

# CSE 421 Intro to Algorithms Autumn 2017

## Homework 3

Due Friday October 20, 4:00 pm

### Problem 1:

A small business—say, a photocopying service with a single large machine—faces the following scheduling problem. Each morning they get a set of jobs from customers. They want to do the jobs on their single machine in an order that keeps their customers happiest. Customer  $i$ 's job will take  $t_i$  time to complete. Given a schedule (i.e., an ordering of the jobs), let  $C_i$  denote the finishing time of job  $i$ . For example, if job  $j$  is the first to be done, we would have  $C_j = t_j$ ; and if job  $j$  is done right after job  $i$ , we would have  $C_j = C_i + t_j$ . Each customer  $i$  also has a given weight  $w_i$  that represents his or her importance to the business. The happiness of customer  $i$  is expected to be dependent on the finishing time of  $i$ 's job. So the company decides that they want to order the jobs to minimize the weighted sum of the completion times,  $\sum_{i=1}^n w_i C_i$ .

Design an efficient algorithm to solve this problem. That is, you are given a set of  $n$  jobs with a processing time  $t_i$  and a weight  $w_i$  for each job. You want to order the jobs so as to minimize the weighted sum of the completion times,  $\sum_{i=1}^n w_i C_i$ .

**Example.** Suppose there are two jobs: the first takes time  $t_1 = 1$  and has weight  $w_1 = 10$ , while the second job takes time  $t_2 = 3$  and has weight  $w_2 = 2$ . Then doing job 1 first would yield a weighted completion time of  $10 \cdot 1 + 2 \cdot 4 = 18$ , while doing the second job first would yield the larger weighted completion time of  $10 \cdot 4 + 2 \cdot 3 = 46$ .

**Problem 2:**

Consider the following variation on the Interval Scheduling Problem. You have a processor that can operate 24 hours a day, every day. People submit requests to run *daily jobs* on the processor. Each such job comes with a *start time* and an *end time*; if the job is accepted to run on the processor, it must run continuously, every day, for the period between its start and end times. (Note that certain jobs can begin before midnight and end after midnight; this makes for a type of situation different from what we saw in the Interval Scheduling Problem.)

Given a list of  $n$  such jobs, your goal is to accept as many jobs as possible (regardless of their length), subject to the constraint that the processor can run at most one job at any given point in time. Provide an algorithm to do this with a running time that is polynomial in  $n$ . You may assume for simplicity that no two jobs have the same start or end times.

**Example.** Consider the following four jobs, specified by (*start-time*, *end-time*) pairs.

(6 P.M., 6 A.M.), (9 P.M., 4 A.M.), (3 A.M., 2 P.M.), (1 P.M., 7 P.M.).

The optimal solution would be to pick the two jobs (9 P.M., 4 A.M.) and (1 P.M., 7 P.M.), which can be scheduled without overlapping.

**Problem 3:**

A group of network designers at the communications company CluNet find themselves facing the following problem. They have a connected graph  $G = (V, E)$ , in which the nodes represent sites that want to communicate. Each edge  $e$  is a communication link, with a given available bandwidth  $b_e$ .

For each pair of nodes  $u, v \in V$ , they want to select a single  $u$ - $v$  path  $P$  on which this pair will communicate. The *bottleneck rate*  $b(P)$  of this path  $P$  is the minimum bandwidth of any edge it contains; that is,  $b(P) = \min_{e \in P} b_e$ . The *best achievable bottleneck rate* for the pair  $u, v$  in  $G$  is simply the maximum, over all  $u$ - $v$  paths  $P$  in  $G$ , of the value  $b(P)$ .

It's getting to be very complicated to keep track of a path for each pair of nodes, and so one of the network designers makes a bold suggestion: Maybe one can find a spanning tree  $T$  of  $G$  so that for *every* pair of nodes  $u, v$ , the unique  $u$ - $v$  path in the tree actually attains the best achievable bottleneck rate for  $u, v$  in  $G$ . (In other words, even if you could choose any  $u$ - $v$  path in the whole graph, you couldn't do better than the  $u$ - $v$  path in  $T$ .)

This idea is roundly heckled in the offices of CluNet for a few days, and there's a natural reason for the skepticism: each pair of nodes might want a very different-looking path to maximize its bottleneck rate; why should there be a single tree that simultaneously makes everybody happy? But after some failed attempts to rule out the idea, people begin to suspect it could be possible.

Show that such a tree exists, and give an efficient algorithm to find one. That is, give an algorithm constructing a spanning tree  $T$  in which, for each  $u, v \in V$ , the bottleneck rate of the  $u$ - $v$  path in  $T$  is equal to the best achievable bottleneck rate for the pair  $u, v$  in  $G$ .

**Problem 4 (Extra Credit):**

Let's go back to the original motivation for the Minimum Spanning Tree Problem. We are given a connected, undirected graph  $G = (V, E)$  with positive edge lengths  $\{\ell_e\}$ , and we want to find a spanning subgraph of it. Now suppose we are willing to settle for a subgraph  $H = (V, F)$  that is "denser" than a tree, and we are interested in guaranteeing that, for each pair of vertices  $u, v \in V$ , the length of the shortest  $u-v$  path in  $H$  is not much longer than the length of the shortest  $u-v$  path in  $G$ . By the *length* of a path  $P$  here, we mean the sum of  $\ell_e$  over all edges  $e$  in  $P$ .

Here's a variant of Kruskal's Algorithm designed to produce such a subgraph.

- First we sort all the edges in order of increasing length. (You may assume all edge lengths are distinct.)
- We then construct a subgraph  $H = (V, F)$  by considering each edge in order.
- When we come to edge  $e = (u, v)$ , we add  $e$  to the subgraph  $H$  if there is currently no  $u-v$  path in  $H$ . (This is what Kruskal's Algorithm would do as well.) On the other hand, if there is a  $u-v$  path in  $H$ , we let  $d_{uv}$  denote the length of the shortest such path; again, length is with respect to the values  $\{\ell_e\}$ . We add  $e$  to  $H$  if  $3\ell_e < d_{uv}$ .

In other words, we add an edge even when  $u$  and  $v$  are already in the same connected component, provided that the addition of the edge reduces their shortest-path distance by a sufficient amount.

Let  $H = (V, F)$  be the subgraph of  $G$  returned by the algorithm.

- (a) Prove that for every pair of nodes  $u, v \in V$ , the length of the shortest  $u-v$  path in  $H$  is at most three times the length of the shortest  $u-v$  path in  $G$ .
- (b) Despite its ability to approximately preserve shortest-path distances, the subgraph  $H$  produced by the algorithm cannot be too dense. Let  $f(n)$  denote the maximum number of edges that can possibly be produced as the output of this algorithm, over all  $n$ -node input graphs with edge lengths. Prove that

$$\lim_{n \rightarrow \infty} \frac{f(n)}{n^2} = 0.$$