


SOI 2P 2021
Solution Booklet



Swiss Olympiad in Informatics

12–16 March 2021



Gummy Bears

Task Idea	Christopher Burckhardt, Daniel Rutschmann
Task Preparation	Johannes Kapfhammer
Description English	Johannes Kapfhammer
Description German	Tobias Feigenwinter
Description French	Florian Gatignon
Solution	Johannes Kapfhammer

In this task you were given a sequence f_0, \dots, f_{n-1} and you had to rearrange it into a $\frac{n}{k}$ groups such that each group consists of exactly k elements and contains a majority element. In other words, inside each group, at least $\lfloor \frac{k}{2} \rfloor + 1$ elements had to be equal.

With the brackets $\lfloor x \rfloor$ we denote the *floor function*, which is operation of rounding down to the next smaller integer. Also see https://en.wikipedia.org/wiki/Floor_and_ceiling_functions. The *ceil function*, written as $\lceil x \rceil$ is rounding up to the next integer.

Let $R = k - (\lfloor \frac{k}{2} \rfloor + 1) = \lceil \frac{k}{2} \rceil - 1$. Then, the output will always have a form similar to this:

$$\begin{array}{cccc|cccc} m_0 & m_0 & \dots & m_0 & r_0 & r_1 & \dots & r_{R-1} \\ m_1 & m_1 & \dots & m_1 & r_R & r_{R+1} & \dots & r_{2R-1} \\ \vdots & & & & & & & \vdots \\ m_{n/k} & m_{n/k} & \dots & m_{n/k} & & & \dots & r_{(n/k)R-1} \end{array}$$

$\underbrace{\hspace{10em}}_{\lfloor \frac{k}{2} \rfloor + 1 \text{ majority elements}} \quad \underbrace{\hspace{10em}}_{R \text{ filler elements}}$

The first $\lfloor \frac{k}{2} \rfloor + 1$ elements are identical and we call them “majority elements”. We don’t care about the value of other elements, so we call them the “filler elements”.

If we can somehow figure out the majority elements $m_0, \dots, m_{n/k}$, we can just use the remaining elements as filler.

For that we can count the number of elements that are the same. Let’s say value x occurs c_x times.

We can use it for a majority element if $c_x \geq \lfloor \frac{k}{2} \rfloor + 1$. We can use it for two majority elements if $c_x \geq 2 \cdot (\lfloor \frac{k}{2} \rfloor + 1)$. So in total, we may use it for $\lfloor \frac{c_x}{\lfloor \frac{k}{2} \rfloor + 1} \rfloor$ majority elements.

50 Points

For 50 points, it was enough to check whether it’s possible or not. This can be done with the above observations: compute how many majorities (the elements m_0, m_1, \dots) we can form, and check whether it will be at least n/k .

```
1 int n, k, s; cin >> n >> k >> s;
2 const int majority_quota = k/2 + 1;
3
4 unordered_map<int, int> flavour_count;
5 for (size_t i = 0; i < n; ++i) {
6     int x; cin >> x;
7     ++flavour_count[x];
8 }
9
10 int majorities = 0;
11 for (auto [f, cnt] : flavour_count) {
12     majorities += cnt / majority_quota;
13 }
14 if (majorities >= n/k)
15     cout << "Yummy!\n";
16 else
17     cout << "Yuck!\n";
```



