


**SOI 2P 2020**  
**Solution Booklet**



Swiss Olympiad in Informatics

13–16 March 2019



## Candy distribution

Task Idea	Monika Steinova
Task Preparation	Monika Steinova
Description English	Monika Steinova
Description German	Jan Schär
Description French	Florian Gatignon
Solution	Monika Steinova

### Subtask 1: 50 Points (simulation)

The task can be solved by a straightforward greedy simulation: Unless the three candy counts are not equal, add a candy to the two smallest counts.

To see why this process terminates and that it yields the minimum number of rounds we observe how candy counts change during the simulation.

Let  $(a, b, c)$  be the candy counts of mice  $A, B, C$ . WLOG let  $a \leq b \leq c$ . In the first  $c - b$  rounds our simulation adds  $c - b$  candy to mouse  $A$  and  $c - b$  candy to mouse  $B$  yielding candy counts  $(a + c - b, b + c - b, c) = (a + c - b, c, c)$ . (Notice that since  $b \geq a$ , the first element  $a - b + c \leq c$ .)

Let's observe how the candy counts change from now on after every two rounds. As long as mouse  $A$  has less candy than her sisters we give 2 candy to  $A$  and 1 candy to each of  $B$  and  $C$ . Hence after  $2k$  such rounds the candy counts are  $(a + c - b + 2k, c + k, c + k)$ .

As the  $A$ 's candy count  $c + (a - b) \leq c$  after  $k := b - a \geq 0$  rounds the candy counts will be equal:  $(a + c - b + (b - a) + (b - a), c + (b - a), c + (b - a)) = (c + k, c + k, c + k)$  and our simulation terminates.

During the whole simulation the candy count of  $A$  was less than candy count of  $C$ . Hence the value  $k := b - a$  is the minimal one needed to reach the equality of candy counts.

Overall we did  $(c - b) + 2k = b + c - 2a$  rounds and hence the time complexity is  $O(\max(a, b, c))$ .

```
1 #include <iostream>
2 #include <vector>
3 #include <algorithm>
4
5 using namespace std;
6
7 int main() {
8     long long round = 0;
9     vector<long long> C(3, 0);
10    cin >> C[0] >> C[1] >> C[2];
11
12    while( !( C[0] == C[1] && (C[1] == C[2]) ) ) {
13        sort(C.begin(), C.end());
14        ++C[0];
15        ++C[1];
16        ++round;
17    }
18    cout << round << endl;
19    return 0;
20 }
```

### Subtask 2: 100 Points in a constant time

In the previous section we have proven that  $b + c - 2a$  rounds are necessary and sufficient. Hence our algorithm can just output this number and terminate.

There exists another approach which doesn't need to analyze rounds. Observe that giving one candy to two different mice is equivalent to taking one candy from the third mouse. You can prove this by observing how the differences between mices change if the two candies are added. Hence the task can be reformulated to "What is the minimum amount of candy that has to be



removed from a mouse to make the candy counts of the three mice equal?" And the answer to this is simple. It is  $(b - a) + (c - a) = b + c - 2a$  ( $a \leq b \leq c$ ).

```
1 #include<iostream>
2 #include<vector>
3 #include<algorithm>
4
5 using namespace std;
6
7 int main(void) {
8     vector<long long> C(3, 0);
9     cin >> C[0] >> C[1] >> C[2];
10    sort(C.begin(), C.end());
11    cout << C[1] + C[2] - 2*C[0] << endl;
12    return 0;
13 }
```

## Workshops

Task Idea	Johannes Kapfhammer
Task Preparation	Florian Gatignon
Description English	Florian Gatignon
Description German	Martin Raszyk
Description French	Florian Gatignon
Solution	Florian Gatignon

### Subtask 1: 20 points (Bruteforce)

In the first subtask, the limits were low, so a simple bruteforce was enough to get those first 20 points. The idea is to compute whether you can organize 1 workshop, then 2, and so on until you find a number of workshops that you can't reach.

