

# CSE 421 Intro to Algorithms Autumn 2017

## Homework 1

Due Friday October 6, 4:00 pm

### Problem 1:

Suppose we have two television networks, whom we'll call  $\mathcal{A}$  and  $\mathcal{B}$ . There are  $n$  prime-time programming slots, and each network has  $n$  TV shows. Each network wants to devise a *schedule*—an assignment of each show to a distinct slot—so as to attract as much market share as possible.

Here is the way we determine how well the two networks perform relative to each other, given their schedules. Each show has a fixed *rating*, which is based on the number of people who watched it last year; we'll assume that no two shows have exactly the same rating. A network *wins* a given time slot if the show that it schedules for the time slot has a larger rating than the show the other network schedules for that time slot. The goal of each network is to win as many time slots as possible.

Suppose in the opening week of the fall season, Network  $\mathcal{A}$  reveals a schedule  $S$  and Network  $\mathcal{B}$  reveals a schedule  $T$ . On the basis of this pair of schedules, each network wins certain time slots, according to the rule above. We'll say that the pair of schedules  $(S, T)$  is *stable* if neither network can unilaterally change its own schedule and win more time slots. That is, there is no schedule  $S'$  such that Network  $\mathcal{A}$  wins more slots with the pair  $(S', T)$  than it did with the pair  $(S, T)$ ; and symmetrically, there is no schedule  $T'$  such that Network  $\mathcal{B}$  wins more slots with the pair  $(S, T')$  than it did with the pair  $(S, T)$ .

The analogue of Gale and Shapley's question for this kind of stability is the following: For every set of TV shows and ratings, is there always a stable pair of schedules? Resolve this question by doing one of the following two things:

(a) give an algorithm that, for any set of TV shows and associated ratings, produces a stable pair of schedules; or

(b) give an example of a set of TV shows and associated ratings for which there is no stable pair of schedules.

## Problem 2:

Gale and Shapley published their paper on the Stable Matching Problem in 1962; but a version of their algorithm had already been in use for ten years by the National Resident Matching Program, for the problem of assigning medical residents to hospitals.

Basically, the situation was the following. There were  $m$  hospitals, each with a certain number of available positions for hiring residents. There were  $n$  medical students graduating in a given year, each interested in joining one of the hospitals. Each hospital had a ranking of the students in order of preference, and each student had a ranking of the hospitals in order of preference. We will assume that there were more students graduating than there were slots available in the  $m$  hospitals.

The interest, naturally, was in finding a way of assigning each student to at most one hospital, in such a way that all available positions in all hospitals were filled. (Since we are assuming a surplus of students, there would be some students who do not get assigned to any hospital.)

We say that an assignment of students to hospitals is *stable* if neither of the following situations arises.

- First type of instability: There are students  $s$  and  $s'$ , and a hospital  $h$ , so that
  - $s$  is assigned to  $h$ , and
  - $s'$  is assigned to no hospital, and
  - $h$  prefers  $s'$  to  $s$ .
- Second type of instability: There are students  $s$  and  $s'$ , and hospitals  $h$  and  $h'$ , so that
  - $s$  is assigned to  $h$ , and
  - $s'$  is assigned to  $h'$ , and
  - $h$  prefers  $s'$  to  $s$ , and
  - $s'$  prefers  $h$  to  $h'$ .

So we basically have the Stable Matching Problem, except that (i) hospitals generally want more than one resident, and (ii) there is a surplus of medical students.

Show that there is always a stable assignment of students to hospitals, and give an algorithm to find one.

### Problem 3:

Set up a stable matching problem that will solve the following problem. How should the preferences be determined? Why will that work?

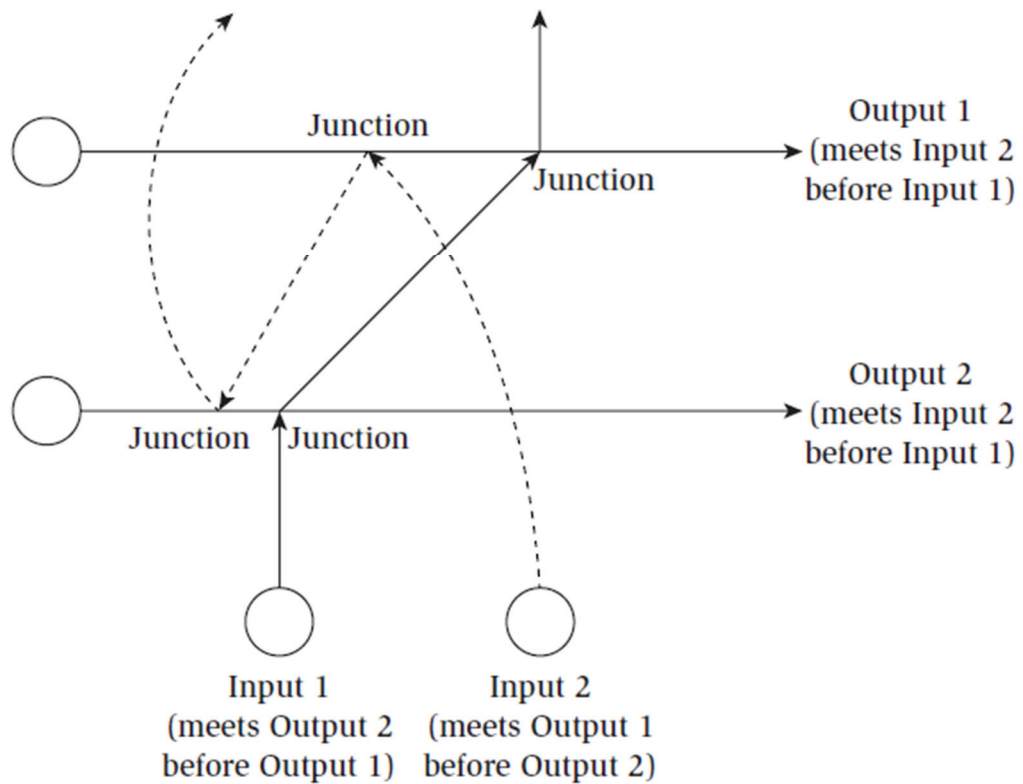
Some of your friends are working for CluNet, a builder of large communication networks, and they are looking at algorithms for switching in a particular type of input/output crossbar.

