Getting Started	
Running Time C++ Debugging	Getting Started COMP4128 Programming Challenges
	School of Computer Science and Engineering UNSW Sydney

Term 2, 2025

	Table of Contents	2
Getting Started		
Running Time C++ Debugging	1 Running Time	
	2 C++	
	3 Debugging	

Complexity and Time Limit

Getting Started

Running Time

Debugging

- Your solution must give the correct output for each possible input, but it must also run within the specified time limit
- If you know your algorithm is not correct or too slow, then there is no point implementing or submitting it
- You can assess whether your algorithm is fast enough using complexity analysis
 - Calculate the number of states your algorithm will enter, and multiply by the amount of work performed in each state
 - Sometimes more sophisticated techniques are required, e.g. recursive algorithms
 - Your solution will not be accepted if it times out on even one test case, so assume the worst case input

Complexity and Time Limit

Getting Started

- Modern computers can handle several hundred million primitive operations per second
 - $\bullet~\text{About}~6\times 10^8$ integer operations is usually safe
 - $\bullet\,$ Maybe only 2×10^8 floating-point operations
- In some easy problems, the naïve algorithm will run in time
- If not, you can use a variety of techniques to reduce the number of states or the amount of work per state
- We'll see more advanced methods in future topics, e.g. data structures

Getting Started

- **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 $(1 \le |S| \le 1,000,000, 1 \le k \le |S|)$
- **Output** A single integer, the maximum sum of a window of size *k*

Getting Started

- 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 operations in the worst case.
- This is way bigger than our 200 million figure from before! We need a way to improve our algorithm.



size k, we are interested in $S_i + S_{i+1} + \ldots + S_{i+k-1}$

Getting Started

- Running Time
- Debugging

- 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

Getting Started

- Running Time C++ Debugging
- Algorithm 2 Instead of computing the sum of each window from scratch, we can modify the sum of the previous window.
- To calculate

$$W_i=S_i+S_{i+1}+\ldots+S_{i+k-1},$$

we can instead evaluate

$$W_i = W_{i-1} - S_{i-1} + S_{i+k-1}.$$

• **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), which we can simplify to O(n) as $n \ge k$.

Getting Started

Running Time

```
Implementation
```

```
#include <iostream>
#include <algorithm>
using namespace std:
const int N = 1001001;
int a[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 \ge k) { sum -= a[i-k]: }
    // add a[i] to the window
    sum += a[i];
    // if a full window is formed, and it's the best so far, update
    if (i \ge k - 1) { ret = max(ret, sum); }
  3
  // output the best window sum
  cout << ret << '\n':
```

10

Getting Started

Running Time

Debugging

• 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	*	*
				*	*	*	
			*		*		*
		*			*		
	*				*		



Running Time

C++

Debugging

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



Getting Started

- **Problem statement** Given a board size *N*, list all the ways of placing *N* queens so that no two attack each other.
- Input An integer $4 \le N \le 12$
- **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, in lexicographic order.
- **Sample** For N = 6, the output should be:

Getting Started

- Running Time C++ Debugging
- Algorithm 1 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.
- There are *N* squares for the queen in each row, so if we simply consider all possibilities, there are *N*^N placements of queens to check.
- Each placement must be checked for duplicates in any column or diagonal (note that we have already assigned exactly one queen per row). 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?

- Running Time C++ Debugging
- 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.
- We could improve by only recursing on *valid* placements, and simply discarding positions that fail before the last row.

Getting Started

Running Time C++ Debugging

- 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

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

states, but in practice most of these are invalid.

Getting Started

Running Time

C++

Debugging

• The true numbers turn out to be as follows:

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

• Each state then requires an O(N) check to ensure that the last queen has been placed legally, by scanning her column and diagonals.



Running Time

Implementation

```
#include <iostream>
using namespace std:
int n, a[12];
void print queens() {
 for (int k = 0; k < n; k++) {
    cout << a[k] + 1 << ' ':
  3
  cout << '\n';
}
bool check_queen (int i, int j) {
  for (int k = 0; k < i; k++) {
   if ((a[k] == j) || (i - k == a[k] - j) || (i - k == j - a[k])) {
      return false;
    3
  }
  return true;
```

void go(int i) {

Getting Started

• Implementation (continued)

```
Running Time
```

```
if (i == n) {
   // we have placed all n queens legally, so print this solution
   print queens();
   return:
 }
 for (int j = 0; j < n; j++) {
   // check whether a queen can be placed at (i,j)
   if (check queen(i,j)) {
     // place queen and recurse
      a[i] = j;
      go(i+1);
   3
 }
int main() {
 cin >> n:
 go(0):
```

Getting Started

- This problem is instructive, but in practice not a great contest problem.
- Why? Because there are only nine possible test cases!
- This allows you to hard-code the answer in your source code.
- But how do you obtain the answer in order to hard-code it?
- You still have to write a solver, but you could let it run locally for several minutes per test case while you work on another problem.