To determine whether it is possible to organize  $x$  workshops, we need to check how many jokers are needed for each job and compare the sum of those values to  $J$ . Another way to do it is to always check whether it is possible to organize one workshop, but removing the needed number of jokers from  $J$  afterwards. This is how the solution below works. For each organizable workshop, we use  $O(N)$  time, so  $O(NW)$  in total if the answer is  $W$ . The answer itself can't be larger than  $J + \min_{0 \leq i < N} b_i$ .

```

1 #include <bits/stdc++.h>
2 #define ll long long
3 using namespace std;
4 ll N, J;
5 vector<ll> a, b;
6
7 int main() {
8     ios_base::sync_with_stdio(false);
9     cin.tie(0);
10    cin >> N >> J;
11    a.resize(N); b.resize(N);
12    for(ll i = 0; i < N; i++) cin >> a[i];
13    for(ll i = 0; i < N; i++) cin >> b[i];
14    ll ans, tot=0;
15    for(ans=-1; tot<=J; ans++) {
16        J--tot;
17        tot = 0;
18        for(ll i = 0; i < N; i++) {
19            if(a[i]>b[i]) {
20                tot += a[i]-b[i];
21                b[i] = 0;
22            } else {
23                b[i] -= a[i];
24            }
25        }
26    }
27    cout << ans << endl;
28 }
```

### Subtask 2: 30 points (No jokers)

In the second subtask, there were more jobs and regular volunteers but there weren't any jokers available. You just had to compute how many workshops could be organized with normal volunteers. To do so was quite easy: the answer is  $\min_{0 \leq i < n} (\lfloor \frac{a_i}{b_i} \rfloor)$ , which one may compute in  $O(N)$ .

Of course, you can combine this solution with that of Subtask 1 to get a total of 50 points.



```
1 #include <bits/stdc++.h>
2 #define ll long long
3 using namespace std;
4 ll N, J;
5 vector<ll> a, b;
6
7 int main() {
8     ios_base::sync_with_stdio(false);
9     cin.tie(0);
10    cin >> N >> J;
11    a.resize(N); b.resize(N);
12    for(ll i = 0; i < N; i++) cin >> a[i];
13    for(ll i = 0; i < N; i++) cin >> b[i];
14    ll ans=1000000000000;
15    for(ll i = 0; i < N; i++) {
16        ans = min(ans,b[i]/a[i]);
17    }
18    cout << ans << endl;
19 }
```

### Subtask 3: 100 points (Binary search)

In the third and last subtask, all limits were quite large. For all  $0 \leq x \leq W$ , it is possible to organize  $x$  workshops, and for all  $x > W$  it is impossible. In other words, we are looking for a *tipping point*. That means that we don't need to test every possible number of workshops until  $W$ , we can binary search for that tipping point.

Basically, we take a large interval that includes all possible answers given the limits. We take the middle of the interval; if we can organize that many workshops, we keep that value in memory as the best answer until now and recurse on the upper half of the interval; else, we look at the lower half. Each time, we can check in  $O(N)$  whether it is possible to organize a certain amount of workshops (in a trivial way), so our total running time is  $O(N \log(J + \min_{0 \leq i < N} b_i))$ .

```
1 #include <bits/stdc++.h>
2 #define ll unsigned long long
3 using namespace std;
4 ll N, J;
5 vector<ll> a, b;
6
7 bool check(ll x) {
8     ll tot=0ll;
9     for(ll i = 0; i < N; i++) {
10        if(x*a[i]>b[i])
11            tot += x*a[i]-b[i];
12    }
13    return tot<=J;
14 }
15
16 int main() {
17     ios_base::sync_with_stdio(false);
18     cin.tie(0);
19     cin >> N >> J;
20     a.resize(N); b.resize(N);
21     for(ll i = 0; i < N; i++) cin >> a[i];
22     for(ll i = 0; i < N; i++) cin >> b[i];
23     ll lo = 0ll, hi = 2134567890, ans = 0ll;
24     while(lo<=hi) {
25         ll mi = (lo+hi)/2;
26         if(check(mi)) {lo = mi+1; ans = mi;}
27         else hi = mi-1;
28     }
29     cout << ans << endl;
30 }
```

