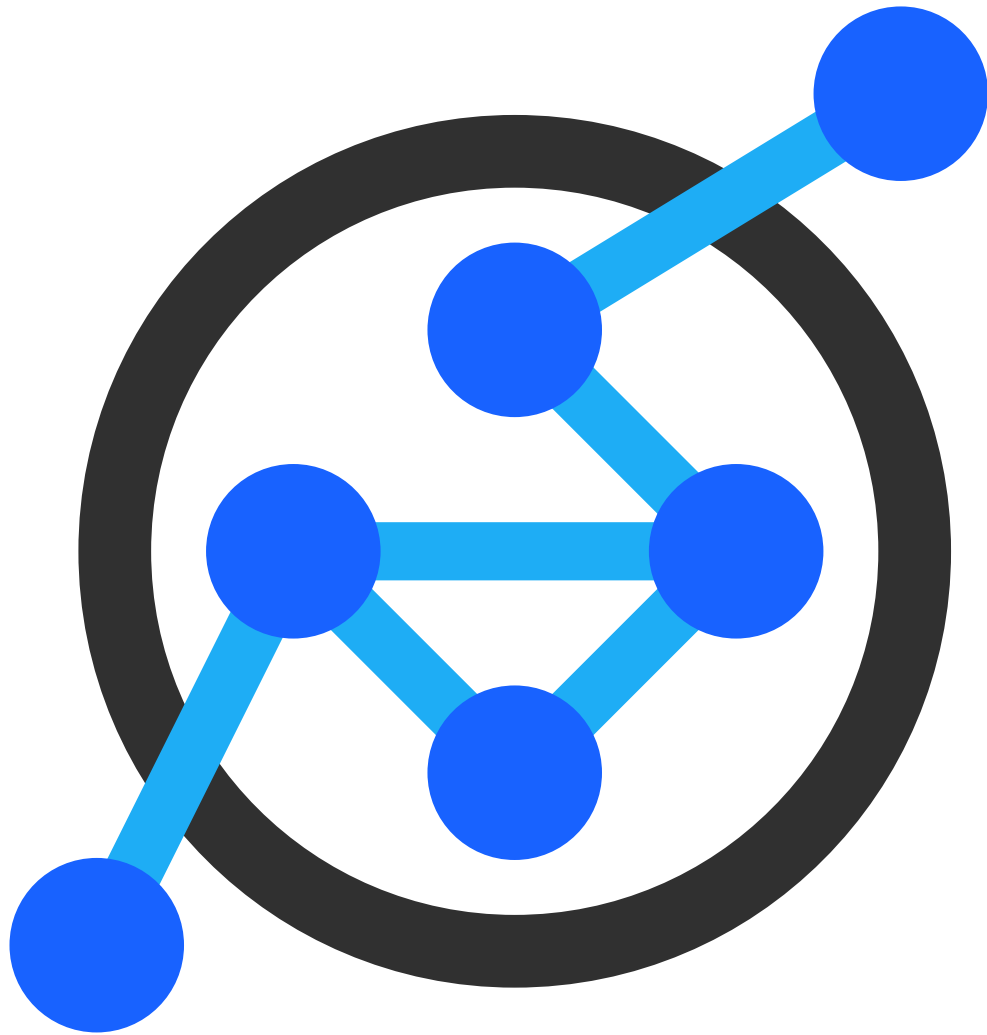


SOI 2P 2024
Solution Booklet



Swiss Olympiad in Informatics

9–10 March 2024

Hidden Keys

Task Idea	Anna Khanova
Task Preparation	Anna Khanova
Description English	Anna Khanova
Description German	Bibin Muttappillil
Description French	Benjamin Faltin
Solution	Anna Khanova

We are given a grid of $n \times m$ numbers where the number in line i and column j is denoted by $c_{i,j}$

Our task is to find i, j, k, l such that $i \neq k$ and $j \neq l$ and $c_{i,j} \neq c_{k,l}$

Subtask 1: $n \cdot m \leq 100$

In this case we can, for each $c_{i,j}$, check all possible k, l to try and find $i \neq k$ and $j \neq l$ and $c_{i,j} \neq c_{k,l}$. If such k, l is found for any i, j we are done and we can output i, j, k, l otherwise no i, j, k, l meet the criteria and we should output -1 . We would have a time complexity of $O(n^2m^2)$ which is sufficient.

```

1 #include <algorithm> // score: 32
2 #include <vector>
3 #include <iostream>
4 using namespace std;
5
6 const int MAXN = 100100;
7
8 int n, m;
9 vector<int> a[MAXN];
10
11 void print_ans(int x1, int y1, int x2, int y2)
12 {
13     cout << x1 << ' ' << y1 << ' ' << x2 << ' ' << y2 << endl;
14     exit(0);
15 }
16
17 void try_at(int x, int y)
18 {
19     for (int i = 0; i < n; ++i)
20         for (int j = 0; j < m; ++j)
21             if (i != x && j != y && a[i][j] != a[x][y])
22                 print_ans(x, y, i, j);
23 }
24
25 int main()
26 {
27     cin >> n >> m;
28     for (int i = 0; i < n; ++i)
29     {
30         a[i] = vector<int>(m);
31         for (int j = 0; j < m; ++j)
32             cin >> a[i][j];
33     }
34     for (int i = 0; i < n; ++i)
35         for (int j = 0; j < m; ++j)
36             try_at(i, j);
37     cout << -1 << endl;
38     return 0;
39 }

```



Subtask 2: $n, m \leq 100$

Here, $O(n^2m^2)$ is too slow however, we can realize that if $c_{i,j}$ and $c_{k,l}$ is a solution, then we know that one of the 8 following cases is also a valid solution:

$c_{i,0}$ and $c_{k,l}$, $c_{0,j}$ and $c_{k,l}$, $c_{i,j}$ and $c_{k,0}$, $c_{i,j}$ and $c_{0,l}$, $c_{i,0}$ and $c_{0,j}$, $c_{i,0}$ and $c_{0,l}$, $c_{k,0}$ and $c_{0,j}$, $c_{k,0}$ and $c_{0,l}$

This is because if $c_{i,j} = c_{i,0}$ then $c_{k,l}$ and $c_{i,0}$ is a solution. The same logic applies to $c_{0,j}$, $c_{k,0}$, $c_{0,l}$. If none of these equalities hold, then either all $c_{i,0}$, $c_{0,j}$, $c_{k,0}$, $c_{0,l}$ are the same and at least one of the 8 inequalities hold since $c_{i,j} \neq c_{k,l}$ or there is a difference between a number on the line 0 and another on the column 0.

Therefore we only need to check if there is a $c_{i,j}$ which doesn't match with a number on the line 0 or the column 0. We have a complexity of $O((n+m)nm)$

```
1 #include <algorithm> // score: 51
2 #include <vector>
3 #include <iostream>
4 using namespace std;
5
6 const int MAXN = 100100;
7
8 int n, m;
9 vector<int> a[MAXN];
10
11 void print_ans(int x1, int y1, int x2, int y2)
12 {
13     cout << x1 << ' ' << y1 << ' ' << x2 << ' ' << y2 << endl;
14     exit(0);
15 }
16
17 void try_at(int x, int y)
18 {
19     for (int i = 0; i < n; ++i)
20         for (int j = 0; j < m; ++j)
21             if (i != x && j != y && a[i][j] != a[x][y])
22                 print_ans(x, y, i, j);
23 }
24
25 int main()
26 {
27     cin >> n >> m;
28     for (int i = 0; i < n; ++i)
29     {
30         a[i] = vector<int>(m);
31         for (int j = 0; j < m; ++j)
32             cin >> a[i][j];
33     }
34     try_at(0, 0);
35     for (int i = 1; i < n; ++i)
36         try_at(i, 0);
37     for (int j = 1; j < m; ++j)
38         try_at(0, j);
39     cout << -1 << endl;
40     return 0;
41 }
```

Subtask 3: $n \cdot m \leq 100\,000$

If there is a mismatch between $c_{i,j}$ and $c_{k,l}$ and $i, j, k, l > 0$, then at least one of them is different than $c_{0,0}$. If there is a mismatch between $c_{i,j}$ and $c_{k,l}$ and i, j, k or l is 0, then either one of them is different than $c_{0,0}$, $c_{1,0}$ or $c_{0,1}$ and the coordinates don't have any common part, or, if they are equal, we have $c_{0,0} \neq c_{0,1}$, $c_{0,0} \neq c_{1,0}$ or $c_{0,1} \neq c_{1,0}$ which has the consequence of having at least one difference either between $c_{0,1}$ and $c_{1,0}$ or between one of the three points and one other point



somewhere in the grid.

To solve the problem we only need to check all $c_{i,j}$ against $c_{0,0}$, $c_{0,1}$ and $c_{1,0}$

```
1 #include <algorithm> // score: 100
2 #include <iostream>
3 #include <vector>
4 using namespace std;
5
6 const int MAXN = 500500;
7
8 int n, m;
9 vector<int> a[MAXN];
10
11 void print_ans(int x1, int y1, int x2, int y2)
12 {
13     cout << x1 << ' ' << y1 << ' ' << x2 << ' ' << y2 << endl;
14     exit(0);
15 }
16
17 void try_at(int x, int y)
18 {
19     for (int i = 0; i < n; ++i)
20         for (int j = 0; j < m; ++j)
21             if (i != x && j != y && a[i][j] != a[x][y])
22                 print_ans(x, y, i, j);
23 }
24
25 int main()
26 {
27     cin >> n >> m;
28     for (int i = 0; i < n; ++i)
29     {
30         a[i] = vector<int>(m);
31         for (int j = 0; j < m; ++j)
32             cin >> a[i][j];
33     }
34     try_at(0, 0);
35     if (m != 1)
36         try_at(0, 1);
37     if (n != 1)
38         try_at(1, 0);
39     cout << -1 << endl;
40     return 0;
41 }
```



Moving Faster

Task Idea	Luca Versari
Task Preparation	Théo von Düring
Description English	Théo von Düring
Description German	Bibin Muttappillil
Description French	Théo von Düring
Solution	Théo von Düring

We are given a value k , a weighted connected graph where each edge has a weight t and a decreasing value d , and two nodes s and e . We want to compute the minimum time needed to go back and forth between s and e a total of k times.

One first observation is that the direction of the path doesn't matter, going from s to e or from e to s on the same set of edges decreases each edge by the same amount and the sum of all edge distances is also the same. Therefore the problem is the same as going k times from s to e and decreasing the edge weights each time. We will use this simplified problem in the solutions.

