

First Round SOI 2018/2019

Solution Booklet



Swiss Olympiad in Informatics

October 1 – November 30, 2018

1 Mergeball

In this task we're given the sizes b_0, \dots, b_{N-1} of N balls. The action we can perform with them is to take 2 adjacent balls with the same size and merge them together, creating a new ball with double the size.

Subtask 1: Final configuration with two balls (20 Points)

In this subtask we're only given pairs of balls which we want to merge.

Because for each test case there are 2 balls and their sizes are the same, it's enough to just print their sum.

An example solution can be found here:

```

1 #include <bits/stdc++.h>
2 using namespace std;
3 int main() {
4     int T; // T is the number of test cases
5     cin >> T;
6     for (int t = 0; t < T; t++) { // A loop to input each test case
7         int N, b0, b1;           // b0 and b1 are the sizes of the 2 balls
8         cin >> N >> b0 >> b1;
9         // For each test case we output the sum of the sizes
10        cout << "Case #" << t << ": " << (b0 + b1) << "\n";
11    }
12 }
```

Subtask 2: All to one (20 Points)

This subtask is building on top of the first one, but this time we have to check if we can merge N balls, all of which have the same size, into one single ball.

In order to solve it, we have to be smart – we can't just take any 2 adjacent balls at each step and merge them together. The reason for that is we might create a ball that divides 2 balls that could have been merged and we end up with multiple of them, instead of a single one, essentially getting stuck.

A way we can make sure this never happens if there is a solution is the following: We take starting from the left the *smallest* pair of balls and we merge them together. We continue doing so, until we get one single ball or we can't merge anymore.

Here's an example:

$(1)(1)(1)(1)(1)(1)(1)(1) \rightarrow (2)(1)(1)(1)(1)(1)(1) \rightarrow (2)(2)(1)(1)(1)(1) \rightarrow (2)(2)(2)(2) \rightarrow (4)(2)(2) \rightarrow (4)(4) \rightarrow (8)$

This is an example of a *bottom-up approach*. The argumentation of why this is indeed correct will be presented in the next subtask.

Another idea that at first sight might seem confusing, but is also correct, is the following. Let's look at the last merge. That merge has to combine 2 balls with the same values. Even more, we know that the size of these 2 balls has to be equal to half of the total sum of all the balls. If we can't split the balls in 2 sides equal to this half-sum, we can't merge everything into one. Otherwise, we



can do the same check, for the 2 sides and do so, until we get to the bottom recursively. This is an example of a *top-down approach*.

There is however a very important observation we have to make. Since we start with N balls, which all have the same size, we merge every consecutive pairs of them, essentially dividing the number of balls by half. Since the new balls we get all have the same size as well, we continue this process until we either have 1 ball or we can't divide the total number of balls we currently have by 2. Therefore in order to check if in the end we'll be left with a single ball or multiple, it's enough to check if the number of balls in the beginning N is a power of 2 or not. An example solution using bittricks to check if a number is a power of 2 can be found here:

```
1 #include <bits/stdc++.h>
2 using namespace std;
3 int main() {
4     int T;
5     cin >> T;
6     for (int t = 0; t < T; t++) {
7         int N;
8         cin >> N; // Number of balls
9         for (int i = 0; i < N; i++) {
10            int bi; // The size of the current ball. In this test case, all sizes are
11                // the same.
12            cin >> bi;
13        }
14        // (N & (N - 1)) is a bit trick to find out if N is a power of 2
15        cout << "Case #" << t << ": " << ((N & (N - 1)) ? "Multiple" : "Single")
16            << "\n";
17    }
18 }
```

Subtask 3: All to one – Part 2 (20 Points)

This subtask is the same as the second one, except in this case we're given balls with different sizes.

The two ideas we used to solve the second subtask can be applied here as well. Again at each step we take the smallest 2 balls and we merge them. An example of this solution is:

$$\textcircled{8} \textcircled{4} \textcircled{2} \textcircled{1} \textcircled{1} \rightarrow \textcircled{8} \textcircled{4} \textcircled{2} \textcircled{2} \rightarrow \textcircled{8} \textcircled{4} \textcircled{4} \rightarrow \textcircled{8} \textcircled{8} \rightarrow \textcircled{16}$$

This is correct because of the following argument: Look at the smallest leftmost ball. If there is *any* solution, this ball has to be merged with some other ball. Balls are either merged to the left or to the right. Because the ball to the left is larger or does not exist (by definition of smallest leftmost), it has to be merged with the one to the right. We can repeat this argument until there is only a single ball left.

Because of this, many modifications will still work - it does not need to be the two smallest balls, it can also be the first two adjacent balls with the same value. A fast solution using this idea uses a stack, which pushes the sizes from left to right and merges the two top values if they are equal. Since we access each value only twice - once when we add it to the stack and once when we take it out to merge it, for a total number of $2N$ times, the running time of this solution is $O(N)$. An example of this solution can be found here:

```

1 #include <bits/stdc++.h>
2 using namespace std;
3 int main() {
4     int T;
5     cin >> T;
6     for (int t = 0; t < T; t++) {
7         int N;
8         cin >> N;
9         stack<int> st; // A stack to store the ball sizes
10        for (int i = 0; i < N; i++) {
11            int bi;
12            cin >> bi;
13            // Check if the previous ball had the same size. We continue doing so,
14            // until we can't
15            while (st.size() >= 1 && st.top() == bi) {
16                st.pop(); // If yes, we merge it with the current ball and remove the
17                           // previous
18                bi *= 2;
19            }
20            st.push(bi);
21        }
22        // If there is only 1 ball in the stack, we successfully merged everything
23        cout << "Case #" << t << ": " << (st.size() == 1 ? "Single" : "Multiple")
24             << "\n";
25    }
26 }

```

Subtask 4: To infinity (20 Points)

In this subtask we have balls of the same size, from which we want to find out what sizes can be reached in the game. The limit here is that we aren't allowed to make balls bigger than size C .

Let's say the size of the balls is b_i . Since the only way to create new balls is by merging existing ones with the same size and Stofl has only balls of one size, in the beginning we can create a ball of size $2 \cdot b_i$. This ball can only be merged with another one of the same type, creating a ball of size $4 \cdot b_i$ etc...

So by just doubling the size at each step and stopping when that value gets bigger than C , we can find all possible sizes that can be achieved using that ball. We have to be careful and use long numbers to make sure there won't be an overflow. An example of this solution can be found here:

```

1 #include <bits/stdc++.h>
2 using namespace std;
3 #define int int64_t
4 signed main() {
5     int T;
6     cin >> T;
7     for (int t = 0; t < T; t++) {
8         int N, C, bi;
9         cin >> N >> C >> bi;
10        vector<int> sizes; // A vector to store the sizes we compute
11        while (bi <= C) { // If the new value is bigger than C, we stop
12            sizes.push_back(bi); // We add the size to the rest
13            bi *= 2; // We double it
14        }

```



```
15     cout << "Case #" << t << ": " << sizes.size();
16     for (int s : sizes)
17         cout << " " << s;
18     cout << "\n";
19 }
20 }
```

Subtask 5: To infinity and beyond! (20 Points)

This subtask is the same as the fourth one, except here we have many balls of different sizes.