# Toll Roads

Task Idea	Daniel Rutschmann
Task Preparation	Martin Raszyk
Description English	Martin Raszyk
Description German	Martin Raszyk
Description French	Erwan Serandour
Solution	Johannes Kapfhammer

Given a graph with edge weights 0 or 1, the task was to answer multiple queries that asked you to compute the distance between a pair of vertices.

Starting from the basic solution (one Dijkstra or one 0-1-BFS), there are two independent observations required for 100 points: precomputing distances to answer the queries faster and compressing 0-edges.

## One Dijkstra per Query (25 points)

The distance between two vertices can be computed using Dijkstra. The graph is the same as the input graph, where the private highway segments have cost 1 and the state highway segments have cost 0. Note that Dijkstra also works for graphs where the edge cost is 0 (Dijkstra only breaks if the costs can be negative). Since the weight corresponds to the cost, the shortest path will be the cheapest one.

```
1 #include <bits/stdc++.h>
2 using namespace std;
3
4 const int INF = 1e9;
5
6 // returns the distances from start to all other vertices
7 vector<int> dijkstra(vector<vector<pair<int, int>>> const& g, int start) {
8     vector<int> dist(g.size(), INF);
9     priority_queue<pair<int, int>, vector<pair<int, int>>, greater<>> pq;
10    pq.emplace(0, start);
11    while (!pq.empty()) {
12        auto [cost, v] = pq.top();
13        pq.pop();
14        if (dist[v] != INF) continue;
15        dist[v] = cost;
16        for (auto [u, c] : g[v])
17            pq.emplace(cost+c, u);
18    }
19    return dist;
20 }
21
22 int main() {
23     ios::sync_with_stdio(false);
24     cin.tie(0);
25
26     int n, m, k, q;
27     cin >> n >> m >> k >> q;
28
29     vector<vector<pair<int, int>>> g(n);
30     for (int i=0; i<m; ++i) {
31         int a, b; cin >> a >> b;
32         g[a].emplace_back(b, 0);
33         g[b].emplace_back(a, 0);
34     }
35     for (int i=0; i<k; ++i) {
36         int a, b; cin >> a >> b;
37         g[a].emplace_back(b, 1);
38         g[b].emplace_back(a, 1);
```



```
39 }
40 for (int i=0; i<q; ++i) {
41     int a, b; cin >> a >> b;
42     cout << dijkstra(g, a)[b] << '\n';
43 }
44 }
```

This solution runs in  $O(q \cdot (m + k) \log(m + k))$ .

## Once Dijkstra per Vertex (50 points)

It turns out we don't need to compute one Dijkstra per query. A run of Dijkstra starting at vertex  $s$  computes the distances  $d(s, v)$  for all other vertices  $v$ . Therefore, we can precompute the distances for all start vertices  $s$  to all other vertices.

The only change required is replacing the query loop with the following:

```
1 vector<vector<int>> dist(n);
2 for (int s=0; s<n; ++s)
3     dist[s] = dijkstra(g, s);
4 for (int i=0; i<q; ++i) {
5     int a, b; cin >> a >> b;
6     cout << dist[a][b] << '\n';
7 }
```

This reduces the running time from  $O(q \cdot (m + k) \log(n))$  to  $O(n \cdot (m + k) \log(m + k) + q)$ .

## Compressing components of 0-edges (50 points)

Until now, we didn't make use of the fact that all distances are 0 or 1. Looking at the limits, most edges have weight 0.

