CS 256: Algorithm Design and Analysis

Assignment 5 (due 04/10/24)

Instructor: Sam McCauley

For this assignment, you are allowed to work in groups of 2 or 3. If you do, please write the names of the students in this group here:

- Student 1

- Student 2

- Student 3

Note on Dynamic Programs. For full credit on a dynamic program, you must clearly state the following parts.

(a) *Subproblem definition:* your subproblem must have an optimal substructure.
(b) *Recurrence:* how should the next subproblem be computed using the previous ones? This is the core of your algorithm and its correctness. A less ideal alternative to a recurrence is clear pseudocode for the final iterative dynamic-programming algorithm.
(c) *Base case(s):* you need to start somewhere!
(d) *Final output:* in terms of your subproblem.
(e) *Memoization data structure:* this is often obvious but should not be skipped.
(f) *Evaluation order:* describes the dependencies between the subproblems.
(g) *Time and space analysis.*

For this assignment, these are sufficient to argue correctness (that is to say, if you explain why the above parts are correct, that's sufficient to show that your algorithm works properly).

Problem 1. (Erickson 3.6) Solving part (a) is sufficient for full credit on this problem; part (b) is extra credit.

A shuffle of two strings X and Y is formed by interspersing the characters into a new string, keeping the characters of X and Y in the same order. For example, the string BANANAANANAS is a shuffle of the strings BANANA and ANANAS in several different ways.

$$\text{BANANA}_\text{ANANAS} \qquad \text{BAN}_\text{ANA}\text{ANA}_\text{NAS} \qquad \text{B}_\text{AN}\text{AN}_\text{A}\text{A}_\text{NA}\text{NA}_\text{S}$$

Similarly, the strings PRODGYRNAMAMMIINCG and DYPRONGARMAMMICING are both shuffles of DYNAMIC and PROGRAMMING:

$$\text{PRO}_\text{D}\text{GY}_\text{R}\text{NAM}\text{AMMI}_\text{I}\text{N}_\text{C}\text{G} \qquad {}_\text{DY}\text{PRO}_\text{N}\text{GA}_\text{R}\text{M}\text{AMM}_\text{IC}\text{ING}$$

(a) Given three strings $A[1..m]$, $B[1..n]$, and $C[1..m+n]$, describe and analyze an algorithm to determine whether C is a shuffle of A and B.

Solution. □

(b) **(Extra credit: 5 pts)** A **smooth** shuffle of X and Y is a shuffle of X and Y that never uses more than two consecutive symbols of either string. For example,

- $\text{PR}_\text{D}\text{O}_\text{Y}\text{G}_\text{NA}\text{RA}_\text{M}\text{MM}_\text{I}\text{C}\text{NG}$ is a smooth shuffle of the strings DYNAMIC and PROGRAMMING.

- ${}_\text{DY}\text{PR}_\text{N}\text{OGR}_\text{A}\text{A}\text{M}_\text{MM}\text{IC}\text{ING}$ is a shuffle of DYNAMIC and PROGRAMMING, but it is not a smooth shuffle (because of the substrings OGR and ING).

- $\text{XX}_\text{X}\text{X}_\text{X}\text{X}_\text{XX}\text{XX}_\text{XX}\text{XX}_\text{X}\text{X}_\text{X}\text{XX}$ is a smooth shuffle of the strings XXXXXXX and XXXXXXXXXXX.

- There is no smooth shuffle of the strings XXXX and XXXXXXXXXXXX.

Describe and analyze an algorithm to decide, given three strings X, Y, and Z, whether Z is a smooth shuffle of X and Y.

> **Hint.** What do you need to change in order to build up a smooth shuffle rather than a normal shuffle? What do you need to keep track of to ensure that you can make this distinction?

Solution. □

Problem 2. (From Steve Skiena's *Algorithm Design Manual*) Consider the problem of storing n books on shelves in a library. The order of the books is fixed by the cataloging system and so cannot be rearraged. Let book b_i have thickness t_i and height h_i, for $1 \leq i \leq n$. Let the length of each bookshelf at this library be L. Suppose we have the freedom to adjust the height of each shelf to fit the tallest book on it. The cost of a particular layout is the sum, over each shelf, of the height of the largest book on that shelf. (So if shelf 1 has books with heights $(1, 5, 3)$ and shelf 2 has books with heights $(2, 4)$, the total cost is $5 + 4 = 9$.)

(a) Give an example to show that the greedy algorithm of stuffing each shelf as full as possible (that is, fill the first shelf with as many books as possible until book b_i does not fit, and then repeat the same process on subsequent shelves) does not always give the minimum overall height.

(b) Give a dynamic programming algorithm that computes the height of the optimal arrangement, and analyze its time and space complexity.

> **Hint.** We have done a similar example in class with a different cost function and constraints.

