

Greedy Algorithms

Charlie Li

Coin Problem

- Suppose you are given infinitely many coins of the following values:
- \$10, \$5, \$2, \$1, \$0.5, \$0.2, \$0.1

- What is the minimum number of coins needed to get \$M where
 - (i) $M = 1.9$
 - (ii) $M = 7.8$
 - (iii) $M = 16.1$

Coin Problem

- Answers:
 - (i) 4
 - (ii) 5
 - (iii) 4

- Why can we get the answer so quickly?

Coin Problem

- Cashier's algorithm
 - While there are outstanding amount, we take one of the coin with largest value but not exceeding the outstanding amount
- (x) means the outstanding amount
- $(\$1.9) = \$1 + (\$0.9) = \$1 + \$0.5 + (\$0.4) = \$1 + \$0.5 + \$0.2 + \0.2
- $(\$7.8) = \$5 + (\$2.8) = \$5 + \$2 + (\$0.8) = \$5 + \$2 + \$0.5 + (\$0.3) = \$5 + \$2 + \$0.5 + \$0.2 + \$0.1$
- $(\$16.1) = \$10 + (\$6.1) = \$10 + \$5 + (\$1.1) = \$10 + \$5 + \$1 + \0.1

Coin Problem

- Does Cashier's algorithm always give us the optimal solution for HK's coins ?
 - {\$10, \$5, \$2, \$1, \$0.5, \$0.2, \$0.1}
- Is there any counter example? If no, why is it optimal?
- If you cannot come up with the answer, never mind, let's look at another example

Coin Problem

- Consider the system of stamps
- \$5, \$3.7, \$3.1, \$2.9, \$2.3, \$2.2, \$2, \$1.7, \$1, \$0.5, \$0.2, \$0.1
- What is the minimum number of stamps needed to get \$M where
 - (i) $M = 1.9$
 - (ii) $M = 7.8$
 - (iii) $M = 16.1$



Coin Problem

- Using the same algorithm,
 - (i) $\$1.9 = \$1.7 + \$0.2$
 - (ii) $\$7.8 = \$5 + \$2.3 + \0.5
 - (iii) $\$16.1 = \$5 + \$5 + \$5 + \$1 + \0.1
-
- The algorithm gives optimal solution for (i) and (ii) but not (iii)
 - $\$16.1 = \$5 + \$3.7 + \$3.7 + \$3.7$
 - We can use 4 stamps only but greedy told us 5

Coin Problem

- Go back to the original problem
- Why the Cashier's algorithm gives us the optimal solution for the coins with value $\{\$10, \$5, \$2, \$1, \$0.5, \$0.2, \$0.1\}$
- Let's see if there are any properties for the optimal solution

Coin Problem

- Property 1
 - The optimal solution contains at most one \$0.1 coin.
 - If there are more than one \$0.1 coin, we can change two \$0.1 into one \$0.2
 - Similarly, the optimal solution contains at most one \$5, \$1, \$0.5
- Property 2
 - The optimal solution contains at most two \$0.2 coins.
 - If there are more than two \$0.2 coins, we can change three \$0.2 into \$0.5 and \$0.1
 - Similarly, the optimal solution contains at most two \$2
- Property 3
 - The optimal solution contains no \$0.1 coins if it contains two \$0.2 coins
 - If there are one \$0.1 and two \$0.2, we can change them to \$0.5
 - Similarly, optimal solution contains no \$1 coins if it contains two \$2 coins

Coin Problem


Value of coin	Constraint	Max. value if we only use coins with less value
\$0.1	At most 1	\$0
\$0.2	At most 2 no \$0.1 when 2	\$0.1
\$0.5	At most 1	\$0.4
\$1	At most 1	\$0.9
\$2	At most 2 no \$1 when 2	\$1.9
\$5	At most 1	\$4.9
\$10	No limit	\$9.9

Coin Problem

- From the table, we cannot get \$10 just using coins less than \$10, so we must use \$10 coin
- This is also true for \$5, \$2, \$1, \$0.5, \$0.2 and \$0.1
- So Cashier's Algorithm can give optimal solution for HK's coin system

Greedy Algorithm

- A greedy algorithm attempt to solve the problem by choosing locally optimal step and reduce the original problem into another sub-problem then solve the sub-problem using the same strategy
- It is not a specific algorithm, it is just an idea
- It is important to know that for a greedy algorithm, the previous choices should not affect the decision
- Unlike brute force and DP(will be taught in April), a greedy algorithm may not always give an optimal solution, eg. The stamp system



Coin Problem – How Greedy?