100 Points

For the reconstruction, you also need to keep track of the values of the majority elements and the values of the filler elements. The following solution builds the output table by first filling out the majority elements and then later pad it with the filler elements.

```
1 int n, k, s; cin >> n >> k >> s;
2 const int majority_quota = k/2 + 1;
3
4 unordered_map<int, int> flavour_count;
5 for (int i = 0; i < n; ++i) {
6     int x; cin >> x;
7     ++flavour_count[x];
8 }
9 vector<vector<int>> blocks;
10 vector<int> filler;
11 for (auto [f, cnt] : flavour_count) {
12     // repeatedly add blocks of size majority_quota
13     while ((int)blocks.size() < n/k &&
14           cnt >= majority_quota) {
15         cnt -= majority_quota;
16         // add majority_quota times the value f
17         blocks.emplace_back(majority_quota, f);
18     }
19     // all remaining elements will be used as filler
20     for (int i=0; i<cnt; ++i)
21         filler.push_back(f);
22 }
23 // not enough majority elements
24 if ((int)blocks.size() < n/k) {
25     cout << "Yuck!\n";
26     return 0;
27 }
28 cout << "Yummy!\n";
29 for (auto& block : blocks) {
30     // fill it up to size k with filler elements
31     while ((int)block.size() < k) {
32         block.push_back(filler.back());
33         filler.pop_back();
34     }
35     copy(block.begin(), block.end(), ostream_iterator<int>(cout, " "));
36     cout << '\n';
37 }
```

Below is also a Python solution:

```
1 from collections import Counter
2
3 n, k, s = map(int, input().split())
4 q = k // 2 + 1
5
6 count = Counter(map(int, input().split()))
7 majority = list(Counter({b:(c - c % q) for b, c in count.items()}).elements())
8 filler = list(Counter({b:(c % q) for b, c in count.items()}).elements())
9
10 if len(majority) // q < n // k:
11     print("Yuck!")
12 else:
13     print("Yummy!")
14     for mouthful in zip(*[iter(majority[:q*n//k])*q, *[iter(filler+majority[q*n//k:])*q]*(k-q)]:
15         print(*mouthful)
```

Water View

Task Idea	Daniel Rutschmann
Task Preparation	Luc Haller
Description English	Luc Haller
Description German	Tobias Feigenwinter
Description French	Elias Boschung
Solution	Luc Haller

Subtask 1: 30 Points ($n \leq 500$)

To solve the first test group, we can iterate over all possible locations $0 \leq i \leq n - 1$, and in an inner loop counting the number of beautiful segments for this choice of i by checking for each segment $0 \leq j \leq n - 1$ if it's beautiful by doing a third nested loop and checking if there's no higher segment between j and the beginning, or the end, or i .

The running time is $O(n^3)$.

```
1 #include <bits/stdc++.h>
2 using namespace std;
3 using vi = vector<int>;
4
5 int main() {
6     int n; cin >> n;
7     vi hs(n);
8     for (int& hi : hs) cin >> hi;
9     int maxbeauty = 0;
10    for (int i=0; i<n; ++i) {
11        int beauty = 0;
12        for (int j=0; j<n; ++j) {
13            bool beautiful = true;
14            for (int k=0; k<=j; ++k) {
15                if (k == i) {
16                    beautiful = true;
17                } else if (hs[k] > hs[j]) {
18                    beautiful = false;
19                }
20            }
21            if (beautiful) {
22                ++beauty;
23                continue;
24            }
25            beautiful = true;
26            for (int k=n-1; k>=j; --k) {
27                if (k == i) {
28                    beautiful = true;
29                } else if (hs[k] > hs[j]) {
30                    beautiful = false;
31                }
32            }
33            beauty += beautiful;
34        }
35        maxbeauty = max(maxbeauty, beauty);
36    }
37    cout << maxbeauty << "\n";
38 }
```

Subtask 2: 60 Points ($n \leq 10^4$)

To improve our solution from above and bring it down to a running time of $O(n^2)$, we observe that it's not really necessary to check each j with a separate loop over the segments: Starting from a lake in one direction, the segments which can see it are exactly those which are at least as



high as everything between themselves and the lake, so we can easily mark all of them in one scan from the lake. So to compute the beautiful segments if we place the third lake at a given i , we'll just do one scan from 0 up to mark all segments which can see the lake on the left, one scan from $n - 1$ down for the lake on the right, and one scan left from $i - 1$ and one right from $i + 1$ for the lake at i . After that, we count the number of marked segments.