Subtask 1: A tree

If the graph is a tree we know that the path between s and e is unique and therefore it will always be used to travel between the two nodes. To solve this task we can just run a DFS between the two nodes and save the path taken. Then compute the time needed to use k times each edge.

```
1 #include <bits/stdc++.h> // score: 17
2
3 using namespace std;
4
5 #define int int64_t
6
7 vector<vector<array<int,2>>> g;
8 stack<array<int,2>> st;
9
10 bool dfs(int u, int target, int i, int p) {
11     st.push({u,i});
12
13     if (u == target)
14         return true;
15
16     for (auto [v,j] : g[u]) {
17         if (v == p)
18             continue;
19         if (dfs(v, target, j, u))
20             return true;
21     }
22     st.pop();
23     return false;
24 }
25
26 signed main() {
27     ios_base::sync_with_stdio(0);
28     cin.tie(0);
29
30     int n,m,k,s,e;
31     cin >> n >> m >> k >> s >> e;
32
33     g.assign(n, {});
34     vector<array<int,2>> edge(m);
35
```



```
36 for (int i = 0; i < m; i++) {
37     int a,b,t,d;
38
39     cin >> a >> b >> t >> d;
40     edge[i] = {t,d};
41     g[a].push_back({b,i});
42     g[b].push_back({a,i});
43 }
44
45 dfs(s, e, -1, s);
46
47 int ans = 0;
48 while (!st.empty()) {
49     auto [u, i] = st.top();
50     auto [t, d] = edge[i];
51     st.pop();
52     if (i == -1)
53         continue;
54     ans += k*t - d*(k*(k-1)/2);
55 }
56 cout << ans;
57 }
```

Subtask 2: $k = 1$

Since $k = 1$, we only need to travel once from s to e and since we never traverse an edge twice, d doesn't have any effect. Thus, to solve this case you need to go from s to e using the shortest path without taking into account d . This problem can be solved using a standard dijkstra implementation on the weighted graph

```
1 #include <bits/stdc++.h> // score: 15
2
3 using namespace std;
4
5 #define int int64_t
6
7 signed main() {
8     ios_base::sync_with_stdio(0);
9     cin.tie(0);
10
11     int n,m,k,s,e;
12     cin >> n >> m >> k >> s >> e;
13
14     vector<vector<array<int,2>>> g(n);
15     for (int i = 0; i < m; i++) {
16         int a,b,t,d;
17         cin >> a >> b >> t >> d;
18
19         g[a].push_back({b, t});
20         g[b].push_back({a, t});
21     }
22
23     priority_queue<array<int,2>> pq;
24     pq.push({0, s});
25     vector<int> dist(n, -1);
26     while (!pq.empty()) {
27         auto [d, a] = pq.top();
28         pq.pop();
29
30         if (dist[a] != -1)
31             continue;
32         dist[a] = -d;
33
34         for (auto [b, t] : g[a]) {
35             if (dist[b] != -1)
36                 continue;
```



```
37     pq.push({d-t, b});
38     }
39 }
40 cout << dist[e] << "\n";
41 }
```

General problem

For the general case it can also be shown that using the same path k times is optimal. The idea of the proof is that if you suppose that the optimal solution uses $n > 1$ paths which we will call p_1, p_2, \dots, p_n each used c_1, c_2, \dots, c_n times respectively, then there exist i, j , such that using p_i $c_i - 1$ times and p_j $c_j + 1$ times is faster which contradicts the hypothesis and shows that we only use one path. To solve the problem we will, for each edge, compute the time it will take to use it k times and use dijkstra on the graph with these new precomputed times as weights. This gives us the shortest weighted path and the time needed to use it k times

```
1 #include <bits/stdc++.h> // score: 100
2
3
4 using namespace std;
5
6 #define int int64_t
7
8 signed main() {
9     ios_base::sync_with_stdio(0);
10    cin.tie(0);
11
12    int n,m,k,s,e;
13    cin >> n >> m >> k >> s >> e;
14
15    vector<vector<array<int,2>>> g(n);
16
17    for (int i = 0; i < m; i++) {
18        int a,b,t,d;
19
20        cin >> a >> b >> t >> d;
21
22        assert(t > 0);
23        assert(d > 0);
24        assert(t*k > 0);
25        assert(d*(k*(k-1)/2) >= 0);
26        assert(t - d*k > 0);
27
28        t = t*k - d*(k*(k-1)/2);
29        assert(t > 0);
30
31        g[a].push_back({b,t});
32        g[b].push_back({a,t});
33    }
34
35    vector<int> dist(n, -1);
36
37    priority_queue<array<int,2>> pq;
38    pq.push({0, s});
39    while(!pq.empty()) {
40        auto [d1, u] = pq.top();
41        pq.pop();
42
43        if (dist[u] != -1)
44            continue;
45
46        dist[u] = -d1;
47        assert(dist[u] >= 0);
48
49        for (auto [v,d2] : g[u]) {
```



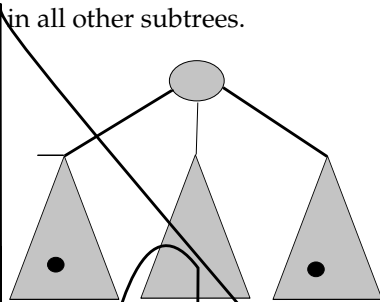
```
50         if (dist[v] != -1)
51             continue;
52         pq.push({d1-d2, v});
53     }
54 }
55
56 assert(dist[e] > 0);
57
58 cout << dist[e] << "\n";
59 }
```


Ruined Chocolates

Task Idea	Andrei Feodorov, Johannes Kapfhammer
Task Preparation	Andrei Feodorov
Description English	Andrei Feodorov
Description German	Charlotte Knierim
Description French	Théo von Düring
Solution	Johannes Kapfhammer

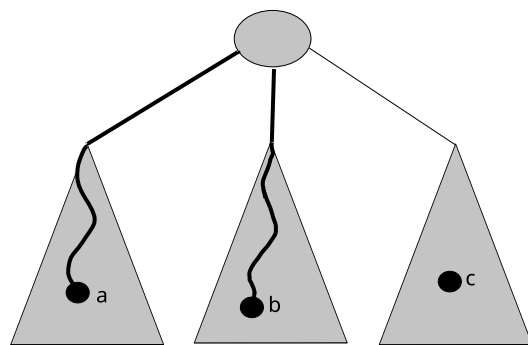
Subtask 1: Observation: All queries go through the root

Root the tree in the vertex that has the maximal value. Then: if a query does not go through the root, the answer is just the value of the root. Else we just have to consider answering the question for two subtrees where the query starts at the root of the subtree. The final answer is the maximum answer of the two subqueries inside the two subtrees, and of the maximum value in all other subtrees.