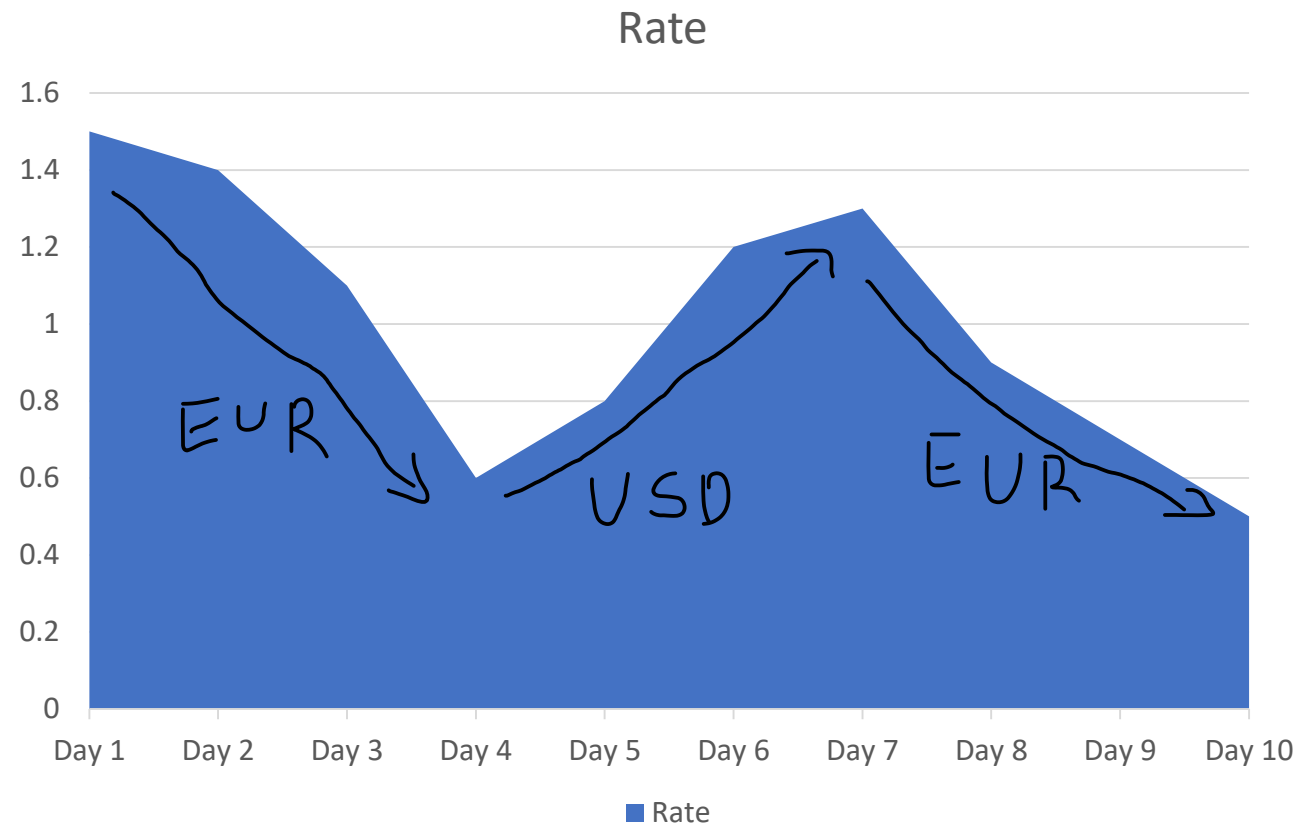
- Answer to main problem =
 - Optimal step's effect to answer + Answer to sub-problem
- Problem: Remaining amount K
 - We will greedily take a coin with largest value but not exceeding K
 - We will take a coin d where $d = \max\{D_i | D_i \leq K\}$
 - The sub-problem is remaining amount $K - d$
 - Answer for remaining amount K : $\text{Ans}(K) = 1 + \text{Ans}(K - d)$
 - Answer to main problem: $\text{Ans}(N)$
 - Base case: $\text{Ans}(0) = 0$

J042 Currency Exchange

- You know the exact exchange rate from USD to EUR for the following N days, you have US\$ M at day 1, how much USD will you have if you trade optimally?