This results in the desired running time of $O(n^2)$.

```
1 #include <bits/stdc++.h>
2 using namespace std;
3 using vi = vector<int>;
4
5 void markgood(const vi& hs, int from, bool reverse, vi& good) {
6     int maxyet = 0;
7     for (int i=from; reverse ? i>=0 : i<good.size(); i += reverse ? -1 : 1) {
8         if (hs[i] >= maxyet) {
9             good[i] = 1;
10            maxyet = hs[i];
11        }
12    }
13 }
14
15 int main() {
16     int n; cin >> n;
17     vi hs(n);
18     for (int& hi : hs) cin >> hi;
19     int maxgood = -1;
20     for (int i=0; i<n; ++i) {
21         vi good(n);
22         markgood(hs, 0, false, good);
23         markgood(hs, n-1, true, good);
24         markgood(hs, i+1, false, good);
25         good[i] = 1;
26         markgood(hs, i-1, true, good);
27         int goodcount = accumulate(good.begin(), good.end(), 0);
28         maxgood = max(maxgood, goodcount);
29     }
30     cout << maxgood << "\n";
31 }
```

Subtask 3: 100 Points ($n \leq 500'000$)

For the full score, we need to consider all possible locations for the third lake in a constant number of scans over the list of heights. First, we do a scan from 0 up and one from $n - 1$ down to mark all segments which can see at least one of the two fixed lakes at the ends, since for these the placement of the third lake won't matter.

We can count the number of beautiful segments to the left of each i and to the right of each i in two separate scans. So now we've essentially reduced the problem to computing a list l_i , where l_i is the number of beautiful segments to the left of i if the third lake is placed at i .

To compute the l_i , consider the lists L_i of beautiful segments left of each i if the third lake is placed at i , or more specifically, consider how L_{i+1} differs from L_i : The only change in the relevant list of heights is that h_i is added at the end. No new segment before i can become visible because of this, but some segments may become hidden. By definition of *visibility*, everything at least as high as h_i stays visible, and everything lower is blocked from view. Also, i becomes visible.

This tells us how to compute L_{i+1} from L_i , but if we implemented this naively by looking at all elements of L_i , the worst case running time of the solution would still be quadratic. But we can combine the idea with the observation we made for the second test set: The heights of the elements of L_i are non-increasing. So to transform L_i into L_{i+1} , we can just keep removing the last element until its height is at least h_i or the list is empty, then add i . Note that we compare at most one element more than we remove, and each element is removed at most once (since it's added at most once), so all of these transformations from L_i to L_{i+1} in total run in $O(n)$.

So far I ignored the issue of how to take into account the two original lakes when computing the



L_i and l_i . We can't just deal with them fully separately, since that might lead to double-counting some beautiful segments. Consider an additional list m_i where we pre-marked the segments which can see an endpoint with 1 and the others with 0. So far l_i would just be $|L_i|$, and L_i may contain some segments j which can also see one of the lakes at the ends. We can easily solve that by just not adding i to L_{i+1} if $m_i = 1$. (Note that we still have to remove the segments which become hidden, though.) Now l_i is just $|L_i| + \sum_{j=0}^{i-1} m_j$.

But what if placing a lake at i and thereby removing h_i means additional $j < i$ can now see the lake at the right end? This would mean that m_j doesn't correctly represent visibility of the endpoints anymore. Yes, but that's not a problem: Any such j can also see the lake at i , since $j < i < n$, so it must be in L_i and we still count it exactly once for l_i . In fact, we don't need to take visibility of the end into account for m_i at all by this argument.

We analogously compute r_i on the reversed height list, and the total beauty if the third lake is placed at i is $l_i + r_i + 1$.