The solution is exactly the same as the previous one, except that we have to pay attention to the fact that we might achieve the same size by using 2 different balls and we don't want to count it twice. The easiest way to do so is by using a set, where we store the values we get during the computation, which deals with the problem of having the same value twice. An example of this solution can be found here:

```
1 #include <bits/stdc++.h>
2 using namespace std;
3 #define int int64_t
4 signed main() {
5     int T;
6     cin >> T;
7     for (int t = 0; t < T; t++) {
8         int N, C;
9         cin >> N >> C;
10        set<int> sizes; // A set to store the sizes we compute
11        for (int i = 0; i < N; i++) {
12            int bi;
13            cin >> bi;
14            while (bi <= C) { // If the new value is bigger than C, we stop
15                sizes.insert(bi); // We add the size to the rest
16                bi *= 2; // We double it
17            }
18        }
19        cout << "Case #" << t << ": " << sizes.size();
20        for (int s : sizes)
21            cout << " " << s;
22        cout << "\n";
23    }
24 }
```

2 Ceremony

Given are the heights h_0, \dots, h_{N-1} of each building in a row of N skyscrapers. For convenience's sake, let us say that if $x < y$, then the x^{th} building is to the west or on the left of the y^{th} building, and, similarly, that if $x > y$, then the x^{th} building is to the east or on the right of the y^{th} building.

For every $0 \leq i, j < N$, fireworks launched from the top of the i^{th} building are visible from the top of the j^{th} building if and only if both $i \neq j$ and $h_k < h_j$ for every k such that $\min(i, j) < k < \max(i, j)$. In that case, we say that the j^{th} building offers an ideal view on the i^{th} building, or that it is ideal for the i^{th} building.

In the three first subtasks, you're given some i and you have to determine how many buildings are ideal for the i^{th} building. In the two last subtasks, your task is to compute the maximal number of ideal buildings if one chooses i optimally.

Subtask 1: Launch the fireworks from the westmost building (10 points)

In this first subtask, the fireworks were always launched from the 0^{th} building and you had to compute from how many buildings they were visible. The limits were still quite low ($1 \leq N \leq 100$), so it was possible to test the visibility from every other building using the definition of visibility presented above without any further trick. For the j^{th} building, one just needs to check for every k such that $0 < k < j$ whether $h_k < h_j$. If there is some case in which this inequality does not hold, then the j^{th} building is not ideal for the i^{th} building. This approach would result into the following code:

```

1 #include <bits/stdc++.h>
2 using namespace std;
3 int main() {
4     int T; cin >> T;
5     for (int t = 0; t < T; t++) {
6         int N, i; cin >> N >> i; // note that i = 0
7         vector<int> h(N);
8         for (int j = 0; j < N; j++)
9             cin >> h[j];
10        int result = 0;
11        for (int j = 1; j < N; j++) {
12            bool visible = true;
13            for (int k = 1; k < j; k++)
14                if (h[k] >= h[j]) // i. e. if it is not the case that h[k] < h[j]
15                    visible = false;
16            if (visible)
17                result++;
18        }
19        cout << "Case #" << t << ": " << result << '\n';
20    }
21 }

```

Overall, we perform $\sum_{k=1}^{N-2} k = \frac{(N-2)^2 + N - 2}{2} = \frac{N^2 - 3N + 2}{2}$ comparisons. Therefore, the program computes the solution to any testcase in $O(N^2)$.



Subtask 2: More available buildings (20 points)

This subtask was the same as the first one, except that N could be a lot larger: the limits were $1 \leq N \leq 1'000'000$. Our previous program was therefore too slow to compute the solution within the time limit of 5 minutes.

The key observation is the following: we were comparing a lot of values that we did not need to compare. For example, let us consider an input with heights 8, 3, 5, 12, 4, and 13. When we want to check whether the last building is ideal for the 0th building, we test four inequalities: $3 < 13$, $5 < 13$, $12 < 13$ and $4 < 13$. It is evident that if the third of these inequalities holds, all of the others do as well. For example, if $12 < 13$, then clearly $3 < 13$, because $3 < 12$ and $12 < 13$ ("being less than" is a transitive relation). More generally, one only needs to check whether $\max_{0 < k < j}(h_k) < h_j$.

Moreover, if one goes through the buildings from left to right, it is easy to keep track of the highest building met so far on the go, so that one can always compute whether the next building is ideal in constant time. This code would be an implementation of that idea:

```
1 #include <bits/stdc++.h>
2 using namespace std;
3 int main() {
4     int T; cin >> T;
5     for (int t = 0; t < T; t++) {
6         int N, i; cin >> N >> i; // note that i = 0
7         vector<int> h(N);
8         for (int j = 0; j < N; j++)
9             cin >> h[j];
10        int result = 0, maxSoFar = 0;
11        for (int j = 1; j < N; j++) {
12            if (maxSoFar < h[j])
13                result++;
14            maxSoFar = max(maxSoFar, h[j]);
15        }
16        cout << "Case #" << t << ": " << result << '\n';
17    }
18 }
```

With that improvement, we get our runtime down to $O(N)$.

Subtask 3: Launching from another building (20 points)

In this subtask, the limits are the same as in the second one, but i , the index of the building from which the fireworks are launched, can vary. Simply reusing our previous solution but starting with the $(i + 1)$ th building would clearly not be enough, because one would not ever look at buildings on the left of the i th building. However, we can apply the same principles to extend our solution with another loop checking for each building on the left of the i th building whether it is ideal for that building. Here is an example of how one would do that:

```
1 #include <bits/stdc++.h>
2 using namespace std;
3 int main() {
4     int T; cin >> T;
5     for (int t = 0; t < T; t++) {
6         int N, i; cin >> N >> i; // i is not always 0 anymore
```



```

7  vector<int> h(N);
8  for (int j = 0; j < N; j++)
9      cin >> h[j];
10 int result = 0, maxSoFar = 0;
11 for (int j = i + 1; j < N; j++) { // right side
12     if (maxSoFar < h[j])
13         result++;
14     maxSoFar = max(maxSoFar, h[j]);
15 }
16 maxSoFar = 0;
17 for (int j = i - 1; j >= 0; j--) { // left side
18     if (maxSoFar < h[j])
19         result++;
20     maxSoFar = max(maxSoFar, h[j]);
21 }
22 cout << "Case #" << t << ": " << result << '\n';
23 }
24 }

```

In total, we look at the visibility of $N - 1$ buildings, and for every one of these, we only need a constant time to compute whether they are visible. Therefore, our runtime is still $O(N)$.

Subtask 4: Find the building with optimal visibility (10 points)

Now the problem is slightly different: you are not given some i as in the previous subtasks, but you are required to find the optimal number of buildings from which the fireworks are visible if i is chosen optimally, that is, you need to find the maximal solution among the solutions for every possible value of i .

In this fourth subtask, the limits are low again: $1 \leq N \leq 100$. That means that we can still use a similar solution to that of the third subtask, but checking the solution for every i such that $0 \leq i < N$ and taking the maximum of these N values. The modification is quite simple and could be implemented in the following manner:

```

1  #include <bits/stdc++.h>
2  using namespace std;
3  int main() {
4      int T; cin >> T;
5      for (int t = 0; t < T; t++) {
6          int N; cin >> N; // there is no given i anymore
7          vector<int> h(N);
8          for (int j = 0; j < N; j++)
9              cin >> h[j];
10         int bestResult = 0;
11         for (int i = 0; i < N; i++) {
12             int result = 0, maxSoFar = 0;
13             for (int j = i + 1; j < N; j++) { // right side
14                 if (maxSoFar < h[j])
15                     result++;
16                 maxSoFar = max(maxSoFar, h[j]);
17             }
18             maxSoFar = 0;
19             for (int j = i - 1; j >= 0; j--) { // left side
20                 if (maxSoFar < h[j])

```



```
21     result++;
22     maxSoFar = max(maxSoFar, h[j]);
23     }
24     bestResult = max(bestResult, result);
25 }
26 cout << "Case #" << t << ": " << bestResult << '\n';
27 }
28 }
```

Since we need $O(N)$ time to compute each of the N partial solutions, we get a runtime in $O(N^2)$.

Subtask 5: Find the building with optimal visibility with many buildings (40 points)

This final subtask was the trickiest: with the higher limit on N ($1 \leq N \leq 1'000'000$), one needed a good idea to get under the time limit.

