NCPC 2023 Presentation of solutions

2023-10-07

NCPC 2023 solutions

Problems prepared by

- Arnar Bjarni Arnarson (Reykjavik University)
- Andreas Björklund (IT University of Copenhagen)
- Atli Fannar Franklín (University of Iceland)
- Nils Gustafsson (KTH Royal Institute of Technology)
- Oskar Haarklou Veileborg (Aarhus University)
- Asger Hautop Drewsen (Aarhus University)

Test solver

• Björn Martinsson

Problem

Convert Roman numerals to regular Arabic numbers.

Solution

Problem

Convert Roman numerals to regular Arabic numbers.

$\operatorname{Solution}$

• Using a map/dictonary to map Roman digits to numbers is convenient.

Problem

Convert Roman numerals to regular Arabic numbers.

Solution

- Using a map/dictonary to map Roman digits to numbers is convenient.
- Sor each digit in a number, check if it is smaller than anything to the right. In that case, subtract it.

Problem

Convert Roman numerals to regular Arabic numbers.

Solution

- Using a map/dictonary to map Roman digits to numbers is convenient.
- Solution of the second seco
- Make sure this runs faster than $\mathcal{O}(L^2)!$

Problem

Convert Roman numerals to regular Arabic numbers.

Solution

- Using a map/dictonary to map Roman digits to numbers is convenient.
- Or each digit in a number, check if it is smaller than anything to the right. In that case, subtract it.
- Make sure this runs faster than $\mathcal{O}(L^2)!$
- This can be done by looping from the end of the string, and keeping track of the largest digit seen so far.

Given the description of a card game and a final state of the game, implement the rules for resolving which player won the match.

- Process each player and location pair individually to apply abilities and add all the power levels of the cards together.
 - If the card is Seraphina and k is the number of friendly cards at the location, add k 1 power to the location.
 - If the card is Thunderheart and the location has 4 friendly cards, add 6 power to the location.
 - If the card is Zenith, and the location is the center location, add 5 power to the location.
- Ount the number of locations where each player wins.
- **③** If one of the players wins more locations than the other, that player wins.

Given the description of three dice, calculate whether any of them will roll higher than both of the remaining two with probability $\geq \frac{1}{2}$.

Given the description of three dice, calculate whether any of them will roll higher than both of the remaining two with probability $\geq \frac{1}{2}$.

Solution

• There are N^2 possible outcomes when rolling two dice with N sides.

Given the description of three dice, calculate whether any of them will roll higher than both of the remaining two with probability $\geq \frac{1}{2}$.

Solution

• There are N^2 possible outcomes when rolling two dice with N sides.

2 N = 6, so just try all combinations of rolls.

Karl has a string of size $2 \le N \le 10^{18}$. It is contained in a buffer of size 2N where the first N bytes are non-zero and the last N bytes are zero. You need to find N by reading from the buffer, without reading at an index which is out of bounds. You can only use $2\lceil \log_2 N \rceil$ reads.

Solution

¹https://www.youtube.com/watch?v=gWexE-YVnOs Author: Oskar Haarklou Veileborg and Asger Hautop Drewsen NCPC 2023 solutions

Karl has a string of size $2 \le N \le 10^{18}$. It is contained in a buffer of size 2N where the first N bytes are non-zero and the last N bytes are zero. You need to find N by reading from the buffer, without reading at an index which is out of bounds. You can only use $2\lceil \log_2 N \rceil$ reads.

Solution

Linear search uses too many queries.

¹https://www.youtube.com/watch?v=gWexE-YVnOs

Karl has a string of size $2 \le N \le 10^{18}$. It is contained in a buffer of size 2N where the first N bytes are non-zero and the last N bytes are zero. You need to find N by reading from the buffer, without reading at an index which is out of bounds. You can only use $2\lceil \log_2 N \rceil$ reads.

Solution

- Linear search uses too many queries.
- ② First idea: Use binary search instead of linear search to reduce number of queries.

¹https://www.youtube.com/watch?v=gWexE-YVnOs

Karl has a string of size $2 \le N \le 10^{18}$. It is contained in a buffer of size 2N where the first N bytes are non-zero and the last N bytes are zero. You need to find N by reading from the buffer, without reading at an index which is out of bounds. You can only use $2\lceil \log_2 N \rceil$ reads.

