

2024 North America Championship Solutions

The Judges

May 27, 2024

Problem

- You are given n tasks and each task must be assigned to one of two interns.
- Based on the task and intern, each task may take a different amount of time. Assign the tasks to minimize T , which is the maximum of the sum of task times assigned to each intern.
- Bounds: $1 \leq n \leq 50$, $1 \leq a, b \leq 10^5$.

Solution

- We can solve this problem with dynamic programming. Specifically, define $f(i, k)$ to be the minimum sum of times that the second intern can be assigned given that the first i tasks are assigned and the first intern's sum of task times is exactly k .
- We can then iterate over $f(n, k)$ and compute $\min_k \max(k, f(n, k))$.
- The time complexity of this solution is $\mathcal{O}(n^2 \max(a, b))$.

Problem

- Consider the infinite string obtained by writing the positive integers in increasing order.
- You are given a regex of length at most 25, find the first match of the regex within the infinite string.

Solution

- Due to the small bounds of the problem, we can solve the problem with brute force on the length of the first number that touches the regex and how many digits of the first number do not appear within the regex.
- There are a lot of tricky cases to handle, so some things that may aid in the implementation include special handling of the following cases:
 - Cases where the first number is less than 100.
 - Cases where a carry is observed to a new power of 10.

Solution, continued

- After handling the above cases, we can reduce the given problem to: Given two regexes of x and $x + 1$ where both numbers have k digits, and x is not $9 \pmod{10}$, compute the minimum value of x that is compatible with both regexes.
- In this reduced problem, all digits are effectively uniquely determined so there is no casework.

Problem

- You are given a series of statements in a programming language called IFFY, used for defining comparators on bitstrings.
- Bounds: $0 \leq n \leq 2 \cdot 10^5$, $1 \leq k \leq 10$.

Solution

- First, parse each expression and convert it to a truth table on x and y .
- There are only 2^4 such tables an expression can evaluate to, so any IFFY program only has at most $16k^2$ useful statements.
- Then evaluate the program on every pair of bitstrings after filtering out the useless statements. In expectation, each invocation will only use 4 statements, so this process takes $4 \cdot 4^k$ time.
- Use bitsets to compute the transitive property quickly.
- Overall runtime is $\mathcal{O}(n + |s| + 4 \cdot 4^k + 8^k)$, but with a low constant factor thanks to the bitsets.

Problem

- You are given a cyclic list which is a permutation of the first n positive integers. Determine if a query list of m integers is a sublist of the permutation or its reverse.
 $1 \leq m \leq n \leq 5 \cdot 10^4$.

Solution

- The first integer in the query list can match exactly one index in the given permutation.
- From there, there are only two different orientations to check, and we can check both naively in $\mathcal{O}(m)$.

Problem

- You are given n people and m houses on a circle.
- Let $f(S)$ be the minimum cost of a matching between the n people and a set S of n houses, where the cost between the person and a house is the distance along the circle.
- Find both the minimum value of $f(S)$ and the number of sets S that achieve that minimum value.
- Note we are only counting the number of sets of houses, not the number of ways to assign people to houses.

Solution

- Draw an arc from each person to their matched house. You can prove the following statements with exchange arguments:
 - ① no arcs will intersect,
 - ② there are no unmatched houses with arcs going over them, and
 - ③ if one arc is nested within another, then both travel in the same direction.
- With these criteria, there are only two potential houses that each person could be matched with.
- We can find these houses using a stack, and build a graph using only these edges.
- Because there are strictly more houses than people, each component in this graph is a line graph.
- All possible costs on a line graph can be computed in linear time using prefix / suffix sums.

Problem

- You are given a puzzle comprising 30 beads, 10 of three different colors, and the ability to rotate beans along one of four cycles.
- Given a puzzle layout, figure out how to get all ten beads of a given color into their assigned cycles with at most 435 moves.

Solution

- We can define a primitive on one of the colored cycles by rotating it once clockwise once, rotating the shared cycle clockwise once, rotating the colored cycle counterclockwise once, and rotating the shared cycle counterclockwise once.
- This has the effect of swapping two beads in specific locations and reordering the other beads within the two cycles that are touched.
- Therefore, with six moves, we can move one bead from an incorrect cycle to a correct cycle, and in at most 180 moves we can use this primitive to solve the puzzle.

Problem

- You are in the top-left corner of an $r \times c$ grid, where each square in the grid has an arrow pointing either down or right along with a timer uniformly at random from $[0, p]$. The arrow changes directions when the timer runs down to zero and resets to p .
- You can only move in the direction specified by the arrow in the square you are in.
- You can only see the arrow in the square you are currently in.
- Compute the expected time you wait for arrows while traveling to the bottom-right square with an optimal strategy.

Solution

- Define $f(x, y)$ to be the answer for a grid with x and y columns. Note that $f(x, y) = f(y, x)$.
- We can see that $f(1, n) = \frac{n-1}{4}$ since with probability $\frac{1}{2}$ we can travel directly to the right, and otherwise we wait $\frac{p}{2}$ seconds in expectation.
- $f(i, j)$ is clearly dependent only on $f(i-1, j)$ and $f(i, j-1)$ otherwise, let the values of those be a and b with $a < b$.
- Clearly, if we are able to transition to a directly, we should. Otherwise, we should wait only if waiting would take less time than going straight to b .
- The rest of the math is left as an exercise to the reader.