Subtask 2: Observation: Two paths and an extra vertex

Let's consider the vertices with the k largest values. Choose the smallest k such that there is no path that contains them all.

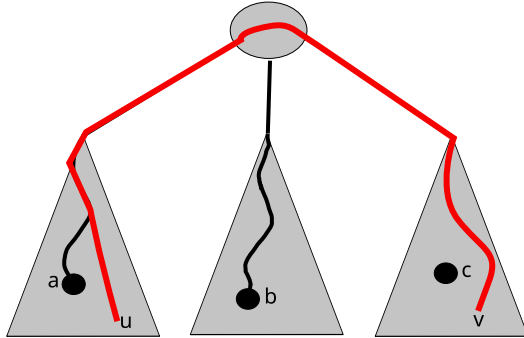


Since the largest value is the root, we can describe this set by giving the two endpoints of the path containing the first $k - 1$ elements, call them a and b , and the extra vertex c .

Note that:

- a and b lie in different subtrees because the root must be part of the path (recall we are considering the largest k values). There is one exception if the root only has a single child.
- c can be in the same subtree as a or b or in a different one.
- Any possible query can never cover all vertices $\text{path}(a, b) \cup \{c\}$ – because queries are paths and by definition there is no path through those vertices.

- Therefore, the answer to any query is at least $\text{value}(c)$.
- As a corollary, the values of vertices outside of $\text{path}(a, b) \cup \{c\}$ don't matter.



So to answer a query (v, w) , we only have check:

- If v and a are in the same subtree, let $q = \text{LCA}(a, v)$ be the vertex where they come together. The answer of the query is the maximum of all values on $\text{path}(a, q) \setminus \{q\}$. Similarly with v and b , w and a , and w and b .
- If vertex c does not lie on the path $v \rightarrow w$, then also consider the value of c a possible answer.
- By the observations above, this is enough – we will have considered at least one value because the set $\text{path}(a, b) \cup \{c\}$ is a strict superset of the query. And all other values are smaller.

Subtask 3: Implementation

We need to do the following precomputation which can be implemented in $O(n)$:

1. Root the tree at the vertex with the maximum value.
2. Iterate over the vertices in decreasing value, and find the two endpoints of the path a and b , and our extra vertex c .
3. For every other vertex, precompute the nearest ancestor that lies on the path $a \rightarrow b$. For each vertex on the path compute the maximum of all vertices strictly below.

And then for each query (v, w) we can answer in $O(1)$:

- Find the two points $p = \text{LCA}(a, v)$ (or w) and $q = \text{LCA}(b, w)$ (or v) where we meet up with the path. Then the maximum value below p and q are possible answers. All of this is already precomputed.
- Check whether c lies on the path $v \rightarrow w$. If not, then the value of c is also a possible answer.

```

1 #include <bits/stdc++.h>
2 using namespace std;
3
4 int main() {
5   cin.tie(nullptr) -> sync_with_stdio(false);
6
7   int n, q; cin >> n >> q;
8   vector<vector<int>>> g(n);
9   for (int i=1; i<n; ++i) {
10    int a, b; cin >> a >> b;
11    g[a].push_back(b);
12    g[b].push_back(a);
13  }

```



```
14 vector<int> value; copy_n(istream_iterator<int>(cin), n, back_inserter(value));
15 vector<int> order(n); iota(order.begin(), order.end(), 0);
16 sort(order.begin(), order.end(), [&](int a, int b) { return value[a]>value[b]; });
17 int root = order[0];
18
19 vector<int> pre(n), post(n);
20 vector<int> subtreerootof(n), parent(n);
21 {
22     int time = 0;
23     auto dfs = [&](auto&& self, int v, int p, int subroot) -> void {
24         pre[v] = time++;
25         subtreerootof[v] = subroot;
26         parent[v] = p;
27         for (auto w : g[v])
28             if (w != p)
29                 self(self, w, v, subroot);
30         post[v] = time;
31     };
32     for (auto w : g[root])
33         dfs(dfs, w, root, w);
34 }
35 auto is_ancestor_of = [&](int a, int b) { return pre[a] <= pre[b] && post[b] <= post[a]; };
36 array<int, 2> subtree_root{-1, -1};
37 array<int, 2> subtree_pathend{};
38 array<int, 2> subtree_maxvalue{};
39 int extra_vertex = -1;
40
41 for (auto v : order) {
42     if (v == root) continue;
43     int i=0;
44     for (; i<2; ++i) {
45         if (subtree_root[i] == -1) {
46             subtree_root[i] = subtreerootof[v];
47             subtree_pathend[i] = v;
48             subtree_maxvalue[i] = value[v];
49             break;
50         } else if (subtree_root[i] == subtreerootof[v]) {
51             if (is_ancestor_of(subtree_pathend[i], v))
52                 subtree_pathend[i] = v;
53             else if (!is_ancestor_of(v, subtree_pathend[i]))
54                 extra_vertex = v;
55             break;
56         }
57     }
58     if (i == 2 && extra_vertex == -1)
59         extra_vertex = v;
60     if (extra_vertex != -1)
61         break;
62 }
63 vector<int> subtree_answer(n, -1);
64 for (int i=0; i<2 && subtree_root[i] != -1; ++i) {
65     int pos = subtree_pathend[i];
66     int submax = 0;
67     auto fill = [&](auto&& self, int v, int p) -> void {
68         subtree_answer[v] = submax;
69         for (int w : g[v])
70             if (w != p && subtree_answer[w] == -1)
71                 self(self, w, v);
72     };
73     while (pos != root) {
74         fill(fill, pos, parent[pos]);
75         submax = max(submax, value[pos]);
76         pos = parent[pos];
77     }
78 }
79
80 auto subtree_index = [&](int v) {
81     if (v == root) return -1;
```



```
82     int subroot = subtreerootof[v];
83     auto it = find(subtree_root.begin(), subtree_root.end(), subroot);
84     return it == subtree_root.end() ? -1 : static_cast<int>(it - subtree_root.begin());
85 };
86 while (q--) {
87     int a, b; cin >> a >> b;
88     int i = subtree_index(a), j = subtree_index(b);
89     int ans = 0;
90     if (i != -1 && j != -1 && i == j) {
91         ans = value[root];
92     } else {
93         for (int k=0; k<2; ++k) {
94             if (k == i) ans = max(ans, subtree_answer[a]);
95             else if (k == j) ans = max(ans, subtree_answer[b]);
96             else ans = max(ans, subtree_maxvalue[k]);
97         }
98         if (extra_vertex != -1 &&
99             !is_ancestor_of(extra_vertex, a) &&
100             !is_ancestor_of(extra_vertex, b))
101             ans = max(ans, value[extra_vertex]);
102     }
103     cout << ans << '\n';
104 }
105 }
```