As we showed above, computing all l_i and r_i runs in $O(n)$, so this solution has linear running time.

```
1 #include <bits/stdc++.h>
2 using namespace std;
3 using vi=vector<int>;
4
5 vi solveleft(const vi& hs) {
6     const int n = hs.size();
7     vi beginvis(n); // m_i in the solution text
8     int maxh = 0;
9     for (int i=0; i<n; ++i) {
10         if (hs[i] >= maxh) {
11             beginvis[i] = 1;
12             maxh = hs[i];
13         }
14     }
15     int leftbeginvis = 0; // running sum over the current prefix of m_i
16     vi leftvis(n); // l_i
17     vi visstack; // L_i
18     for (int i=0; i<n; ++i) {
19         leftvis[i] = leftbeginvis + visstack.size();
20         if (beginvis[i]) ++leftbeginvis;
21         while (!visstack.empty() && visstack.back() < hs[i])
22             visstack.pop_back();
23         if (!beginvis[i])
24             visstack.push_back(hs[i]);
25     }
26     return leftvis;
27 }
28
29 int main() {
30     int n; cin >> n;
31     vi hs(n);
32     for (int& hi : hs) cin >> hi;
33     vi leftvis = solveleft(hs);
34     reverse(hs.begin(), hs.end());
35     vi rightvis = solveleft(hs);
36     reverse(rightvis.begin(), rightvis.end());
37     int maxvis = -1;
38     for (int i=0; i<n; ++i)
39         maxvis = max(maxvis, leftvis[i] + rightvis[i] + 1);
40     cout << maxvis << '\n';
41 }
```



Sliding mouse

Task Idea	Daniel Rutschmann
Task Preparation	Jan Schär
Description English	Jan Schär
Description German	Jan Schär
Description French	Florian Gatignon
Solution	Jan Schär

Subtask 1: 25 Points ($n, m \leq 100$ and there are no ice tiles)

In the first test group, we just need to find the length of the shortest path in a graph where all edges have length 1. This can be done with a simple BFS.

This solution runs in $O(nm)$.

```
1 #include <bits/stdc++.h>
2 using namespace std;
3 struct state {
4     int x, y, distance;
5 };
6 struct point {
7     int x, y;
8 };
9 vector<point> offset = {{-1, 0}, {1, 0}, {0, -1}, {0, 1}};
10 int main () {
11     int n, m;
12     cin >> n >> m;
13     vector<string> field(m);
14     for (string &s : field) cin >> s;
15
16     vector<vector<bool>> > visited(m, vector<bool>(n, false));
17     queue<state> q;
18     for (int y = 0; y < m; ++y) {
19         for (int x = 0; x < n; ++x) {
20             if (field[y][x] == 'd') q.push({x, y, 0});
21         }
22     }
23
24     while (!q.empty()) {
25         state now = q.front();
26         q.pop();
27         if (visited[now.y][now.x]) continue;
28         visited[now.y][now.x] = true;
29         char type = field[now.y][now.x];
30         if (type == '#') continue;
31         if (type == 'g') {
32             cout << now.distance << "\n";
33             return 0;
34         }
35         for (point off : offset) {
36             q.push({now.x + off.x, now.y + off.y, now.distance + 1});
37         }
38     }
39     cout << "IMPOSSIBLE\n";
40 }
```

Subtask 2: 75 Points ($n, m \leq 1000$)

We can compute for each node and direction, where mouse Daniel stops when going from that node in that direction. In this interpretation, edges can be longer than one (when Daniel slides over the ice), which means that we need to use Dijkstra's algorithm.



Computing where mouse Daniel stops can be done by a linear scan, which takes $O(n + m)$ and leads to an overall running time of $O(nm(n + m))$. This solution achieves 75 points (or if the inner loop is very efficient, 100 points).

The solution can be improved by precomputing the ice slide stops. This can be done by a simple linear scan in all four directions, which takes $O(nm)$. The overall solution then has running time $O(nm \log(nm))$, which achieves 100 points.