The idea was the following: if one knows which buildings on the left of the k^{th} building offer an ideal view of it, then it is easy to compute which buildings on the left of the $(k + 1)^{\text{th}}$ building offer an ideal view of it. The heights of the buildings on the left of the k^{th} building that are ideal for it make up a strictly decreasing sequence (for it is a direct consequence of the definition of visibility that if this sequence was not strictly decreasing, then some of the buildings therein would not be ideal). The ideal buildings on the left of the $(k + 1)^{\text{th}}$ building are then the same as those on the left of the k^{th} one, without the buildings that were ideal for the k^{th} building but are not for the $(k + 1)^{\text{th}}$ building because their height is smaller than h_k , and with the addition of the k^{th} building itself. But the buildings that we have to remove are always a suffix of the sequence of the ideal buildings on the left of the k^{th} building, because of the fact that this sequence is strictly decreasing. In other words, we can remove the last element of that sequence, and then the next one and so on, until we meet a building whose height is greater than h_k or we reach the end of the sequence. At that point, we get a new sequence with all the buildings on the left of the $(k + 1)^{\text{th}}$ building offering an ideal view of it, and we can later use that sequence to compute the sequence of all the buildings offering an ideal view from the left of the $(k + 2)^{\text{th}}$ building, and so on until we reach the last building. We also know that no building on the left of the first building offers an ideal view of fireworks launched from the first building, because there is no building on the left of the first building.

By maintaining such a sequence (which is more easily represented by a stack) by starting from the leftmost building and iterating through all buildings from left to right, we can find the number of buildings on the left of every building offering an ideal view of fireworks launched from the top of that building. The procedure would go as follows:

1. We create an empty stack S .
2. We start considering the 0^{th} building. S contains all the buildings left of the 0^{th} which offer an ideal view of fireworks launched from the 0^{th} building.
3. S 's size is 0, so we memorize that 0 is the number of ideal buildings on the left of the 0^{th} building.
4. The 0^{th} building is taller than no element in S , so we don't remove anything from S , and we add the 0^{th} building to the stack.

... (We continue by considering the next buildings)

5. We're now considering the i^{th} building. S contains all the buildings left of the i^{th} which offer an ideal view of fireworks launched from the i^{th} building.
6. S 's size is now X , so we memorize that X is the number of ideal buildings on the left of the i^{th} building.
7. The i^{th} building is of at least the same height as the first Y items on top of S , so we remove these from S , and we add the i^{th} building on top of S .

... (We continue by considering the next buildings)

8. We've now found the number of ideal buildings on the left of the $(N - 1)^{\text{th}}$ building, so we can stop.

It is possible to implement a similar process to compute the number of buildings offering an ideal view from the right of any building. Then, after we have computed both the number of ideal buildings on the left and on the right of some building, we know that the sum of these two numbers is the number of buildings offering an ideal view on fireworks launched from the top of that building. Hence, to get the overall solution, it suffices to iterate through every i between 0 and $N - 1$ and remember what the maximal solution was among all of these.

Here is some code implementing this algorithm:

```

1 #include <bits/stdc++.h>
2 using namespace std;
3 int main() {
4     int T; cin >> T;
5     for (int t = 0; t < T; t++) {
6         int N, result = 0; cin >> N;
7         vector<int> h(N), left(N), right(N);
8         stack<int> leftStack, rightStack;
9         for (int j = 0; j < N; j++)
10            cin >> h[j];
11        for (int i = 0; i < N; i++) { // ideal buildings on the left
12            left[i] = leftStack.size();
13            while (!leftStack.empty() && h[i] >= leftStack.top())
14                leftStack.pop();
15            leftStack.push(h[i]);
16        }
17        for (int i = N - 1; i >= 0; i--) { // ideal buildings on the right
18            right[i] = rightStack.size();
19            while (!rightStack.empty() && h[i] >= rightStack.top())
20                rightStack.pop();
21            rightStack.push(h[i]);
22        }
23        for (int i = 0; i < N; i++) { // find optimal value
24            result = max(result, left[i] + right[i]);
25        }
26        cout << "Case #" << t << ": " << result << '\n';
27    }
28 }

```

The runtime analysis of this algorithm is a little bit more complicated than for the earlier subtasks, because it is at first glance unclear how much time the inner loops on the 13th and 19th lines need. The key observation is that every building will be added once to each queue and removed at most once from each queue. Overall, our solution runs in $O(N)$.



3 Touristtrap

Given is a rectangular map and a starting spot, your task is to determine whether or not it is possible to reach any of the border tiles from the starting point.

The four different subtasks had differing representations of the map and increasing limits on map complexity and size.

Subtask 1: ASCII map (20 points)

In the first subtask the map is represented as a grid where each tile is either accessible (`_`) or not (`#`). This case can be solved by starting a depth-first-search (<https://soi.ch/wiki/dfs/>) from Mouse Stofl's starting tile on the graph where each tile is connected to its four neighboring tiles if they are accessible. If the DFS visits any border tiles, it's possible to escape, otherwise not. Since for a graph with N nodes and M edges DFS runs in $O(N + M)$ time, and we have one node for each accessible tile and at most four edges for each tile, the solution runs in $O(WH)$.

```
1 #include <bits/stdc++.h>
2 using namespace std;
3
4
5 signed main() {
6     int T; cin >> T;
7     for(int t=0; t<T; ++t) {
8         int w, h, x, y; cin >> w >> h >> x >> y;
9         vector<string> map;
10        copy_n(istream_iterator<string>(cin), h, back_inserter(map));
11        cout << "Case #" << t << ": ";
12        vector<array<int,2>> stack={{x,y}};
13        bool poss = false;
14        while(!stack.empty()) {
15            auto pos = stack.back(); stack.pop_back();
16            if (pos[0]<0 || pos[0]>=w || pos[1]<0 || pos[1]>=h)
17                continue;
18            if (map[pos[1]][pos[0]] != '_')
19                continue;
20            if (pos[0]==0 || pos[1]==0 || pos[0]==w-1 || pos[1]==h-1) {
21                poss = true;
22                break;
23            }
24            map[pos[1]][pos[0]] = '.';
25            stack.push_back({pos[0]-1,pos[1]});
26            stack.push_back({pos[0]+1,pos[1]});
27            stack.push_back({pos[0],pos[1]-1});
28            stack.push_back({pos[0],pos[1]+1});
29        }
30        cout << ("IMPOSSIBLE\n"+2*poss);
31    }
32 }
```

Subtask 2: Vector map (20 points)

In this subtask, the map is now given as a list of rectangles, whose contained tiles should be considered accessible. Since the limits on the size is still there, any solution for Subtask 1 should still work, we just need to convert to the old format. We can start with a grid filled with # characters and fill each rectangle with _ characters. This will take $O(RWH)$ time in the worst case (most rectangle covering most of the grid). The test data didn't contain many cases with large-area rectangles, but even in the worst case the code is fast enough for $R \leq 1000$ because constants are very small.

```

1 #include <bits/stdc++.h>
2 using namespace std;
3
4
5 signed main() {
6     int T; cin >> T;
7     for(int t=0; t<T; ++t) {
8         int w, h, x, y, r; cin >> w >> h >> x >> y >> r;
9         vector<string> map(h,string(w,'#'));
10        for(int i=0; i<r; ++i) {
11            int w, h, x, y; cin >> x >> y >> w >> h;
12            for_each(map.begin()+y, map.begin()+y+h, [x,w](auto &row) {
13                fill_n(row.begin()+x, w, '_');
14            });
15        }
16        cout << "Case #" << t << ": ";
17        vector<array<int,2>> stack={{x,y}};
18        bool poss = false;
19        while(!stack.empty()) {
20            auto pos = stack.back(); stack.pop_back();
21            if (pos[0]<0 || pos[0]>=w || pos[1]<0 || pos[1]>=h)
22                continue;
23            if (map[pos[1]][pos[0]] != '_')
24                continue;
25            if (pos[0]==0 || pos[1]==0 || pos[0]==w-1 || pos[1]==h-1) {
26                poss = true;
27                break;
28            }
29            map[pos[1]][pos[0]] = '.';
30            stack.push_back({pos[0]-1,pos[1]});
31            stack.push_back({pos[0]+1,pos[1]});
32            stack.push_back({pos[0],pos[1]-1});
33            stack.push_back({pos[0],pos[1]+1});
34        }
35        cout << ("IMPOSSIBLE\n"+2*poss);
36    }
37 }

```

Subtask 3: Bigger map (30 points)

Now that we don't have a significant restriction on W and H (up to 10^9), our $O(WH)$ algorithm won't be fast enough.

