1, p(3) = 5$. Then you should return “buy” on 2, sell on “3”; i.e. buying on day 2 and selling on day 3 means they would have made \$4 per share, the maximum possible for that period. Clearly, there is a simple algorithm that takes time $O(n^2)$: try all possible pairs of buy/sell days and see which makes them the most money. Your investment friends were hoping for something a little better. Show how to find the correct numbers i and j in time $O(n)$.