

CS358: Applied Algorithms

Mini-Midterm 3: MIPs and ILPs (due 11/17/21)

Instructor: Sam McCauley

Instructions

All submissions are to be done through github, as with assignments. 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.” You will also be given one `.lp` file to fill in: `clerks.lp`. The mini-midterm will not have a leaderboard or any automated testing.

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

This is a mini-midterm (as defined on the syllabus). This means that **all work must be done alone**. Please do not discuss solutions with any other students, even at a high level. You should not look up answers, hints, or even code libraries on the internet. This midterm was designed to be completed with no external resources, beyond those explicitly linked in the midterm. (Class resources, such as slides, notes, the GLPK manual, and your previous assignment submissions, are of course acceptable, as are basic resources such as looking up debugging information.)

For this mini-midterm, each question begins a new page. (It got very busy without this change.) Make sure you answer all six questions!

A comment on grading: remember that we have four tests over the course of the semester, and (while important) this midterm is only worth 20% of your final grade.

Furthermore, some problems may seem a bit harder than others. As always, if you're unsure of a 100% correct solution, make sure to write down the ideas you have for partial credit.

Problem 1 (15 points). Let's generalize the router assignment problem we saw in class to the case with multiple routers.

There are n roommates living in a single-floor house; roommate i is located at coordinates (x_i, y_i) . You want to place k routers in the house. Let the distance between roommate i and a router at (x, y) be defined as $|x_i - x| + |y_i - y|$.

You want to place the routers under the following constraints:

- Each roommate must be assigned to exactly one router
- The router i assigned to a roommate r must satisfy that the distance between router i and roommate r is at most 10.

Your goal is to minimize the sum, over all roommates, of the distance between the roommate and the router they are assigned to.

To be clear, you receive as input n , k , and the coordinates of all n roommates: $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$.

Give an ILP or MIP for this problem and prove that it is correct.

Solution.

Problem 2 (15 points). The following is a fairly classic bin packing problem. Let’s solve it using a branch and bound approach. This question will not have any LP, ILP, or MIP in it—it’s just about how branch and bound itself works.¹

Problem Statement: You have n items with weights $W = w_1, \dots, w_n$. You have m bins with capacities $C = c_1, \dots, c_m$; these are given in *decreasing* order. Your job is to assign items to bins, with the restriction that the total weight of all items assigned to bin i is at most c_i .

The goal is to minimize the number of bins that have any item assigned to them. (To rephrase in case it helps with clarity: your goal is to maximize the number of empty bins with no item assigned.)

You may assume that, without loss of generality, a solution that uses k bins uses bins $1, \dots, k$ (you don’t want to “skip over” bins to use a later bin—that wouldn’t make sense since the later bins are smaller).

A Recursive Approach: The following recursive algorithm solves this problem optimally by considering every possible solution (recall that $V \setminus \{v\}$ means the set resulting when removing v from V):

```

BINPACKING( $V, i, w$ ):
  1. If  $|V| = 0$  or  $i > m$  return  $i$ 
  2.  $\text{bestSoFar} \leftarrow m + 1$ 
  3. for all  $v \in V$ :
      • if  $w + w_v \leq c_i$ , then  $\text{bestSoFar} \leftarrow \min\{\text{bestSoFar}, \text{BINPACKING}(V \setminus \{v\}, i, w + w_v)\}$ .
  4.  $\text{bestSoFar} \leftarrow \min\{\text{bestSoFar}, \text{BINPACKING}(V, i + 1, 0)\}$ .
  5. return  $\text{bestSoFar}$ 

```

For any set of items V , any $i \in \{1, \dots, m\}$, and any w , $\text{BINPACKING}(V, i, w)$ returns the smallest j such that all items in V can be stored in bins $\{i, i + 1, \dots, j\}$ (where bin i already has weight w of items stored in it); if no such j exists then BINPACKING returns $m + 1$.²

Therefore, we can solve this problem by calling $\text{BINPACKING}(W, 1, 0)$.

Developing Branch and Bound: The remainder of this question is in two parts.

(a) Give a *simple* method³ $\text{LOWERBOUND}(V, i, w)$ that gives a lower bound on the number of bins necessary from $i, i + 1, \dots, m$ in order to store all items from V given that bin i already has w worth of items in it. Explain your solution: why is this a lower bound? (A brief explanation is sufficient.)

Solution.

¹That said—yes, in general one could absolutely solve this with an ILP.

²I won’t include the proof of correctness—you may take as given that this is true.

³The “simple” here is to emphasize that I’m not looking for a complicated approach. I just want a way to immediately get some lower bound on the value of the solution. In particular, I don’t anticipate taking off points for a solution not being simple enough.

(b) Use $\text{LOWERBOUND}(V, i, w)$ to create an algorithm $\text{BINPACKINGWITHBOUNDS}$ that avoids some of the recursive subcalls made by BINPACKING . Your algorithm should be a modification of the BINPACKING algorithm (to this end, the code for BINPACKING is included below; modify it to obtain your solution).

