ICPC Europe Regionals

icpc global sponsor
programming tools

icpc.foundation

icpc diamond
multi-regional sponsor

The 2022 ICPC Southwestern Europe Regional Contest

Mirror Solutions

A Walking Boy

Author: Federico Glaudo
Preparation: Andrea Ciprietti

Notice that you may assume that the judge has sent a message at minute 0 and at minute 1440 and this does not change the answer (but it simplifies the reasoning and the implementation).

Let us consider two consecutive messages sent by the judge, at times $s < t$.

If $t - s < 120$, then the judge cannot have walked the dog between the two messages.

If $120 \leq t - s < 240$, then the judge may have walked at most once between the two messages.

If $240 \leq t - s$, then the judge may have walked the dog two times between the two messages.

Hence, if $a_{i+1} - a_i \geq 240$ for some i, then the answer is `YES`. If $a_{i+1} - a_i \geq 120$ for two distinct values of i, then the answer is `YES`. Otherwise the answer is `NO`.

B Vittorio Plays with LEGO Bricks

Author: Giovanni Paolini
Preparation: Alex Danilyuk

Each purple brick needs to be at the top of a chain of h bricks at heights $0, 1, \dots, h$ such that the x coordinates of consecutive bricks differ by at most 1. In an optimal structure, we may assume that every brick belongs to at least one of the n chains, and that any two chains that differ at some height h' also differ at all heights $\geq h'$.

We say that two chains diverge at height h' if their bricks coincide up to height $h' - 1$ but not at height $\geq h'$. In particular, two chains with different bricks at height 0 are said to diverge at height 0. For now, let's pretend that bricks can partially overlap with each other, so that there are no further constraints on how chains can be formed; at the end we will show how to account for overlapping bricks.

For $1 \leq l \leq r \leq n$, denote by $f(l, r)$ the minimum number of additional bricks needed to support the purple bricks $l, l + 1, \dots, r$. We will recursively compute $f(l, r)$. For $l = r$, a single chain is needed and so $f(l, r) = h$.

Suppose now that $l < r$. For any m with $l \leq m < r$, we can build a structure as follows. First, place $f(l, m)$ blocks to support the purple bricks $l, l + 1, \dots, m$. It is easy to see that we can modify such a structure so that the chain supporting the leftmost purple brick always climbs to the left, i.e., it consists of the bricks at positions $(x_l + h, 0, 0), (x_l + h - 1, 0, 1), \dots, (x_l, 0, h)$. Similarly, add $f(m+1, r)$ bricks to support the purple bricks $m + 1, m + 2, \dots, r$, while ensuring that the chain supporting the rightmost purple brick always climbs to the right. If we set $h_{lr} := \max(0, h + 1 - \lceil (x_r - x_l)/2 \rceil)$, then it is possible to further change the leftmost and rightmost chains so that they coincide at heights $0, 1, \dots, h_{lr} - 1$ (and in fact diverge at height h_{lr}). We obtained a valid structure to support the purple bricks $l, l + 1, \dots, r$ consisting of $f(l, m) + f(m + 1, r) - h_{lr}$ additional bricks.

Conversely, suppose to have an optimal structure to support the purple bricks $l, l + 1, \dots, r$. Let m be the index of any purple brick such that the m-th and the $(m + 1)$-th chain diverge at the lowest possible height h'. In particular, all chains coincide at heights $0, 1, \dots, h' - 1$. Additionally, the first m chains share no bricks with the other chains at heights $\geq h'$. Note that $h' \leq h_{lr}$, otherwise it would not be possible to support both the l-th and the r-th purple bricks. Therefore, the number of non-purple bricks is at least $f(l, m) + f(m + 1, r) - h' \geq f(l, m) + f(m + 1, r) - h_{lr}$.

This allows us to use dynamic programming to calculate $f(l, r)$ through the following recursive formula:

$$f(l, r) = \min_{l \leq m < r} f(l, m) + f(m + 1, r) - h_{lr}.$$

In particular, we can compute the answer $f(1, n)$ in $O(n^3)$ time.

To conclude, we now show that any optimal structure with overlapping bricks can be modified into an optimal structure with the same number of bricks which do not overlap. Suppose to have two overlapping bricks with x coordinates $\overline{x}$ and $\overline{x} + 1$ (and same height). There can't be any brick at the same height and with x coordinate equal to $\overline{x} - 2$, $\overline{x} - 1$, $\overline{x} + 2$, or $\overline{x} + 3$, otherwise we could remove one of the two overlapping bricks (it would not be needed to support the bricks above). Then we can move the left brick to position $\overline{x} - 1$ or the right brick to position $\overline{x} + 2$ (at least one of the two works, based on the position of the bricks in the row immediately below).

C Library game

Author: Andrea Ciprietti
Preparation: Andrea Ciprietti

Sort the numbers $x_1, x_2, \dots, x_n$ in decreasing order, i.e. $x_1 \geq x_2 \geq \cdots \geq x_n$. For a real number x, let $\lfloor x \rfloor$ and $\lceil x \rceil$ denote, respectively the floor function of x (that is, the largest integer which does not exceed x) and the ceil function of x (that is, the smallest integer which is not less than x).

Lemma. Bernardo has a winning strategy if and only if there exists an index $1 \leq k \leq n$ such that $x_k > \lfloor m/k \rfloor$.

Proof. We will show the correctness of the criterion above by exhibiting a winning strategy for both Alessia (when the condition is not satisfied) and Bernardo (when the condition is satisfied).

- First, suppose that such an index k does not exist. Alessia will play by choosing the number x_i in her i-th turn (recall that the x_i's are sorted decreasingly). Let us show that, at the beginning of Alessia's i-th turn, there is necessarily an interval of length x_i that does not contain any number selected by Bernardo. Indeed, Bernardo has so far selected $i-1$ numbers, which form i "gaps" among the numbers $1, 2, \dots, m$. By the pigeonhole principle, one of these gaps contains at least $\lceil \frac{m-i+1}{i} \rceil = \lfloor \frac{m}{i} \rfloor \geq x_i$, where the last inequality follows from our hypothesis. Then, Alessia can safely choose this interval.

- Now suppose that $x_k > \lfloor m/k \rfloor$ for some k. Bernardo will play every turn as follows: if the interval chosen by Alessia contains at least one multiple of x_k, he selects one of those multiples; otherwise, he selects any number in the interval. Note that, every time Alessia chooses an x_i with $i \leq k$, whatever interval she chooses next will contain a multiple of x_k (because $x_i \geq x_k$). Since there are $\lfloor m/x_k \rfloor$ multiples of x_k in $[1, m]$, and the condition $x_k > \lfloor m/k \rfloor$ is equivalent to $k > \lfloor m/x_k \rfloor$, again by the pigeonhole principle there will necessarily be a turn where the multiple of x_k selected by Bernardo was already selected previously.

Implementing the strategies described is not difficult, due to the generous constraints that allow for implementations with $\mathcal{O}(nm)$ runtime.

Remark. In fact, if a "bad" k exists it is necessarily greater than 1, since $a_1 \leq m$ by the problem assumptions.

Remark. In the proofs above, we took for granted that $\lceil \frac{m-i+1}{i} \rceil = \lfloor \frac{m}{i} \rfloor$ and that $x_k > \lfloor m/k \rfloor$ if and only if $k > \lfloor m/x_k \rfloor$. The former is the well-known $\lceil \frac{a}{b} \rceil = \lfloor \frac{a+b-1}{b} \rfloor$ when a, b are integers (just replace a with $m-i+1$ and b with i). The latter can be proved by observing that both inequalities are equivalent to $k \cdot x_k > m$.

Understanding the game's criterion

How does one come up with Bernardo's winning condition $x_k > \lfloor m/k \rfloor$? While everyone has their own combination of methods, intuition and luck, in this problem it can be particularly useful to focus on small values of n, and specifically on the case $n = 2$ ($n = 1$ is not interesting at all: Alessia always wins!).

