

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

Introduction

COMP4128 Programming Challenges

School of Computer Science and Engineering
UNSW Australia

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

1 Admin

2 Classes

3 Assessment

4 Competitions and Practice

5 Solving Problems

6 Example Problems

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- Ray Li (z3424810)
- Declan McDonnell (z5163949)
- Michael Chen (z5118285)
- Sahan Fernando (z5113187)
- LIC: Aleks Ignjatovic (ignjat)

Consultation: feel free to email, or catch me after lectures. I don't have a room :(

If you want a longer consultation, best to email me to set up a time first.

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- To learn algorithms and data structures
- To practice fundamental problem solving ability
- To practice your implementation and general programming skills
- To prepare for programming competitions

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- Significant programming experience in a C-like programming language
- Understanding of fundamental data structures and algorithms:
 - Arrays, structs, heaps, merge sort, BSTs, graph search, etc...
- COMP3121/3821 (coreq)
- Enthusiasm for problem solving

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- 1 Introduction
- 2 Data structures
- 3 Graph algorithms
- 4 Dynamic programming
- 5 Network flow
- 6 Strings
- 7 Computational geometry
- 8 A bit of math

There is a tentative course schedule on the website.

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- 1 Admin
- 2 Classes**
- 3 Assessment
- 4 Competitions and Practice
- 5 Solving Problems
- 6 Example Problems

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- Lectures:
 - Tuesday, 12:00 - 14:00 in Ainsworth 202
 - Wednesday, 16:00 - 17:00 in Ainsworth 102
- Labs:
 - Thursday, 12:00 - 15:00 in Drum Lab (K17 B8)
 - Thursday, 15:00 - 18:00 in Drum Lab (K17 B8)
 - Friday, 14:00 - 17:00 in Oboe Lab (J17 304)

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- Lectures for each topic will present the theory, and apply this to some example problems
- Any code in lectures will be in C++
- Slides will be available before each lecture
- Please ask questions at any time if anything is unclear

Introduction

Admin

Classes

Assessment

Competitions
and PracticeSolving
ProblemsExample
Problems

- There are three hours of lab time assigned a week, to work on problem sets
- You may take this opportunity to ask for help, or just work on the problem sets by yourself
- If you get stuck, your tutors are your best first resource
- In weeks 5 and 10 (tentatively) we will hold 2.5 hour contests during lab time

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- 1 Admin
- 2 Classes
- 3 Assessment**
- 4 Competitions and Practice
- 5 Solving Problems
- 6 Example Problems

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- Weekly problem sets: 45%
- Assignment: 15%
- Contests: 16%
- Final: 24%
- Bonus Assignment :O 4%

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- A set of 5-7 problems will be released each week (except the last)
- Links will be posted on the course website
- You have ~ 3 weeks to complete each set
- Worth 5% each, for a total of 45%
- Each problem in a set is weighted equally
- For full marks in a set, you need to complete all but one. Completing the entire set counts for 1 bonus mark

Introduction

- Solve problems with multiple steps and harder difficulty than the rest of the course.
- However, you will be given steps leading to a solution and hints.
- The goal is to help you break down and tackle a harder problem. Also cover some interesting ideas I won't have time to explicitly cover.
- Individual or in pairs
- Worth 15%
- More details will be available in the spec which will be released around week 4

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- Form a group of 3, attempt a past ACM style contest. Afterwards, with the help of the editorial, solve most of the problems and write a brief summary.
- Should be fun! Hopefully gives you some exposure to contests outside this course.
- Worth 4 bonus marks.
- More details will be on the website later in the course.

Introduction

1 Admin

Admin

Classes

2 Classes

Assessment

Competitions
and Practice

3 Assessment

Solving
Problems

4 Competitions and Practice

Example
Problems

5 Solving Problems

6 Example Problems

- ACM-ICPC
 - South Pacific Programming Competition. Divisionals September 21! Registration due *tomorrow*!
 - ANZAC League
- Big companies
 - Google Code Jam
 - Facebook Hacker Cup
 - Microsoft Coding Competition (Probably around T1 at UNSW)
 - AmazAlgo (Probably around May at UNSW)
 - Will be announced on CSESoc Facebook page.
- Online competitions
 - Atcoder
 - Codeforces
 - topcoder
 - CodeChef