Problem

- A mountain is defined by a peak at point (x, y) , and by line segments that are parallel to the lines $y = x$ and $y = -x$ that go down to the y -axis, possibly clipped on the left by $x = 0$ and the right by $x = w$.
- Given a collection of mountains, we wish to compute the sum of the lengths of segments of mountains that are visible. If two mountains overlap in area, the affected segments are not visible.
- Operations consist of adding and removing mountains. After each operation, output the sum of lengths of visible segments.

Solution

- If we project all segments down to the x -axis, we see that the answer is equal to $\sqrt{2}$ times the length of segments that are covered from above by some mountain.
- To solve this problem, we can maintain a segment tree. On a given range, we maintain the minimum number of segments observed at some point within the segment, along with the sum of lengths of segments covered within.
- Though it may seem like lazy propagation is necessary, it suffices to take a given segment and just break it up into a maximal collection of disjoint segments on the segment tree, and to track how many segments fully cover a given node.
- This is the same data structure used when computing the union of area of axis-aligned rectangles that may overlap.

Problem

- You are given a partial string. Change each question mark to some letter such that the resulting string has exactly k instances of NAC as a subsequence.

Not Another Constructive!

Solution

- Naively, we can recursively backtrack on the string, keeping track of the prefix that has been fully determined, the number of N, NA, and number of NAC subsequences.
- This is too slow, but with memoization there are at most $\mathcal{O}(n^6)$ states.
- Most of these states are actually not reachable, the true number of reachable states is less than $5 \cdot 10^7$. With some pruning, the number of reachable states is small enough that recursive backtracking will also pass.

Problem

- You have a passport with n pages. You will take a trip where on each trip, t_i contiguous pages are stamped and unusable in the future. Compute the first trip where you are at risk of not having enough contiguous pages.

Solution

- To check if the k -th trip is at risk, we see that $n - \sum_{i=1}^{k-1} t_i$ pages will be available, and there are at most k sections that they can be in. If $\left\lfloor \frac{n - \sum_{i=1}^{k-1} t_i}{k} \right\rfloor \geq t_k$, then it is guaranteed to be fine.
- We can therefore check k in increasing order.
- Note that the condition is not binary searchable!

Problem

- A point light source is placed somewhere on the negative x -axis. n vertical segments cast shadows from the light source on a wall. Compute the sum of lengths of intervals where the shadows formed on the wall form a single contiguous interval. $1 \leq n \leq 3 \cdot 10^3$.

Solution

- When the light source is at negative infinity, this problem reduces to checking if the projection of the vertical segments on the wall forms a single contiguous interval.
- As the light source moves towards the origin, if we consider an angle sweep of the light from $y = -\infty$ to $y = \infty$, the ordering of the endpoints of the segments changes. Specifically, two segment endpoints change their relative ordering at the intersection of the negative x -axis and the line connecting the two endpoints.
- A given x coordinate is therefore valid if and only if the number of “bottom” endpoints exceeds the number of “top” endpoints at every point in the angle sweep until every segment has been processed.

Solution, continued

- Naively, we can check the validity at all critical points in $\mathcal{O}(n^3)$.
- To optimize this to $\mathcal{O}(n^2 \log n)$, we maintain the prefix sum of count of bottom segment endpoints minus count of top segment endpoints. When we swap two points in the relative ordering, one prefix sum either increases by 2 or decreases by 2.
- It is also possible to solve this problem with offline dynamic connectivity techniques, though the constant factor of this solution may need to be optimized.

Square of Triangles

Problem

- You are given four triangles. Determine if they can be arranged to form a square.

Solution

- This problem can be solved either by exhaustive search on arranging the triangles, or by exhaustively listing all arrangements of a square into four triangles and doing some tedious casework.
- A couple cases that teams commonly missed include:
 - Having one triangle take one side of the square and part of an adjacent side, and partitioning the remaining trapezoid into three triangles.
 - Having one triangle be half the square and partitioning the other half of the square into three triangles.

Problem

- Ashley is solving n problems.
- Her skill can be modeled with two integers (s, t) .
- A problem can only be solved if $l_i \leq s \leq r_i$ and $l_t \leq t \leq r_t$.
- After solving a problem, she may increase one of her skill points by 1.
- What is the maximum number of problems she can solve?

Solution

- Keep a set of (s, t) values that Ashley can achieve before solving a problem.
- For each such value, if she can solve the problem, then add $(s + 1, t)$ and $(s, t + 1)$ to the set.
- Implemented naively, this runs in $\mathcal{O}(n^3)$, which is too slow for the given bounds.
- This complexity can be improved by removing elements from the set after processing them.
- Since the bounds are large, it is recommended to use a bitset to store the states and find elements in the set quickly.
- With proper implementation to only iterate over set bits, this reduces the runtime to $\mathcal{O}(n^2)$.