J042 Currency Exchange

- When thinking greedily, we should be holding EUR when the rate decrease and holding USD when the rate increases



J042 Currency Exchange

- Try this input:

10 100

1.5 1.4 1.1 0.6 0.8 1.2 1.3 0.9 0.7 0.5

- Between day 1 and day 2, we will change to EUR at day 1 and change back at day 2 since the rate decreases
- Between day 2 and day 3, we will change to EUR at day 2 and change back at day 3 since the rate decreases
- ...
- Between day 4 and day 5, we will do nothing since the rate increases ...

J042 Currency Exchange

- The code looks like this:

```
1 #include<stdio>
2
3 int n;
4 double m, c[10004];
5
6 int main() {
7     scanf("%d %lf", &n, &m);
8     for (int i = 0; i < n; i++) scanf("%lf", &c[i]);
9     for (int i = 1; i < n; i++)
10         if (c[i - 1] > c[i])
11             m = m/c[i]*c[i - 1];
12     printf("%.2lf\n", m);
13 }
```

J042 Currency Exchange

- Why is this optimal?
- What if we hold EUR when the rate increase?
- What if we hold USD when the rate decrease?
- How are these compare to our greedy solution?

Fractional Knapsack

- Suppose you have a cup with volume V ml.
- There are N flavors of drinks, for each flavor, there are only a_i ml available and contain a total of s_i g sugar
- At most how many grams of sugar can you consume if you fill up the cup optimally?

Fractional Knapsack

Flavor	Volume available (ml)	Total sugar (g)
A	320	35
B	220	0
C	240	40
D	200	30

Fractional Knapsack

- We can just sort the flavors according to sugar per unit volume, and then fill up our cup according to the order

Flavor	Volume available (ml)	Total sugar (g)	Sugar per volume	Used
C	240	40	0.167	240
D	200	30	0.15	200
A	320	35	0.109	160
B	220	0	0	0

Fractional Knapsack

- Why is this optimal?
- Try to answer seeing what will happen if we change some of C into D, A or B

Flavor	Volume available (ml)	Total sugar (g)	Sugar per volume	Used
C	240	40	0.167	240
D	200	30	0.15	200
A	320	35	0.109	160
B	220	0	0	0

0-1 Knapsack

- What will happen if the question is changed to “If you used some flavor X, you must use up all flavor X”
- Can we still use greedy? Can you think of counter example?

Flavor	Volume available (ml)	Total sugar (g)	Sugar per volume	Choice
C	240	40	0.167	0 or 240
D	200	30	0.15	0 or 200
A	320	35	0.109	0 or 320
B	220	0	0	0 or 220

Longest Increasing Subsequence

- Given an integer array, find the length of its longest strictly increasing subsequence.
- For example, the length of LIS of the following array is 4

8 1 2 7 5 6 3 4

Longest Increasing Subsequence

- For an array $a[1..n]$ to have an increasing subsequence of length k , we need the following two statements to be true
 - The array $a[1..n-1]$ has an increasing subsequence of length $k-1$
 - One of such increasing subsequences has the $(k-1)^{\text{th}}$ element smaller than $a[n]$

Longest Increasing Subsequence

- So it maybe useful if we find the length of LIS of the whole array first by finding the one of its prefixes'
- However, instead of recording the LIS, it would be more useful if we record the smallest k^{th} element in any increasing subsequence
 - We can check which increasing subsequence $a[n]$ can be appended to
- It is easy to show that if we find the answer in this way, it will be correct

Longest Increasing Subsequence

- For example, use $a[] = 8\ 1\ 2\ 7\ 5\ 6\ 3\ 4$
- We record the smallest k^{th} element in any increasing subsequence

For []:

i	1	2	3	4
v[i]	+inf	+inf	+inf	+inf

Longest Increasing Subsequence

- For example, use $a[] = 8\ 1\ 2\ 7\ 5\ 6\ 3\ 4$
- We record the smallest k^{th} element in any increasing subsequence