For $n = 2$, it might be easier to visualize what is going on, and to figure out that Bernardo wants to select the first number so that it is as close as possible to $m/2$ (indeed, he wants to minimize the largest of the two gaps that the selected number creates). After understanding this case, one can try to generalize to other values of n.

D Teamwork

Author: Giovanni Paolini
Preparation: Giovanni Paolini

To start, we determine whether it is possible to solve a easy problems, b medium problems, and c hard problems in a contest that lasts l time units. Once we are able to do this, it is easy to find the optimal number of problems that can be solved; we will do it in the "Full solution" section below.

We are going to prove that it is possible to solve a easy problems, b medium problems, and c hard problems if and only if the following two constraints are satisfied:

$$(1) \qquad l \geq a+b+c+\begin{cases} 0 & \text{if } a=b=c=0 \\ 1 & \text{if } a \geq 1 \\ 2 & \text{if } a=0 \text{ and at least one of } b \text{ and } c \text{ is } \geq 1 \end{cases}$$

$$(2) \qquad 3l \geq 2a+3b+4c+\begin{cases} 0 & \text{if } a=b=c=0 \\ 3 & \text{if } a,b,c \geq 1 \\ 4 & \text{if at least one of } a,b,c \text{ is } 0 \text{ but at least one of } a,b \text{ is } \geq 1 \\ 6 & \text{if } a=b=0 \text{ and } c \geq 1 \end{cases}$$

These two constraints might seem complicated at first glance, but for the most part, they are quite easy to come up with. They have a very natural interpretation in terms of available computer time and contestant time.

Proof that the constraints are necessary

In this section, we show that constraints (1) and (2) above need to be satisfied for it to be possible to solve the given problems in l time units.

Any problem takes 1 time unit of computer time and there are l time units in total, so $l \geq a+b+c$. We only need to strengthen this inequality a little bit in order to obtain constraint (1). Notice that the computer is necessarily idle during the first time unit of the contest, because any problem's solution starts with at least 1 time unit of non-computer time. Therefore, if at least one of a, b, c is ≥ 1, we must have $l \geq a+b+c+1$. Similarly, if $a = 0$ and at least one of b and c is ≥ 1, then any problem's solution starts with at least 2 time units of non-computer time; thus the computer is idle during the first 2 time units and we must have $l \geq a+b+c+2$. This proves constraint (1). Note that we could further strengthen the constraint to $l \geq a+b+c+3$ if $a = b = 0$ and $c \geq 1$, but such improvement is not needed (constraint (2) is stronger than (1) in this case).

Apart from computer time, our other limited resource is contestant time, of which there is a total of $3l$ time units (l time units for each of the 3 contestant). Each easy problem requires 2 units of contestant time, each medium problem requires 3 units, and each hard problem requires 4 units. Therefore, we must have $3l \geq 2a + 3b + 4c$. We are going to strengthen this inequality to obtain constraint (2).

If $a = b = c = 0$ there is nothing more to do, so suppose from now on that at least one of a, b, c is ≥ 1. Then the contest must last at least 2 time units for it to be possible to solve all problems. The last time unit of the contest can be used by at most 1 contestant, who is completing a problem's solution using the shared computer. Similarly, the last 2 time units can be used by at most 2

contestants: one contestant can complete a problem's solution at time l (using the computer in the last time unit of the contest) and one can complete a problem's solution at time $l-1$ (using the computer in the second to last time unit). In total, there are at least 3 units of contestant time that are necessarily wasted. Therefore, we must have $3l \geq 2a + 3b + 4c + 3$.

The contest can last 2 time units only if $a = 1$ and $b = c = 0$, and in this case we have $3l = 2a + 3b + 4c + 4$ (so constraint (2) holds). Assume from now on that the contest lasts at least 3 time units. If at least one of a, b, c is 0, then it is not possible for all 3 contestants to start solving a problem at time 0, because at least two contestants would finish solving a problem at the same time; therefore, at least one additional unit of contestant time is wasted, so $3l \geq 2a + 3b + 4c + 4$. This improvement is useful for instance in the case $a = 1$, $b = 0$, $c = 1$, where we get $3l \geq 10$ and so $l \geq 4$ (in this case, constraint (1) only yields $l \geq 3$).

Finally, we need to strengthen the inequality once more when $a = b = 0$ and $c \geq 1$. In this case, the contest necessarily lasts at least 4 time units. At most one contestant can start solving a problem at time 0, and at most one other contestant can start solving a problem at time 1. This means that at least 2 units of contestant time are wasted between time 0 and time 1, and at least 1 additional unit is wasted between time 1 and time 2. If we add the 3 units of contestant time that are wasted at the end of the contest, we obtain that $3l \geq 2a + 3b + 4c + 6$.

Proof that the constraints are sufficient & construction of a strategy

In this section, we inductively show that constraints (1) and (2) are sufficient for it to be possible to solve all problems. We do so by describing an explicit greedy strategy built by induction on $a + b + c$.

We first observe that constraint (1) implies constraint (2) whenever $a \geq c + 1$. Indeed, in this case we have $a + b + c + 1 \geq \frac{1}{3}(2a + 3b + 4c + 4)$.

Suppose that $b \geq 2$, or $b = 1$ and $a \geq c + 1$. Note that constraint (1) implies $l \geq 3$. By induction, it is possible to solve a easy problems, $b - 1$ medium problems, and c hard problems in $l - 1$ time units, because both constraints remain satisfied (if $b = 1$ and $a \geq c + 1$, constraint (2) is implied by constraint (1) which remains satisfied). Given any strategy to do so, at least one contestant is idle between time $l - 3$ and time $l - 1$, because otherwise, the computer would need to be used by all 3 contestants within these 2 time units. Any such idle contestant can then solve the remaining medium problem between time $l - 3$ and time l; this gives us a strategy to solve all problems in l time units.

Suppose now that $a \geq 2$ and $c \geq 1$. Note that constraint (1) implies $l \geq 4$. By induction, it is possible to solve $a - 1$ easy problems, b medium problems, and $c - 1$ hard problems in $l - 2$ time units, because both constraints remain satisfied (if $c = 1$, constraint (2) is implied by constraint (1) which remains satisfied). Given any strategy to do so, at least one contestant (say, x) is idle between time $l - 4$ and time $l - 2$, and at least one other contestant (say, y) is idle between time $l - 3$ and time $l - 2$. Contestant x can then solve a hard problem between time $l - 4$ and time l, and contestant y can solve an easy problem between time $l - 3$ and time $l - 1$.

We are left with the following four small cases (leaving out $a = b = c = 0$ where there is nothing to solve).

- $b = c = 0$. By constraint (1), we have $l \geq a + 1$. A strategy is to have the i-th easy problem solved between time $i - 1$ and time $i + 1$, by the first contestant if i is odd and by the second contestant if i is even. The third contestant gets bored and leaves the contest hall to get some pizza.

- $a \leq 1$, $b = 0$. By constraint (2), we have $3l \geq 4c + 6$ (regardless of whether $a = 0$ or $a = 1$), so

$l \geq \lceil \frac{4}{3}c + 2 \rceil$. If $a = 1$, the easy problem can be solved by the first contestant between time 0 and time 2. Then, the i-th hard problem can be solved by contestant $(i \bmod 3) + 1$ between time $\lceil \frac{4}{3}i - 2 \rceil$ and time $\lceil \frac{4}{3}i + 2 \rceil$.