	Table of Contents	21
Getting Started		
Running Time ++ Debugging	Running Time	
	2 C++	
	3 Debugging	



Getting Started

- Rule of thumb: whatever workflow you used for C programming will (almost) work for this course
- Write code in an editor or IDE of your choice
- Compile and run locally
 - Save samples as 1.in/1.ans etc
 - Then run ./a < 1.in > 1.out; diff -w 1.out 1.ans
 - Of course, diff is only helpful if the output has to match exactly
- Or work from your CSE account using SSH or VLAB

Compilers

Getting Started

Running Tim

Debugging

- Use g++, the GNU C++ Compiler
 - Includes some libraries (e.g. Policy-Based Data Structures) that we make use of
 - Any semi-recent version is fine
 - Allows #include<bits/stdc++.h> to catch all standard libraries
 - Sometimes throws in headers "for free"; if your code doesn't compile on the judge this might be why
- Makefile: see tips page of the course website.
- \bullet DCC Help and DCC Sidekick work for C++, not just C

Compiler Installation

Getting Started

Running Tim

Debugging

• Linux users: you probably know what you're doing

• Mac users: default C++ compiler is Clang, so to work locally you will want to install GCC (e.g. via Homebrew)

• Windows users: probably set up WSL, works nicely enough with VSCode

C++ Documentation

Running Ti C++ Most of the C programs you wrote in COMP1511 and COMP2521 are valid C++ as is

• You can find C++ documentation at cppreference.com

• This is the only online resource that you can use during the contests and exam, other than the course website and associated platforms



Getting Started

- Running Time
- Debugging

• The original C++ standard was C++98

• All problems in the weekly problem sets will be judged using C++11 or more recent versions

 \bullet In the contests, your submissions will be compiled using the C++20 standard

C++11 Features for Competitive Programming 27

Getting Started

Running Time C++

- C++11 has lots of nice things for us
- Very important:
 - long long width
- Situational
 - unordered_set and unordered_map
 - random
 - tuple
- Conveniences
 - auto type
 - range-based for loops
 - various functions in <algorithm>

C++14/17/20 Features for Competitive Programming

28

Getting Started

- Running Time
- Debugging

- $\bullet\,$ There's much less in the later C++ standards that affects us
- C++14: nothing
- C++17:
 - tuple
 - gcd and lcm
- C++20:
 - <compare> for three-way comparison <=>
 - <bit> for bitmask manipulation
 - <numbers> for numerical constants
 - .contains() in common containers
- Almost all are only minor conveniences

Naming conflicts

Getting Started

- You might want to use count as a function name, but the <algorithm> header also defines a function named count.
- When you call count(), how does the compiler know which one you are calling?
- Two C++ features are relevant here.
- Unlike C, C++ allows overloading of function names. If two functions with the same name have different signatures (e.g. different number of parameters, differently typed parameters), then the compiler may be able to resolve the conflict automatically.

Getting Started

- Running Time C++ Debugging
- A namespace defines a named scope, within which you can define variables, functions and classes.
- You can then refer to those entities by prepending the name of the namespace, differentiating them from other entities with the same name that might be defined elsewhere in the project.
- In the example from the previous slide, the <algorithm> header is part of the C++ standard library, so all functions in this header belong to the std namespace.
- Therefore there is no conflict; count() refers to the local function, whereas std::count() refers to the one from <algorithm>.

Getting Started

- Running Time C++ Debugging
- A common theme in competitive programming is that in competitive programming, you can (and often should) use lot of shortcuts and hacks that would be considered bad style or worse in other settings.
- Namespaces are useful, but they make us type five extra characters whenever we want to use something from the standard library, which is a nuisance.
- Instead, we include the global directive using namespace std; immediately after our #include directives. This tells the compiler to try looking in the std namespace whenever it tries to resolve a variable or function name.
 - Global using directives are a bad idea in most contexts! Don't do this in larger projects.