Solution.   □

Problem 3. Two friends Rosa and Beth are planning a hike to the top of a mountain. There are n lookout points situated at varying points on the mountain (including one at the top), and an extensive network of trails connecting these lookout points.[1]

Rosa and Beth want to find the shortest path from the base of the mountain to the top of the mountain under two constraints. First, Rosa and Beth want to (between the two of them) visit all n lookout points: any lookout point i should either be on Rosa's route or Beth's route. Second, neither wants to backtrack: if Rosa visits lookout point i, Rosa's next lookout point should be at a higher elevation. Similarly, when Beth visits some lookout point j, Beth's next lookout point should be at a higher elevation.

Assume that the n lookout points are given in sorted order of elevation and that there are no ties (so lookout point i is higher than lookout point k when $i > k$).

You are given the following input: for each pair of lookout points x, y, you are given the distance $d(x, y)$ to hike from lookout point x to lookout point y. Give an algorithm to find the shortest *total* path length (i.e. the sum of Rosa's path length and Beth's path length) through the mountain such that (1) Rosa or Beth visit all lookout points, and (2) the path Rosa takes, and the path Beth takes, each visit lookout points in increasing order. Both paths start at lookout point 1, and end at lookout point n (lookout point n is at the top of the mountain).[2]

Give an $O(n^2)$ dynamic programming algorithm to find the optimal pair of paths, minimizing the total distance hiked in sum by Rosa and Beth (In the "time and space" section please explain why your algorithm is $O(n^2)$.)

Solution.

(a) **Subproblem definition:**

> **Hint.** Use the following subproblem definition. You do not need to create your own subproblem for this problem. (If you want to change this that's OK—but there is a correct answer that uses this subproblem definition.)

Entry (i, j) in the dynamic programming table represents two paths such that Rosa's path is from 1 to i; Beth's path is from 1 to j, and every vertex from 1 to $\max\{i, j\}$ is visited by one of the two paths.

Let $S(i, j)$ be the smallest total length (sum of the two paths) of Rosa and Beth's path satisfying these constraints.

(We do not use $S(i, i)$.)

(b) **Recurrence:**

---

[1] Note that we don't use the specific trail layout in this problem: we just use the distance between any two lookout points. Basically: assume there is some path from any lookout point to any other.

[2] This means that both Rosa and Beth visit lookout points 1 and n. Any other lookout point will only be visited by one of them.

> **Hint.** First, let's say that either $i > j + 1$, or $j > i + 1$. What is the recurrence for $S(i, j)$? (This recurrence is unusually short.)
>
> Then, let's say that $i = j + 1$ or $j = i + 1$. What is the recurrence for $S(i, j)$?

(c) **Base case(s):**

(d) **Final output:**

(e) **Memoization data structure:**

(f) **Evaluation order:**

(g) **Time and space analysis**

> **Hint.** For the time analysis, treat the total time of the two cases in your recurrence separately.

□

> **Note.** Problem 4 relies on material we will not see until Monday, 04/08/24.

Problem 4 (KT 6.13). Consider a firm that trades shares in n different companies. For each pair $i \neq j$, they maintain a trading rate r_{ij}, meaning that one share of i trades for r_{ij} shares of j. Here we allow the rate r to be fractional; that is, $r_{ij} = \frac{2}{3}$ means that you can trade three shares of i to get two shares of j.

A trading cycle for a sequence of shares $i_1, i_2, \ldots, i_k, i_1$ consists of successively trading shares in company i_1 for shares in company i_2, then shares in company i_2 for shares i_3, and so on, finally trading shares in i_k back to shares in company i_1.[3]

After such a sequence of trades, one ends up with shares in the same company i_1 that one starts with. Trading around a cycle, in practice, is usually a bad idea, as you tend to end up with fewer shares than you started with. But occasionally, for short periods of time, there are opportunities to increase shares.

We will call such a cycle an *opportunity cycle*, if trading along the cycle *increases* the number of shares. This happens exactly if the product of the ratios along the cycle is above 1. In particular, a sequence of trades $i_1, i_2, \ldots, i_k, i_1$ form an opportunity cycle if $r_{i_1,i_2} \cdot r_{i_2,i_3} \cdots r_{i_k,i_1} > 1$.

In analyzing the state of the market, a firm engaged in trading would like to know if there are any opportunity cycles. Give an efficient algorithm that determines whether or not such an opportunity cycle exists.

> **Hint.** Reduce this problem to a problem we have discussed in class. A reduction from a problem A to a problem B means you take an input instance of A and turn it into an input instance of problem B.
>
> **This means that your solution does not need to have a dynamic program.** Instead, describe how to use a dynamic programming algorithm we already know from class to describe this problem. How can you transform an instance of the above trading problem into an instance of a problem we've already seen?

Solution. □

---

[3]If it helps, you can think about trading different currences: suppose that 1 U.S. (item i) dollar buys 0.7 British pound (item j), then $r_{ij} = 1/0.7 = 10/7$.