If two vertices  $a$  and  $b$  are connected with an edge of weight 0, we can just pretend they are the same vertex. Why? Because we have  $d(a, v) = d(b, v)$  for all vertices  $v$ . (Proof: in any graph we have the triangle inequality  $d(a, b) \leq d(a, x) + d(x, b)$ . Applying this to  $a, b$  and  $v$  and noting  $d(a, b) = 0$ :  $d(a, v) \leq d(a, b) + d(b, v) = d(b, v)$  and  $d(b, v) \leq d(b, a) + d(a, v) = d(a, v)$ , so  $d(a, v) = d(b, v)$ .)

More generally, this applies to all pairs of vertices that have distance 0. So we look for "clumps" of vertices that have distance 0 and compress them to one mega-vertex.

This is fast because we have to compress the vertices just once, and then then all remaining edges are the ones with distance 1.

To compress, we create a graph containing only the edges of cost 0. We identify components using DFS and save for each vertex to which component it belongs. Then, we build a second graph where the vertices are components of the first graph, containing only edges of cost 1.

Precomputing can be done in  $O(n + m)$ . Then, to answer queries one could run one BFS per query and get  $O(n + m + q \cdot k)$

```
1 #include <bits/stdc++.h>
2 using namespace std;
3
4 // sets component[w]=comp_id for all vertices w that are reachable from v
5 // assuming none of those vertices have been visited yet
6 void dfs(vector<vector<int>> const& g, int v, int comp_id, vector<int>& component) {
7     if (component[v] != -1) return;
8     component[v] = comp_id;
9     for (auto w : g[v])
10         dfs(g, w, comp_id, component);
11 }
12
13 // returns the number of components and the component id of each vertex
14 pair<int, vector<int>> compress(vector<vector<int>> const& g) {
15     vector<int> component(g.size(), -1);
16     int comp_id = 0;
```



```
17     for (int i=0; i<(int)g.size(); ++i)
18         if (component[i] == -1)
19             dfs(g, i, comp_id++, component);
20     return {comp_id, component};
21 }
22
23 // returns the distances from start to all other vertices
24 vector<int> bfs(vector<vector<int>> const& g, int start) {
25     vector<int> dist(g.size(), -1);
26     queue<int> q;
27     dist[start] = 0;
28     q.push(start);
29     while (!q.empty()) {
30         auto v = q.front();
31         q.pop();
32         for (auto w : g[v]) {
33             if (dist[w] == -1) {
34                 dist[w] = dist[v] + 1;
35                 q.push(w);
36             }
37         }
38     }
39     return dist;
40 }
41
42 int main() {
43     ios::sync_with_stdio(false);
44     cin.tie(0);
45
46     int n, m, k, q;
47     cin >> n >> m >> k >> q;
48
49     vector<vector<int>> g0(n); // graph of 0-edges
50     for (int i=0; i<m; ++i) {
51         int a, b; cin >> a >> b;
52         g0[a].push_back(b);
53         g0[b].push_back(a);
54     }
55     auto [num_components, component] = compress(g0);
56     vector<vector<int>> g1(num_components); // graph of 1-edges
57     for (int i=0; i<k; ++i) {
58         int a, b; cin >> a >> b;
59         a = component[a];
60         b = component[b];
61         if (a != b) {
62             g1[a].push_back(b);
63             g1[b].push_back(a);
64         }
65     }
66     for (int i=0; i<q; ++i) {
67         int a, b; cin >> a >> b;
68         a = component[a];
69         b = component[b];
70         cout << bfs(g1, a)[b] << '\n';
71     }
72 }
```

## Compressing and Precomputing (100 points)

Both ideas can be combined for 100 points.

Again, the only part that changes is that we precompute the table of all distances before answering the queries:

```
1 vector<vector<int>> dist(num_components);
2 for (int s=0; s<num_components; ++s)
3     dist[s] = bfs(g1, s);
4 for (int i=0; i<q; ++i) {
```





```
5  int a, b; cin >> a >> b;  
6  cout << dist[component[a]][component[b]] << '\n';  
7  }
```

With this, the running time is  $O(n + m + n \cdot k + q)$ .

