

2023 South Conference Division 1 Solutions

The Judges

Feb 24, 2024

Problem

- You are given a string with n A's, B's, and C's. Partition the string into the minimum number of subsequences such that each subsequence is the concatenation of several permutations of ABC, not necessarily the same permutation in each instance.

Initial Observations

- Note that the first three letters of a string are some permutation of ABC, then if we partition the remainder of the string optimally, we can always prepend the first three characters to an arbitrary subsequence.
- Therefore, we can greedily assign characters to subsequences looping over the characters in order.

Solution

- If there is an active subsequence of length 2 that does not contain the character we are currently considering, assign the current character to that subsequence and set that subsequence to be empty.
- Otherwise, if there is some subsequence of length 1 that does not contain the character we are currently considering, assign the current character to that subsequence.
- Otherwise, create a new subsequence consisting of just that character.
- Since we always reset the subsequence lengths when they reach 3, this algorithm runs in $\mathcal{O}(|s|)$.

Balanced Tree Path

Problem

- You are given a tree where each node has an opening or closing bracket, brace, or parenthesis. A path along the tree generates a string by concatenating all characters on the path. The string is balanced if the characters match up properly. Compute the number of paths that yield balanced strings.

Initial Observations

- There are $\mathcal{O}(n^2)$ distinct paths in the tree, so if we can check them all in $\mathcal{O}(1)$ time, that would run in time.
- Naively, it takes $\mathcal{O}(n)$ time to check if a string is balanced, but due to the nature of the paths, a lot of redundant work is performed.

Solution

- We can DFS through the tree, maintaining a stack of characters seen.
- If the current character is a closing character, make sure it matches properly, otherwise terminate the recursion. If it matches properly, pop the character from the stack.
- Otherwise, the current character is an opening character and should be pushed on to the stack.
- If the stack is empty, we have observed a balanced path.
- This optimizes the $\mathcal{O}(n^3)$ solution to $\mathcal{O}(n^2)$.

Problem

- You are given a binary grid where n entries are 1. Compute, over all entries that are zero, the minimum number of 1 entries that are Euclidean distance at most d away from the given zero entry.

Initial Observations

- n and d are small relative to the grid size. If we can get an answer of zero, we can immediately short-circuit.

Solution

- We start by doing a line sweep. For a fixed row, compute for each 1 entry which columns in the given row it affects. Since these columns are contiguous, we can efficiently compute this by first maintaining a difference array and then computing a running sum over the difference array, taking the minimum over the difference array over all entries that are zero.
- A given circle has $\mathcal{O}(d)$ interesting points along its circumference, so this line sweep only needs to consider $\mathcal{O}(nd)$ intervals. Due to the difference array, some sorting is required internally, so this algorithm runs in $\mathcal{O}(nd \log n)$.

Ordered Problem Set

Problem

- You are given a permutation of the first n positive integers. Compute all values of k such that k divides n and the first k integers in the array are a permutation of 1 through k , the next k integers are a permutation of $k + 1$ through $2k$, and so on.

Solution

- Because n is small, we can check this with brute force for all possible values of k .
- Care should be exercised to confirm that k is divisible by n .

Problem

- You are given a tree of n vertices. For a given k , compute the lexicographically smallest permutation of vertices such that exactly $n - k - 1$ pairs of adjacent vertices in the permutation are connected by an edge.

Initial Observations

- For a tree, if l is the fewest number and r is the largest number such that exactly l and r pairs of adjacent vertices in some permutation could be connected by an edge, then all intermediate values are attainable.
- l is almost always 0. Otherwise, it is 1.

Solution

- We will incrementally build the permutation element by element. Given the permutation so far, we check to see if appending the current element would still be valid. To compute r , there are multiple ways to compute the maximum matching. This can be done with a tree DP or a greedy assignment with a precomputed postorder traversal of the tree.
- We need to carefully check if $l = 1$, which requires some casework based on the previous node and whether the rest of the tree is a star.
- Since checking that a node is valid takes $\mathcal{O}(n)$ time, this algorithm runs in $\mathcal{O}(n^3)$ time.