We notice that any two tiles within the same rectangle are connected (possibly with some intermediate tiles). Because of the transitivity property of connectivity (if both tiles A and B



and tiles B and C are connected, A is also connected to C), if any two tiles of two rectangles are connected, all tiles in the rectangles are connected. Because all tiles inside a rectangle have the same connectivity, we can simplify our graph to only have one node for each rectangle. Now we just need a good way to figure out if two rectangles are connected.

We have to consider four cases for a pair of rectangles:

Image	Description of extents	Connected?
	Disjoint in at least one dimension	×
	Intersecting in both dimensions.	✓
	Intersecting in one and touching in one dimension	✓
	Only touching in both dimensions	×

We can then create a node for each rectangle, and check every pair of rectangles to add an edge if they are connected. We then run DFS on this graph, starting from the rectangle containing Mouse Stofl's starting point (any of them if there are several). If we visit any rectangle that touches the border, we have an escape route, otherwise not.

Checking connectivity between rectangles is a constant-time operation, so we spend $O(R^2)$ time on generating the graph. The DFS also takes at most $O(R^2)$ to run in the worst case (many connections), giving a total running time of $O(R^2)$.

```
1 #include <bits/stdc++.h>
2 using namespace std;
3
4 typedef array<int, 4> int4;
5
6 namespace std {
7     istream& operator>>(istream& stream, int4 &r) {
8         return stream >> r[0] >> r[1] >> r[2] >> r[3];
9     }
10 }
11
12 int connectivity(int xs, int xe, int ys, int ye) {
13     if (xs > ye || ys > xe)
14         return 0;
15     if (xs == ye || ys == xe)
16         return 1;
17     return 2;
18 }
19 bool connected(int4 &a, int4 &b) {
20     return connectivity(a[0], a[0]+a[2], b[0], b[0]+b[2])
21         + connectivity(a[1], a[1]+a[3], b[1], b[1]+b[3]) > 2;
22 }
23 bool has(int4 &a, int x, int y) {
24     return a[0] <= x && x <= a[0]+a[2]
25         && a[1] <= y && y <= a[1]+a[3];
```

```

26 }
27 bool border(int4 &a, int w, int h) {
28     return a[0] == 0 || a[1] == 0 || a[0]+a[2] == w || a[1]+a[3] == h;
29 }
30
31 signed main() {
32     int T; cin >> T;
33     for(int t=0; t<T; ++t) {
34         int w, h, x, y, r, s; cin >> w >> h >> x >> y >> r;
35         vector<int4> rects;
36         copy_n(istream_iterator<int4>(cin), r, back_inserter(rects));
37
38         vector<vector<int>> >g(r);
39         for (int i = 0; i < r; ++i) {
40             if (has(rects[i], x, y))
41                 s = i;
42             for (int j = 0; j < i; ++j) {
43                 if (connected(rects[i], rects[j])) {
44                     g[i].push_back(j);
45                     g[j].push_back(i);
46                 }
47             }
48         }
49
50         cout << "Case #" << t << ": ";
51         vector<int> stack={s};
52         vector<bool> vis(r,false);
53         bool poss = false;
54         while(!stack.empty()) {
55             int pos = stack.back(); stack.pop_back();
56             if (vis[pos])
57                 continue;
58             vis[pos] = true;
59             if (border(rects[pos], w, h)) {
60                 poss = true;
61                 break;
62             }
63             copy(g[pos].begin(), g[pos].end(), back_inserter(stack));
64         }
65         cout << ("IMPOSSIBLE\n"+2*poss);
66     }
67 }

```

Subtask 4: Even bigger map (30 points)

For this subtask, we need a solution that runs in $O(n \log n)$ in the worst case. Even if we have a smarter way of finding connecting rectangles, we can still have up to $O(R^2)$ connections (consider many long vertical and horizontal rectangles).

Dealing with rectangles touching only at corners

We first deal with rectangles that touch only at a corner. If we replace $[a, b] \times [c, d]$ by the “cross” $[3a, 3b] \times [3c + 1, 3d - 1] \cup [3a + 1, 3b - 1] \times [3c, 3d]$, then two rectangles are connected if and only if the corresponding crosses touch or intersect. We also note that two crosses never touch only at a



corner, so this transformation got rid of that special case.

An elementary $O(n \log n)$ solution

We perform scanline on x and use a union-find datastructure to store the components of rectangles based on all intersection to the left of the sweepline. We maintains two sets of pairs (y -coordinates, rectangle-id):

- The set S stores the lower and upper coordinate of every rectangle that intersects the sweepline.
- The set T stores all elements of S that are in a different component than their predecessor in S or in a different component than their successor in S .
- Whenever we insert or delete something in S , we need to check whether its predecessor or successor needs to be added to or removed from T .

The intuitive idea behind the second set is the following: If many adjacent y -coordinates b_1, \dots, b_k belong to the same component, then a new rectangle $[y_1, y_2]$ that intersects some of them is either fully contained in $[b_1, b_k]$, or it contains b_1 or b_k .

At the current x -position, we first process add all rectangles that start here and then we remove all rectangles that end here. To add a rectangle R with y -range $[y_1, y_2]$, we do the following.

1. We first seek to find all components that R intersects. Let A be the set of all coordinates in T that are in $[y_1, y_2]$. We unite the R with all elements of A . Next, we remove all but the first and last element of A from T . This removal step makes the solution fast, as whenever we need a lot of time, there are a lot of elements in A , so also remove a lot of elements from T . Finally, we check whether we need to remove the first and last elements of A from T .
2. If in the first step, A was empty, then we're in the special case where all elements of S in $[y_1, y_2]$ belong to the same component. In this case, we just unite A with any element of S in $[y_1, y_2]$.
3. Insert y_1 and y_2 into S (and then as always check if we need to add / remove elements to / from T).

To remove a rectangle, simply remove its endpoints from S and T .

After the scanline, we know how the connected components of the rectangles look like, so we just iterate over all components and check if they contain both the starting point and a rectangle touching the border.

Note that this solutions uses only requires some scanline with "std::set". No fancy datastructures are used here.

Runtime analysis

This solution runs in $\Theta(r \log r)$. Sorting the events and the union-find take $\Theta(r \log r)$ time. The $2r$ insert and remove operations performed on S take $O(\log r)$ time each, so that's also fine. The non-obvious part is that maintaining the set T takes $O(r \log r)$ time.

We first note, that T can be efficiently updated when we insert or erase some element from S , as we only check if the predecessor and successor should belong to T , so we insert or remove at most two elements from T . Hence this doesn't increase the asymptotic runtime.