Solution

- Linear search uses too many queries.
- Sirst idea: Use binary search instead of linear search to reduce number of queries.
 - Problem 1: Uses too many queries if N is small.

¹https://www.youtube.com/watch?v=gWexE-YVnOs

Karl has a string of size $2 \le N \le 10^{18}$. It is contained in a buffer of size 2N where the first N bytes are non-zero and the last N bytes are zero. You need to find N by reading from the buffer, without reading at an index which is out of bounds. You can only use $2\lceil \log_2 N \rceil$ reads.

Solution

- Linear search uses too many queries.
- **②** First idea: Use binary search instead of linear search to reduce number of queries.
 - Problem 1: Uses too many queries if N is small.
 - **2** Problem 2: Will read out of bounds if N is small.

¹https://www.youtube.com/watch?v=gWexE-YVnOs

• First idea: Use binary search instead of linear search to reduce number of queries.

- First idea: Use binary search instead of linear search to reduce number of queries.
- Second idea: First do exponential search (with base 2) to "safely" find an index *i*, where $N \le i < 2N$. Then do binary search to find N in the range $\left[\frac{i+1}{2}, i\right]$.

- First idea: Use binary search instead of linear search to reduce number of queries.
- Second idea: First do exponential search (with base 2) to "safely" find an index *i*, where $N \le i < 2N$. Then do binary search to find N in the range $\left[\frac{i+1}{2}, i\right]$.
- This can be done with $2\lceil \log_2 N \rceil$ reads in total.

- First idea: Use binary search instead of linear search to reduce number of queries.
- Second idea: First do exponential search (with base 2) to "safely" find an index *i*, where $N \le i < 2N$. Then do binary search to find N in the range $\left[\frac{i+1}{2}, i\right]$.
- This can be done with $2\lceil \log_2 N \rceil$ reads in total.
- You need to be careful to avoid off-by-one errors and to not use too many queries when N is small.



















- First idea: Use binary search instead of linear search to reduce number of queries.
- Second idea: First do exponential search (with base 2) to "safely" find an index *i*, where $N \le i < 2N$. Then do binary search to find N in the range $\left[\frac{i+1}{2}, i\right]$.
- This can be done with $2\lceil \log_2 N \rceil$ reads in total.
- You need to be careful to avoid off-by-one errors and to not use too many queries when N is small.

Problem

Partition N items among M scouts so that no one has more than two items and the maximum total size is minimised.

Solution

Problem

Partition N items among M scouts so that no one has more than two items and the maximum total size is minimised.

Solution

• Sort the N items after decreasing size.

Problem

Partition N items among M scouts so that no one has more than two items and the maximum total size is minimised.

Solution

• Sort the N items after decreasing size.

2 If $N \leq M$, the answer is the size of largest item.

Problem

Partition N items among M scouts so that no one has more than two items and the maximum total size is minimised.

Solution

- Sort the N items after decreasing size.
- 2 If $N \leq M$, the answer is the size of largest item.
- If N > M, pair up the M + 1:th largest item with the M:th largest, the M + 2:th largest item with the M 1:th largest, and so on until we have M groups of pairs and single elements. Report the maximum size.

Problem

Partition N items among M scouts so that no one has more than two items and the maximum total size is minimised.

Solution

- Sort the N items after decreasing size.
- 2 If $N \leq M$, the answer is the size of largest item.
- If N > M, pair up the M + 1:th largest item with the M:th largest, the M + 2:th largest item with the M 1:th largest, and so on until we have M groups of pairs and single elements. Report the maximum size.
- This greedy strategy works since for any four sizes $a \le b \le c \le d$, the best way to divide them in two pairs minimising the maximum is a + d and b + c.

A string is generated by repeatedly appending a 1 if the last characters contain a pattern that repeats K times, and a 0 otherwise. Count the number of ones.
A string is generated by repeatedly appending a 1 if the last characters contain a pattern that repeats K times, and a 0 otherwise. Count the number of ones.

Solution

A good idea is to write out the string for K = 3 or 4 on paper, and look for patterns.

$\operatorname{Solution}$

• Guess: the string is generated by first letting s = 0, and then applying s = s * k + 1 k times, and then append infinitely many ones.

$\operatorname{Solution}$

- Guess: the string is generated by first letting s = 0, and then applying s = s * k + 1 k times, and then append infinitely many ones.
- **2** There are now some recursive ways of counting the number of ones in $\mathcal{O}(\log(N))$.