### Alternative idea: 0-1 BFS

Instead of using Dijkstra, one could do a so-called 0-1-BFS. This is a BFS that works if the graph has edges of weights 0 and 1 and runs in  $O(n + m)$  (where  $m$  is the number of total edges) instead of  $O(n + m \log n)$  of a normal Dijkstra (and  $O(n \log n + m)$  for a Dijkstra with a Fibonacci Heap).

The algorithm is explained here: [https://cp-algorithms.com/graph/01\\_bfs.html](https://cp-algorithms.com/graph/01_bfs.html).

### Alternative idea: Union Find

In the code above the components are identified using DFS. Of course, one could also use union find to handle this. Union find is more general because it can handle dynamic updates. This was not needed in this task, however.

# Pizzacut

Task Idea	Daniel Rutschmann
Task Preparation	Daniel Rutschmann
Description English	Daniel Rutschmann
Description German	Benjamin Schmid
Description French	Florian Gatignon
Solution	Daniel Rutschmann

## Greedy doesn't work

There are various wrong greedy solutions, for instance based on giving each mouse the minimum number of grams of pizza. In fact, most of the already fail on the samples. While one can show that some these solutions use at most one pizza too many (see 90p/100p solution), there is no easy way of fixing this. In this task, the limits for  $A$  and  $B$  are quite small. This is a strong indicator against greedy solutions, as those would be fast enough even for large  $A$  and  $B$ .

## 30 Points

Let's use "A-slices" for slices of  $A$  grams and "B-slices" for slices of  $B$  grams. The small limits for  $A$  and  $B$  suggest some solution with dynamic programming. In this subtask, we can use  $DP[i][x]$  = "minimum number of B-slices we need to use to feed the first  $i$  mice if we use exactly  $x$  A-slices". When considering the  $i$ -th mouse, we try all possible ways of efficiently feeding this mouse. More precisely, we iterate over the number of A-slices this mouse gets and compute the minimum number of B-slices this mouse then needs.

To bound the runtime, note that each mouse needs at most  $2B$  A-slices, as  $g_i \leq 2AB$ , so  $x$  is at most  $2NB$ . Hence the above DP has  $O(N^2B)$  states. Each state takes  $O(B)$  times to compute, as there are  $O(B)$  ways of efficiently feeding a fixed mouse, so we get a runtime of  $O(N^2B^2)$ .

```

1 #include <bits/stdc++.h>
2 using namespace std;
3
4 constexpr int inf = 1e9;
5
6 // ceil(a/b)
7 int cdiv(int a, int b){
8     if(a > 0) return (a + b-1)/b;
9     return a/b;
10 }
11 template<typename T>
12 void xmin(T&a, T const&b){
13     a = min(a, b);
14 }
15 signed main(){
16     int n, A, B;
17     cin >> n >> A >> B;
18     vector<int> v(n);
19     for(auto &e:v) cin >> e;
20     const int X = 2*n*N+1;
21     vector<vector<int>> > dp(n+1, vector<int> (X, inf));
22     dp[0][0] = 0;
23     for(int i=0; i<n; ++i){
24         const int e = v[i];
25         for(int a=0; a <= e/A+1; ++a){
26             const int b = cdiv(max(0, e-a*A), B);
27             for(int x = 0; x < X-a; ++x){
28                 xmin(dp[i+1][x+a], dp[i][x] + b);
29             }
30     }

```



```
31 }
32 int ans = inf;
33 for(int i=0; i<X; ++i){
34     xmin(ans, cdiv(dp[n][i], A) + cdiv(i, B));
35 }
36 cout << ans << "\n";
37 }
```

## 45 Points

To get 45 Points, we need to handle larger values of  $g_i$ . We make the following observation: If a mouse eats more than  $2AB$  grams of pizza, then it eats at least  $AB$  grams worth of A-slices or at least  $AB$  grams worth of B-slices, i.e. a full pizza of A-slices or a full pizza of B-slices. It makes no difference whether this full pizza is cut into A-slices or B-slices, so we can just give this mouse a full pizza and reduce  $g_i$  by  $AB$ . This allows us to assume that  $g_i \leq 2AB$ , at which point we can use the previous DP solution.