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- The best practice is to solve lots of interesting problems
- Live contests
 - UNSW ACM Training. Tentatively 6pm Thursday.
 - ANZAC League
- Online problem sets and competitions
 - AtCoder, Codeforces, TopCoder, CodeChef
 - USACO, ORAC
 - Project Euler
- Or ask me or your tutor

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- 1 Admin
- 2 Classes
- 3 Assessment
- 4 Competitions and Practice
- 5 Solving Problems**
- 6 Example Problems

- Read the problem statement
 - Check the input and output specification
 - Check the constraints
 - Check for any special conditions which might be easy to miss
 - Check the sample input and output
- Reformulate and abstract the problem away from the flavour text

- Design an algorithm to solve the problem
 - It must be *both* correct and fast enough - complexity analysis
 - If you know your algorithm is not correct or too slow, then there is no point implementing or submitting it
 - Modern computers can handle about 200 million primitive operations per second
- Implement the algorithm
 - Debug the implementation
- Submit!

- **Problem statement** Alice and Bob are two friends who are visiting a milk bar. The milk bar is owned by the crotchety old Mr Humphries. If Alice buys A dollars worth of items and Bob buys B dollars, how much must they pay in total?
- **Input** Two integers, A and B ($0 \leq A, B \leq 10$)
- **Output** A single integer, the total amount Alice and Bob must pay.

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- **Problem** Output $A + B$
- **Algorithm** Calculate $A + B$, and then print it out.

- **Complexity** $O(1)$ time and $O(1)$ space
- **Implementation**

```
#include <iostream>

int main() {
    // read input
    int a, b;
    cin >> a >> b;

    // compute and print output
    cout << (a + b) << '\n';
    return 0;
}
```


- **Problem statement** Given an array of positive integers S and a window size k , what is the largest sum possible of a contiguous subsequence (a *window*) with exactly k elements?
- **Input** The array S and the integer k
($1 \leq |S| \leq 1,000,000$, $1 \leq k \leq |S|$)
- **Output** A single integer, the maximum sum of a window of size k

- **Algorithm 1** We can iterate over all size k windows of S , sum each of them and then report the largest one
- **Complexity** There are $O(n)$ of these windows, and it takes $O(k)$ time to sum a window. So the complexity is $O(nk)$. So we will need roughly around 1,000,000,000,000 operations in the worst case.
- This is way bigger than our 200 million figure from before! We need a way to improve our algorithm.

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- What are we actually computing?
- For some window beginning at position i with a window size k , we are interested in $S_i + S_{i+1} + \dots + S_{i+k-1}$

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- Let's look at an example with $k = 3$
- We compute:
 - $S_0 + S_1 + S_2$
 - $S_1 + S_2 + S_3$
 - and so on

- **Algorithm 2** Instead of computing the sum of each window from scratch, we can use the sum of the previous window and just subtract off the first element, then add our new element to obtain the correct sum.
- To calculate $W_i (= S_i + S_{i+1} + \dots + S_{i+k-1})$, we can instead just do $W_{i-1} - S_{i-1} + S_{i+k-1}$
- **Complexity** After the $O(k)$ computation of the sum of the first window, each subsequent sum can be computed in $O(1)$ time. Hence the total complexity of the algorithm is $O(k + n) = O(n)$

● Implementation

```
#include <iostream>
#include <algorithm>
using namespace std;

const int N = 1e6 + 5;
int s[N];

int main() {
    // read input
    int n, k;
    cin >> n >> k;
    for (int i = 0; i < n; i++) cin >> a[i];

    long long ret = 0, sum = 0;
    for (int i = 0; i < n; i++) {
        // remove a[i-k] if applicable
        if (i >= k) sum -= s[i-k];
        // add a[i] to the window
        sum += s[i];

        // if a full window is formed, and it's the best so far, update
        if (i >= k - 1) ret = max(ret, sum);
    }

    // output the best window sum
    cout << ret << '\n';
    return 0;
}
```

- **Problem statement** In chess, a queen is allowed to move any number of squares horizontally, vertically or diagonally in a single move. We say that a queen *attacks* all squares in her row, column and diagonals.