Given a rooted tree, label the vertices with positive integers so that the divisor graph equals the rooted tree.

Given a rooted tree, label the vertices with positive integers so that the divisor graph equals the rooted tree.

Solution

O Put the value 1 on the root, there is no reason not to.

Given a rooted tree, label the vertices with positive integers so that the divisor graph equals the rooted tree.

- O Put the value 1 on the root, there is no reason not to.
- ② If the tree is a line, then the best we can do is to use powers of 2. The largest number becomes $2^{N-1} \le 2^{59} < 10^{18}$.

Given a rooted tree, label the vertices with positive integers so that the divisor graph equals the rooted tree.

- O Put the value 1 on the root, there is no reason not to.
- ② If the tree is a line, then the best we can do is to use powers of 2. The largest number becomes $2^{N-1} \le 2^{59} < 10^{18}$.
- A useful reduction: instead of x_v on vertices, consider the values $y_e = x_v/x_u$ that we put on edges. These should be small primes.

Given a rooted tree, label the vertices with positive integers so that the divisor graph equals the rooted tree.

- O Put the value 1 on the root, there is no reason not to.
- ② If the tree is a line, then the best we can do is to use powers of 2. The largest number becomes $2^{N-1} \le 2^{59} < 10^{18}$.
- A useful reduction: instead of x_v on vertices, consider the values $y_e = x_v/x_u$ that we put on edges. These should be small primes.
- It seems like a good idea to put 2s on the deepest leaf, then 3s on the second deepest, and so on.

F — Factor-Full Tree



Solution

- O Put the value 1 on the root, there is no reason not to.
- 3 If the tree is a line, then the best we can do is to use powers of 2. The largest number becomes $2^{N-1} \le 2^{59} < 10^{18}$.
- A useful reduction: instead of x_v on vertices, consider the values $y_e = x_v/x_u$ that we put on edges. These should be small primes.
- It seems like a good idea to put 2s on the deepest leaf, then 3s on the second deepest, and so on.
- This turns out to work. There are a few other ways to get numbers of size at most 2^{N-1} as well.

There are N types of components, each with f_i copies and placement time t_i . In one move you can place two components of different types i and j in max (t_i, t_j) nanoseconds. Find the minimum total time to place all components.

There are N types of components, each with f_i copies and placement time t_i . In one move you can place two components of different types i and j in max (t_i, t_j) nanoseconds. Find the minimum total time to place all components.

- Greedily place two component types with maximum t_i? Doesn't work on second sample.
- 2 If t_i values are similar, then it is more important to match all components.

Solution

• Let's try DP instead. Sort the types with respect to time.

Author: Nils Gustafsson and Bengt Evertsson NCPC 2023 solutions

- Let's try DP instead. Sort the types with respect to time.
- 2 DP(i, u) = answer of first *i* components, if there are *u* unmatched components that we already "paid" for.
- **③** Number of states is $N^2 F$, transitions take $\mathcal{O}(F)$. Running time $\mathcal{O}(N^2 F^2)$.





Author: Nils Gustafsson and Bengt Evertsson NCPC 2023 solutions

Solution

- DP(i, u) = answer of first *i* components, if there are *u* unmatched components that we already "paid" for.
- 2 Number of states is $N^2 F$, transitions take $\mathcal{O}(F)$. Running time $\mathcal{O}(N^2 F^2)$.

Solution

- DP(i, u) = answer of first *i* components, if there are *u* unmatched components that we already "paid" for.
- ② Number of states is $N^2 F$, transitions take $\mathcal{O}(F)$. Running time $\mathcal{O}(N^2 F^2)$.
- \bigcirc The value of u never exceeds F in an optimal solution.

Solution

- DP(i, u) = answer of first *i* components, if there are *u* unmatched components that we already "paid" for.
- 3 Number of states is $N^2 F$, transitions take $\mathcal{O}(F)$. Running time $\mathcal{O}(N^2 F^2)$.
- \bigcirc The value of u never exceeds F in an optimal solution.
- Write out the DP transitions carefully. They can be reduced to RMQ:s on a sliding window.