Prove that your $\text{BINPACKINGWITHBOUNDS}$ algorithm will give the correct answer. (You may assume what was already mentioned above: that LOWERBOUND correctly lower bounds the subproblem solution, and that BINPACKING is correct.)

Hint: There must be some optimal solution to each recursive $\text{BINPACKINGWITHBOUNDS}$ call. Prove that your algorithm will always make a recursive call containing this solution. The remainder of the proof follows from this.

Solution.

$\text{BINPACKING}(V, i, w)$:

1. If $|V| = 0$ or $i > m$ return i
2. $\text{bestSoFar} \leftarrow m + 1$
3. for all $v \in V$:
 - if $w + w_v \leq c_i$, then $\text{bestSoFar} \leftarrow \min\{\text{bestSoFar}, \text{BINPACKING}(V \setminus \{v\}, i, w + w_v)\}$.
4. $\text{bestSoFar} \leftarrow \min\{\text{bestSoFar}, \text{BINPACKING}(V, i + 1, 0)\}$.
5. return bestSoFar

Problem 3 (20 points). Let's go back to a problem you may have seen in CS 136.⁴

You are given a set of n students, each of whom are taking four of m classes. Your goal is to assign each of the m classes to one of k time slots, under the following constraint:

- if any student is taking two classes i and j , then i and j may not be assigned to the same time slot.

Under this constraint, you want to minimize the number of times a student has an exam scheduled in three consecutive time slots. (For example, if a student has an exam scheduled in time slot 6, another in time slot 7, and another in time slot 8, that student would count towards this total.) Note that I said *number of times*: if a student's courses are scheduled in four consecutive time slots, that would add 2 to this total. (So a student with exams in time slots 2, 3, 4, and 5, there are two sequences of three slots: 2, 3, 4 and 3, 4, 5. These two sequences each add 1 to the total, for an increase of 2 overall.)

You are given as input a sequence of n lists of length 4; list i gives the classes taken by student i . You are also given n , m , and k .

Give an MIP or ILP to solve this problem. You do not need to prove that it is correct (though it may be a good exercise to verify your solution).

Solution.

⁴It was extra credit; it's quite possible you didn't encounter it. Note that we're solving a slightly different version of the problem presented there.

Problem 4 (15 points). A bank is open from 9:00 to 15:00. During each hour of the day, the number of clerks required is shown in the following table:

Time Period	No. of Clerks
9:00–10:00	4
10:00–11:00	3
11:00–12:00	4
12:00–13:00	6
13:00–14:00	5
14:00–15:00	6

The bank can hire full-time and part-time clerks. Full-time clerks work from 9:00 to 15:00 except for a one-hour lunch break, which is from 12:00–13:00 or from 13:00–14:00 (the bank decides the time at which each clerk takes their lunch break). The clerks are paid \$8 per hour (and receive payment for their lunch break). Part-time clerks work for three consecutive hours and the bank specifies the start time for each of them. Part-time clerks are paid \$6 per hour. No more than five part time clerks can be hired.

Solve this problem using GLPK, in the file `clerks.lp`. Write below: how many full-time employees are hired in your solution, and what are their lunch breaks?

Solution.

Problem 5 (20 points). The following is a variant of the Travelling Salesman Problem.

Let's say you have a collection of n points; for each pair of points i and j there is a cost c_{ij} to travel from i to j .

In addition, k *target points* are identified (they are a subset of the n points above).

Your goal is to find the shortest cycle that visits all k of the target points. Your cycle may also visit points that are not target points—the only requirement is that at least the k target points *must* be visited. To be clear, the solution must be a “cycle”: the cycle must start and stop in the same place, and the cycle cannot visit any point more than once. (In particular, each of the target points must be visited *exactly* once; each of the $n - k$ points that are not target points can be visited *at most* once.)

The input consists of: n , k , the n^2 values for c_{ij} , and a list of k target points.

Give an MIP or ILP for this problem and prove that it is correct.

Solution.

Problem 6 (15 points). You are given a set of n strings, each of length m .

Your goal is to find the longest string s such that s is a subsequence of all n input strings. That is to say: for every input string i , the characters of s should all appear in i (in order, but not necessarily consecutively).

For example, if the input is:

tomorrownamaste
tomatosteerings
antimasterheads
tastetastetaste

Then the solution is **taste**. (The red highlighted characters in the below can be used to verify that this is, in fact, a solution; I believe there is no longer solution.)

tomorrownam aste
tomatosteering s
anti ma ste re head s
taste taste taste

Give an MIP or ILP to solve this problem. You do not need to prove that it is correct (though it may be a good exercise to verify your solution).

Hint: if the solution has length k , then there are exactly k characters from each input string that are matched to the solution string (here I'm referring to the red characters in the above). This may simplify your objective function.

Solution.