

1 Dynamic Programming

1.1 Max independent set in a vertex-weighted path

Let $G = (V, E)$ be an undirected graph with n nodes. A subset S of V is called an *independent set* if no two vertices of S are joined by an edge. Finding large independent sets is difficult in general but not for paths. A graph is a *path* if its nodes can be written as v_1, v_2, \dots, v_n , with an edge between v_i and v_j if and only if i and j differ by exactly 1. With each node v_i , we associate a positive integer weight w_i . The goal in this question is to find an independent set in a path G whose total weight is as large as possible.

1.2 Optimum plan

Given a sequence of n weeks containing the value ℓ_i for “low-stress” job and h_i for “high-stress” job for every week i . A *plan* is specified by a choice of “low-stress”, “high-stress”, or “none” for each week, with the constraint that if high-stress is chosen for week $i > 1$, then none has to be chosen for week $i - 1$. (It’s okay to choose a high-stress job in week 1.) The value of the plan is determined in the natural way: for each week i , you add ℓ_i or h_i or 0 to the value if you choose low-stress or high-stress or none in that week, respectively. Your task is to find a plan of maximum value.

Example: Suppose $n = 4$, and the values are as follows.

	Week 1	Week 2	Week 3	Week 4
ℓ	10	1	10	10
h	5	50	5	1

Then the plan of maximum value would be to choose none in week 1, a high-stress job in week 2, and low-stress jobs in weeks 3 and 4. The value of this plan would be $0 + 50 + 10 + 10 = 70$.

1.3 Longest path in ordered graph

Let $G = (V, E)$ be a directed graph with nodes v_1, \dots, v_n . We say that G is an *ordered graph* if it has the following properties.

- (i) Each edge goes from a node with a lower index to a node with a higher index. That is, every directed edge has the form (v_i, v_j) with $i < j$.
- (ii) Each node except v_n has at least one edge leaving it. That is, for every node v_i , $i = 1, \dots, n - 1$, there is at least one edge of the form (v_i, v_j) with $i < j \leq n$.

Given an ordered graph G , find the length of the longest path that begins at v_1 and ends at v_n . (The *length of a path* is the number of edges in it.)

(a) Show that the following greedy algorithm does not correctly solve this problem.

```
Set  $w = v_1$ 
Set  $L = 0$ 
While there is an edge out of the node  $w$ 
    Choose the edge  $(w, v_j)$  for which  $j$  is as small as possible
    Set  $w = v_j$ 
    Increase  $L$  by 1
end while
Return  $L$  as the length of the longest path
```

(b) Give an efficient algorithm that takes an ordered graph G and returns the length of the longest path that begins at v_1 and ends at v_n .

1.4 Minimum operating cost

Suppose you're running a lightweight consulting business. Your clients are distributed between the East Coast and the West Coast, and this leads to the following question.

Each month, you can either run your business from an office in New York (NY) or from an office in San Francisco (SF). In month i , you'll incur an operating cost of N_i if you run the business out of NY or an operating cost of S_i if you run it out of SF. However, if you run the business out of one city in month i , and then out of the other city in month $i + 1$, then you incur a fixed moving cost of M to switch base offices.

Given a sequence of n months, a plan is a sequence of n locations—each one equal to either NY or SF—such that the i -th location indicates the city in which you will be based in the i -th month. The cost of a plan is the sum of the operating costs for each of the n months, plus a moving cost of M for each time you switch cities. The plan can begin in either city.

Given a value for the moving cost M , and sequences of operating costs N_1, \dots, N_n and S_1, \dots, S_n , find a plan of minimum cost.

Example: Suppose $n = 4$, $M = 10$, and the operating costs are given by the following table.

	Month 1	Month 2	Month 3	Month 4
NY	1	3	20	30
SF	50	20	2	4

Then the plan of minimum cost would be the sequence of locations (NY, NY, SF, SF), with a total cost of $1 + 3 + 2 + 4 + 10 = 20$, where the final term of 10 arises because you change locations once.

(a) Show that the following algorithm does not correctly solve this problem, by giving an instance on which it does not return the correct answer.

```

for  $i = 1$  to  $n$ 
  if  $N_i < S_i$  then
    choose NY in Month  $i$ 
  else
    choose SF in Month  $i$ 
end for

```

In your example, say what the correct answer is and also what the algorithm above finds.

(b) Give an example of an instance in which every optimum plan must move (i.e., change locations) at least three times. Provide a brief explanation, saying why your example has this property.

(c) Give an efficient algorithm that finds an optimal plan.

1.5 Maximum-quality segmentation

Given a string of letters $y = y_1 y_2 \dots y_n$, its segmentation means a partition into blocks of contiguous letters. If a block is a dictionary word, then its *quality* is +1, else -1. The quality for any block is available from a black box in constant time.

The *total quality* of a segmentation is determined by adding up the qualities of its blocks. For example, for **meetateight**, the total quality for the segmentation **meet** + **ate** + **ight** is $1 + 1 - 1 = 1$, whereas that for **meet** + **at** + **eight** is $1 + 1 + 1 = 3$.

Give an efficient algorithm that takes a string and computes a segmentation of maximum total quality.

1.6 Minimum-slack text formatting

Suppose our text consists of a sequence of words, $W = \{w_1, \dots, w_n\}$, where w_i consists of c_i characters. Assume that all characters have same fixed-width font and ignore issues of punctuation or hyphenation. We have a maximum line-length of L . A formatting of W consists of a partition of the words in W into lines. In the words assigned to a single line, there should be a space after each word except the last; and so if w_j, \dots, w_k are assigned to one line, then we should have

$$k - j + \sum_{i=j}^k c_i \leq L.$$

An assignment of words to a line is valid if it satisfies the above constraint. The amount by which the sum (left-hand side of the above equation) falls short of L , is called the *slack* of the line—that is, the number of spaces left at the right margin.

Give an efficient algorithm to find a partition of a set of words W into valid lines, so that the sum of the squares of the slacks of all lines (including the last line) is minimized.

The goal of “minimum-slack text formatting” is to take text with a ragged right margin, like this:

```
Call me Ishmael.
Some years ago,
never mind how long precisely,
having little or no money in my purse,
and nothing particular to interest me on shore,
I thought I would sail about a little
and see the watery part of the world.
```

and turn it into text whose right margin is as “even” as possible, like this:

```
Call me Ishmael. Some years ago, never
mind how long precisely, having little
or no money in my purse, and nothing
particular to interest me on shore, I
thought I would sail about a little
and see the watery part of the world.
```