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| | | ★ | | | ★ | | |
| | | | ★ | | ★ | | ★ |
| | | | | ★ | ★ | ★ | |
| ★ | ★ | ★ | ★ | ★ | Q | ★ | ★ |
| | | | | ★ | ★ | ★ | |
| | | | ★ | | ★ | | ★ |
| | | ★ | | | ★ | | |
| | ★ | | | | ★ | | |

- For $N \geq 4$, it is always possible to place N queens on an N -by- N chessboard so that no pair attack each other.

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| | | | | | Q | | |
| | | | Q | | | | |
| | | | | | | Q | |
| Q | | | | | | | |
| | | | | | | | Q |
| | Q | | | | | | |
| | | | | Q | | | |
| | | Q | | | | | |

- **Input** The integer $4 \leq N \leq 12$
- **Output** For each valid placement of queens, print out the sequence of column numbers, i.e. the column of the queen in the first row, the column of the queen in the second row, etc., separated by spaces and on a separate line.
- **Sample** For $N = 6$, the output should be:

```
2 4 6 1 3 5
3 6 2 5 1 4
4 1 5 2 6 3
5 3 1 6 4 2
```

- **Algorithm 1** Try placing queens one row at a time. The easiest way to do this is through recursion (sometimes this is called "recursive backtracking").
- We place queens one row at a time, by simply trying all columns, and then recurse on the next row. When N queens have been placed, we check whether the placement is valid.
- Complexity? Naively there are N^N placements of queens to check. We need to check if this queen duplicates any column or diagonal. This check takes $O(N)$ time.
- Thus the naïve algorithm takes $O(N^{N+1})$ time, which will run in time only for N up to 8.
- How can we improve on this?

- We need to cut down the search space; N^N is simply too large for $N = 12$.
- Many of the possibilities considered earlier fail because of conflicts within the first few rows — indeed, a single pair of conflicting queens in the first two rows could rule out N^{N-2} of the possibilities.
- Add pruning! Only recurse on *valid* placements, and simply discarding positions that fail before the last row.

- **Algorithm 2** We place queens one row at a time, by trying all *valid* columns, and then recurse on the next row. When N queens have been placed, we print the placement.
- Unfortunately, as is typical of backtracking algorithms like this, it is difficult to formulate a tight bound for the number of states explored; there are theoretically up to

$$\frac{N!}{N!} + \frac{N!}{(N-1)!} + \dots + \frac{N!}{0!} < N \times N!$$

states, but in practice most of these are invalid. The true numbers turn out to be as follows:

| N | 8 | 9 | 10 | 11 | 12 |
|--------|-------|-------|--------|---------|----------|
| states | 15720 | 72378 | 348150 | 1806706 | 10103868 |

- Each state requires an $O(N)$ check to ensure that the last queen does not share her column or diagonal.

● Implementation Gist

```
#include <iostream>
using namespace std;

int n, a[12];

void go(int i) {
    if (i == n) {
        // we have placed all n queens legally, so print this solution
        for (int k = 0; k < n; k++) cout << a[k] + 1 << ' ';
        cout << '\n';
        return;
    }

    for (int j = 0; j < n; j++) {
        // TODO: check whether a queen can be placed at (i,j)
        if (can_place(j)) {
            a[i] = j;
            go(i+1);
        }
    }
}

int main() {
    cin >> n;
    go(0);
}
```

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- 1 Admin
- 2 Classes
- 3 Assessment
- 4 Competitions and Practice
- 5 Solving Problems
- 6 Example Problems

- **Problem statement** You are playing a 2-player game with $2 \leq N \leq 1000$ rounds. You and your opponent have N different cards numbered from 1 to N . In round i , each player picks an unplayed card from their hand. The player with the higher card wins i points (no points are given for draws).
Through "psychology" you know exactly what cards your opponent will play in each round. What is your maximum possible margin of victory?
- **Input** An integer N and a permutation of 1 to N , the i -th value is the card your opponent plays in the i -th round.
- **Output** A single integer, your maximum margin of victory assuming optimal play.
- **Source** Orac

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- **Example Input**

3
3 1 2

- **Example Output 4**

- **Explanation:** Play 1 2 3. You lose the first round (-1) but win the second and third (+2, +3).

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- Brute force? There are $N!$ possible play orders.
- But maybe we can eliminate many of these play orders as suboptimal.
- For this, it helps to imagine what a possible play order could look like.