- $a = 0$ and $b = 1$. By constraint (2), we have $3l \geq 4c + 7$, so $l \geq \lceil \frac{4}{3}c + \frac{7}{3} \rceil$. We can have the first contestant solve the only medium problem between time 0 and time 3. Then, the i-th hard problem can be solved by contestant $(i \bmod 3) + 1$ between time $\lceil \frac{4}{3}i - \frac{5}{3} \rceil$ and time $\lceil \frac{4}{3}i + \frac{7}{3} \rceil$.
- $a = b = 1$, $c \geq 1$. By constraint (2), we have $3l \geq 4c + 8$, so $l \geq \lceil \frac{4}{3}c + \frac{8}{3} \rceil$. We can have the first contestant solve the only small problem between time 0 and time 2. and the second contestant solve the only medium problem between time 0 and time 3. Then, the i-th hard problem can be solved by contestant $(i + 1 \bmod 3) + 1$ between time $\lceil \frac{4}{3}i - \frac{4}{3} \rceil$ and time $\lceil \frac{4}{3}i + \frac{8}{3} \rceil$.

Alternative construction of a strategy

The following is a simpler greedy strategy to solve a easy problems, b medium problems, and c hard problems in the minimum amount of time.

Consider the time units one by one, from the beginning of the contest. At the i-th time unit (for $i = 1, 2, \dots$), do at most one of the following:

- if at least one hard problem remains, $i \geq 4$, and at least one contestant has been idle between time $i - 4$ and time i, then let one such contestant solve a hard problem between time $i - 4$ and time i;
- if at least one medium problem remains, $i \geq 3$, and at least one contestant has been idle between time $i - 3$ and time i, then let one such contestant solve a medium problem between time $i - 3$ and time i;
- if at least one easy problem remains, $i \geq 2$, and at least one contestant has been idle between time $i - 2$ and time i, then let one such contestant solve an easy problem between time $i - 2$ and time i.

If more than one option applies, choose the first one that applies (i.e., prioritize harder problems over easier problems). If no option applies at the i-th time unit, then the computer will be idle between time $i - 1$ and time i (this surely happens for $i = 1$). Note that this strategy coincides with the one from the previous section in the four "small" cases.

For simplicity, we only prove that the constructed strategy is optimal for $a, b, c \geq 1$; the remaining cases are left to the reader. The strategy begins by solving an easy problem between time 0 and time 2, a medium problem between time 0 and time 3, and a hard problem between time 0 and time 4. Then, it proceeds by solving all remaining medium problems up to time $b + 3$. Afterwards, it alternatively solves an easy and a hard problem up to time $b + 2\min(a, c) + 1$. Finally, the remaining easy or hard problems are solved. If $a \geq c$, then only 1 unit of computer time is wasted, so constraint (1) is tight and the strategy is optimal. On the other hand, if $a \leq c$, then at most 5 units of contestant time are wasted, so constraint (2) is tight and the strategy is optimal.

Full solution

Let us now go back to the original task, which requires us to determine the maximum number of problems n that can be solved in time l. To do so, we run a binary search on n in the interval $[0, a + b + c]$. For a given candidate value of n, we greedily decide the types of problems to be solved: first up to a easy problems, then up to b medium problems, then up to c hard problems, until we meet the desired total number of problems n (obviously, there is no advantage in planning to solve a problem while skipping a strictly easier one). Thanks to the previous sections, we are

able to determine in $O(1)$ time whether it is possible to solve the n chosen problems. Once the optimal value of n is found, we can print an explicit strategy in $O(n)$ time. Overall, the complexity of this solution is $O(\log(a + b + c) + n)$.

To determine whether it is possible to solve the n chosen problems, we can also greedily construct an optimal strategy and check whether it takes at most l time units. Such a solution does not explicitly rely on checking constraints (1) and (2).

E Crossing the Railways

Author: Cesc Folch
Preparation: Cesc Folch

First of all, let us transform the given problem to a geometry problem. We consider the plane where the x coordinate represents the distance of Isona from the first validating machine (remember that Isona is always in the straight segment between the two validating machines), and the y coordinate represents time.

Thus, the point (a, b) of the plane corresponds to being a meters away from the first validating machine after b seconds. If we represent trains in this plane they are vertical segments from (r_i, a_i) to (r_i, b_i).

Let us look at how is represented the fact that Isona is m_j meters away from the validating machine and starts moving at constant speed v_j at time t_j, and does so for s_j seconds. We get a straight segment with endpoints (m_j, t_j) and $(m_j + v_j \cdot s_j,\, t_j + s_j)$.

So, the movement of Isona at constant speed for a period of time is represented by a straight segment.

The constraint on the speed of isona corresponds to a constraint on the slope of the segment. Let us suppose that a segment passes through the points (m_1, t_1) and (m_2, t_2) with $t_1 < t_2$. Since Isona cannot run backward, it must hold $m_1 \leq m_2$. Isona's constant speed is $v_I = \frac{m_2 - m_1}{t_2 - t_1}$. From the statement we know that $\frac{1}{v} \geq v_I \geq 0$, so we can write $\frac{t_2 - t_1}{m_2 - m_1} \geq v$. That means that the slope of the segment must be greater or equal then v.

From now on we will consider a segment to be *valid* if it does not cross any train segment (take into account that train segments are open, so the endpoints are not considered part of the segment) and has a valid slope (that is $m_1 = m_2$ or $\frac{t_2 - t_1}{m_2 - m_1} \geq v$).

The problem is equivalent to finding the minimum number of segments that constitute a path from $x = 0$ to $x = m + 1$ that does not intersect any train segment. This statement is similar to a BFS where the nodes are the valid segments and the edges the intersection points. The source node is the segment $(0, 0) - (0, s)$ and the target node is $(m + 1, 0) - (m + 1, s)$. The issue is that we have infinitely many valid segments. The following lemma allows us to consider only finitely many segments.

Lemma. A *special* segment is a valid segment such that at least one of the following holds:

- It contains $(0, 0)$ and has slope equal to v.
- It contains the upper endpoint of a train (i.e., (r_i, b_i)) and has slope equal to v.
- It contains two points (r_i, a_i) and (r_j, b_j) with $r_i < r_j$.

We can construct an optimal solution using only special segments.

From the previous lemma we see that we only need to consider $O(n^2)$ segments, and we can construct the full set of segments from the lemma in time $O(m \cdot n^2 \cdot \log(n))$, as for each valid segment we only need to check if it crosses any train segment, which can be done using binary search in each rail in total time $O(m \cdot \log(n))$.

But there may be $O(n^4)$ intersection points, so running a BFS on this graph is too slow. So let us see when Isona shall change speed in an optimal solution. She may change speed at a given railway or

between two railways. It is never convienient to change speed between the first validating machine and the first railway or between the last railway and the second validating machine.

Lemma. Any solution using only special segments changes speed at most once between any two railways.

Proof. Any solution changing speed twice between two rails is using a segment which is strictly contained between two railways, hence it cannot be a special segment because, by definition, any special segment has at least one point with coordinate x that is integer.

With this new observation, we can now proceed to describe an efficient algorithm. For each $i = 1, \dots, m$, and for each segment j, we keep a value $d_j^{(i)}$ that denotes the minimum number of changes of speed that are necessary to reach this segment, starting from $x = 0$, considering only the region $0 \leq x \leq i$. If we know $d_j^{(m)}$ then we know the answer to the problem.

Given $d_j^{(i)}$ for all segments j, we will efficiently compute $d_j^{(i+1)}$ for all segments j.

Let us see how to deal with the intersections (the changes of speed) between two railways (or on one railway).

Let us consider all the special segments that intersect both railway i and railway $i+1$; these segments will be indexed by an integer j (from here on, we consider only these segments, as the other ones are not important for the changes that happen between railway i and railway $i+1$). Let us identify the j-th segment with a pair of integers (a_j, b_j) where a_j and b_j are the positions that the segment s_j has if we order the segments by the y coordinate of the crossing point with $x = i$ and $x = i + 1$ respectively. From now on we will only use the pair (a_j, b_j) to work with the segments.

Notice that having the pairs (a_j, b_j) is sufficient to tell if two segments intersect (taking care of intersections on the railways is a technical detail we skip in this explanation). If we have two segments j, k with $a_j > a_k$ and $b_j < b_k$ it means that they cross. Also if $a_j < a_k$ and $b_j > b_k$, then they cross.