Tram Tickets

Task Idea	Bibin Muttappillil
Task Preparation	Charlotte Knierim
Description English	Charlotte Knierim
Description German	Charlotte Knierim
Description French	Benjamin Faltin
Solution	Johannes Kapfhammer

Step 1: Without Binna (51 Points)

We sort by t_i then we do the following DP approach:

$$\begin{aligned} DP[i] &= \text{cheapest assignment for } t_0, \dots, t_{i-1} \\ &= \min(\text{buy single ticket at } t_i, \text{ buy multiride ticket at time } t_i - d) \\ &= \min(DP[i - 1] + a, DP[\text{lower_bound}(t_i - d)] + b) \end{aligned}$$

where $\text{lower_bound}(t_i - d)$ is the largest index j such that $t_j < t_i - d$. This index can be found via binary search for a $O(n \log n)$ solution, or via two pointer which runs in $O(n)$.

```

1 #include <bits/stdc++.h>
2 using namespace std;
3 #define int int64_t
4
5 signed main(){
6     int n, _, singleticket, multiride, d;
7     cin >> n >> _ >> singleticket >> multiride >> d;
8
9     vector<int> time; copy_n(istream_iterator<int>(cin), n, back_inserter(time));
10    sort(time.begin(), time.end());
11
12    vector<int> dp(n+1, 0);
13    for (int i=0, j=0; i < n; ++i) {
14        while(j+1 < n && time[j] <= time[i] - d)
15            ++j;
16        dp[i+1] = min(dp[i] + singleticket,
17                    dp[j] + multiride);
18    }
19    cout << dp[n] << '\n';
20 }
```

Step 2: With Binna (75 Points)

$$\begin{aligned} DP[i][j] &= \text{cheapest assignment for } t_0, \dots, t_{i-1} \\ &\text{with Binna having multiride ticket remaining for } j \text{ more time units} \end{aligned}$$

Let t_0, \dots, t_{n-1} be the unique times in which we have tickets. Define $\text{asum}(l, r)$ as the number of distinct times t_i with $l \leq t_i < r$ multiplied by a (the cost for single tickets).

We call tickets which are "doubled" if there exactly two tickets at the same time. Let's define $\text{bsum}(l, r)$ as number of distinct times t_i with $l \leq t_i < r$ where we have doubled tickets, again multiplied by a (so this will be the cost for doubled single tickets). Let's define bsum to be 0 if $r \leq l$. With prefix sums precomputation this can be evaluated in constant time.



$$DP[i][j] = \min \begin{cases} \min_{0 \leq k \leq d} (DP[i-d][k] + bsum(i-d+k, i) + b) & \text{Stofl: buy M at } i-b, \text{ Binna: only S} \\ \min_{0 \leq k \leq d} (DP[i-d][k] + bsum(i-d+k, i-d+j) + 2 \cdot b) & \text{Stofl: M at } i-b, \text{ Binna: M at } i-(d-j) \\ DP[i-1][j+1] + bsum(i-1, i) & \text{if } j+1 \leq d \\ DP[i-1][0] + asum(i-1, i) + bsum(i-1, i) & \text{if } j=0 \\ DP[i][0] + b & \text{if } j=d \end{cases}$$

Use Binna's M, Stofl buys doubled S
Both buy S
Binna: buy M at i

gives a DP that runs in $O(n \cdot d^2)$.

This can be improved by noticing that the first case does not depend on j so it can be cached, and the last three cases can be computed in constant time. So we only need to speed up the second case:

$$\begin{aligned} A[j][k] &= DP[i-d][k] + bsum(i-d+k, i-j) + 2 \cdot b \\ &= \underbrace{DP[i-d][k] + 2 \cdot b}_{\text{independent of } j} + \underbrace{bsum(i-d+k, i-j)}_{\text{only a prefix changes}} \end{aligned}$$

What happens when we go from $A[j+1]$ to $A[j]$?

$$A[j][k] = A[j+1][k] + \begin{cases} bsum(i-j-1, i-j) & \text{if } i-j-1 > i-d+k \\ 0 & \text{otherwise} \end{cases}$$

This means we can store array $A[j]$ as segment tree, which supports the two operations we need: increment a prefix, get the global minimum.