Shown here is the solution without precomputation:

```
1 #include <bits/stdc++.h>
2 using namespace std;
3 struct state {
4     int x, y, distance;
5 };
6 bool operator<(const state &a, const state &b) {
7     return a.distance > b.distance;
8 }
9 struct point {
10    int x, y;
11 };
12 vector<point> offset = {{-1, 0}, {1, 0}, {0, -1}, {0, 1}};
13 int main () {
14    ios_base::sync_with_stdio(false);
15    int n, m;
16    cin >> n >> m;
17    vector<string> field(m);
18    for (string &s : field) cin >> s;
19
20    vector<vector<bool>> > visited(m, vector<bool>(n, false));
21    priority_queue<state> pq;
22    for (int y = 0; y < m; ++y) {
23        for (int x = 0; x < n; ++x) {
24            if (field[y][x] == 'd') pq.push({x, y, 0});
25        }
26    }
27
28    while (!pq.empty()) {
29        auto [x, y, distance] = pq.top();
30        pq.pop();
31        if (visited[y][x]) continue;
32        visited[y][x] = true;
33        if (field[y][x] == 'g') {
34            cout << distance << "\n";
35            return 0;
36        }
37        for (point off : offset) {
38            if (field[y + off.y][x + off.x] != '#') {
39                int d = 1;
40                while (field[y + d * off.y][x + d * off.x] == '+' &&
41                    field[y + (d + 1) * off.y][x + (d + 1) * off.x] != '#') {
42                    d++;
43                }
44                pq.push({
45                    x + d * off.x, y + d * off.y,
46                    distance + d});
47            }
48        }
49    }
50    cout << "IMPOSSIBLE\n";
51 }
```

Subtask 3: 100 Points ($n, m \leq 2000$)

The intended solution for 100 points is to use BFS on a state graph. This means that each tile corresponds to 4 nodes – one for each direction. Both the BFS queue and the `visited` vector need to be extended with this extra dimension. For a dirt (.) tile, all 4 nodes have a directed edge to each non-wall neighbor. But for an ice (+) tile, each of the 4 nodes only has a single outgoing edge,



corresponding to its direction (unless there is a wall in that direction, in that case it connects to all non-wall neighbors). It is not necessary to explicitly construct the state graph, instead the edges can be easily computed from the input.

The running time is $O(nm)$.

A common mistake was to either not add the direction dimension to the visited vector, or to forget to check if a node was already visited.

```
1 #include <bits/stdc++.h>
2 using namespace std;
3 struct state {
4     int x, y, dir, distance;
5 };
6 struct point {
7     int x, y;
8 }
9 vector<point> offset = {{-1, 0}, {1, 0}, {0, -1}, {0, 1}};
10 int main () {
11     ios_base::sync_with_stdio(false);
12     int n, m;
13     cin >> n >> m;
14     vector<string> field(m);
15     for (string &s : field) cin >> s;
16
17     vector<vector<vector<bool>>> visited
18         (m, vector<vector<bool>>(n, vector<bool>(4, false)));
19     queue<state> q;
20     for (int y = 0; y < m; ++y) {
21         for (int x = 0; x < n; ++x) {
22             if (field[y][x] == 'd') q.push({x, y, 0, 0});
23         }
24     }
25
26     while (!q.empty()) {
27         auto [x, y, dir, distance] = q.front();
28         q.pop();
29         if (visited[y][x][dir]) continue;
30         visited[y][x][dir] = true;
31         char type = field[y][x];
32         if (type == 'g') {
33             cout << distance << "\n";
34             return 0;
35         }
36         if (type == '+' &&
37             field[y + offset[dir].y][x + offset[dir].x] != '#') {
38             q.push({
39                 x + offset[dir].x, y + offset[dir].y,
40                 dir, distance + 1});
41         } else {
42             for (int ndir = 0; ndir < 4; ++ndir) {
43                 if (field[y + offset[ndir].y][x + offset[ndir].x] != '#') {
44                     q.push({
45                         x + offset[ndir].x, y + offset[ndir].y,
46                         ndir, distance + 1});
47                 }
48             }
49         }
50     }
51     cout << "IMPOSSIBLE\n";
52 }
```



Ladder Balcony

Task Idea	Bibin Muttappillil
Task Preparation	Daniel Rutschmann
Description English	Daniel Rutschmann
Description German	Tobias Feigenwinter
Description French	Florian Gatignon
Solution	Daniel Rutschmann

40 Points $O(NWH^4)$

This task can be solved with dynamic programming: Let $DP[x][n][y]$ be the maximum amount of sunlight of any ladder balcony that is contained in the first x columns, uses exactly n ladders and has a ladder in the y -th position of the x -th column.