What remains to show is that the first step when adding a rectangle is fast. Finding the set A and removing its interior elements from T takes $O(\log r + |A|)$ time. While $|A|$ might be large, we

then also remove $|A| - 2$ elements from T . As all other steps run in $\Theta(r \log r)$ time, we add at most that many elements to T , so we can also remove at most that many elements from T . Hence the combined size of all A is at most $O(r \log r)$ and this part is also fast.

```

1 #include <bits/stdc++.h>
2 using namespace std;
3
4 using ll = int64_t;
5 constexpr int inf = 1e9;
6
7 struct Rectangle{
8     ll x1, x2, y1, y2;
9     int id;
10    bool contains(ll x, ll y) const {
11        return x1 <= x && x < x2 && y1 <= y && y < y2;
12    }
13    bool touches_border(ll w, ll h) const {
14        return x1 == 0 || y1 == 0 || x2 == w || y2 == h;
15    }
16 };
17 struct Event{
18     ll x;
19     ll y1, y2;
20     int id;
21     bool is_deletion;
22     bool operator<(Event const&o) const {
23         return make_pair(x, is_deletion) < make_pair(o.x, o.is_deletion);
24     }
25 };
26 struct Union_Find{
27     vector<int> p;
28     Union_Find(int n) : p(n) {
29         iota(p.begin(), p.end(), 0);
30     }
31     int f(int x){
32         return p[x] == x ? x : p[x] = f(p[x]);
33     }
34     void u(int a, int b){
35         p[f(a)] = f(b);
36     }
37 };
38
39 signed main(){
40     int T;
41     cin >> T;
42     for(int cas=0; cas<T; ++cas){
43         cout << "Case #" << cas << ": ";
44         int r;
45         ll W, H, X, Y;
46         cin >> W >> H >> X >> Y >> r;
47         vector<Rectangle> recs;
48         // read rectangles and transform them into crosses
49         for(int i=0; i<r; ++i){
50             ll x1, y1, x2, y2, w, h;
51             cin >> x1 >> y1 >> w >> h;

```



```
52     x2 = x1+w;
53     y2 = y1+h;
54     recs.push_back(Rectangle{3*x1+1, 3*x2-1, 3*y1, 3*y2, i});
55     recs.push_back(Rectangle{3*x1, 3*x2, 3*y1+1, 3*y2-1, i});
56 }
57 // also transform grid
58 X=3*X+1; Y=3*Y+1; W*=3; H*=3;
59 // compute events
60 vector<Event> evs;
61 for(auto const&rec:recs){
62     // insertion event
63     evs.push_back(Event{rec.x1, rec.y1, rec.y2, rec.id, false});
64     // deletion event
65     evs.push_back(Event{rec.x2, rec.y1, rec.y2, rec.id, true});
66 }
67 sort(evs.begin(), evs.end());
68 // set up datastructures
69 set<pair<ll, int> > S, T;
70 Union_Find uni(r);
71 auto should_be_in_T = [&](set<pair<ll, int> >::iterator it){
72     // check if first or last
73     if(it == S.begin()) return true;
74     auto it2 = next(it);
75     if(it2 == S.end()); return true;
76     // check if successor is in different component
77     if(uni.f(it2->second) != uni.f(it->second)) return true;
78     // check if predecessor is in different component
79     auto it3 = prev(it);
80     if(uni.f(it3->second) != uni.f(it->second)) return true;
81 };
82 auto update_in_T = [&](set<pair<ll, int> >::iterator it){
83     if(should_be_in_T(it)){
84         T.insert(*it);
85     } else {
86         T.erase(*it);
87     }
88 };
89 auto check_pred_succ = [&](pair<ll, int> const&p){
90     auto it_suc = S.upper_bound(p);
91     if(it_suc != S.end()) update_in_T(it_suc);
92     auto it = S.lower_bound(p);
93     if(it != S.begin()) update_in_T(prev(it));
94 };
95 auto remove_point = [&](pair<ll, int> const&p){
96     S.erase(p);
97     T.erase(p);
98     check_pred_succ(p);
99 };
100 auto add_point = [&](pair<ll, int> const&p){
101     auto it = S.insert(p).first;
102     update_in_T(it);
103     check_pred_succ(p);
104 };
105 // scanline
106 for(auto const&ev:evs){
107     if(ev.is_deletion){
```

```

108     remove_point(make_pair(ev.y1, ev.id));
109     remove_point(make_pair(ev.y2, ev.id));
110 } else { // insertion
111     auto it_l = T.lower_bound(make_pair(ev.y1, -inf)), it_r =
112     ↪ T.upper_bound(make_pair(ev.y2, inf));
113     // The set A of intersections is now [it_l, ..., prev(it_r)]
114     if(it_l == it_r){
115         // special case: [y1, y2] is inside endpoints belonging to the same
116         ↪ component
117         // find the <=1 component [y1, y2] intersects
118         auto it = S.lower_bound(make_pair(ev.y1, -inf));
119         if(it != S.end() && it->first <= ev.y2){
120             uni.u(it->second, ev.id);
121         }
122     } else {
123         // unite with elements in A
124         for(auto it = it_l; it != it_r; ++it){
125             uni.u(it->second, ev.id);
126         }
127         // remove elements in the interior of A
128         auto erase_l = next(it_l), erase_r = prev(it_r);
129         if(erase_l != it_r){ // don't erase if A has size 1
130             T.erase(erase_l, erase_r);
131             // check if first and last element of A still belong into T
132             update_in_T(S.find(*it_l));
133             update_in_T(S.find(*erase_r));
134         }
135     }
136     add_point(make_pair(ev.y1, ev.id));
137     add_point(make_pair(ev.y2, ev.id));
138 }
139 // check if a component contains (X, Y) and touches the border.
140 vector<bool> contains(r,0), touches(r, 0);
141 for(auto &e:recs){
142     if(e.contains(X, Y)) contains[uni.f(e.id)] = true;
143     if(e.touches_border(W, H)) touches[uni.f(e.id)] = true;
144 }
145 bool can_escape = false;
146 for(int i=0; i<r; ++i){
147     if(contains[i] && touches[i]) can_escape = true;
148 }
149 if(can_escape){
150     cout << "POSSIBLE\n";
151 } else {
152     cout << "IMPOSSIBLE\n";
153 }
154 }

```



4 Pipeline

Given are N spots to place pumps ($x_0 < x_1 < \dots < x_{N-1}$). For any subset of points including the start and the end, that is all $\{x_{i_0}, x_{i_1}, \dots, x_{i_P}\}$ with $0 = i_0 < i_1 < \dots < i_P = N - 1$, we define the LTC as the maximum distance between any two adjacent points distance, i.e. $LTC = \max_{0 \leq j \leq P-2} \{x_{i_{j+1}} - x_{i_j}\}$.

The three subtasks were:

1. Find the minimal P for a fixed LTC.
2. Find the minimal LTC for a fixed P .
3. Find the minimal P for all $LTC \in \{1, 2, \dots, M\}$ with an algorithm that has minimal space usage.

Subtask 1: Fixed LTC (15 points)

To find the minimal number of pumps for a fixed LTC, we do the following algorithm: Go from left to right and only take those pumps that are absolutely necessary. A pump is necessary if taking the next pump has a distance larger than LTC from the previous pump that has been taken.