Solution

- DP(i, u) = answer of first *i* components, if there are *u* unmatched components that we already "paid" for.
- 3 Number of states is $N^2 F$, transitions take $\mathcal{O}(F)$. Running time $\mathcal{O}(N^2 F^2)$.
- \bigcirc The value of u never exceeds F in an optimal solution.
- Write out the DP transitions carefully. They can be reduced to RMQ:s on a sliding window.
- **③** Use segment tree, RMQ, or priority queue in the transition: $\mathcal{O}(NF \log(F))$.
- Bonus: can you solve it in $\mathcal{O}(NF)$ (without fancy RMQ)?

Problem

Find a proper three-coloring of the $N \le 1000$ vertices in an undirected graph on $M \le 100000$ edges, or report that none exists.

Problem

Find a proper three-coloring of the $N \le 1000$ vertices in an undirected graph on $M \le 100000$ edges, or report that none exists.

Solution - first part

1 Three-coloring is NP-hard - we need backdoor.

Problem

Find a proper three-coloring of the $N \le 1000$ vertices in an undirected graph on $M \le 100000$ edges, or report that none exists.

- Three-coloring is NP-hard we need backdoor.
- ② The backdoor here is a vertex that has a subset of few neighbors, say at most $B \le 12$, through which it reaches every other vertex in at most two steps.

Problem

Find a proper three-coloring of the $N \le 1000$ vertices in an undirected graph on $M \le 100000$ edges, or report that none exists.

- Ihree-coloring is NP-hard we need backdoor.
- ② The backdoor here is a vertex that has a subset of few neighbors, say at most $B \le 12$, through which it reaches every other vertex in at most two steps.
- Such a vertex is guaranteed to exist (e.g., the CEO even has $B \le 8$). In case there are many, anyone will do!

Problem

Find a proper three-coloring of the $N \le 1000$ vertices in an undirected graph on $M \le 100000$ edges, or report that none exists.

- Three-coloring is NP-hard we need backdoor.
- ② The backdoor here is a vertex that has a subset of few neighbors, say at most $B \le 12$, through which it reaches every other vertex in at most two steps.
- Such a vertex is guaranteed to exist (e.g., the CEO even has $B \le 8$). In case there are many, anyone will do!
- To find backdoor, first find candidate "CEO" vertices of degree ≤ 15 that reaches every vertex in at most two steps in $O(n^2)$ time.

Problem

Find a proper three-coloring of the $N \le 1000$ vertices in an undirected graph on $M \le 100000$ edges, or report that none exists.

- Three-coloring is NP-hard we need backdoor.
- ② The backdoor here is a vertex that has a subset of few neighbors, say at most $B \le 12$, through which it reaches every other vertex in at most two steps.
- Such a vertex is guaranteed to exist (e.g., the CEO even has $B \le 8$). In case there are many, anyone will do!
- To find backdoor, first find candidate "CEO" vertices of degree ≤ 15 that reaches every vertex in at most two steps in $O(n^2)$ time.
- Then for subsets of each candidate CEO's neighbors from larger to smaller, see if they still cover the graph. $O(\sum_{k=B}^{15} {\binom{15}{k}n})$ time per candidate CEO.

Find a proper three-coloring of the $N \le 1000$ vertices in an undirected graph on $M \le 100000$ edges, or report that none exists.

Solution - second part

Find a proper three-coloring of the $N \le 1000$ vertices in an undirected graph on $M \le 100000$ edges, or report that none exists.

Solution - second part

Pick an arbitrary color for the CEO, and test all 2^B ways to color the at most B spanning neighbors.

Find a proper three-coloring of the $N \le 1000$ vertices in an undirected graph on $M \le 100000$ edges, or report that none exists.

Solution - second part

- Pick an arbitrary color for the CEO, and test all 2^B ways to color the at most B spanning neighbors.
- **2** Reduce to 2SAT since every vertex now have at most two possible colors. O(N + M) time per tried coloring of spanning neighbors.

You are given n points in the plane along with a rope of length d. As it wraps around the origin you want to find out which point it ends up wrapping around.

Solution

Imagine drawing a line through the rope and taking the convex hull of the points on the left side of the rope.

You are given n points in the plane along with a rope of length d. As it wraps around the origin you want to find out which point it ends up wrapping around.

- Imagine drawing a line through the rope and taking the convex hull of the points on the left side of the rope.
- Then as you rotate the line, the next point you would touch (aside from when the rope is too short) is on this hull.

You are given n points in the plane along with a rope of length d. As it wraps around the origin you want to find out which point it ends up wrapping around.