## 70 points

As in the solution for 45 Point, we reduce to the case where  $g_i \leq 2AB$ . In the 30 points solution, we used  $DP[i][x]$  where  $x$  was the total number of A-slices we used, so  $x$  could get as large as  $2BN$ . Since we only care about the number of pizzas and not about the exact number of A-slices or B-slices, we could replace a whole pizza of A-slices by a whole pizza of B-slices. In other words, if  $x \geq B$ , then we can reduce  $x$  by  $B$  if we increase the number of B-slices of A. This allows us to only consider  $0 \leq x < B$ , which reduces the number of DP states to  $O(NB)$ . This solution runs in  $O(NB^2)$ .

```
1 #include <bits/stdc++.h>
2 using namespace std;
3
4 constexpr int inf = 1e9;
5
6 int cdiv(int64_t a, int64_t b){
7     if(a > 0) return (a + b-1)/b;
8     return a/b;
9 }
10
11 template<typename T>
12 void xmin(T&a, T const&b){
13     a = min(a, b);
14 }
15
16 signed main(){
17     int n, A, B;
18     cin >> n >> A >> B;
19     vector<int> v(n);
20     int64_t whole = 0;
21     for(auto &e:v) {
22         cin >> e;
23         if(e > 2*A*B){
24             int x = e/(A*B)-2;
25             e -= x*A*B;
26             whole += x;
27         }
28     }
29     vector<vector<int64_t>> dp(n+1, vector<int64_t>(B, inf));
30     dp[0][0] = 0;
31     for(int i=0; i<n; ++i){
32         const int e = v[i];
33         for(int a=0; a <= min(B, e/A+1); ++a){
34             const int b = (max(0, e-a*A)+B-1)/B;
35             for(int x = 0; x < B-a; ++x){
36                 xmin(dp[i+1][x+a], dp[i][x]+b);
37             }
38         }
39     }
40     cout << dp[n][0] << "\n";
41 }
```



```
37     }
38     // whole pizza of A-slices -> whole pizza of B-slices
39     for(int x = B-a; x < B; ++x){
40         xmin(dp[i+1][x+a-B], dp[i][x]+b + A);
41     }
42 }
43 }
44 int ans = inf;
45 for(int i=0; i<B; ++i){
46     xmin(ans, cdiv(dp[n][i], A) + cdiv(i, B));
47 }
48 cout << whole+ans << "\n";
49 }
```

## 90 / 100 points

Getting more than 70 points is very difficult. The basic idea is to first start with some greedy solution and then do a DP to see if this greedy solution can be improved. Starting with a greedy solution greatly reduces the number of values the DP can attain. This allows us to use some data structure to quickly decide which options improve our DP and which don't.

Concretely, we first precompute all numbers of grams we can reach with A-slices and B-slices. Then for each mouse, we give that mouse the fewest number of grams of pizza possible. This gives us some number of A-slices and some number of B-slices we use, so we order enough pizzas to get that many slices.

We claim that this uses at most one pizza too many. To show this claim, note that this greedy solution uses the minimum number of grams of pizza. The only inefficiency comes from the fact that number of A-slices we use might be not divisible by B, in which case we have some A-slices that we don't use. We have at most  $(B - 1)$  leftover A-slices and at most  $(A - 1)$  leftover B-slices, so the number of grams of pizza we order but don't use is less than  $2AB$ . Hence any other solution can save strictly less than  $2AB$  grams of pizza, so it can save at most one whole pizza. Based on this observation, there are a few different ways of getting an  $O(n + AB^2)$  solution. We will first discuss the original solution and then an easier solution that we only discovered after the contest.