For all segments j, it is trivial that $d_j^{(i+1)} \leq d_j^{(i)}$. Moreover, if segment j and segment k cross, then necessarly $d_j^{(i+1)} < d_k^{(i)} + 1$. These two observations allow us to compute $d_j^{(i+1)}$, but how can we do it efficiently?

We shall use a Fenwick tree (or a segment tree) to efficiently compute the minimum over an interval.

First, for each j, we consider the crossings such that $a_j > a_k$ and $b_j < b_k$. We process them in increasing order of a_j. When processing the j-th segment, we update $d_j^{(i+1)}$ with the minimum over the interval $[b_j + 1, \infty]$ (if such value is smaller than the current value of d'_j). Then, we update the Fenwick with the value $d_j^{(i)} + 1$ at position b_j.

The other type of crossings, those such that $a_j < a_k$ and $b_j > b_k$, can be processes analogously by considering the segments in increasing order of b_j.

This allow us to process all the crossings between two railways in $O(n^2 \cdot \log(n))$. So, at the end, the total time complexity is $O(m \cdot n^2 \cdot \log(n))$.

F Train Splitting

Author: Alex Danilyuk
Preparation: Alex Danilyuk

The problem is equivalent to:

Formal statement: Given a connected graph, paint its edges with several colors so that the edges of any single color do not make the graph connected, but any 2 colors together make the graph connected.

This is a constructive problem and there can be a lot of different approaches. We left the limitations small on purpose, so that you could let your imagination run wild. We will describe a simple solution which works in time proportional to the size of the input.

How can we make sure the graph is not connected? Take a vertex and remove all incident edges. So if we paint edges incident to one vertex with color 1 and all other edges with color 2, then color 2 will not connect the graph, while colors 1 and 2 together will connect. And what about color 1 alone? It will connect the graph if and only if the chosen vertex was connected to all of the other vertices, so if we can find a vertex that is not connected to all of the other vertices, we are done.

Unless the graph is complete, i.e., it contains all possible edges, we can find such a vertex. If, on the other hand, the graph is complete, we can paint the edges from one vertex with two different colors (at least one of such edges with one color and at least one of such edges with the other), and all the other edges with the third color. You can verify that all the conditions are satisfied by this coloring.

The complexity of this solution is $O(n + m)$.

G Another Wine Tasting Event

Author: Andrea Ciprietti
Preparation: Andrea Ciprietti

For an interval $[l, r]$ $(1 \leq l \leq r \leq 2n-1)$, let $w(l, r)$ be the number of `W`'s in the substring $s_l \dots s_r$.

We are going to show that $x = \max_{1 \leq l \leq n} w(l,\, l+n-1)$ is a solution (that is, there are at least n intervals of length $\geq n$ containing exactly x `W`'s).

Let k be an index such that $w(k,\, k+n-1) = x$. The idea is to find $n-1$ other intervals with the same number of `W`'s by "sliding" the original interval $[k,\, k+n-1]$ in a clever way.

For each $l = 1, 2, \dots, k-1$, let L_l be the shortest interval starting at l with $w(L_l) = x$. Such interval exists because $w(l,\, k+n-1) \geq w(k,\, k+n-1) = x$ and moreover the length of L_l cannot be smaller than n by definition of x. Similarly, for each $r = k+n, k+n+1, \dots, 2n-1$, let R_r be the shortest interval ending at r with $w(R_r) = x$.

We have constructed a family of n intervals: $L_1,\, L_2,\, \dots,\, L_{k-1},\, [k,\, k+n-1],\, R_{k+n},\, R_{k+n+1},\, \dots,\, R_{2n-1}$. All of them have length $\geq n$ and all of them contain exactly x `W`'s. Are they all distinct? Clearly $L_i \neq L_j$ and $R_i \neq R_j$ whenever $i \neq j$. Moreover, $[k,\, k+n-1]$ is different from all the other intervals. Could it be that $L_l = R_r$? No, because that would mean that such interval is $[l, r]$ with $l < k$ and $k+n-1 < r$, but then $[k,\, k+n-1]$ would be strictly contained inside $[l, r]$, violating the minimality of L_l and R_r.

H Beppa and SwerChat

Author: Andrea Ciprietti
Preparation: Lifu Jin

Instead of finding the minimum number of members that must have been online at least once between 9:00 and 22:00, let us study the complement of that set: What is the maximum number of members that have never been online? We have the following observations:

- Members that have never been online must be a suffix of the sequence b. It cannot happen that member b_i has been online but b_{i-1} has not, for $2 \leq i \leq n$.
- Consider two members x and y who have never been online. Their relative order in a and b is the same. Indeed, if someone else goes online the relative order of x and y in the list does not change.

From the above observations, we deduce that the members who have not been online form a suffix of b that is also a subsequence of a. Let $p : \{1, 2, \dots, n\} \to \{1, 2, \dots, n\}$ be the function such that $a_{p(x)} = x$; we say that a sequence s of length m is a subsequence of a if $p(s_{i-1}) < p(s_i)$, for $2 \leq i \leq m$.

It turns out that any "suffix of b that is a subsequence of a" can be a valid subset of members who have never been online as the following construction shows. For any such suffix $b_t, b_{t+1}, \dots, b_n$, it could be that members $b_{t-1}, b_{t-2}, \dots, b_1$ went online *in this order* between 9:00 and 22:00, leading to a valid list b of last seen online at 22:00.

Thus, the answer to the problem is the length of the *longest* suffix of b that is also a subsequence of a. Let us explain how to find such length quickly.

We first preprocess the sequence a to find the corresponding function p. Then, we iterate from b_n to b_1. If, for some i we discover that $p(b_i) < p(b_{i-1})$, then we know that b_{i-1} must have been online, and we denote this position i as r. In this way, $b_r, b_{r+1}, \dots, b_n$ is the sought longest suffix of b. The minimum number of members that must have been online at least once between 9:00 and 22:00 is then $r - 1$.

I Spinach Pizza

Author: Gerard Orriols
Preparation: Gerard Orriols

This problem has a greedy solution. The idea is that, at any moment, if the player who has to eat picks the triangle with smallest area among those that can be eaten in that moment, then any triangle chosen later will have area not less than it.

Therefore if n is even, Alberto can win by choosing the smallest possible slice at each turn, since the quantity eaten by Beatrice in the next turn will always be greater or equal and thus, since they eat the same number of slices, Alberto will eat no more in total. On the other hand, for n odd, no matter what slice Alberto chooses at the beginning, Beatrice can apply the explained strategy for the remaining number of turns, which is even. In this case Alberto will eat the initial amount plus at least the quantity Beatrice eats in the remaining turns, and therefore a total area strictly bigger than hers.

Now we prove the above claim. More precisely, given a convex polygon with vertices $P_1, \dots, P_n$, the minimum area of any triangle with vertices P_i, P_j, P_k is attained by a triangle determined by two consecutive edges in the polygon (we will call such triangles *edible*). Given a non-edible triangle, it is clear that we can label its vertices P, Q and R in such a way that neither of the edges PR and QR belong to the polygon. Then it is enough to show that we can change R by a vertex adjacent to P or Q and the remaining triangle will have at most the initial area. Indeed, by repeating this argument at most twice we can start with any triangle and obtain an edible triangle without increasing its area.

Let A and B be the vertices in the same side of the line PQ as R which are adjacent to P and Q, respectively. Then one of the triangles PQA or PQB has less area than PQR: since they all have the common base PQ, the area is proportional to the height with respect to the line PQ. However, if the height of R were less than that of A and B, then we would be able to write R as a convex combination of A, B, P and Q, which contradicts the strict convexity of the polygon.