An algorithm that solves a problem by making locally optimal decision is called *greedy*. In this case, we take the optimal pump location at each step, ignoring the possible spots after that point or before the last station that was taken. While this gives us the best solution locally, it does not necessarily give us a globally optimal solution. A proof of why this is indeed correct for this task is given in the last subtask.

```
1 #include <bits/stdc++.h>
2 using namespace std;
3
4 vector<int> greedy(int ltc, vector<int> const& xs) {
5     vector<int> take = { xs[0] }; // always take the first
6     int x_prev = xs[0];
7     for (auto x : xs) {
8         if (x - take.back() > ltc)
9             take.push_back(x_prev); // gap too large, need to take the last one
10        if (x - take.back() > ltc)
11            return {}; // impossible to cover, return empty vector
12        x_prev = x;
13    }
14    if (x_prev != take.back()) // take the last, if not already taken
15        take.push_back(x_prev);
16    return take;
17 }
18
19 int main() {
20     int t; cin >> t;
21     for (int tc=0; tc<t; ++tc) {
22         int n, ltc; cin >> n >> ltc;
23         vector<int> xs;
24         copy_n(istream_iterator<int>(cin), n, back_inserter(xs));
25         cout << "Case #" << tc << ": ";
26         auto answer = greedy(ltc, xs);
```

```

27     if (answer.empty())
28         cout << "Impossible";
29     else if (answer.size() > 9000)
30         cout << answer.size();
31     else
32         copy(answer.begin(), answer.end(), ostream_iterator<int>(cout<<answer.size()<<' ', " "));
33     cout << '\n';
34 }
35 }

```

Subtask 2: Fixed Pumps Count (15 points)

This time we want to minimize the LTC given a fixed number of pumps. The key idea was to not construct the solution directly, but to use the algorithm of the previous subtask.

Say we want to use 4 pumps. We can run the previous algorithm for multiple values of LTC, for example:

- Solve with LTC = 1 \rightsquigarrow impossible
- Solve with LTC = 2 \rightsquigarrow impossible
- Solve with LTC = 3 \rightsquigarrow $P = 9$
- Solve with LTC = 4 \rightsquigarrow $P = 6$
- Solve with LTC = 5 \rightsquigarrow $P = 5$
- Solve with LTC = 6 \rightsquigarrow $P = 4$
- Solve with LTC = 7 \rightsquigarrow $P = 3$
- Solve with LTC = 8 \rightsquigarrow $P = 3$
- Solve with LTC = 9 \rightsquigarrow $P = 3$
- Solve with LTC ≥ 10 \rightsquigarrow $P = 2$

And then we see that we can get an LTC of 6 using exactly 4 pumps. A smaller LTC, say LTC=5 requires at least 5 pumps.

What if we have $P = 3$ pumps? Because we want to minimize the LTC, we take the *smallest*, in this case LTC=7. Also what if we have $P = 7$ pumps? We can't ever select spots with an LTC of 3, so we have to do an LTC=4 with 6 pumps. As we're required by the task description to actually use 7 pumps, we just add another pump that has been previously unused.

In general, the algorithm would be to take the smallest LTC that uses at most P pumps and then fill up with remaining pumps until it uses exactly P pumps.

If we were to code just that, trying all possible LTC's until we're good is too slow. However, we can do it smarter. We notice that the values of the P are decreasing when the LTC is increasing. This allows us to make use of binary search: We start with some value, say LTC = 6. If we use at least P pumps, we decrease it, say to LTC = 3. If we use less than P pumps, we increase it, say to LTC = 12.

More formally, we have two variables: l (for left) and r (for right). Those two variables define our search range $\{l + 1, l + 2, \dots, r\}$ (meaning we know that the optimal LTC lies in this range) and satisfy the following conditions:

- if we have LTC = l , it's either impossible or we need $\geq P$ pumps



- if we have $LTC = r$, it's possible and we need $< P$ pumps

Initially, we set $l = 0$ (then it's never possible) and $r = x_{N-1} - x_0$ (then we only need 2 pumps). We then take a value in the middle, $m = \lfloor \frac{l+r}{2} \rfloor$. We check what happens with $LTC = m$. If it's possible and we need less than P pumps, we set the variable r to m . If not, it's either impossible or we need at least P pumps, so we can set l to m . Note that our search range decreased from $r - l$ to $l - m$ or $r - m$ and in both cases it has been halved. So if originally we had $l = 0$ and $r = 2^{30}$, we need exactly 30 steps to have $r = l + 1$. In this case, we have found our optimal LTC of r . As the distance between the last and the first pump station are at most $10^9 < 2^{30}$, we will use at most $\lceil \log_2 10^9 \rceil = 30$ steps. So this algorithm runs in $O(n \log C)$ (where $C = x_{N-1} - x_0$).

Note that the initial values as given above are incorrect for $N = 1$. In this case, solving it with $l = 0$ is always possible. We can fix that by starting with $l = -1$.

There are many ways to "fill up" the remaining pumps. One is to put all pump locations selected by the greedy algorithm in a `std::set<int>` and insert values until its size is P . A set keeps each value only once, so this works and the running time would be $O(n \log(\max n, C))$ is fast because inserting into a set is logarithmic time and we do it at most n times.

Another way to do it is to make use of the fact that the values are already sorted, so we can compute the unused values in linear time and then take as many of them as we need. The function `set_difference` computes the difference between two sorted lists ($A \setminus B = \{x \in A \mid x \notin B\}$) and the function `inplace_merge` sorts a list that consists of two sorted halves. Doing it this way gets rid of a log-factor in the last step. The total running time thus is $O(n \log C)$.

```
1 vector<int> greedy(int ltc, vector<int> const& xs) {
2     ... // same as in previous subtask
3 }
4
5 pair<int, vector<int>> binary_search(size_t max_pumps, vector<int> const& xs) {
6     // invariant:
7     // - impossible with <=max_pumps and LTC=left
8     // - possible with <=max_pumps and LTC=right
9     int left=-1, right=xs.back() - xs.front();
10    while (right - left > 1) {
11        int mid = (left + right)/2;
12        auto ans = greedy(mid, xs);
13        if (ans.empty() || ans.size() > max_pumps)
14            left = mid; // impossible
15        else
16            right = mid; // possible
17    }
18    auto take = greedy(right, xs);
19    vector<int> unused;
20    set_difference(xs.begin(), xs.end(), take.begin(), take.end(), back_inserter(unused));
21    int necessary = take.size();
22    copy_n(unused.begin(), max_pumps - necessary, back_inserter(take));
23    inplace_merge(take.begin(), take.begin() + necessary, take.end());
24    return {right, take};
25 }
26
27 int main() {
28     int t; cin >> t;
29     for (int tc=0; tc<t; ++tc) {
30         int n, p; cin >> n >> p;
```

```

31     vector<int> xs;
32     copy_n(istream_iterator<int>(cin), n, back_inserter(xs));
33     auto [ltc, answer] = binary_search(p, xs);
34     cout << "Case #" << tc << ": " << ltc << '\n';
35     if (answer.size() <= 9000)
36         copy(answer.begin(), answer.end(), ostream_iterator<int>(cout << '\n', " "));
37     else
38         cout << "Over 9000";
39     cout << '\n';
40 }
41 }

```

Subtask 3: Walking down the BTC (10 points)

To find the minimal P for all $LTC \in \{1, 2, \dots, M\}$, it was enough to run the solution of subtask 1 all different values of LTC. The goal of this task was to prepare for the last subtask.

Subtask 4: Walking down the ETH in $O(N \cdot M)$ (20/60 points)

This section shows the intended solution for 20 out of 60 points. We will look at the solution that score more than that in the following section. We structure the proof as follows. First, we show that greedy works and is indeed correct. Then, we start with the trivial solution that runs greedy on all values and transform the solution into another one that computes the same results but uses less memory.

So, why does greedy work? Looking again at the source code for subtask 1:

```

1 vector<int> greedy(int ltc, vector<int> const& xs) {
2     vector<int> take = { xs[0] }; // always take the first
3     int x_prev = xs[0];
4     for (auto x : xs) {
5         if (x - take.back() > ltc)
6             take.push_back(x_prev); // gap too large, need to take the last one
7         if (x - take.back() > ltc)
8             return {}; // impossible to cover, return empty vector
9         x_prev = x;
10    }
11    if (x_prev != take.back()) // take the last, if not already taken
12        take.push_back(x_prev);
13    return take;
14 }

```

The idea was to always select the last pump station we can take without violating the LTC. By construction, we always take the first and the last element and if we find a solution, we always find one that is viable (not violating any of the constraints). What remains to be shown is that what we select is optimal, meaning it has the minimal number of pump stations.