### Original solution

Starting from the greedy solution, we can try to get a better solution by give some mice more A-slices and fewer B-slices or vice-versa. This results in the following:  $DP[i][\Delta a]$  = "minimum number of additional B-slices we need if we use  $\Delta a$  additional A-slices for the first  $i$  mice" (this will be zero or negative). Initially,  $DP[0][\Delta a] = 0$ . When considering the  $i$ -th mouse, we again try all efficient ways of feeding it and check how this would change the greedy solution. As in the 70p solution, we again use the trick of replacing whole pizzas of A-slices by whole pizzas of B-slices to assume  $0 \leq \Delta a < B$ .

This solution still runs in  $O(nB^2)$ , so it will only give 70 points. To speed it up, note that  $DP[i+1][\Delta a] \leq DP[i][\Delta a]$ , so this is never positive and that  $DP[i][\Delta a] > -3A$  as we can't improve the greedy solution by more than one pizza. Hence  $DP[i][\Delta a]$  only decreases as  $i$  increases and it can decrease at most  $O(A)$  times.

When considering the  $i$ -th mouse, each change will be of the form: If we use  $x$  additional A-slices, then we need  $y$  additional B-slices, where either  $x > 0, y < 0$  or  $x < 0, y > 0$ . The key idea is to use some data structure to maintain for each  $x$  ( $-B \leq x \leq B$ ) the minimum  $y$  that would cause a decrease in the DP. This data structure should maintain a multiset of  $B$  values (one value for each possible value of  $\Delta a$ ), each between  $-2A$  and  $+2A$  and it should support finding the maximum and updating a value in the set in amortized  $O(1)$ . This can be achieved by storing a vector of counts and the current maximum. Whenever we compute a new entry of our DP table, we need to update this datastructure.

This data structure allows us to quickly (in  $O(1)$ ) decide whether a change  $(x, y)$  will improve the DP or not. As the DP can be improved at most  $O(AB)$  times, we spend  $O(AB^2)$  time updating the DP in total. We need  $O(nB)$  time to consider all possible changes, so we get a total runtime of  $O(nB + AB^2)$ . This scores 90 points.

To get 100 points, we note that there are at most  $2AB$  different values of  $g_i$ . If some value of  $g_i$  doesn't change the DP, we can discard all mice with that value. Hence for every Mouse we consider, we either throw away some  $g_i$ , or we find some improvement in the DP. Hence we consider only  $O(AB)$  mice, so the runtime is  $O(n + AB^2)$ .

### Easier solution (Thanks Fabian!)

We again start with the greedy solution and try to improve it by some changes of the form: Give some mouse  $x$  additional  $A$ -slices and  $y$  additional  $B$ -slices where one of  $x$  and  $y$  is positive and the other is negative. We make two observations.

1. Whether or not some change  $(x, y)$  with  $x < 0$  is possible for some mouse only depends on the the number of  $A$ -slices this mouse has in the greedy solution. More precisely, this mouse needs to have at least  $x$   $A$ -slices. The same also holds for changes with  $y < 0$  and  $B$ -slices.
2. There is an optimal solution that does at most  $B$  such changes.

To see (2), we use the following easy lemma: Given  $B$  integers  $x_1, \dots, x_B$ , there is a non-empty subset with sum divisible by  $B$ . The proof of this is an easy exercise (Hint: look at prefix sums modulo  $B$ ). Suppose now that we do  $B$  changes  $(x_1, y_1), (x_2, y_2), \dots, (x_B, y_B)$ . By the lemma, there is a subset of changes whose  $x_i$  is divisible by  $B$ . Let's remove these changes. This does not change the number of  $A$ -slices modulo  $B$  and it reduces the number of changes. As a change can only increase the total number of grams of pizza we use, removing these changes only improves our solution. Hence we can always get an optimal solution with fewer than  $B$  changes.