To compute $DP[x+1][n][y]$ from $DP[x]$, we iterate over all ranges $[l, r]$ with $0 \leq l \leq y \leq r < n$ and consider placing ladders at positions $l, l+1, \dots, r$ in the $(x+1)$ -st column. When placing ladders this way, we should pick the best option for the first x columns, which is $\max_{l \leq y' \leq r} DP[x][n][y']$, add the sunlight for the $(x+1)$ -st column, i.e. $a_{l,x} + \dots + a_{r,x}$ and consider this as an option for $DP[x+1][n+(r-l+1)][y]$.

Doing this naively takes $O(NWH^4)$ time, as we loop over x, n and y , consider $O(H^2)$ ranges and process each such range in $O(H)$ time.

```
1 #include <bits/stdc++.h>
2 using namespace std;
3
4 template<typename T>
5 void xmax(T&a, T const&b){
6     a = max(a, b);
7 }
8 constexpr int inf = 1e9;
9
10 signed main(){
11     int H, W, n;
12     cin >> H >> W >> n;
13
14     vector<vector<int>> > v(H, vector<int>(W));
15     for(auto &e:v){
16         for(auto &f:e){
17             cin >> f;
18         }
19     }
20
21     // Only store last 2 rows of table, re-use memory
22     vector<vector<int>> dp(n+1, vector<int>(H, -inf));
23     auto dp2 = dp;
24     auto swap_dp = [&]() {
25         dp.swap(dp2);
26         for(auto &f:dp2) {
27             fill(f.begin(), f.end(), -inf);
28         }
29     };
30
31     int ans = -inf;
32     for(int y=0; y<H; ++y){
33         dp[0][y] = 0;
34     }
35     for(int x=0; x<W; ++x){
36         for(int i=0; i<n; ++i){
37             for(int y=0; y<H; ++y){
38                 // iterate over vertical ranges [l, r]
```



```
39         for(int l=0; l<=y; ++l){
40             for(int r=y; r<H; ++r){
41                 const int d = r-l+1;
42                 if(i+d > n) continue;
43                 // find best option and sum
44                 int best = -inf;
45                 int sum = 0;
46                 for(int y2=l; y2<=r; ++y2){
47                     xmax(best, dp[i][y2]);
48                     sum += v[y2][x];
49                 }
50                 best += sum;
51                 xmax(dp2[i+d][y], best);
52             }
53         }
54     }
55 }
56 // update answer
57 swap_dp();
58 for(int y=0; y<H; ++y){
59     xmax(ans, dp[n][y]);
60 }
61 }
62 cout << ans << "\n";
63 }
```

80 Points $O(NWH^3)$

We can speed the previous solution up by a factor of H if we loop over ranges $[l, r]$ and then consider all y for which this range is valid. This way, we can separate the loops for y and $y2$, so we don't don't recompute the sum and maximum over $[l, r]$ for every y .

```
1     ...
2     int ans = -inf;
3     for(int y=0; y<H; ++y){
4         dp[0][y] = 0;
5     }
6     for(int x=0; x<W; ++x){
7         for(int i=0; i<n; ++i){
8             // iterate over vertical ranges [l, r]
9             for(int l=0; l<H; ++l){
10                for(int r=l; r<H; ++r){
11                    const int d = r-l+1;
12                    if(i+d > n) continue;
13                    // find best option and sum
14                    int best = -inf;
15                    int sum = 0;
16                    for(int y=l; y<=r; ++y){
17                        xmax(best, dp[i][y]);
18                        sum += v[y][x];
19                    }
20                    best += sum;
21                    // consider the range [l, r]
22                    // for all valid y
23                    for(int y=l; y<=r; ++y){
24                        xmax(dp2[i+d][y], best);
25                    }
26                }
27            }
28        }
29        swap_dp();
30        for(int y=0; y<H; ++y){
31            xmax(ans, dp[n][y]);
32        }
33    }
34    ...
```