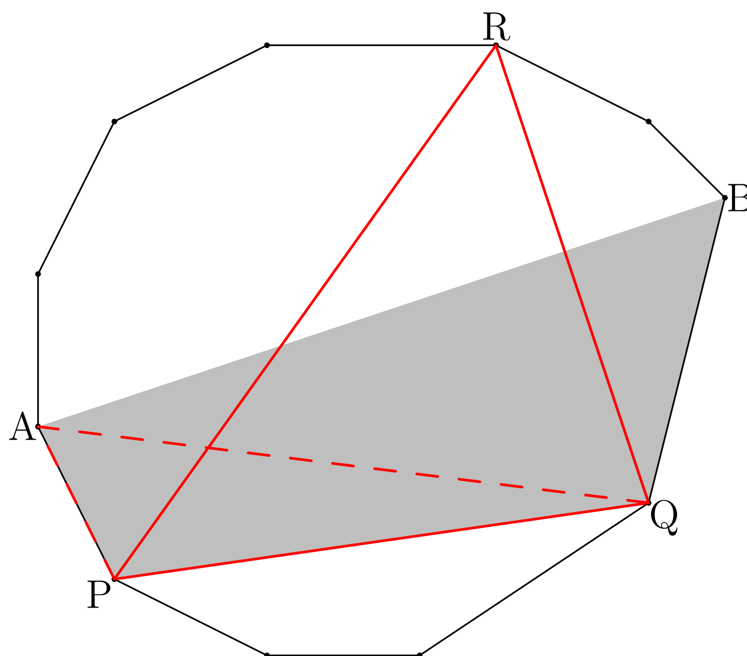

J Italian Data Centers

Author: Pedro Paredes
Preparation: Pedro Paredes

This problem is ultimately about graphs, so let's start by translating the statement into a more formal language using graphs. Let $G = (V, E)$ be a graph with $|V| = n$ vertices and $|E| = m$ edges, where each vertex is colored in one of three colors. The problem statement describes an operation we shall call *doubling*, since in a way it "doubles" the graph. The double of G, denoted by $d(G)$, is a graph with $2n$ vertices and $2m + n$ edges defined as such:

1. For each $v \in V$, the double graph $d(G)$ contains two copies of v of the same color, v_1 and v_2. There is also an edge connecting v_1 and v_2 in $d(G)$.

2. For each $\{v, u\} \in E$: if v and u have the same color, then $d(G)$ contains an edge between v_1 and u_1, and an edge between v_2 and u_2; otherwise, $d(G)$ contains an edge between v_1 and u_2, and an edge between v_2 and u_1.

Now consider applying this doubling operation k times to G. Our goal is to find the diameter of the resulting graph, i.e. the largest distance (shortest path) between any two vertices in the graph $\underbrace{d(d(\cdots d(}_{k \text{ times}}G)\cdots))$. Let G_i be the graph after i doubling operations, so $G_i = \underbrace{d(d(\cdots d(}_{i \text{ times}}G)\cdots))$.

Dissecting the problem

At first glance the problem sounds really hard, one has to compute the diameter of an exponentially sized graph ($2^k n$ vertices), so we have to study the structure of the problem to gain some intuition.

Observation 1. Each vertex of G is copied 2^k times in G_k, so we can describe each vertex in G_k by a vertex in G and a length-k bitstring. For example, $(v, \texttt{011})$ would represent $((v_1)_2)_2$, so the second copy of the second copy of the first copy of v.

Before we continue with our study of the doubling process let's make one definition.

Definition 1. The *hypercube graph* of dimension n is a graph with 2^n vertices given by all length-n bitstrings such that there is an edge between two vertices if their associated bitstrings differ in exactly one element (i.e. if the Hamming distance between them is 1).

Observation 2. The edges in G_k between copies of a vertex v from G form an hypercube graph of dimension k, in other words, each subgraph of G_k induced by all of the copies of a single vertex forms an hypercube graph of dimension k.

This is easy to see inductively, but for completeness here is a proof: consider the subgraph in G_{i-1} induced by all the copies of $v \in V$ and assume this forms an hypercube graph of dimension $i - 1$. When forming G_i notice that since all of the vertices have the same color, step 2 of the doubling operation creates two copies of subgraph induced by the copies of v in G_{i-1}. The first copy has all the vertices whose bitstring starts with a `0` in the G_i representation, and the second copy has are all the vertices whose bitstring starts with a `1`. All the vertices in the first copy differ in the first bit from the ones on the second copy, so to form an hypercube graph of dimension i we need to add an edge between vertices with the same $i - 1$ last bits, which is exactly what step 1 of the doubling operation does.

Now that we understand the structure of copies of a single vertex, let's describe the edges between

copies of different vertices of G. To make our notation easier to read, we refer to a vertex from G_i as (v, b), where $v \in V$ and b is a length-i bitstring.

Observation 3. Assume that $\{u, v\} \in E$. Consider a vertex (u, b) from G_k, where b is a length-k bitstring. (u, b) is connected to exactly one copy of v which is (v, b) if v and u have the same color, or $(v, \overline{b})$ if they have different colors, where $\overline{b}$ is the bitwise-negation of b.

We can see why this is true by noting that if v and u have the same color, then all the copies of u will be connected to the corresponding copy of v. If v and u have different colors, then we can see the observation by thinking inductively. Let $b^{(i)}$ be the length-i suffix of i. When creating G_i from G_{i-1} we go from $(u, b^{(i-1)})$ to $(u, b^{(i)})$ by adding b_i to the beginning of $b^{(i-1)}$. Since u and v have different colors, we connect $(u, b^{(i)})$ to $(v, \overline{b_i}\overline{b^{(i-1)}}) = (v, \overline{b^{(i)}})$.

Observation 4. Let $(u_1, b_u^1), (u_2, b_u^2), \ldots, (u_\ell, b_u^\ell)$ be a path in G_k. Then there is a valid path of the same length $(v_1, b_v^1), (v_2, b_v^2), \ldots, (v_\ell, b_v^\ell)$ and an index i such that $v_1 \neq v_2$ and $v_2 \neq v_3$, $\ldots$, $v_{i-1} \neq v_i$ and $v_i = v_{i+1} = \ldots = v_\ell$, so we first take all steps that move between vertices of G and then we take only steps in copies of v_i.

We obtain this by "rearranging" the steps in the first path. Given the symmetry of the graph, if we take one step between two copies of the same vertex too early, we can take that step at the end instead.

Computing the diameter

Let's start by describing the distance between two vertices (u, b_u) and (v, b_v) from G_k, where b_u and b_v are length-k bitstrings and $u, v \in V$ (not necessarily connected or distinct). But first, a quick definition:

Definition 2. Let P be a path in G between u and v. We call P an *even* path if the number of multicolored edges it uses is even, i.e. the number of times it walks to a vertex of a different color from the previous one is even. We analogously define *odd* paths.

Observation 5. The length of the shortest path between (u, b_u) and (v, b_v) is given by the shortest of the following:

- The length of the shortest even path between u and v plus $\Delta(b_u, b_v)$, where $\Delta(x, y)$ represents the Hamming distance between x and y, i.e. the number of positions at which the corresponding bits differ.
- The length of the shortest odd path between u and v plus $\Delta(b_u, \overline{b_v})$.

Proof. Consider a path in G_k of length ℓ from (u, b_u) to (v, b_v), say $(u_1, b_u^1), (u_2, b_u^2), \ldots, (u_\ell, b_u^\ell)$, where each pair of consecutive vertices is connected and $(u_1, b_u^1) = (u, b_u)$ and $(u_\ell, b_u^\ell) = (u_v, b_v)$. Without loss of generality, we can assume that there is an i such that $u_1 \neq u_2$ and $u_2 \neq u_3$, $\ldots$, $u_{i-1} \neq u_i$ and $u_i = u_{i+1} = \ldots = v$, by Observation 4.