With these observations, we see that we only need to do changes on the  $B$  mice with a largest number of  $A$ -slices and the  $B$  mice with a largest number of  $B$ -slices. All other mice will get exactly the slices they would get in the greedy solution. Hence we can run the 70p DP on these  $2B$  mice and the greedy on all other mice. This reduces  $n$  to  $2B$ , so the runtime is  $O(n + B^3)$ .

```

1 #include <bits/stdc++.h>
2 using namespace std;
3
4 constexpr int inf = 1e9;
5 template<typename T>
6 void xmin(T&a, T const&b){
7     a = min(a, b);
8 }
9 int cdiv(int a, int b){
10     if(a > 0) return (a + b-1)/b;
11     return a/b;
12 }
13 struct Greedy_Decomp{
14     Greedy_Decomp(int A_, int B_) : A(A_), B(B_), U(2*A*B+1), pre(U, {-1, -1}){
15         for(int i=0; i <= min(B-1, U/A); ++i){
16             for(int j=0; j <= U/B; ++j){
17                 const int r = i*A + j*B;
18                 if(r < U){
19                     pre[r] = make_pair(i, j);
20                 }
21             }
22         }
23         for(int i=U-1; i >= 0; --i){
24             if(pre[i].first == -1) pre[i] = pre[i+1];
25         }
26     }
27     pair<int, int> decompose(int x) const {
28         return pre[x];
29     }
30     int A, B, U;
31     vector<pair<int, int> > pre;
32 };
33 struct Mouse{
34     int a, b, hunger;

```



```
35 };
36 signed main(){
37     cin.tie(0);
38     ios_base::sync_with_stdio(false);
39     int n, A, B;
40     cin >> n >> A >> B;
41     Greedy_Decomp greed(A, B);
42     vector<int> v(n);
43     int64_t whole = 0;
44     for(auto &e:v) {
45         cin >> e;
46         if(e > 2*A*B){
47             int x = e/(A*B)-1;
48             e -= x*A*B;
49             whole += x;
50         }
51     }
52     vector<Mouse> mice(n);
53     for(int i=0; i<n; ++i){
54         auto [a,b] = greed.decompose(v[i]);
55         mice[i].hunger = v[i];
56         mice[i].a = a;
57         mice[i].b = b;
58     }
59     vector<Mouse> used_mice;
60     if(n <= 2*B){
61         used_mice.swap(mice); // use all mice
62     } else {
63         // use B mice with largest number of A-slices
64         nth_element(mice.begin(), mice.begin()+B, mice.end(),
65             [](Mouse const&x, Mouse const&y){return x.a > y.a;});
66         used_mice.insert(used_mice.end(), mice.begin(), mice.begin()+B);
67         mice.erase(mice.begin(), mice.begin()+B);
68
69         // use B mice with largest number of B-slices
70         nth_element(mice.begin(), mice.begin()+B, mice.end(),
71             [](Mouse const&x, Mouse const&y){return x.b > y.b;});
72         used_mice.insert(used_mice.end(), mice.begin(), mice.begin()+B);
73         mice.erase(mice.begin(), mice.begin()+B);
74     }
75     // now run 70p dp on the 2B best mice
76     const int k = used_mice.size();
77     vector<vector<int>> dp(k+1, vector<int>(B, inf));
78     dp[0][0] = 0;
79     for(int i=0; i<k; ++i){
80         const int e = used_mice[i].hunger;
81         for(int a=0; a <= min(B, e/A+1); ++a){
82             const int b = (max(0, e-a*A)+B-1)/B;
83             for(int x = 0; x < B-a; ++x)
84                 xmin(dp[i+1][x+a], dp[i][x]+b);
85             // whole pizza of A-slices -> whole pizza of B-slices
86             for(int x = B-a; x < B; ++x)
87                 xmin(dp[i+1][x+a-B], dp[i][x]+b + A);
88         }
89     }
90     // run greedy on other mice
91     int total_a = 0, total_b = 0;
92     for(auto const&e:mice){
93         total_a+=e.a;
94         total_b+=e.b;
95     }
96     int ans = inf;
97     for(int i=0; i<B; ++i){
98         xmin(ans, cdiv(total_b + dp[k][i], A) + cdiv(total_a + i, B));
99     }
100     cout << whole+ans << "\n";
101 }
```