This runs in $O(n \cdot d \log d)$ if we only look at times i where there is a ticket. Instead of segment trees we can also use sets: We have $A[j][k] \leq A[j][k+1]$ (having Binna's ticket last longer is always more expensive) so the array $A[k]$ is non-decreasing. Let's just keep the differences in a map with entries $(k, A[j][k] - A[j][k-1])$. When we increment a prefix we increment the delta of the first entry and decrement the delta of some entry in the middle. Let's say the new element is (k_i, d_i) and the entry before is (k_{i-1}, d_{i-1}) . Then we replace merge both entries into a new one $(k_i, d_i + d_{i-1})$ or remove it if $d_i + d_{i-1} = 0$. This gets rid of the irrelevant entries and we only keep the prefix minima. The minimum is always at the beginning of the set.

And instead of segtree/sets we can also just use vectors: since the prefixes are always getting smaller, and the array is non-decreasing initially, we pop all values that are no longer getting changed from the end of the vector and keep one minimum of the values that have been removed. We also store a global offset that we want to add to all elements of the vector. Then incrementing a prefix just means popping the last value of the vector, add offset to it and update the global minimum. Then increment the offset. The overall minimum is the minimum between the first entry of the vector and the popped minimum.

This runs in $O(n \cdot d)$.

```
1 #include<bits/stdc++.h>
2 using namespace std;
3
4 #define int int64_t
5
6 const int INF = 1e9;
7
8 void xmin(int& x, int v) {
9     if (v < x)
10         x = v;
11 }
12
```



```
13 signed main() {
14     cin.tie(nullptr)->sync_with_stdio(false);
15     cin.exceptions(istream::badbit | istream::failbit | istream::eofbit);
16
17     int num_stofl_tickets, num_binna_tickets, singleticket, multiride, d;
18     cin >> num_stofl_tickets >> num_binna_tickets >> singleticket >> multiride >> d;
19
20     vector<int> time;
21     vector<bool> both;
22     {
23         array<vector<int>, 2> xs;
24         xs[0].resize(num_stofl_tickets);
25         xs[1].resize(num_binna_tickets);
26         for (auto& v : xs) {
27             for (auto& x : v) cin >> x;
28             sort(v.begin(), v.end());
29         }
30         time.push_back(-d-2);
31         set_union(xs[0].begin(), xs[0].end(), xs[1].begin(), xs[1].end(), back_inserter(time));
32         vector<int> bs;
33         set_intersection(xs[0].begin(), xs[0].end(), xs[1].begin(), xs[1].end(), back_inserter(bs));
34         for (auto t : time)
35             both.push_back(binary_search(bs.begin(), bs.end(), t));
36     }
37     const int n = time.size();
38
39     vector<vector<int>> dp(n+1, vector<int>(d+1, INF));
40     dp[0][0] = 0;
41     vector<int> dp_i_d(d+1);
42     for (int i=1, i_d=0; i<n; ++i) {
43         int time_i_d = time[i] - d;
44         while (time[i_d+1] <= time_i_d)
45             ++i_d;
46         // dp_i_d is the interpolated DP state at time_i_d.
47         int padding = time_i_d - time[i_d];
48         assert(padding >= 0);
49         if (padding == 0) {
50             dp_i_d = dp[i_d];
51         } else {
52             fill(dp_i_d.begin(), dp_i_d.end(), INF);
53             for (int k=d-padding; k>=0; --k)
54                 dp_i_d[k] = min(dp_i_d[k+1], dp[i_d][k+padding]);
55             dp_i_d[0] = min(dp_i_d[0], dp[i_d][0]);
56         }
57     }
58
59     int inbetween = 0;
60     for (int i_k=i_d + (time[i_d] < time[i]-d); i_k <= i; ++i_k)
61         inbetween += both[i_k];
62
63     // buy Stofl a double ticket at time t
64     vector<int> aj(d+1, INF);
65     int aj_offset=0, aj_popmin=INF;
66     for (int k=0, i_k=i_d + (time[i_d] < time[i]-d); k<=d; ++k) {
67         assert(time[i_k] >= time[i] - d + k);
68
69         // cover [i-d,i] with a multiride but don't buy a second one
70         xmin(dp[i][0], dp_i_d[k] + singleticket*inbetween + multiride);
71
72         // prepare array a[j] for the next case:
73         // cover [i-d,i] with a multiride, buy new multiride at time i-d+j
74         xmin(aj[k], dp_i_d[k] + singleticket*inbetween + 2*multiride);
75
76         // Needed to make a[j] decreasing
77         if (k) xmin(aj[k], aj[k-1]);
78
79         if (i_k+1 < n && time[i_k] < time[i] && time[i_k] == time[i]-d+k)
80             inbetween -= both[i_k++];

```



```
81     }
82     for (int j=d, i_k=i; j>=0; --j) {
83         assert(time[i_k] <= time[i]-d+j);
84         if (time[i_k] == time[i]-d+j) {
85             if (both[i_k])
86                 aj_offset -= singleticket;
87             --i_k;
88         }
89         assert((int)aj.size() == j+1);
90         xmin(aj_popmin, aj.back() + aj_offset);
91         aj.pop_back();
92         xmin(dp[i][j], min(aj_popmin, aj.empty() ? INF : aj.back()+aj_offset));
93     }
94
95     // use up existing multiride and buy singles
96     int td = time[i] - time[i-1];
97     for (int j=1; j + td <= d; ++j)
98         xmin(dp[i][j], dp[i-1][j + td] + (both[i] ? singleticket : 0));
99
100    // just buy singletickets
101    xmin(dp[i][0], dp[i-1][0] + (both[i] ? 2 : 1)*singleticket);
102
103    // buy a new multiride
104    xmin(dp[i][d], dp[i-1][0] + multiride + (both[i] ? singleticket : 0));
105
106    // make monotone (not needed)
107    for (int j=d-1; j>=0; --j)
108        xmin(dp[i][j], dp[i][j+1]);
109 }
110 cout << dp[n-1][0] << '\n';
111 }
```

Step 3: Happy subtask merging

Submit both solutions to step 1 and step 2 and subtask merging will do its job giving you well-deserved 100 points.