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- Consider the round where the opponent plays card N .
- In such a round, we can either draw (play card N too) or lose.
- If we lose, which card should we play?
- May as well play our worst card, 1.
- But now we can win every other round!

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- Okay, how about the play patterns where we play card N and draw?
- Then it's like we're repeating the problem with $N - 1$ in place of N .
- Unrolling this recursion, we now see, we can assume our play pattern is:
 - Pick a number i .
 - Draw all rounds with opponent card $> i$.
 - Lose the round with card i .
 - Win all rounds with cards $< i$.
- Only N play patterns! Can simulate each in $O(N)$. Total $O(N^2) = O(1,000,000)$.

● Implementation

```
#include <bits/stdc++.h>
using namespace std;

const int MAXN = 1005;
int N, opp[MAXN];

int main() {
    cin >> N;
    for (int i = 0; i < N; i++) cin >> opp[i];
    int ans = 0;
    for (int i = 1; i <= N; i++) {
        // draw > i, lose round i, win rounds < i
        int cur = 0;
        for (int j = 0; j < N; j++) {
            if (opp[j] == i) cur -= j+1;
            if (opp[j] < i) cur += j+1;
        }
        ans = max(ans, cur);
    }
    cout << ans << '\n';
    return 0;
}
```

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- **Moral:** One way to eliminate states is figure out conditions "good" states must satisfy. For this, it helps to consider a problem from different angles.
- Other angles would have worked too. E.g: one could have considered the round the opponent plays card 1, or the round you played card N , etc...

- **Problem statement** You have a list of closed intervals, each with an integer start point and end point. For reasons only known to you, you want to stab each of the intervals with a knife. To save time, you consider an interval stabbed if you stab any position that is contained within the interval. What is the minimum number of stabs necessary to stab all the intervals?
- **Input** The list of intervals, S . $0 \leq |S| \leq 1,000,000$ and each start point and end point have absolute values less than 2,000,000,000.
- **Output** A single integer, the minimum number of stabs needed to stab all intervals.

Introduction

Admin

Classes


Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- **Example**



- The answer here is 3.

Introduction

Admin

Classes

Assessment

Competitions and Practice


Solving Problems

Example Problems

- How do we decide where to stab? State space is again laughably big.
- Again let's ask ourselves if we can eliminate many of the stab possibilities.
- Focus on a single stab for now.

- **Observation 1:** We can move it so it is an end point of an interval without decreasing the set of intervals we stab.
- **Proof:** Consider any solution where there is a stab *not* at the endpoint of an interval. Then we can create an equivalent solution by moving that stab rightwards until it hits an end point.

- Now let's try drawing sample data and consider moving from left to right. Where do we put our first stab?



- **Observation 2:** By Observation 1, we may assume it is at the first endpoint.

- **Algorithm 1** Stab everything that overlaps with the first end point. Then, remove those intervals from the intervals to be considered, and recurse on the rest of the intervals.
- **Complexity** There are a few different ways to implement this idea, since the algorithm's specifics are not completely defined. But there is a simple way to implement this algorithm as written in $O(|S|^2)$ time.

- If we look closely at the recursive process, there is an implicit order in which we will process the intervals:
ascending by end point
- If we sort the intervals by their end points and can also efficiently keep track of which intervals have been already stabbed, we can obtain a fast algorithm to solve this problem.

- Given all the intervals sorted by their end points, what do we need to keep track of? **The last stab point**
- Is this enough? How can we be sure we haven't missed anything?
- Since we always stab the next unstabbed end point, we can guarantee that there are *no unstabbed intervals* that are *entirely* before our last stab point.
- For each next interval we encounter (iterating in ascending order of end point), that interval can start before or on/after our last stab point.
- If it starts before our last stab point, it is already stabbed, so we ignore it and continue.
- If it starts after our last stab point, then it hasn't been stabbed yet, so we should do that.

- **Algorithm 2** Sort the intervals by their end points. Then, considering these intervals in increasing order, we stab again if we encounter a new interval that doesn't overlap with our right most stab point.
- **Complexity** For each interval, there is a constant amount of work, so the main part of the algorithm runs in $O(|S|)$ time, $O(|S| \log |S|)$ after sorting.