100 Points $O(NWH^2)$

To get full points, we need shave off another factor of H . When going from $DP[x]$ to $DP[x + 1]$, let's try to simultaneously consider all ranges $[l, r]$: Fix x and let

$$\text{Best}[n][l][r] = \sum_{y=l}^r a_{y,x} + \max_{l \leq y \leq r} dp[x][n - (r - l + 1)][y]$$

We can compute Best with dp as follows:

$$\begin{aligned} \text{Best}[n][y][y] &= dp[x][n - 1][y] + a_{y,x} \\ \text{Best}[n][l][r] &= \max(\text{Best}[n - 1][l + 1][r] + a_{x,l}, \text{Best}[n - 1][l][r - 1] + a_{x,r}) \end{aligned}$$

This avoids the first loop over y in the previous solution. We can similarly avoid the second loop: Fix n and x and let

$$\text{Up}[l][r] = \max_{l' \leq l, r' \geq r} \text{Best}[n][l'][r']$$

This will then allows us to compute $DP[x + 1]$ via

$$DP[x + 1][n][y] = \text{Up}[y][y]$$

We can compute Up (with some care to avoid out of bounds errors) via

$$\text{Up}[l][r] = \max(\text{Up}[l - 1][r], \text{Up}[l][r + 1], \text{Best}[n][l][r])$$

The total running time is $O(NWH^2)$.

Note: When implementing this, you should avoid excessive memory allocations and re-use the same vectors for every x . This makes the solution up to $5\times$ faster. (But the time limit was lenient and the model solution only uses $1/3$ of the time limit.)

```
1 #include <bits/stdc++.h>
2 using namespace std;
3
4 template<typename T>
5 void xmax(T&a, T const&b){
6     a = max(a, b);
7 }
8 constexpr int inf = 1e9;
9
10 signed main(){
11     int H, W, n;
12     cin >> H >> W >> n;
13
14     vector<vector<int>> > v(H, vector<int>(W));
15     for(auto &e:v){
16         for(auto &f:e){
17             cin >> f;
18         }
19     }
20     // re-use memory
21     vector<vector<int>> > dp(n+1, vector<int>(H, -inf));
22     vector<vector<vector<int>>> > best(n+1, vector<vector<int>> > (H, vector<int>(H, -inf)));
23     vector<vector<int>> > up(H, vector<int>(H, -inf));
24
25     int ans = -inf;
26     for(int y=0; y<H; ++y){
27         dp[0][y] = 0;
28     }
29     for(int x=0; x<W; ++x){
30         // consider all ranges [l, r] in parallel
31         for(int i=1; i<=n; ++i){
32             for(int y=0; y<H; ++y){
33                 best[i][y][y] = dp[i-1][y] + v[y][x];
34             }
35             for(int l=0; l<H; ++l){
```



```
36         for(int r=l+1; r<H; ++r){
37             best[i][l][r] = max(best[i-1][l][r-1]+v[r][x], best[i-1][l+1][r]+v[l][x]);
38         }
39     }
40 }
41 for(int i=0; i<=n; ++i){
42     for(int l=0; l<H; ++l){
43         for(int r=H-1; r>=1; --r){
44             up[l][r] = best[i][l][r];
45             if(l>0) xmax(up[l][r], up[l-1][r]);
46             if(r+1<H) xmax(up[l][r], up[l][r+1]);
47         }
48     }
49     for(int y=0; y<H; ++y){
50         dp[i][y] = up[y][y];
51     }
52 }
53 for(int y=0; y<H; ++y){
54     xmax(ans, dp[n][y]);
55 }
56 }
57 cout << ans << "\n";
58 }
```