Algorithms

Example Pre-Sheet

Petar Veličković

Lent Term 2017

Introduction

The Algorithms course is arguably the most fundamental Computer Science course you will have in Part IA, with ideas and methodologies covered therein constantly reappearing throughout virtually every other Tripos course.

It is also the course that has by far the most variability with respect to the level of relevant background that people possess prior to taking up the course. In order to leverage this as much as possible, I would like you to have a look at the questions outlined below, and send your answers to me anytime *before Lent term starts*. I will then use this information to produce optimal pairings (so that everyone is put in the position to make the most out of these supervisions, regardless of prior background).

Before you begin

A couple of *important points* to note:

- The provided examples all correspond to well-defined algorithmic problems. The inputs, required outputs and constraints will be clear.
- I only require you to provide an **informal discussion** of the possible idea(s) that could be used to solve the problem, keeping the execution time and memory usage as low as possible. Justifying your decisions in a convincing way is also highly appreciated. It might be useful to know that a typical computer may execute $\sim 3 \cdot 10^8$ commands per second.
- You are **not required** to provide any code, pseudocode, or rigorous analysis of your proposed ideas (although you may freely apply ideas already covered in, say, Foundations of Computer Science).
- Finally, you aren't in any way *expected*, at this time, to have the optimal solutions for any of the problems given below—by the end of the course, you will all be capable of optimally reasoning about solutions to all of them. Please submit partial or inefficient solutions whenever necessary.

Examples

1. You want to become the vice-chancellor of the University of Algobridge. There are N ($1 \leq N \leq 1000$) constituent Colleges of the University, and you require the support of more than 50% of all Colleges to become elected. The i-th College has f_i Fellows, and you will be supported by this College if you are supported by more than 50% of its Fellows. What is the minimal number of Fellows that need to support you in order to become elected?
   Example input: $N = 3$, Fellow numbers: $\{6, 5, 6\}$.
   Example output: 7.
   Explanation: Gaining the support of four Fellows of College 1 and three Fellows of College 2 is sufficient to win the elections. There does not exist a more efficient way.

2. Your friend wants to play a game of memorising sequences of numbers. He will give you N queries ($5 \leq N \leq 10^5$), where each query is either:

   - Your friend gives you an integer. It is known that your friend has a set of 400 favourite integers, and he will always give integers out of this set (repeats of the same number are possible). You do not know which numbers these are upfront, however.
   - Your friend asks you the kth smallest number out of all the numbers given so far (repeated numbers are counted separately)!

   Example input: $N = 7$.
   Example queries and replies:
   Add 0; Add 2; Add 5;
   1st smallest? (your reply: 0);
   3rd smallest? (your reply: 5);
   Add 2;
   3rd smallest? (your reply: 2).

3. A Formula 1 race has taken place, but you have missed it! There were N ($1 \leq N \leq 10^6$) drivers competing, and you know the starting and final positions in which they have finished the race. As the races are usually quite dull, you would like to infer how many overtakes had to have taken place, to decide whether you should watch the re-run.
   Example input: $N = 5$, final positions: $\{4, 3, 2, 1, 5\}$.
   Example output: 6.
   Explanation: In the most pessimistic case, the fifth car stayed in its position throughout, then each of the other cars overtook all the cars that were initially ahead. This gives the total number of overtakes as $0 + 1 + 2 + 3 + 0 = 6$.

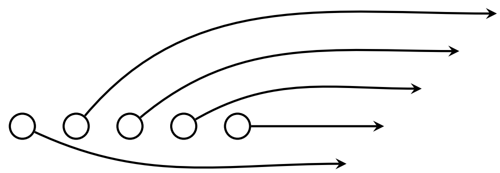

4. You own a field of apple trees, arranged in a rectangular $N \times M$ ($1 \leq N, M \leq 250$) grid. For each of the trees in the grid, you know how many apples are ripe to be eaten. You want to choose a sub-rectangle of this grid from which you will take all the apples, such that you can evenly distribute all the apples across your K ($1 \leq K \leq 10^6$) children. You need to determine all such sub-rectangles (by noting their top-left and bottom-right corners).
Example input: $N = 3, M = 3,\ K = 5$. Grid:
$$\begin{bmatrix} 2 & 9 & 3 \\ 10 & 8 & 6 \\ 1 & 4 & 12 \end{bmatrix}$$
Example output: $\{(1,1)-(3,3), (2,1)-(2,1), (3,1)-(3,2), (2,2)-(3,3)\}$.
Explanation: The sums of the numbers of ripe apples in the aforementioned sub-rectangles are: $\{55, 10, 5, 30\}$, all divisible by 5.