Problem

- You are given a binary matrix and the ability to toggle individual entries. Compute the minimum number of entries to toggle such that each row and each column have a 1, and within each row and each column, all 1 are contiguous.

Initial Observations

- If we iterate over the rows in order, we can maintain a DP that maps the exact interval that contains 1's to the minimum number of operations needed to get exactly that interval to contain 1's.
- $r \times c \leq 2 \cdot 10^5$, which means that at least one of r and c must be less than or equal to $\sqrt{2 \cdot 10^5}$.

DP formulation

- Our DP tracks $f(r_i, c_l, c_r)$ which is equal to the minimum number of operations needed such that the matrix up to row r_i has at least one 1 in every column from column 1 to column c_r and one 1 in every row from row 1 to row r_i , and row r_i contains 1's exactly on the interval $[c_l, c_r]$.
- We can transition from $f(r_i, c_l, c_r)$ to $f(r_i + 1, c'_l, c'_r)$ if and only if $c'_l \geq c_l$ and $c'_r \geq c_r$. There are $\mathcal{O}(rc^2)$ states and $\mathcal{O}(c^2)$ transitions per state, so the DP currently runs in $\mathcal{O}(rc^4)$.

DP formulation

- Assume without loss of generality that $r \geq c$ - transposing the matrix and solving this problem on the transposed matrix is equivalent.
- This observation gets us down to $\mathcal{O}(r^2 c^2 \min(r, c))$ which is still too slow.
- If we look harder at the transition for $f(r_i + 1, c'_l, c'_r)$, we see that the transition uses the 2D prefix minimum for $f(r_i, c_l, c_r)$. If we compute these prefix minima, we can optimize the transition down to $\mathcal{O}(1)$. The number of states is now $\mathcal{O}(rc \min(r, c))$, and since $\min(r, c) \leq \sqrt{2 \cdot 10^5}$, this is fast enough.

Problem

- You have a team of n engineers, each of who is familiar with some of m services. A team has a *robustness level* of k if, no matter which k services go down, k different engineers can each be assigned to one of the services such that each engineer is familiar with their assigned service. Compute the team's robustness level.

Initial Observations

- m is much smaller than n .
- Hall's Theorem gives us a simple way to check if a bipartite graph has a perfect matching.
- By Hall's Theorem, if the robustness level of a team is k , then there exists some collection of x engineers where $x > k$ and the union of the services they are familiar with is some subset of size k .

Solution

- For each subset of x engineers, we compute the size of the union of the services they are familiar with. If the size of the union, y , is greater than or equal to x , we get no information on the robustness level. Otherwise, the robustness of the team must be at most y by Hall's Theorem.
- Naively, this runs in $\mathcal{O}(n2^m)$. We can optimize this to $\mathcal{O}(nm + m2^m)$ by short-circuiting when we observe that the size of the union must be at least x .

Problem

- You are given an array of n integers. In a single operation, you can take a subarray and change all entries in the subarray to some other integer. Compute the minimum number of operations needed to get to a specific configuration.

Initial Observations

- If the outermost element of the array has the correct value, it is never optimal to change it.
- This gives us an $\mathcal{O}(n^4)$ DP where, for a fixed interval and the current color of the leftmost element in the interval, we can maintain the minimum number of operations needed to get that state.

Solution

- Our DP tracks $f(l, r, c)$ which is equal to the minimum number of operations needed to get the subarray $[l, r]$ correctly assigned, given that l currently has a value of c .
- If l is correctly assigned, then we compute $f(l + 1, r, c)$. Otherwise, we have to assign l to the correct color and also determine which prefix it gets colored with. This DP therefore has $\mathcal{O}(n^3)$ states and $\mathcal{O}(n)$ transitions, which gives an $\mathcal{O}(n^4)$ algorithm which is too slow.
- To optimize this, note that there is no reason to explicitly maintain DP states where c does not match element l , since we know that we are forced to set that element properly.
- Therefore, we do not need to maintain c in our state, reducing the number of states to $\mathcal{O}(n^2)$. This DP runs in $\mathcal{O}(n^3)$ and is fast enough.