For [8]:

i	1	2	3	4
v[i]	8	+inf	+inf	+inf

Longest Increasing Subsequence

- For example, use $a[] = 8\ 1\ 2\ 7\ 5\ 6\ 3\ 4$
- We record the smallest k^{th} element in any increasing subsequence

For [8 1]:

i	1	2	3	4
v[i]	1	+inf	+inf	+inf

Longest Increasing Subsequence

- For example, use $a[] = 8\ 1\ 2\ 7\ 5\ 6\ 3\ 4$
- We record the smallest k^{th} element in any increasing subsequence

For [8 1 2]:

i	1	2	3	4
v[i]	1	2	+inf	+inf

Longest Increasing Subsequence

- For example, use $a[] = 8\ 1\ 2\ 7\ 5\ 6\ 3\ 4$
- We record the smallest k^{th} element in any increasing subsequence

For $[8\ 1\ 2\ 7]$:

i	1	2	3	4
$v[i]$	1	2	7	+inf

Longest Increasing Subsequence

- For example, use $a[] = 8\ 1\ 2\ 7\ 5\ 6\ 3\ 4$
- We record the smallest k^{th} element in any increasing subsequence

For $[8\ 1\ 2\ 7\ 5]$:

i	1	2	3	4
$v[i]$	1	2	5	+inf

Longest Increasing Subsequence

- For example, use $a[] = 8\ 1\ 2\ 7\ 5\ 6\ 3\ 4$
- We record the smallest k^{th} element in any increasing subsequence

For [8 1 2 7 5 6]:

i	1	2	3	4
v[i]	1	2	5	6

Longest Increasing Subsequence

- For example, use $a[] = 8\ 1\ 2\ 7\ 5\ 6\ 3\ 4$
- We record the smallest k^{th} element in any increasing subsequence

For $[8\ 1\ 2\ 7\ 5\ 6\ 3]$:

i	1	2	3	4
v[i]	1	2	3	6

Longest Increasing Subsequence

- For example, use $a[] = 8\ 1\ 2\ 7\ 5\ 6\ 3\ 4$
- We record the smallest k^{th} element in any increasing subsequence

For $[8\ 1\ 2\ 7\ 5\ 6\ 3\ 4]$:

i	1	2	3	4
v[i]	1	2	3	4

Longest Increasing Subsequence

- Finally, we just need to check the length of v , it will be the answer.
- Note, the array v may not be an LIS (see the second last step in the example), so if you are asked to output an LIS, don't just output v directly
- How can we find which cell to update quickly?
 - Use binary search

Longest Increasing Subsequence

- In C++, it will be very handy using `std::vector` and `std::lower_bound` or `std::upper_bound` (try and look for the differences yourself)
- The code can be as short as follows:

```
for (int i = 0; i < n; i++) {
    auto it = lower_bound(v.begin(), v.end(), a[i]);
    if (it == v.end()) {
        v.push_back(a[i]);
    } else {
        *it = a[i];
    }
}
printf("%d\n", v.size());
```

S011 Activities

- You have N activities to join.
- The i^{th} activity start at S_i and end at E_i
- For any two activities, you can join both of them iff they do not overlap
- Find the max. no. activities you can join

S011 Activities

- 4 ways to greedy:
 - (1) Attend the activity with smallest starting time
 - (2) Attend the activity with smallest ending time
 - (3) Attend the activity with smallest conflicts
 - (4) Attend the activity with shortest interval
- Three of them are wrong, do you know which three?
- What are the counter examples?