We can "project" this path down to G, i.e. drop all bitstring labels to obtain $u_1, u_2, \ldots, u_\ell$. Note that this projected path is a path between u and v. If the projected path is an even path, then using Observation 3 we conclude that the ith vertex is (v, b_u), otherwise it is $(v, \overline{b_u})$. In the even case, the distance between (v, b_u) and (v, b_v) is given by $\Delta(b_u, b_v)$, and in the odd case the distance between $(v, \overline{b_u})$ and (v, b_v) is $\Delta(\overline{b_u}, b_v)$. This concludes the proof. □

Now, we can use this observation to describe the solution to the problem. First, because of the symmetries of the doubling operation, note that given a path $(u_1, b_u^1), (u_2, b_u^2), \ldots, (u_\ell, b_u^\ell)$, there is another valid path $(u_1, \mathtt{00\ldots0}), (u_2, (b_u^2)'), \ldots, (u_\ell, (b_u^\ell)')$, which is obtained by inverting all the `1` bits of b_1. This means that to find the diameter of G_k we only need to look at paths that start on

vertices of the form $(v, \texttt{00}\ldots\texttt{0})$.

Given u, v from G, let's try to determine the b that maximizes the distance between $(u, \texttt{00}\ldots\texttt{0})$ and (v, b). From Observation 4 and 5, we know that we either first take an even path from $(u, \texttt{00}\ldots\texttt{0})$ to $(v, \texttt{00}\ldots\texttt{0})$, or an odd path from $(u, \texttt{00}\ldots\texttt{0})$ to $(v, \texttt{11}\ldots\texttt{1})$. Let $|b| = x$, i.e. the number of 1s in b is x. Then, the distance from $(v, \texttt{00}\ldots\texttt{0})$ to (v, b) is given by x and the distance from $(v, \texttt{11}\ldots\texttt{1})$ to (v, b) is given by $k - x$. Let's denote $\delta_e(u, v)$ be the shortest even path between u and v, and define $\delta_o(u, v)$ analogously. Then the maximum distance between $(u, \texttt{00}\ldots\texttt{0})$ and (v, b) for any b, is given by $\max_{0 \le x \le k}\{\delta_e(u, v) + x, \delta_o(u, v) + k - x\}$.

So here is a possible algorithm:

Algorithm. For every pair of vertices u, v from G, compute the length of the shortest even path ($\delta_e(u, v)$), and the shortest odd path between them ($\delta_e(u, v)$). We can do this by running a Breadth-First Search starting at u that also keeps track of the parity of the number of times we've traversed multicolored edges. Then we compute $\max_{0 \le x \le k}\{\delta_e(u, v) + x, \delta_o(u, v) + k - x\}$ and output the maximum of this over all pairs.

Note that this algorithm runs in time $O(nm + n^2k)$, since we need one BFS per vertex (which takes $O(nm)$ time) and then for each pair of vertices we need to compute $\max_{0 \le x \le k}\{\delta_e(u, v) + x, \delta_o(u, v) + k - x\}$ (which takes $O(n^2k)$ time). This can be improved to $O(nm + n^2)$, but that isn't necessary for this problem since $n, k \le 100$.

K Uniform Chemistry

Author: Federico Glaudo and Petr Mitrichev
Preparation: Petr Mitrichev

Small constraints

The small version of this problem, which was used in the onsite competition, had constraints $n \leq 100$, $m \leq 10$.

The standard approach for this type of problem would be to use dynamic programming to compute the probability that each researcher wins the prize for each possible state of the process — the set of elements that each researcher has. However, even for the small constraints the number of such sets can be as high as 100^{10}, which is clearly too big for us to process them all one by one.

Therefore the key idea is to use the independence of the processes that each researcher follows. Suppose we compute for each researcher i and each year j the probability p_{ij} that the i-th researcher discovers element n in the j-th year. In order for them to win the SWERC prize when doing so, the other researchers must have not discovered element n yet. The probability that researcher k has not yet discovered element n by the j-th year can be computed as $1 - \sum_{t<j} p_{kt}$, and because of the independence of different researchers, the probability that all other researchers have not yet discovered element n is equal to $\prod_{k \neq i}(1 - \sum_{t<j} p_{kt})$, and the probability that the i-th researcher wins the SWERC prize is equal to

$$\sum_j p_{ij} \prod_{k \neq i} \left(1 - \sum_{t<j} p_{kt}\right)$$

Now we just need to compute the probability p_{ij} that the i-th researcher discovers element n in the j-th year for all i and j. This can be done using dynamic programming that computes q_{ij}: the probability that a researcher that has element $n - i$ in the beginning (in other words, is i elements away from the goal) discovers element n in the j-th year. The probabilities that we want are then found as $p_{ij} = q_{n-s_i,j}$.

From the fusion experiment definition we directly obtain:

$$q_{ij} = \frac{1}{i} \sum_{0 \leq k < i} q_{k,j-1}$$

We can then apply this formula in increasing order of i, or in increasing order of j, to find all those probabilities, starting from $q_{00} = 1$, $q_{i0} = 0$ for $i > 0$, and then apply the formula that combines those values into our answer.

The dynamic programming has $O(n^2)$ states, and each state is processed in $O(n)$, so its running time is $O(n^3)$. The final formula is computed in $O(n^2 m)$ for each of the m researchers, so the overall running time is $O(n^2(n + m^2))$, which is fast enough for $n \leq 100$, $m \leq 10$.

Large constraints

The above solution can be optimized to run in $O(n(n+m))$ using relatively standard techniques, but that still won't be enough to solve the large version that was given in the mirror round: $n \leq 10^{18}$, $m \leq 100$. Here we need to come up with more radical improvements.

The first big improvement is to notice that the distance to n decreases twice every year on average, and therefore the typical number of years needed to reach n is likely to be proportional to $\log n$. Moreover, for each ε there likely exists a constant C such that the number of years is below $C \log n$ with probability $1 - \varepsilon$.

To prove this formally, we can notice that each year the distance to n decreases at least twice with probability of at least $\frac{1}{2}$, so after k years the distance to n decreases at least twice with probability at least $1 - \frac{1}{2^k}$. We can choose k such that $\frac{1}{2^k} < \frac{\varepsilon}{log_2 n}$, and then the number of years will not exceed $k \log_2 n$ with probability at least $1 - \varepsilon$. This is a very simple estimate, using Bernstein inequalities one can prove much more.

But we don't really need a formal proof during the round, instead the intuition about the distance decreasing twice on average can be coupled with an experiment that measures how quickly the numbers q_{ij} described above become very small. In practice, it turns out that we can always assume $j \leq 100$ and get the required precision.

This observation improves the running time to $O((n + m) \log n)$, which is still not good enough. We will describe two independent ways to proceed further.

The exact approach. Let us examine the recurrence for q_{ij} more closely. If we apply it repeatedly, we end up having the following sum:

$$q_{ij} = \sum_{0 < a_1 < \cdots < a_j = i} \prod_{k=1}^{j} \frac{1}{a_k}$$

This prompts us to consider the following polynomial:

$$f(x) = \frac{1}{i} \prod_{k=1}^{i-1} \left(1 + \frac{1}{k} x\right)$$

The coefficient next to x^{j-1} in this polynomial is equal to q_{ij}, because to get it we can use any of the ways to pick $j - 1$ terms of the product from which we take the part with x and then we get a product of the reciprocals of their indices, which is precisely what the sum for q_{ij} has.

So now we need to compute the first $O(\log n)$ coefficients of this polynomial quickly. Let us take its natural logarithm (in the formal power series sense), and substitute the Taylor series for $\log(1 + y)$:

$$\log f(x) = -\log i + \sum_{k=1}^{i-1} \log \left(1 + \frac{1}{k} x\right) = -\log i - \sum_{k=1}^{i-1} \sum_{t=1}^{\infty} \frac{(-1)^t}{tk^t} x^t = -\log i - \sum_{t=1}^{\infty} \frac{(-1)^t}{t} x^t \sum_{k=1}^{i-1} \frac{1}{k^t}$$

Since the first $O(\log n)$ coefficients of $f(x)$ can be found by exponentation of the formal power series and depend only on the first $O(\log n)$ coefficients of $\log f(x)$, all we need now is a way to find the sums of inverse powers $\sum_{k=1}^{i-1} \frac{1}{k^t}$ for t up to $C \log n$.

