

CS 473: Fundamental Algorithms, Spring 2013

Discussion 7

March 6, 2013


7.1. WEIGHTED SCHEDULING.

We have n jobs J_1, J_2, \dots, J_n which we need to schedule on a machine. Each job J_i has a processing time t_i and a weight w_i . A schedule for the machine is an ordering of the jobs. Given a schedule, let C_i denote the finishing time of job J_i . For example, if job J_j is the first job in the schedule, its finishing time C_j is equal to t_j ; if job J_j follows job J_i in the schedule, its finishing time C_j is equal to $C_i + t_j$. The weighted completion time of the schedule is $\sum_{i=1}^n w_i C_i$.

- (A) For the case when $w_i = 1$ for all i , show that choosing the shortest job first is optimal.
- (B) Give an efficient algorithm that finds a schedule with minimum weighted completion time given arbitrary weights.

7.2. MINIMUM SPANNING TREE.

Consider the following graph:



- (A) Draw the edges in the Minimum Spanning Tree for the following graph.
- (B) Given G and MST T , suppose you decrease the weight of an edge e not in T . Give an algorithm to recompute the MST in $O(n)$ time.

7.3. BORŮVKA'S ALGORITHM.

Borůvka's algorithm computes the MST of a graph $G = (V, E)$, by repeatedly picking for every vertex in the graph the cheapest edge adjacent to it. Let $F \subseteq E$ be the set of edges picked by this process. The Borůvka's algorithm then collapse every connected component of (V, F) into a single vertex. It continues this process iteratively till remaining with a single vertex. The set of edges picked formed the required MST.

Formally, the collapsing of the graph is done as follows: An edge in the original graph that connects two vertices in the same connected component disappears in the new graph. And edge of the original graph that connects two different connected components, now connects the two respective connected components. Naturally, if there are several edges connecting the same pair of connected components, we remember only the cheapest one.

- (A) Show how to compute the collapsed graph in linear time (i.e., $O(|V| + |E|)$), for any set of edges $F \subseteq E$.
- (B) Show that Borůvka's algorithm decreases the number of vertices by two at each iteration.
- (C) Conclude that Borůvka's algorithm takes $O((n + m) \log n)$ time in the worst case. Why is the running time not $O(n \log n + m)$?

7.4. MST UNDERSTANDING

Let G be an undirected graph with m edges and n vertices with weights on the edges.

- (A) You are given a minimum spanning tree T of G . The weight of one edge e of the graph had changed (the edge might be an edge of T). Describe an $O(n + m)$ time algorithm that computes the MST of G with the updated weights.
- (B) You are given a minimum spanning tree T of G . For social reasons that are still not well understood, Vogon children just broke into your house and stole k edges of T . Describe an algorithm to compute an MST for the graph G without these k edges. The running time of your algorithm should be $O(k \log k + m)$.
- (C) You are given a minimum spanning tree T of G . The weight of k edges of G (that are not in T) had been suddenly decreased. Describe an $O((n + k) \log n)$ time algorithm that computes the MST of the new graph.
- (D) You are given a graph G and a minimum spanning tree T . The *max-price* of a path π is the price of the most expensive edge on π . Describe an algorithm, as efficient as possible, for computing the minimum max-price path between two given vertices x and y of G . (For full credit, your algorithm should work in $O(n)$ time.)
Prove the correctness of your algorithm.
(And no, you can not use hashing in the solution for this question.)

7.5. STOCK PICKING.

You have a group of investor friends who are looking at n consecutive days of a given stock at some point in the past. The days are numbered $i = 1, 2, \dots, n$. For each day i , they have a price $p(i)$ per share for the stock on that day.

For certain (possibly large) values of k , they want to study what they call *k-shot strategies*. A *k-shot strategy* is a collection of m pairs of days $(b_1, s_1), \dots, (b_m, s_m)$, where $0 \leq m \leq k$ and

$$1 \leq b_1 < s_1 < b_2 < s_2 \cdots < b_m < s_m \leq n.$$

We view these as a set of up to k nonoverlapping intervals, during each of which the investors buy 1,000 shares of the stock (on day b_i) and then sell it (on day s_i). The *return* of a given *k-shot strategy* is simply the profit obtained from the m buy-sell transactions, namely,

$$1000 \cdot \sum_{i=1}^m (p(s_i) - p(b_i)).$$

- (A) Design an efficient algorithm that determines, given the sequence of prices, the *k-shot strategy* with the maximum possible return. Since k may be relatively large, your running time should be polynomial in both n and k .

(B) Now, modify your algorithm to only use $O(n)$ space.

7.6. SET COVER AND THE GREEDY ALGORITHM.

Given a set U , and a family of subsets \mathcal{F} , the *set cover* problem asks for the minimum number of sets in \mathcal{F} that fully cover U . For example, here is an instance of set cover, with the ground set being

$$U = \{1, 2, 3, 4\}.$$

and the family of sets being

$$\mathcal{F} = \left\{ A = \{1\}, B = \{1, 3\}, C = \{2, 4\}, D = \{2\}, E = \{3\}, F = \{4\} \right\}.$$

Valid covers of U might be A, D, E, F as $A \cup D \cup E \cup F = U$, and it is in this case of size 4. In this specific case, the best set cover is B, C as it is made of two sets of \mathcal{F} .

The greedy algorithm for set cover always picks the set covering the largest number of elements not covered yet into the set cover. Show an example, where this greedy algorithm fails. Show an example where for a set of size n , the greedy algorithm outputs a set of size $\Theta(\log n)$, by the optimal cover is made out of two sets!