S011 Activities

- The correct answer is (2) attend the activity with smallest ending time
- We first sort the activities with ascending end time.
- Then we will attend the first one. (local optimal step)
- So we cannot join any other activities before the one end, we just eliminate those we cannot join. (Reduce to sub-problem)
- Repeat until there are no activities left. (base case)

S011 Activities

- The code looks like this:

```
13 bool cmp(activity x, activity y) {
14     return x.e < y.e;
15 }
16
17 bool cmp2(activity x, activity y) {
18     return x.i < y.i;
19 }
20
21 int main() {
22     cin >> n;
23     for (int i = 1; i <= n; i++) {
24         cin >> a[i].s >> a[i].e;
25         a[i].i = i;
26     }
27     sort(a + 1, a + n + 1, cmp);
28     for (int i = 1, j = 0; i <= n; i++) {
29         if (a[i].s < j) continue;
30         j = a[i].e;
31         a[i].u = true;
32         c++;
33     }
34     sort(a + 1, a + n + 1, cmp2);
35     cout << c << endl;
36     for (int i = 1; i <= n; i++) if (a[i].u) cout << i << endl;
37 }
```

S011 Activities

- Why is this optimal?
- Consider only 2 activities, if we cannot join both, which should we join?
- We should join the one which ends earlier, because this can help us reserve more time to join other activities
- By applying the same argument to all pairs of activities, we know that always attending the one with smallest ending time is the best.

M1713 Biscuit Clicker

- When the production rate is P , Alice will get a biscuit every $1/P$ seconds
- The basic production rate is 1
- There are N upgrades where every upgrade can be bought at most once
- For the i^{th} upgrade, the cost is C_i and it multiply the production rate by P_i
- Find **a** fastest way to collect K biscuits in the bank

M1713 Biscuit Clicker

- If we consider only two upgrades
- We will buy upgrade i before j if

$$C_i + \frac{C_j}{P_i} \leq C_j + \frac{C_i}{P_j}$$

- If we consider only one upgrade
- We will buy upgrade i if

$$C_i + \frac{K}{P_i} \leq K$$

- Actually, if we sort according to the first inequality and pick upgrades according to the second inequality, this will give us the full solution

M1713 Biscuit Clicker

- The code looks like this:

```
10 bool cmp(int x, int y) {
11     return c[x] + 1. * c[y] / p[x] < c[y] + 1. * c[x] / p[y];
12 }
13
14 int main() {
15     cin >> n >> k;
16     for (int i = 0; i < n; i++) f[i] = i;
17     for (int i = 0; i < n; i++) cin >> c[i] >> p[i];
18     for (int i = 0; i < n; i++)
19         if (c[i] + 1. * k / p[i] < k) {
20             u[i] = true;
21             tt++;
22         }
23     sort(f, f + n, cmp);
24     cout << tt << endl;
25     for (int i = 0; i < n; i++) if (u[f[i]]) cout << f[i] + 1 << " ";
26     cout << endl;
27 }
```

M1713 Biscuit Clicker

- The most important part is to see the transverse property of the first inequality

$$C_i + \frac{C_j}{P_i} \leq C_j + \frac{C_i}{P_j}$$

- If we buy upgrade i before upgrade j, and if we buy upgrade j before upgrade k, then we should buy upgrade i before upgrade k
- If this property holds, then we may try to use greedy to solve the problem
 - First, this tell us that we can sort the upgrades
 - Second, we can easily show if there are any inversions, it will not be the optimal solution

Conclusion

- When can we use greedy?
 - First, when you think you can (TRUE!!!)
 - Second, you cannot find counter example (You should try to do so)
 - The system gives you full feedback (it worth trying)
 - As you can see, greedy algorithm is not long, so it worth trying if you can code fastly
 - points are given per subtask and you have no idea (it worth trying)
 - Greedy solution sometimes give extremely good approximation for the optimal solution
- It is not realistic to prove the correctness during the competition
 - Just prove in mind for a little bit is enough

Suggested Tasks

- 01014 Stamps
- D109 Giving changes
- M0632 Machine Scheduling
- M0633 Children's Game
- J044 Amazing Robot
- J064 Cave Adventures
- J154 Father's Will

Extra: M0633 Children's Game

- Idea similar to M1713 Biscuit Clicker, think of what will happen if $N = 2$
- The code looks like this:

```
1 #include<iostream>
2 #include<string>
3 #include<algorithm>
4
5 using namespace std;
6
7 int n;
8 string a[55];
9
10 bool cmp(string x, string y) {
11     return [REDACTED]
12 }
13
14 int main() {
15     cin >> n;
16     for (int i = 0; i < n; i++) cin >> a[i];
17     sort(a, a + n, cmp);
18     for (int i = 0; i < n; i++) cout << a[i];
19     cout << endl;
20 }
```