As with many greedy algorithm, we can do a proof by contradiction followed by an *exchange argument*. Consider an arbitrary input. Our algorithm will produce the solution $A = \{a_0, a_1, \dots, a_{P-1}\}$. Look at an optimal solution $O = \{o_0, o_1, \dots, o_{Q-1}\}$. Assume that this other solution is *better* than ours, that is $Q < P$. Without loss of generality, let's further assume that in case there are multiple optimal solutions, we take the one that agrees with ours on the longest possible prefix. In other words take that O such that there is an k where both mismatch ($a_i = o_i$ for $0 \leq i \leq k$ and $a_k \neq o_k$)



and this k is maximal. Such an k always exists because the last element of both solutions must be equal (as it was required by the task statement that start and end are in there) and as they have different length, they must differ at some element. Furthermore, we have $2 \leq k < Q$ because the start is the same in both elements.

Now, by construction of A , we must have $a_k \geq o_k$. Why? Firstly, a_k is not the last element ($a_k \neq x_{N-1}$) because only a_{Q-1} can be equal to x_{N-1} and we know that $Q < P$. So we can only have chosen a_k on line 6. The condition on line 5 tells us that the spot following a_k is larger than $a_{k-1} + \text{LTC}$, so it would be illegal to be chosen. Therefore, $o_k > a_k$ is impossible (illegal), $o_k = a_k$ is impossible (by definition of k), so we must have $o_k < a_k$.

As $o_k \neq x_{N-1}$, we must have $k < Q - 1$ and thus o_{k+1} exists. Now comes the clue: We construct a new solution $O' = \{o_0, o_1, \dots, o_{k-1}, a_k, o_{k+1}, \dots, o_{Q-1}\}$. This solution is valid, because (a) $a_k - o_{k-1} \leq \text{LTC}$ (because $o_{k-1} = a_{k-1}$), (b) $o_{k+1} - a_k \leq \text{LTC}$ (because $a_k > o_k$), (c) $o_{k+1} - a_k > 0$ (otherwise, $O \setminus \{o_k\}$ would be a valid solution and thus O would not be optimal). However, this solution agrees with A on the first k elements, and a mismatch only happens at $k' \geq k + 1$, contradicting the assumption that k was maximal. We have a contradiction, thus we can only conclude that such an O can not exist. \square

The running time of this solution is clearly $O(n)$ as we iterate through each value exactly once.

Doing this naively for every LTC we end up with the following algorithm:

```
1 vector<int> serial_greedy(int max_ltc, vector<int> const& xs) {
2     const int INF = 1e9; // sentinel
3     vector<int> last(max_ltc+1, xs[0]);
4     vector<int> need(max_ltc+1, 1);
5     for (int ltc=1; ltc <= max_ltc; ++ltc) { // ltc-loop
6         for (int i=0; i<xs.size(); ++i) { // xs-loop
7             if (xs[i] - last[ltc] > ltc) {
8                 last[ltc] = xs[i-1]; // gap too large, need to take the last one
9                 need[ltc] += 1;
10            }
11            if (xs[i] - last[ltc] > ltc)
12                need[ltc] = INF; // impossible
13        }
14        if (xs.back() != last[ltc]) // take the last, if not already taken
15            need[ltc] += 1;
16    }
17    need.erase(need.begin()); // ltc=0 is uninteresting
18    replace_if(need.begin(), need.end(), [](int n) { return n >= INF; }, -1);
19    return need;
20 }
```

This code is slightly different to the one shown before in order to prepare for what's about to come. We don't return early if we see something is impossible, we set the number of required pumps to $+\infty$ instead. We also store our temporary variables inside the arrays `last` and `need`. Functionally the code remains equivalent after these changes.

The memory required for this solution is $O(N + M)$. We need an array of size N for all values $\{x_0, \dots, x_{N-1}\}$. We also need two arrays of size $M + 1$ for our temporary variables.

We can't get rid of this $O(N)$ easily because we need all the values at different points in time. What we need to do first is to swap the two for loops. Doing so will not change *what's* computed, it will only change *when* it's computed.


```

1 vector<int> parallel_greedy(int max_ltc, vector<int> const& xs) {
2     const int INF = 1e9; // sentinel
3     vector<int> last(max_ltc+1, xs[0]);
4     vector<int> need(max_ltc+1, 1);
5     for (int i=0; i<xs.size(); ++i) { // xs-loop
6         for (int ltc=1; ltc <= max_ltc; ++ltc) { // ltc-loop
7             if (xs[i] - last[ltc] > ltc) {
8                 last[ltc] = xs[i-1]; // gap too large, need to take the last one
9                 need[ltc] += 1;
10            }
11            if (xs[i] - last[ltc] > ltc)
12                need[ltc] = INF; // impossible
13        }
14    }
15    for (int ltc=1; ltc <= max_ltc; ++ltc)
16        if (xs.back() != last[ltc]) // take the last, if not already taken
17            need[ltc] += 1;
18    need.erase(need.begin()); // ltc=0 is uninteresting
19    replace_if(need.begin(), need.end(), [](int n) { return n >= INF; }, -1);
20    return need;
21 }

```

Now we don't need to store all N values of x 's anymore at the same time! We only ever look at $xs[i]$ and $xs[i-1]$. We can thus read the values on the fly, only storing x (replacing $xs[i]$) and x_{prev} (replacing $xs[i-1]$). We've reached a streaming algorithm that only keeps $O(M)$ values at the same time, but still computes the same as the correct greedy algorithm.

```

1 vector<int> streaming_greedy(int max_ltc, int n) {
2     const int INF = 1e9; // sentinel
3     int x, x_prev;
4     cin >> x;
5     vector<int> last(max_ltc+1, x);
6     vector<int> need(max_ltc+1, 1);
7     for (int i=1; i<n; ++i) {
8         x_prev = x;
9         cin >> x;
10        for (int ltc=1; ltc <= max_ltc; ++ltc) {
11            if (x - last[ltc] > ltc) {
12                last[ltc] = x_prev; // gap too large, need to take the last one
13                need[ltc] += 1;
14            }
15            if (x - last[ltc] > ltc)
16                need[ltc] = INF; // impossible
17        }
18    }
19    for (int ltc=1; ltc <= max_ltc; ++ltc)
20        if (x != last[ltc]) // take the last, if not already taken
21            need[ltc] += 1;
22    need.erase(need.begin()); // ltc=0 is uninteresting
23    replace_if(need.begin(), need.end(), [](int n) { return n >= INF; }, -1);
24    return need;
25 }

```

The solution that runs in $(N \cdot M^{0.5}(\log M)^2)$ will be released in an update to this booklet in the online version.



5 Storytelling

In this task you were given n mice sitting around a circular table about to tell their stories. You were given rules how they were allowed to talk and your goal was to calculate the minimum time in which all mice could finish storytelling.

Subtasks 1 to 4

The first four subtasks had increasing input size which distinguished among different solutions. The Subtask 1 was solvable through a bruteforce (e.g., try all $n!$ permutations to find the answer) and Subtask 2 by a binary search.