Problem

- You are given a UI with some items initially selected. You can click between pages, toggle individual selections, select all items on a page, or deselect all items on a page. Compute the minimum number of clicks needed to get to a final configuration of items selected.

Initial Observations

- In terms of moving between pages, it is never optimal to change direction more than once. Therefore, one should either move left and then move right, or move right and then move left. It is possible that the movements are empty.
- On a given page, we should only use at most one of the select all and deselect all options.

Solution

- For each page, we do some casework to compute the minimum number of button clicks based on whether we use select all, deselect all, or neither.
- We then keep track of all pages that require changes. We compute the minimum number of clicks needed to navigate pages by either going left and then going right, or going right and then going left.

Sequence Guessing

Problem

- You are asked to generate a sorted sequence of integers starting with 0 ending with 10^5 where adjacent entries differ by either 1 or 2. You tell an adversary how long your sequence is, then the adversary will guess integers in $[0, 10^5]$ and you must tell the adversary either the index of that integer in your list or assert that it is not present. Generate and maintain such a sequence while forcing the adversary to guess an integer that is not present 33333 times.

Initial Observations

- $\left\lfloor \frac{10^5}{3} \right\rfloor = 33333$.
- If you always include every multiple of 3, then there are 33333 pairs of adjacent integers and you can always force the adversary to miss by selecting the other integer in the pair.

Solution

- The sequence will take the form $0, a_2, 3, a_4, \dots, 99996, a_{66666}, 99999, 10^5$.
- Tell the adversary that your sequence has length 66668.
- If the adversary guesses $3x + 1$, and the adversary has not guessed $3x + 2$, report it is not present and include $3x + 2$ in your sequence.
- Otherwise, if the adversary guesses $3x + 2$, and the adversary has not guessed $3x + 1$, report it is not present and include $3x + 2$ in your sequence.
- Otherwise, the entry must already be in your sequence, report its index accurately.

Problem

- You have n serious runners and k casual runners. If you get x serious runners and y casual runners to run together in a group, and $\frac{x}{x+y} \geq \frac{a}{b}$, the y runners will all become serious runners. Compute the minimum number of runs to schedule such that the entire group only has serious runners.

Initial Observations

- In a scheduled run, all serious runners should join. It is then optimal to maximize the number of casual runners that join while still maintaining a sufficient proportion of serious runners.
- Naively simulating this is too slow.

Solution

- Given n and k , we start by computing $\lfloor \frac{nb-na}{a} \rfloor$. If this value is zero, then it is impossible. Otherwise, this is the number of casual runners we can convert in one run.
- We can also compute the number of times we need to schedule training runs until this value changes.
- This value can change at most $\sqrt{n+k}$ times, since the sum of the first x positive integers is $\mathcal{O}(x^2)$.
- The complexity of this solution is therefore $\mathcal{O}(t\sqrt{n+k})$.

Problem

- You are given n labeled points p_i and line segments $[l_i, r_i]$ on the x -axis, and you must connect each point to its segment with a newly drawn segment. The newly drawn segments may touch but cannot strictly intersect. Compute the minimum sum of lengths of the newly drawn segments.

Initial Observations

- The collection of points p_i can connect to on the line segment $[l_i, r_i]$ form a contiguous subsegment.
- Therefore, it suffices to compute the leftmost and rightmost points each p_i can connect to.

Solution

- Scan over the points from left to right in line segment order, and maintain a stack of points. While the current point is not below the topmost point in the stack, pop that point off the stack and check if the popped element's segment possibly interferes with the current point's segment. If so, update the new leftmost point that the current point can connect to. When done, push the current point on the stack.
- Repeat this process in the reverse order, updating the rightmost point each point can connect to.
- If any point now has an empty segment to connect to, it is impossible. Otherwise, the optimal arrangement follows directly from the given intervals - either draw straight down if possible, otherwise connect to the closest of the boundary points.
- Be careful to use integer arithmetic everywhere.

Problem

- Determine if an integer n is composed of the first k integers written in increasing order.

Solution

- There are multiple possible approaches. For example, given the length of n , generate the integer $123\dots k$ and see if it matches.
- Since there are only nine valid entries, one could also explicitly handle all entries with nine if statements.