This can be done using the Euler-Maclaurin formula, or using the following trick which is likely equivalent to it in some sense. We denote $S_t(a, b) = \sum_{k=a}^{b} k^{-t}$, and notice that using the Taylor series for $\log(1 + y)$ we get:

$$S_1(a, b) = \sum_{k=a}^{b} \log \left(1 + \frac{1}{k}\right) + \sum_{k=a}^{b} \left(\frac{1}{k} - \log \left(1 + \frac{1}{k}\right)\right) = \log(b + 1) - \log a + \sum_{k=a}^{b} \sum_{p=2}^{\infty} \frac{(-1)^p}{pk^p} =$$

$$= \log(b+1) - \log a + \sum_{p=2}^{\infty} \frac{(-1)^p}{p} S_p(a,b)$$

Similarly for $t > 1$ using the Taylor series for $(1+y)^{-(t-1)}$ we get:

$$S_t(a,b) = \sum_{k=a}^{b} \frac{1}{t-1}\left(\frac{1}{k^{t-1}} - \frac{1}{(k+1)^{t-1}}\right) + \sum_{k=a}^{b}\left(\frac{1}{k^t} - \frac{1}{t-1}\left(\frac{1}{k^{t-1}} - \frac{1}{(k+1)^{t-1}}\right)\right) =$$

$$= \frac{1}{t-1}\left(\frac{1}{a^{t-1}} - \frac{1}{(b+1)^{t-1}}\right) + \sum_{k=a}^{b}\left(\frac{1}{k^t} - \frac{1}{(t-1)k^{t-1}}\left(1 - \frac{1}{(1+\frac{1}{k})^{t-1}}\right)\right) =$$

$$= \frac{1}{t-1}\left(\frac{1}{a^{t-1}} - \frac{1}{(b+1)^{t-1}}\right) + \sum_{k=a}^{b} \frac{1}{(t-1)k^{t-1}} \sum_{p=2}^{\infty} (-1)^p \binom{t-2+p}{p} \frac{1}{k^p} =$$

$$= \frac{1}{t-1}\left(\frac{1}{a^{t-1}} - \frac{1}{(b+1)^{t-1}}\right) + \frac{1}{t-1} \sum_{p=t+1}^{\infty} (-1)^{p-t+1} \binom{p-1}{p-t+1} S_p(a,b)$$

These formulas express $S_t(a,b)$ using $S_p(a,b)$ where $p > t$. Since for sufficiently large t $S_t(a,b)$ can be computed with good enough precision directly because $\frac{1}{k^t}$ decreases very quickly, this allows to compute $S_t(a,b)$ fast for all t up to $C \log n$, and therefore to obtain the polynomial we need and to solve the problem.

The approximate approach. Alternatively, we can approximate our problem with its continuous version: suppose the result of the fusion experiment was not an integer, but a uniformly distributed real number, and we were interested in the probability to get a number less than 1 after exactly j years if we start with i.

This yields the following formulas:

$$q'_{i1} = \frac{1}{i}$$

$$q'_{ij} = \frac{1}{i} \int_1^i q'_{k,j-1} dk$$

The integral is actually not very hard to compute, and we get

$$q'_{ij} = \frac{\log^{j-1} i}{i(j-1)!}$$

This approach does not give good enough approximation by itself, but we can improve on it in the following way: suppose we want to find some number q_{ij}. Let us choose some boundary b, for example $b = 10^6$, and for all values of $i < b$ we can just compute q_{ij} using the original recurrence. When i is larger than or equal to b, we will separate the process into two stages: while i is larger than or equal to b, and when it becomes smaller, and we will use the continuous approximation only for the first stage.

Getting from i to b in the continuous setting is the same as getting from i/b to 1, so we can use the above formula for q' for this.

To glue the two stages together, we can notice that whenever i becomes less than b, because of the uniform choice all values of i between 0 and $b-1$ are equally likely, and therefore this situation is the same as if we had $i = b$ on the previous step. Therefore we can use the following formula to compute our approximate answer:

$$q_{ij} \approx \sum_{k=1}^{j} q'_{\frac{i}{b},k} q_{b,j-k+1}$$

This (barely) gives good enough precision for $b = 10^6$.

Precision issues. The solutions to this problem, especially the exact approach involving formal power series exponentiation and Taylor expansions, can easily suffer from floating-point precision issues. We are very confident that our reference solutions are correct because we have several different approaches, implemented in different languages, using different floating-point types, all agree on the answers.

To achieve this, we had to make some changes that we would like to mention to make it easier to follow in our footsteps. In the exact solution, we used the above formulas to compute $S_t(a, b)$ only for sufficiently large values of a (at least $a \geq 2$), and added the first few terms directly, as that allowed for better convergence. We used those formulas only for small values of t ($t \leq 5$, or even just $t \leq 2$), as otherwise the large binomial coefficients have amplified the precision errors too much. When computing the exponential of a formal power series using an $O(n \log n)$ approach (which was not necessary in this problem), we found that using the identity $exp(f(x)) = (exp(\frac{f(x)}{8}))^8$ can improve the precision dramatically.

In addition, we noticed that if we use the above formulas to compute $S_t(a, b)$ only for $a \geq 10^6$, then we need them only for $t = 1$ and $t = 2$, and even for those we need only a couple of terms, making this part of the solution really short:

$$S_1(a, b) \approx \log(b + 1) - \log a + \frac{1}{2} S_2(a, b)$$
$$S_2(a, b) \approx \frac{1}{a} - \frac{1}{b + 1}$$

L Controllers

Author: Stefanie Zbinden
Preparation: Stefanie Zbinden

Denote by p the number of + in s and by m the number of - in s and by tot the value $p - m$. We want to figure out whether we can win the game for a single controller with values x and y on the two buttons. If we only press the button with value x written on it, then our sum at the end is $(p - m) \cdot x = tot \cdot x$. So if $tot = 0$ we can always win the game and from here on we assume that $tot \neq 0$.

Assume we press the button with value x k_1 times when the symbol + comes up and k_2 times when the symbol - comes up. This means we press the button with value y $p - k_1$ times when the symbol + comes up and $m - k_2$ times when the symbol - comes up. Hence our total sum in the end is $k_1 \cdot x - k_2 \cdot x + (p - k_1)y - (m - k_2)y$.

We can define $k = k_1 - k_2$ and rewrite this as $k \cdot x + (tot - k)y$. Rewriting this even further gives that the sum is 0 if and only if $k(y - x) = tot \cdot y$. So if $x = y$ we cannot win the game (recall that we are assuming here that $tot \neq 0$). If $tot \cdot y$ is not divisible by $(y - x)$ we cannot win the game either.

Further, if $x \neq y$ we win if and only if $k = tot \cdot y/(y - x)$.

Can we get such value of k? Since $0 \leq k_1 \leq p$ and $0 \leq k_2 \leq m$, $k = k_1 - k_2$ can have any integer value in the range $[-m, p]$. Or in other words, we can win the game if and only if $tot \cdot y/(y - x)$ is an integer between $-m$ and p.

M Parmigiana With Seafood

Author: Simon Mauras
Preparation: Simon Mauras

First, as for many tree problem, let us solve it first on chains, then generalize the solution for trees.

Game on a path

First, start with several observations on Alessandro's strategy.

Definition (bad pair). We say that two ingredients form a bad pair if they are at odd distance of each other.

Lemma 1. Alessandro can choose the best of the following strategies:

- First, he can include any ingredient which is a terminal ingredient at the begining of the game.
- Second, if n is even, then he can make sure to include ingredient n.
- Third, if n is odd and given a bad pair, he can always make sure that one of the two ingredient is included in the recipe.