Here is the setup. There are  $n$  *input wires* and  $n$  *output wires*, each directed from a *source* to a *terminus*. Each input wire meets each output wire in exactly one distinct point, at a special piece of hardware called a *junction box*. Points on the wire are naturally ordered in the direction from source to terminus; for two distinct points  $x$  and  $y$  on the same wire, we say that  $x$  is *upstream* from  $y$  if  $x$  is closer to the source than  $y$ , and otherwise we say  $x$  is *downstream* from  $y$ . The order in which one input wire meets the output wires is not necessarily the same as the order in which another input wire meets the output wires. (And similarly for the orders in which output wires meet input wires.) Figure 1.8 gives an example of such a collection of input and output wires.

Now, here's the switching component of this situation. Each input wire is carrying a distinct data stream, and this data stream must be *switched* onto one of the output wires. If the stream of Input  $i$  is switched onto Output  $j$ , at junction box  $B$ , then this stream passes through all junction boxes upstream from  $B$  on Input  $i$ , then through  $B$ , then through all junction boxes downstream from  $B$  on Output  $j$ . It does not matter which input data stream gets switched onto which output wire, but each input data stream must be switched onto a *different* output wire. Furthermore—and this is the tricky constraint—no two data streams can pass through the same junction box following the switching operation.

Finally, here's the problem. Show that for any specified pattern in which the input wires and output wires meet each other (each pair meeting exactly once), a valid switching of the data streams can always be found—one in which each input data stream is switched onto a different output, and no two of the resulting streams pass through the same junction box. Additionally, give an algorithm to find such a valid switching.

(see the following page for a picture)



**Figure 1.8** An example with two input wires and two output wires. Input 1 has its junction with Output 2 upstream from its junction with Output 1; Input 2 has its junction with Output 1 upstream from its junction with Output 2. A valid solution is to switch the data stream of Input 1 onto Output 2, and the data stream of Input 2 onto Output 1. On the other hand, if the stream of Input 1 were switched onto Output 1, and the stream of Input 2 were switched onto Output 2, then both streams would pass through the junction box at the meeting of Input 1 and Output 2—and this is not allowed.

#### Problem 4 (Extra Credit)

You're doing some stress-testing on various models of glass jars to determine the height from which they can be dropped and still not break. The setup for this experiment, on a particular type of jar, is as follows. You have a ladder with  $n$  rungs, and you want to find the highest rung from which you can drop a copy of the jar and not have it break. We call this the *highest safe rung*.

It might be natural to try binary search: drop a jar from the middle rung, see if it breaks, and then recursively try from rung  $n/4$  or  $3n/4$  depending on the outcome. But this has the drawback that you could break a lot of jars in finding the answer.

If your primary goal were to conserve jars, on the other hand, you could try the following strategy. Start by dropping a jar from the first rung, then the second rung, and so forth, climbing one higher each time until the jar breaks. In this way, you only need a single jar—at the moment it breaks, you have the correct answer—but you may have to drop it  $n$  times (rather than  $\log n$  as in the binary search solution).

So here is the trade-off: it seems you can perform fewer drops if you're willing to break more jars. To understand better how this trade-off works at a quantitative level, let's consider how to run this experiment given a fixed "budget" of  $k \geq 1$  jars. In other words, you have to determine the correct answer—the highest safe rung—and can use at most  $k$  jars in doing so.

- (a) Suppose you are given a budget of  $k = 2$  jars. Describe a strategy for finding the highest safe rung that requires you to drop a jar at most  $f(n)$  times, for some function  $f(n)$  that grows slower than linearly. (In other words, it should be the case that  $\lim_{n \rightarrow \infty} f(n)/n = 0$ .)
- (b) What is the asymptotic rate of growth of  $f(n)$  as a function of  $n$  in your solution? What are  $f(15)$  and  $f(16)$  using your solution?

#### Problem 5 (Extra Credit)

In the previous problem, determine the best possible solution and prove that nobody can do better.