5. You are tasked with maintaining a dictionary of words. This dictionary needs to be capable of supporting the following operations as efficiently as possible:

    - Inserting new words (represented as lowercase strings consisting of up to 100 characters of the English alphabet) into the dictionary;
    - Determining the number of words in the dictionary that are lexicographically *smaller* than a given word (which is, itself, in the dictionary). For example, if the dictionary currently contains: {`"joke", "steve", "peter", "stephen", "maya", "chris", "maria"`} and we are interested in the number of words lexicographically smaller than `"stephen"`, then the answer is 5 (all except `"stephen"` and `"steve"`).

6. You are given a power grid with N nodes ($1 \leq N \leq 10^4$), and some pairs of nodes are linked together by M ($1 \leq M \leq 10^5$) power lines. Each power line $i \leftrightarrow j$ has a *teardown time*, w_{ij}, specifying the number of hours it needs to be neglected to stop being operational. You also have a list of K pairs of nodes that **should not** be connected (either directly or via several power lines) in order for the network to operate safely. Determine the minimal amount of hours you need to wait before the grid can be safely transporting electricity.
Example input: $N = 7, M = 8,\ K = 3$.
Power grid (with teardown times highlighted on each link):

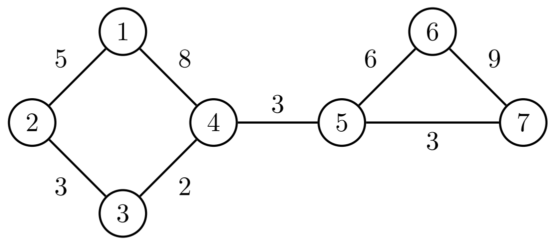

Forbidden connections: $2 \leftrightarrow 5$, $1 \leftrightarrow 3$, $5 \leftrightarrow 7$.
Example output: 6.
Explanation: After waiting for 3 hours, the links $2 \leftrightarrow 3$, $3 \leftrightarrow 4$, $4 \leftrightarrow 5$ and $5 \leftrightarrow 7$ will tear down, effectively disconnecting 2 and 5, as well as 1 and 3. However, there is still a link between 5 and 7 (through node 6), which requires a waiting time of 6 hours to fully tear down.

7. After a failed DNA experiment, N ($1 \leq N \leq 300$) dinosaurs are sleeping in a rectangular laboratory of size $W \times H$ ($2 \leq W, H \leq 10^6$)—for each dinosaur, an (x, y) coordinate is known. A scientist (considered as a point within this rectangle) starts at $(0, 0)$ and wants to get to (W, H) (the opposite corner). He does not want to awake the dinosaurs, so he wants to take the path that will be as far as possible from them. What is the shortest distance to any dinosaur that the scientist will ever have if he takes this optimal path?
Example input: $N = 1$, $W = H = 2$, $(x_1, y_1) = (1, 1)$.
Example output: 1.0.
Explanation: The optimal path for the scientist to take is exactly along the (west&north or south&east) edges of the lab, since the sole dinosaur is exactly in the centre of the room. That way, he will be the closest to the dinosaur when he is exactly halfway through each edge, when the dinosaur will be a distance 1 away from him.
Example input 2: $N = 2$, $W = 5$, $H = 4$,
$(x_1, y_1) = (1, 3)$, $(x_2, y_2) = (4, 1)$.
Example output 2: ≈ 1.581.

8. A captain is exploring an archipelago of N ($3 \leq N \leq 200$) islands. He does this by going from island i to island j for all possible pairs (i, j), and recording how many islands were on his *right-hand side*. No three islands are on the same line. These counts are stored in a matrix, isl_{ij} ($isl_{ii} = 0$ by definition). After doing this task, he wants to know: for every tuple (A, B, C) of islands, *was island C on his right-hand side when he travelled from A to B?*
Example input: $N = 3$.

$$isl = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$$

Example output: The "YES" tuples are $(1, 2, 3)$, $(2, 3, 1)$ and $(3, 1, 2)$.
Explanation: As there are only three islands, if an island was on the RHS while going from i to j, it is exactly the third island. Therefore in this case we can infer the "YES" tuples directly from the nonzero entries of isl.