Baguette Magique

Task Idea	Anna Khanova, Johannes Kapfhammer
Task Preparation	Anna Khanova
Description English	Johannes Kapfhammer
Description German	Bibin Muttappillil, Johannes Kapfhammer
Description French	Benjamin Faltin
Solution	Jan Schär

Subtask 1: Cut single centimeters

In this subtask, we can only cut out parts of length 1 cm from the baguette. We need to pick some flavor which should remain at the end, and cut out all centimeters which have a different flavor. We can just try this for all possible flavors and pick the best option.

```
1 #include<bits/stdc++.h> // score: 12
2 using namespace std;
3
4 signed main () {
5     int n, k;
6     string baguette;
7     cin >> n >> k >> baguette;
8
9     int best_cuts = n - 1;
10    for (char keep = 'a'; keep <= 'z'; keep++) {
11        int count = 0;
12        for (char c : baguette) if (c != keep) count++;
13        best_cuts = min(best_cuts, count);
14    }
15    cout << best_cuts << "\n";
16 }
```

This solution runs in $O(n)$ time (when treating the size of the alphabet as a constant).

Subtask 4: 10 meter baguette

The baguette can be at most 1'000 cm in length.

Once we have decided which flavor we want to keep, the problem becomes simpler: We need to remove all the centimeters which are not this flavor.

There are still two tricky things: It can be that it is more efficient to remove a larger part which also includes centimeters of the flavor we want to keep. For example, let's say we have a baguette which contains the pattern `aba`, and we can remove parts of length at most 3 cm, and want to end up with only flavor `b`. Then it's better to cut out the entire pattern at once, rather than cutting out the two centimeters of flavor `a` individually.

But then we need to be careful that we don't end up with an empty baguette.

One possible solution is to brute-force a centimeter of the baguette which we definitely keep. Once we know which centimeter to keep, we can then scan the baguette left to right, and when we see a centimeter of the wrong flavor, we remove a part of the baguette which is as long as possible, stopping before the centimeter which we definitely keep.

It's optimal to remove a part which is as long as possible, because we only need to optimize the number of cuts, not their length, and this way we maximize the chance of removing as many centimeters of the wrong flavor as possible.



```
1 #include<bits/stdc++.h> // score: 67
2 using namespace std;
3
4 signed main () {
5     int n, k;
6     string baguette;
7     cin >> n >> k >> baguette;
8
9     int best_cuts = n;
10    for (int keep = 0; keep < n; keep++) {
11        int cuts = 0;
12        int cut_until = 0;
13        for (int i = 0; i < n; i++) {
14            if (i == keep) {
15                cut_until = 0;
16            } else if (cut_until <= i && baguette[i] != baguette[keep]) {
17                cuts++;
18                cut_until = i + k;
19            }
20        }
21        best_cuts = min(best_cuts, cuts);
22    }
23    cout << best_cuts << "\n";
24 }
```

This solution runs in $O(n^2)$ time, which is enough for this subtask.

Subtask 5: 100 km baguette

The baguette can be at most 10^7 cm in length. The previous solution is too slow for this.

We can improve it by making the observation that we can actually ignore the constraint that the baguette must be non-empty. We can only violate this constraint if we remove the entire baguette. In this case, we can just add back the last centimeter of the baguette, it doesn't matter what flavor it has.

The solution is to try out all flavors, and for each, scan the baguette from left to right and remove a maximal part whenever we see a centimeter of the wrong flavor.

```
1 #include<bits/stdc++.h> // score: 100
2 using namespace std;
3
4 signed main () {
5     int n, k; cin >> n >> k;
6     string s; cin >> s;
7
8     int ans = n;
9
10    for (char keep = 'a'; keep <= 'z'; ++keep) {
11        int cuts = 0;
12        for (int i = 0; i < s.size(); i) {
13            if (s[i] != keep) {
14                i += k;
15                cuts += 1;
16            } else {
17                i += 1;
18            }
19        }
20        ans = min(ans, cuts);
21    }
22    cout << ans << "\n";
23 }
```

This solution runs in $O(n)$ time.

Catering

Task Idea	Johannes Kapfhammer
Task Preparation	Amirkeivan Mohtashami
Description English	Amirkeivan Mohtashami
Description German	Bibin Muttappillil, Johannes Kapfhammer
Description French	Benjamin Faltin
Solution	Amirkeivan Mohtashami

We are given a weighted connected graph with n nodes and m edges representing a network of roads and the price needed to cover each road with cheese. Our goal is to cover enough roads with cheese such that between any two nodes there exists a path between them which is completely covered in cheese. A sponsor also accepts to cover k of the roads we chose for free. We want to compute the minimum cost needed to cover the remaining roads we chose with cheese.

Observation

The roads we want to cover with cheese must form a spanning tree. If the solution wasn't a tree then there would exist a cycle and the number of roads we must pay for ourselves could be reduced by one by removing one edge in the cycle and the solution wouldn't be optimal. Additionally, it must be a spanning tree since all two nodes must be connected by a path.