Proof. The first observation is easy as Alessandro plays first. For the second observation, Alessandro's strategy is to include ingredient n whenever allowed to, and otherwise to select an ingredient which does not make n a terminal. Observe that when only 3 ingredients remain, such that n is the middle one, then it must be Bianca's turn, and Alessandro will be able to pick n on the next turn. For the third observation, the same argument hold, and it must be Bianca's turn when the only possible moves make either x or y terminal. An alternative explaination is to look at all the ingredients on the path between x and y (including x and y) as one very large ingredient (changing the parity of the number of items), and to apply the second observation.

Now, we have to check whether or not Alessandro can do even better, by looking at Bianca's strategy.

Lemma 2. Assuming that n is odd, consider a set S of non-terminal vertices which contain no bad pair of ingredients. Bianca can make sure that each ingredient of S is discarded.

Proof. Bianca's strategy is (i) to select and discard ingredients from S whenever possible, and (ii) to only select ingredients which do not make any ingredient of S terminal. Observe that ingredients of S are not adjacent and thus (i) does not contradict (ii). To prove (ii), observe that whenever it is Bianca's turn the number of remaining ingredients is even. If only two ingredient remains then at most one can be in S and (ii) holds. If the chain has four or more ingredients then it has the form

$$a - x - \cdots - y - b$$

where x and y are at odd distance. In particular, either x or y is not is S, and Bianca can select and discard the corresponding neighbour and (ii) holds. Finally, we can prove by induction that whenever Alessandro plays, none of the terminal ingredients is in S, concluding the proof.

Game on a tree

We now move on to the same question, when the graph of ingredients is a tree. One can check that Lemma 1 still apply, with the exact same proof. However, the proof of Lemma 2 relied on the chain structure and does not hold for trees, as the following example shows.

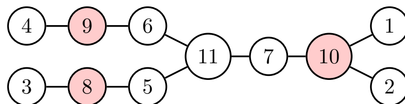

Nodes in S colored in red. If Alessandro selects either 1 or 2, then Bianca will have to make one of the red ingredients terminal. This is caused by the fact that there is an even number of ingredients between red ingredients. Equivalently, this happens because there is a *bad triple*, as formalized in the following Definition and Lemma.

Definition (bad triple). We say that three ingredients x, y and z form a bad triple if there exist one indredient m whose removal would disconnect x, y and z (m is sometimes called the median), and such that all three ingredients x, y and z are at even distance of m.

In the example above, ingredients $x = 8$, $y = 9$ and $z = 10$ form a bad triple, as they are all three at distance two of their median $m = 11$.

Lemma 3. If n is odd and a bad triple is given, Alessandro can make sure one of the three ingredients of the triple is included in the recipe.

Proof. If one player is forced to make one of the three ingredients a terminal, then an even number of ingredient remains, and thus it is Bianca's turn to play.

Finally, we show that Alessandro cannot do better than combining strategies from Lemmas 1 and 3.

Lemma 4. Assuming that n is odd, consider a set S of non-terminal vertices which contain no bad pair and no bad triple of ingredients. Bianca can make sure that each ingredient of S is discarded.

Proof. Once again, Bianca's strategy is (i) to select and discard ingredients from S whenever possible, and (ii) to only select ingredients which do not make any ingredient of S terminal. To prove that (ii) holds, one can use an induction on $|S|$, which is left as an exercise.

Therefore if both players play optimally, the largest index included by Alessandro is characterized by Lemma 1 and 3. More precisely:

- if n is even then the answer is n,
- otherwise, it is the maximum between
  - the index of any terminal ingredient,
  - the index of any ingredient at odd distance to n,
  - $\min(x, y)$ for any bad triple (x, y, n),
  - $\min(x, y, z)$ for any bad triple (x, y, z) whose median is n.

In particular, this can be computed in linear time with a dfs in the tree (rooted in ingredient n).

N Count Permutations

Author: Federico Glaudo and Giovanni Paolini
Preparation: Federico Glaudo

This was meant to be the hardest problem of the mirror. It required advanced knowledge in modern combinatorics, together with some clever ideas to reach the desired precision. The solution splits in two parts: first we obtain a nontrivial (actually, pretty hard) formula for the exact answer, then we explain how to approximate it quickly enough.

Exact formula for the answer via Hook-length formula for skew diagrams

We will give for granted the definition of Young Diagram and Young Tableau.

As a first step, we associate to the string s a Young diagram as follows.

We construct a path, that will be the border of the Young diagram. Start with a single cell (that will be the bottom left corner of the diagram), then iterate over the characters of s. Each time you read >, go up. Each time you read <, go right. The path we have constructed is the *border* (apart from the two long sides) of the Young diagram associated to s. See the figure, which refers to the string $s = ' <> < < > ><>< <> ><'$.

In this Young diagram, we say that a cell is special if it belongs to the path we have constructed.

Notice that each way to fill the special cells according to the usual law of a Young tableaux coincides with a permutation satisfying the constraints given by s. So, counting the permutations is the same as counting the number of Young tableaux restricted to the special cells.

If we were to count the number of Young tableaux on the whole Young diagram, we could apply the hook-length formula. It turns out that there is a new formula, due to Naruse, that is valid also for skew diagrams, that are diagrams obtained as differences of two Young diagrams (notice that the special cells can be seen as the difference of two Young diagrams).

Refer to the original slides by Naruse and to this more recent paper by Morales, Pak, Panova for a presentation of the formula in the general case (with proofs).

In the case we need, i.e., for border strips (this is how the paper by Morales, Pak, Panova call the special cells) the formula simplifies a lot. Instead of stating the general formulas, which would require the definition of excited diagram, let us state here the formula for the special case of border strips.

Summing up these results we get the following formula.

Number of Tableaux on Border Strips (consequence of Naruse's hook length formula): Let λ be a Young diagram and let λ' be its border strip (i.e., its *nontrivial* border, as described above). A valid path in the λ is a sequence of adjacent cells in λ going from the bottom left corner to the top right corner that goes only up and right. The weight of a valid path is the product of the inverses of the hook-lengths of the cells of the path.

Let $W(\lambda)$ be the sum of the weights of all valid paths on λ. The number of tableaux on λ' is $|\lambda'|!W(\lambda)$ (here $|\lambda'|$ denotes the number of cells of λ').

This formula can be used to compute the answer in $O(n^2)$, which is too slow for this problem (notice that there is also a much simpler solution with complexity $O(n^2)$).

How to optimize the solution

Notice that the hook-length of a cell is, in general, large for cells that are far from the border. Hence, we may expect that the contribution to the result given by path which are not *close* to the border at all times is negligible. This claim is highly nontrivial and we, the judges, have no idea of how to prove it. But experimentally it is true; for the precision required in this problem it is sufficient to consider path that stay all the time at a *diagonal distance* from the border $\leq D = 500$.

So, we must compute the sum, over all paths that stay at a distance $\leq D$ from the border of the Young diagram. This can be done with a simple dynamic programming.

It remains to understand how to take care of the huge numbers that we are going to get. There are fundamentally two ways:

- Instead of keeping the numbers, we keep their logarithms. But if we are given $\log(x)$ and $\log(y)$, how do we compute the sum without computing explicitly x and y? Assume that $x < y$, then we compute
$$\log(x + y) = \log(y) + \log(1 + \exp(\log(x) - \log(y))).$$
- We keep the numbers as long double. Still, during our dynamic programming, the exponent may become larger then the largest exponent admissible for long doubles. During our dynamic programming we will keep a vector of values, denoting the sum of the weights of paths starting from the bottom left corner and ending at one point on a fixed diagonal. At each step we will keep D values and if we divide all of them by 2, then the final result will change by 2. Therefore, we divide all of them by 2, until the largest one is ≤ 1. Doing this, we do not lose too much precision and everything fits in a long double. This solution is much faster than the previous one.