To earn a full score one had to come up with a linear (both in time and space) solution which uses just a few minimums and maximums. One such implementation is listed below. The proof why this works is discussed in the following section and would earn you full score for the theoretical part.

```
1 #include <iostream>
2 #include <vector>
3 #include <algorithm>
4
5 using namespace std;
6
7 void schedule(int n, vector<long long>& V, vector<long long>&ret) {
8     ret.resize(n);
9
10    if (n % 2 == 0) { // even case
11        for (int i=0; i<n; ++i)
12            if (i % 2 == 0)
13                ret[i] = V[i]; // even mice speak from the beginning
14            else // odd mices speak when they can
15                ret[i] = V[i] + max(V[(i-1+n)%n], V[(i+1)%n]);
16    } else { // odd case
17        long long min_3 = V[n-2] + V[n-1] + V[0];
18        int min_i = 0; // find the mice min_i which will wait for min_i-2 and min_i-1
19        for (int i=0; i<n; ++i) {
20            long long s = V[(n+i-2)%n] + V[(n+i-1)%n] + V[i];
21            if (s < min_3) {
22                min_3 = s;
23                min_i = i;
24            }
25        }
26
27        for (int i = min_i + 1; i < min_i + n - 1; ++i) {
28            if (i % 2 != min_i % 2)
29                ret[i%n] = V[i%n]; // odd mices relative to min_i speak first
30            else
31                ret[i%n] = V[i%n] + max(V[(n+i-1)%n], V[(n+i+1)%n]);
32        }
33        // min_i-1 speaks once min_i-2 is finished
34        ret[(min_i-1+n)%n] = V[(min_i-2+n)%n] + V[(min_i-1+n)%n];
35        // in the end once min_i-1 and min_i+1 are finished min_i speaks
36        ret[min_i] = max(min_3, V[min_i] + V[(min_i+1)%n]);
37    }
```

```

37 }
38 }
39
40 int main() {
41     int T, N;
42
43     cin >> T;
44     for (int t=0; t<T; ++t) {
45         cin >> N;
46         vector<long long> V(N);
47         for (int i=0; i<N; ++i) cin >> V[i];
48
49         vector<long long> res; // endtimes of speaking
50         schedule(N, V, res);
51         cout << "Case #" << t << ": " << *max_element(res.begin(), res.end()) << '\n';
52         for (int i=0; i<N; ++i)
53             cout << res[i]-V[i] << (i == N-1 ? "\n" : " ");
54     }
55 }

```

Subtask 5

In order to have simpler formulas, all mouse indices in this solution are implicitly considered modulo n . For example, mouse $n + 3$ is the same as mouse 3, and $y_{3-n} = y_{3+n} = y_3$.

First, let's assume that n is even, i.e., $n = 2k$ for some integer k .

We'll start by showing how to find a lower bound: for any instance y_1, \dots, y_n we will find a time Y such that no solution can finish earlier than at Y . For each i we can make the following argument: Mice i and $i + 1$ cannot speak at the same time, so we know that mouse i cannot be done speaking and listening sooner than at the time $y_i + y_{i+1}$. Hence, the entire group cannot be done sooner than at the time $Y = \max_{1 \leq i \leq n} (y_i + y_{i+1})$.

Next, we will show that this Y is also an upper bound. In other words, we will find an algorithm that will always construct a schedule such that everybody will finish talking by the time Y . One such algorithm looks as follows:

- All mice with odd indices start telling their stories at time 0.
- Each mouse with an even index starts telling its story as soon as both of its neighbours finish telling theirs.

Clearly, the last mouse to finish talking will be an even-numbered mouse. Mouse $2j$ will finish talking precisely at the time $\max(y_{2j-1} + y_{2j}, y_{2j} + y_{2j+1})$. It is obvious that the maximum of all these times is precisely Y .

Now let's assume that n is odd, i.e., $n = 2k + 1$ for some integer k .

The lower bound shown in the even case still holds, as the argument is valid for any number of mice. However, in some situations this lower bound cannot be achieved. For example, if we have three mice with $y_1 = y_2 = y_3 = 1$, we get the lower bound $Y = 2$, but the actual optimal solution needs 3 seconds. Below, we will show a second lower bound that will cover such situations.

Suppose that we have any valid schedule telling each mouse when it should start speaking. As the number of mice is odd, there has to be a mouse who starts speaking **after** one of its neighbors but **before** the other one. (The opposite is impossible, because if you only have mice who start



speaking before both neighbors and mice who start speaking after both neighbors, these two type of mice have to alternate, which is impossible if the total number of mice is odd.)

Without loss of generality, suppose that mice q , $q + 1$, and $q + 2$ have the property that mouse $q + 1$ speaks after mouse q but before mouse $q + 2$. Then, mouse $q + 2$ cannot finish its own story sooner than at the time $y_q + y_{q+1} + y_{q+2}$.

Note that for the lower bound we now need to take the **minimum** (and not the **maximum**) of all such triples. This is because all we know is that each valid schedule has to contain at least one such triple, and it is possible that the optimal solution will contain the cheapest of all such triples. Thus, the second lower bound we get is $Y' = \min_{1 \leq i \leq n} (y_i + y_{i+1} + y_{i+2})$.

To finish this task, we will now show an algorithm for $n = 2k + 1$ that always produces an optimal solution – i.e., a solution that ends at the time $\max(Y, Y')$. This new algorithm looks as follows:

- Rotate the mice around the table until you get the situation where $y_{n-2} + y_{n-1} + y_n = Y'$.
- All mice with odd indices, except for mouse n , can start telling their stories at time 0.
- Each mouse with an even index, except for mouse $n - 1$, starts telling its story once all its neighbours from the first group finish telling theirs.
- Mouse $n - 1$ starts telling its story after mouse $n - 2$.
- Finally, mouse n can tell its story once both its neighbors are done.

Why is this algorithm optimal? The last mouse to finish talking is either an even-numbered mouse, or mouse n . An even-numbered mouse $2j$ will finish telling its story at the time $\max(y_{2j-1} + y_{2j}, y_{2j} + y_{2j+1}) \leq Y$. Mouse n will finish its story at the time $\max(y_n + y_1, y_{n-2} + y_{n-1} + y_n)$. The first of these two is at most Y , while the second one is exactly Y' . As a result, all mice will finish their stories by the time $\max(Y, Y')$.

Note that the entire algorithm can be implemented in linear time and space – each of the first three steps can be implemented as a single pass through the array of talking times.

6 Tank Golf

Scoring

Bugs

10 points were awarded for submitting running code. This means that the submission should compile and be able to handle all matches without crashing.

Elo

After fixing all crashes, the bots were tested against each other. Each bot was pitted against every other bot multiple times. An Elo rating was calculated for each bot using the Bayesian Elo algorithm¹. This means that Elo scores were chosen such that they maximize the likelihood of correctly predicting every match. The Elo ratings of the bots are shown in figure 6.1.

Elo scores r can be converted to playerstrength γ using

$$\gamma = 10^{\frac{r}{400}}$$

The win chance of a player E_1 can then be calculated using

$$E_1 = \frac{\gamma_1}{\gamma_1 + \gamma_2}$$

90 points were awarded proportionally to the playerstrength of the submission, i.e. $\text{score} = 90 \frac{\gamma}{\gamma_{\text{ref}}}$. The maximum score γ_{ref} was set at the third place submission by Nicolas Camenisch.

Strategies

Condition checking

By simulating multiple possible shooting directions or spawn locations, each possible outcome can be tested for a few conditions. This way the bot can recognize which shots will throw the enemy off the map. Advanced conditions such as trying to return ones own tank back into an upright position if it was flipped over can be evaluated as alternative options. The main weakness with checking for conditions directly is that there are too many cases to efficiently cover all of them.

State scoring

Instead of checking for conditions strictly, each condition can be converted into a score. These scores are added up to give a numeric value to every possible game state. Then, from all possible actions the highest scoring one is chosen. In addition to the previous conditions, things like distance to the nearest ledge and distance to the enemy tank can be included in the score.

¹<https://www.remi-coulom.fr/Bayesian-Elo/#theory>




Figure 6.1: Distribution of Elo ratings depending on the rank.

Targeted search

Once a scoring function has been established, possible shooting positions can be refined. As only a limited number of requests are available, it makes sense to focus the search near high scoring actions. Variations of gradient ascent can be used to find local maxima.