Solution

Since the roads must form a tree with $n - 1$ edges we must pay for $\max(0, n - k - 1)$ of these edges. The claim is that these $n - k - 1$ edges are the $n - k - 1$ smallest edges contained in a minimum spanning tree of the weighted graph. Let's prove this. The lemma 1 isn't needed to prove that the algorithm is correct but is needed to prove that the $n - k - 1$ smallest edges of any spanning trees form a solution.

Lemma 1.

All minimum spanning trees have edges of same weight, alternatively, the i -th smallest edge of two different spanning trees are of same weight.

Proof. We will proceed by contradiction. Let's suppose we have two minimum spanning trees T_1, T_2 with different edge weights.

Then T_1 and T_2 must have a different number of edges with weight w for some weight w in the possible edge weights. Lets call $n_{w,1}$ the number of edges of weight w in T_1 and $n_{w,2}$ the number in T_2 .

Without loss of generality, lets suppose $n_{w,1} > n_{w,2}$

Now, for each edge e_w of weight w in T_1 we can add it to T_2 and it will create a cycle containing the edge e_w provided e_w isn't already in T_2 .

Since T_2 is a MST, all edges on that created cycle must have weight smaller or equal to w otherwise we could remove one of the longer edges and have a smaller MST than T_2 which is absurd

Since T_1 is a MST, there must be an other edge on the created cycle in T_2 which doesn't belong to T_1 with weight w . If that wasn't the case, we could remove e_w from T_1 and add one of the smaller edge in that created cycle to connect the two newly disconnected subtrees in T_1 which would imply that we can have a smaller MST than T_1 which is absurd.

We will now, for each edge e_w not already in T_2 , remove one edge of weight w not in T_1 from T_2 and add e_w to T_2 .

We conclude that we can remove all edges with weight w from T_2 and insert all the ones from T_1 and know that T_2 is still a MST after the changes. But then T_2 has more edges than T_1 since $n_{w,1} > n_{w,2}$ hence it is no longer a MST which is absurd
Therefore, T_1 and T_2 must have edges with the same weights. □

Lemma 1. shows us that all MST have the same set of edge weights and so taking the $n - k - 1$ smallest edges isn't impacted by the found MST. In this solution Kruskal's algorithm was used to compute the edges but because of Lemma 1., any other method to get a valid MST gives a correct answer.

Let's now show that the $n - k - 1$ smallest edges of an MST are a set without cycles with minimum weight sum.

(We will use the $n - k - 1$ first edges obtained using Kruskal's algorithm.)

Proof. Let S_1 be a set of $n - k - 1$ edges which is solution to the problem.
We will proceed using induction.

(I.H.) The set S_2 of l edges obtained using Kruskal on the first m edges is optimal. i.e. the l edges in S_2 are the l first edges of some solution S_1 and all edges between the l -th edge of S_2 and the m -th edge of the graph are not in S_1 .

($m = 0$) The set containing edges present in the 0 first smallest edges of the graph is the empty set which is contained in S_1 . Therefore the algorithm is correct for a $m = 0$.

($m \implies m + 1$) Suppose the algorithm works for m edges. Let's show it works for $m + 1$ edges.

If at any point the subset S_2 of edges obtained with Kruskal applied on the m first edges is of size $n - k + 1$, we are done and the proof is finished.

Consider e_{m+1} the $m + 1$ -th smallest edge of the graph, we have two cases

- if e_{m+1} is in S_1 :

We are done and Kruskal is optimal on $m + 1$ edges.

- if e_{m+1} isn't in S_1 :

If when we insert e_{m+1} in S_2 (optimal by (I.H)) it create a cycle, not adding it to S_2 is correct as adding it would not create a valid solution.

If it doesn't create a cycle, then adding it to S_1 must create a cycle with no edges greater than e_{m+1} as otherwise we could remove a longer edge f from S_1 and $S_1 \cup \{e\} \setminus \{f\}$ would be a better solution than S_1 which is absurd.

Additionally there must exist an edge with weight greater or equal to e_{m+1} on the cycle since otherwise it would create a cycle in S_2 which is not the case by hypothesis.

Therefore there exists an edge f with weight equal to e_{m+1} such that $S'_1 = S_1 \cup \{e_{m+1}\} \setminus \{f\}$ is a valid solution. We can now replace S_1 with S'_1 while keeping all our induction steps true since S'_1 works like S_1 for all previous induction steps.

Since Kruskal gives us a tree with size $n - 1 \geq n - k - 1$, the set S_2 obtained is always a solution by induction. □

Again, by Lemma 1., this shows us that any MST algorithm works since Kruskal's algorithm works for this problem.

```

1 #include <bits/stdc++.h> // score: 100
2
3 using namespace std;
4
5 #define int int64_t
6 #define all(x) x.begin(), x.end()
7 #define sz(x) (int)x.size()
8
9 vector<int> dsu;
10

```



```
11 int find(int u) {
12     if (dsu[u] == u)
13         return u;
14     return dsu[u] = find(dsu[u]);
15 }
16
17 signed main() {
18     ios_base::sync_with_stdio(0);
19     cin.tie(0);
20
21     int n,m,k;
22     cin >> n >> m >> k;
23
24     dsu.resize(n);
25     for (int i = 0; i < n; i++)
26         dsu[i] = i;
27
28     vector<array<int,3>> e(m);
29     for (int i = 0; i < m; i++) {
30         int u,v,c;
31         cin >> u >> v >> c;
32
33         e[i] = {c, u, v};
34     }
35
36     sort(all(e));
37
38     int cost = 0;
39     int cnt = k;
40     for (int i = 0; cnt < n-1; i++) {
41         auto [c,u,v] = e[i];
42
43         u = find(u);
44         v = find(v);
45         if (u == v)
46             continue;
47         cost += c;
48         cnt++;
49         dsu[u] = v;
50     }
51     cout << cost << "\n";
52 }
```