- Imagine drawing a line through the rope and taking the convex hull of the points on the left side of the rope.
- Then as you rotate the line, the next point you would touch (aside from when the rope is too short) is on this hull.
- Thus we can solve this by maintaining convex hulls. When we miss a point due to the rope being too short we can delete that point instead.

I — Intertwined

Solution

• Using a decremental convex hull data structure we can maintain the hulls in $O(n \log(n))$ with $O(\log(n))$ deletion.

I — Intertwined

Solution

- Using a decremental convex hull data structure we can maintain the hulls in O(n log(n)) with O(log(n)) deletion.
- This data structure consists of two halves, each maintaining an upper and lower hull. They consist of a tree where the leaves are the hull points and when a point is deleted, the change is propagated to maintain the hull.
I — Intertwined

Solution

- Using a decremental convex hull data structure we can maintain the hulls in $O(n \log(n))$ with $O(\log(n))$ deletion.
- This data structure consists of two halves, each maintaining an upper and lower hull. They consist of a tree where the leaves are the hull points and when a point is deleted, the change is propagated to maintain the hull.
- We maintain one data structure for the upper half plane and iterate until we touch the hull for all points, then move to maintaining the data structure for all points instead.

Statistics at 4-hour mark: 6 submissions, 0 accepted, first after HH:MM

I — Intertwined

Solution

- Using a decremental convex hull data structure we can maintain the hulls in O(n log(n)) with O(log(n)) deletion.
- This data structure consists of two halves, each maintaining an upper and lower hull. They consist of a tree where the leaves are the hull points and when a point is deleted, the change is propagated to maintain the hull.
- We maintain one data structure for the upper half plane and iterate until we touch the hull for all points, then move to maintaining the data structure for all points instead.
- Edge case: If we travel around the entire hull without deleting any points we have to modulo the rope length with the distance traveled so we don't go in $\approx 10^9$ circles.

Statistics at 4-hour mark: 6 submissions, 0 accepted, first after HH:MM

Problem

There are N/2 berries randomly placed in a sequence of length N. You play a game where the players take turn picking all berries in 4 adjacent positions. Your opponent makes the first move, and follows a simple greedy strategy. Your goal is to get at least as many berries.

Problem

There are N/2 berries randomly placed in a sequence of length N. You play a game where the players take turn picking all berries in 4 adjacent positions. Your opponent makes the first move, and follows a simple greedy strategy. Your goal is to get at least as many berries.

Solution

It is impossible to pick more berries than grandpa.

Problem

There are N/2 berries randomly placed in a sequence of length N. You play a game where the players take turn picking all berries in 4 adjacent positions. Your opponent makes the first move, and follows a simple greedy strategy. Your goal is to get at least as many berries.

Solution

- It is impossible to pick more berries than grandpa.
- 2 The only way to tie the game is if our scores mirror grandpa's scores.

44433222211

Author: Nils Gustafsson NCPC 2023 solutions



• The game consists of four phases. In each phase, we want to make the last move.

- The game consists of four phases. In each phase, we want to make the last move.
- Main strategy: Each phase is an impartial game. We will play each game perfectly by calculating Grundy numbers.

Why and how

Grandpa plays the four impartial games randomly. Each game consists of a few thousand moves. If he has a winning strategy, then it is very likely that he will make a mistake.

Statistics at 4-hour mark: 17 submissions, 0 accepted, first after HH:MM

Why and how

- Grandpa plays the four impartial games randomly. Each game consists of a few thousand moves. If he has a winning strategy, then it is very likely that he will make a mistake.
- ² The Grundy numbers can be calculated in $\mathcal{O}(N^3)$ with DP, which is too slow. We need various tricks to speed it up.

Statistics at 4-hour mark: 17 submissions, 0 accepted, first after HH:MM

Why and how

- Grandpa plays the four impartial games randomly. Each game consists of a few thousand moves. If he has a winning strategy, then it is very likely that he will make a mistake.
- ² The Grundy numbers can be calculated in $\mathcal{O}(N^3)$ with DP, which is too slow. We need various tricks to speed it up.
- Main speedup: The sequence has many gaps with no berries, so the game can be partitioned into many small independent games. Find the Grundy number of each and XOR them.

Statistics at 4-hour mark: 17 submissions, 0 accepted, first after HH:MM

Results!

NCPC 2023 solutions