

CS358: Applied Algorithms

Homework 5: Linear Programming (due 11/07/24 10PM EDT)

Instructor: Sam McCauley

Instructions

All submissions are to be done through github. This process is detailed in the handout “Handing In Assignments” on the course website. Answers to the questions below should be submitted by editing this document. All places where you are expected to fill in solution are marked in comments with “FILL IN.”

Please contact me at srm2@cs.williams.edu if you have any questions or find any problems with the homework materials.

This homework will not have any coding component. (Note that this means that normal CS collaboration rules apply: you’re encouraged to discuss high-level concepts with other students, but should not be sharing the text of solutions. This is often referred to as “hands-in-pockets” collaboration, as you’re encouraged to collaborate if you’re not writing anything down.) Your handin will consist of two parts: solutions filled in below (in the tex file), and lp files for use with GLPK (`grocery.lp` and `investment.lp`).

Problems

The following question comes from (the extended version of) “Algorithms” by Jeff Erickson.

Problem 1 (20 points). Suppose you are given a rooted tree T , where every edge e has two associated values: a non-negative length $\ell(e)$, and a cost $c(e)$ (which may be positive, negative, or zero). Your goal is to add a non-negative stretch $s(e) \geq 0$ to the length of every edge e in T , subject to the following conditions:

- Every root-to-leaf path π in T has the same total stretched length $\sum_{e \in \pi} \ell(e) + s(e)$
- The total weighted stretch $\sum_e s(e)c(e)$ is as small as possible.

Give an LP for this problem and prove that it is correct.

Solution.

Problem 2 (20 points). You have a problem that is nearly an LP:

Objective:

$$\max x + y$$

Constraints:

$$x^2 + y^2 \leq 9$$

$$x - y \leq 2$$

$$x \geq 0$$

$$y \geq 0$$

Will the simplex algorithm give a correct solution to this problem? Please explain briefly why it always will, or why it may not. Note that I am looking for a reason why simplex will work or fail for *this problem in particular*, not just a statement that it may not work because the problem is not an LP.

Hint: Drawing a picture of the search space may help.

Solution.

Problem 3 (30 points). You are in charge of scheduling employees at a grocery store. This question is in two parts. The first part asks for a general “recipe” of how to turn the problem into a linear program. The second part asks for the linear program of a specific instance (along with a couple of additional restrictions).

(a) The grocery store has m employees that must be assigned to n days. Each employee is assigned a certain number of hours to work per day (this number of hours does not have to be integral). Each employee $i \in \{1, \dots, m\}$ is paid s_i per hour that they work (they are paid more than this for overtime, as mentioned below).

Some employees fall in special categories. Let M be the set of employees who are under 18. Let C be the set of employees who are able to work the cash register.

Any employee who is under 18 can only work 6 hours a day. The remaining employees can work up to 10 hours a day. But, any time that an over-18 employee works beyond 8 hours a day is paid at time and a half (employee i makes $1.5s_i$ per hour after their 8th hour in a day). To be clear: minors cannot work overtime.

During the day, there must be at least 8 hours assigned (in order to operate the cash register), all from employees who are both over 18 and are able to work the cash register.¹ In addition, each day $j \in \{1, \dots, n\}$ has a minimum amount of work ℓ_j (this is in addition to the aforementioned 8 hours) that must be assigned; at least half of this must be from employees who can work the cash register (they do not need to be over 18). This means that in total on day j , there must be at least $8 + \ell_j/2$ hours assigned to employees who can work the cash register.

Finally, no employee under 18 can work more than 30 hours per week; the remaining employees cannot work more than 40 hours per week. A week is a 7-day period in the n days; for example, days $\{1, \dots, 7\}$ constitute a week; days $\{8, \dots, 14\}$ constitute the second week.

¹Employees who are under 18 can't sell alcohol so it's important that one over-18 employee is present.

Your goal is to find a solution that meets all of these constraints while minimizing the cost of running the store.

Give an LP to solve this problem. (You do not need to prove that it is correct—that said, it may be a useful exercise to verify that you didn’t make any mistakes.) You should briefly explain any variables and constraints that are used.

Solution.

(b) Let’s look at an actual instance of the above problem.

Employee	s_i	Minor?	Cash register?
Sarah	9.50	yes	yes
Sun	7.10	yes	yes
Abed	11.10	no	no
Brian	15.00	no	yes
Ruhi	20	no	yes

The employees need to be assigned to the following days; these days are all in the same week:

Day:	Monday	Tuesday	Wednesday	Thursday	Friday
ℓ_j :	12	12	14	14	24

There is one additional restriction: Brian and Abed are siblings, and they can only work 10 hours total (total between the two of them) on Friday.

Write the LP for this instance. Then, solve it using GLPK, putting your solution in `grocery.lp`. (You will almost certainly want to use descriptive variable names to keep track of what you’re doing— x_{ij} and the like probably isn’t a good idea when it comes to an `.lp` file of this size.) Record below: what is the cost of the optimal solution?

Side note: You may notice from your solution that while the schedule you obtain is optimal mathematically, it may not be very workable for real people being scheduled. You should think to yourself (you don’t need to write anything about this): how could you modify the LP to avoid some of these cases? Are there issues that you can’t avoid with linear constraints—issues that may require a human to tweak the LP output? Again: no need to actually respond to these questions below, but these considerations are crucially important when applying these techniques in the real world.

Solution.

Problem 4 (30 points). You have \$1000 to invest over the next ten years. You can invest the money in the following ways.

- A savings account gives 1% interest per year.

- A 2-year CD gives 2% interest per year.
- A 5-year CD gives 3% interest per year.

In addition, at the end of years 3, 7, and 9 you will have a \$300 bill you will need to pay (you must have enough money, not in any CD, to pay it—and, to be clear, you must pay it, losing \$300 each time). However, you'll receive an additional \$200 at the end of year 6.

How can you invest this money to maximize your ten-year return? (Note that the money at the end of the ten years must be immediately available—it cannot be tied up in a CD.)

Give an LP to solve this problem. **You do not need to prove that it is correct.** Then, use GLPK to find an optimal solution; put your linear program in `investment.lp`. Record below: what is the most amount of money you can end up with at the end of the ten years?

A couple things to note about this problem:

- You can reinvest immediately: I can buy a 2 year CD at the beginning of year 1, and then invest the proceeds into a 5-year CD at the beginning of year 3.
- Careful about off-by-one errors! Be very clear on what your variables mean. For example: “after ten years” in the problem statement can be interpreted as “money at the end of year ten” or “money at the beginning of year eleven.” You can even create separate variables for the money coming in at the end of a year and the money going in at the beginning of a year (though you do not need to and it's likely one is redundant).
- I recommend, once you have a solution, tracing through it and making sure that it makes sense
- The correct solution is nontrivial (it's not a series of 5 2-year CDs or anything like that)

Solution.