● Implementation

```
#include <iostream>
#include <utility>
#include <algorithm>
using namespace std;

const int N = 1001001;
pair<int, int> victims[N];

int main() {
    // scan in intervals as (end, start) so as to sort by endpoint
    int n;
    cin >> n;
    for (int i = 0; i < n; i++) cin >> victims[i].second >> victims[i].first;
    sort(victims, victims + n);

    int last = -2000000001, res = 0;
    for (int i = 0; i < n; i++) {
        // if this interval has been stabbed already, do nothing
        if (victims[i].second <= last) continue;
        // otherwise stab at the endpoint of this interval
        res++;
        last = victims[i].first;
    }

    cout << res << '\n';
    return 0;
}
```

Introduction

Admin

Classes

Assessment

Competitions
and Practice

Solving
Problems

Example
Problems

- **Moral:** Sorting into a sensible order is often helpful. As is drawing pictures.
- I often find it helpful to play with a problem on paper and see how I would solve it manually.

- **Problem statement** There are $N \leq 2000$ countries, the i -th has $a_i \leq 20$ delegates.
There are $M \leq 2000$ restaurants, the i -th can hold $b_i \leq 100$ delegates.
For "synergy" reasons, no restaurant can hold 2 delegates from the same country.
What's the minimum number of delegates that need to starve?
- **Input** An integer N , N integers a_i . An integer M , M integers b_i .
- **Output** A single integer, the minimum number of delegates that need to starve.
- **Source** Orac

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- **Example Input**

3

4 3 3

3

5 2 3

- **Example Output 2**

- **Explanation:** Someone from the first country starves.
Furthermore, the second restaurant has too few seats.

Introduction

Admin

Classes

Assessment

Competitions and Practice


Solving Problems

Example Problems

- Yet again, trying all assignments is laughably slow. So again, let us try to think about what conditions a good assignment may have?
- Makes sense to consider all delegates of a country at once so we don't have to keep track of who has been assigned where.
- Consider the countries in any arbitrary order. Suppose "Australia" is the first country we are considering.

- **Observation 1:** We should assign as many delegates as possible.
- **Proof:** In any solution that does not, there is some restaurant with no Australian delegates and there is a starving Australian delegate.
- We can then kick out any delegate for an Australian delegate without making the solution any worse.
- But where should we assign the Australian delegates?
- Our main objective is to make it easier to seat the other country's delegates.
- From some extreme examples, the bottleneck seems to be the restaurants with few seats.

- **Observation 2?** We should assign delegates to the restaurants with the most seats remaining.
- **Proof:** Again, consider a solution that does not.
- Then we skip restaurant i for a restaurant j where $b_i > b_j$.
- But this means we can swap some delegate from restaurant i with the Australian delegate in j while preserving uniqueness.



- By repeating these swaps, we obtain a solution just as optimal except **Observation 2** was obeyed.

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- Hence we may consider just solutions where Australia's delegates are assigned to the restaurants with the most seats remaining.
- Now repeat all other countries in the same manner.
- One easy way to implement: Sweep through the countries one by one. For each country, sort the restaurants in decreasing capacity order and assign to them in that order.

● Implementation

```
#include <bits/stdc++.h>
using namespace std;

const int MAXN = 2005, MAXM = 2005;
int N, numDelegates[MAXN], M, numSeats[MAXN];

int main() {
    cin >> N;
    for (int i = 0; i < N; i++) cin >> numDelegates[i];
    cin >> M;
    for (int i = 0; i < M; i++) cin >> numSeats[i];
    int starved = 0;
    for (int i = 0; i < N; i++) {
        int delegatesRemaining = numDelegates[i];
        sort(numSeats + i, numSeats + i + M, greater<int>());
        for (int j = 0; j < M; j++) {
            if (numSeats[j] > 0 && delegatesRemaining > 0) {
                numSeats[j]--;
                delegatesRemaining--;
            }
        }
        starved += delegatesRemaining;
    }
    cout << starved << '\n';
    return 0;
}
```

Introduction

Admin

Classes

Assessment

Competitions and Practice

Solving Problems

Example Problems

- **Complexity?** $O(N)$ countries. For each we sort a M length list and then a linear sweep.
- $O(NM \log M) \approx O(4\text{mil} \cdot 11)$, fast enough.
- **Moral:** One way to make observations is think abstractly about what should hold. Often this is guided by examples.
- Once you have some guess, you can try to prove it after.