Getting Started

Running Tim C++

- Now, we can write cout << ans << endl; rather than std::cout << ans << std::endl;!
 - Surely the precious seconds saved will give us the edge to win a contest.
- The downside of this is that we reintroduce naming conflicts.
 - If we name a variable or function as count, rank or various other common words which name entities in std, the compiler might complain.
 - No matter, we will just use names like cnt and rnk instead if necessary.

```
Getting
  Started
                #include <algorithm>
                #include<cassert>
                #include <iostream>
                using namespace std;
C++
                // count nonzero array entries
                const int N = 100100:
                int a[N];
                int main (void) {
                  int n:
                  cin >> n;
                  int count = 0;
                  for (int i = 0; i < n; i++) {
                    cin >> a[i];
                    if (a[i] != 0) { count++; }
                  3
                  // check our answer
                  assert(count == n-count(a, a+n, 0)); // conflict!
                  cout << count << '\n':
```



```
Getting
                #include <algorithm>
  Started
                #include<cassert>
                #include<iostream>
                // count nonzero array entries
C++
                const int N = 100100;
                int a[N]:
                int main (void) {
                  int n:
                  std::cin >> n:
                  int count = 0;
                  for (int i = 0; i < n; i++) {
                    std::cin >> a[i];
                    if (a[i] != 0) { count++; }
                  }
                  // check our answer
                  std::assert(count == n-std::count(a,a+n,0)); // no conflict
                  std::cout << count << '\n';</pre>
                }
```

• This is the 'good' way to resolve the naming conflict, by omitting the using directive that caused the problem.

#include <algorithm>

Getting Started

C++

```
#include<cassert>
#include <iostream>
using namespace std;
// count nonzero array entries
const int N = 100100;
int a[N]:
int main (void) {
  int n;
  cin >> n:
 int cnt = 0;
 for (int i = 0; i < n; i++) {
    cin >> a[i];
   if (a[i] != 0) { cnt++; }
  }
 // check our answer
  assert(cnt == n-count(a,a+n,0)); // no conflict
  cout << cnt << '\n';</pre>
}
```

 This is the lazy way to resolve the naming conflict, by renaming our local variable.

Input and Output

Getting Started

- In every problem in the contests and almost every problem in the weekly problem sets, input is from the standard input stream stdin and output is from the standard output stream stdout.
 - We will provide specific instructions for any problem that requires you to read from and write to files.
- C-style I/O (scanf, printf, getline, and so on) is supported in the <cstdio> header.
- C++-style I/O (cin, cout) is supported in the <iostream> header.

Input

Getting Started

```
C++
```

- If we know the format of the input (as is the case in contest problems), we can reliably extract it from cin using the >> operator.
- This is very convenient; it's fast to type, and doesn't require us to think about pointers and addresses.
- We can read an entire line into a string s using getline(cin,s).
- There is a trade-off in speed. Reading input from cin is quite a bit slower than using scanf().
 - If you need to read 10MB per second, consider including the following statements:

```
|| cin.tie(nullptr); // prevents cout from flushing on every cin read
cin.sync_with_stdio(false); // unsyncs iostream from cstdio
```

• or just revert to scanf().

Getting Started

Running Time

• We can similarly print output to cout using the << operator.

- Almost always print newline characters as '\n' rather than endl, since the latter flushes the output buffer, making it slower.
 - Flushing is only necessary for interactive problems.

Floating-Point Output

Getting Started

- Suppose x is a double, and we run cout << x << '\n'. The default behaviour is to print to six digits of precision. This is often insufficient.
- Worse still, if x exceeds 1,000,000 in absolute value, it will be printed in scientific notation (yes, really).
- To insist on fixed-point (not scientific) notation, use the manipulator std::fixed from the <ios> header, which is included within <iostream>.
- To change the number of decimal places, use the manipulator std::setprecision() from the <iomanip> header.
- Therefore, to print to 9 decimal places, we write cout << fixed << setprecision(9) << x << '\n';. Consider just using printf() instead!

Memory

Getting Started

Running Ti C++

- Debugging
- We will (almost) never need dynamic memory in programming contests, so you don't need to know about malloc or free (or the analogous C++ operators new and delete).
- Instead, we will use the input size defined in the problem to allocate as much (or more) static memory than we could ever need.
- Good habit to declare arrays slightly larger than necessary.
 - If the length is up to 100,000, we will usually declare it with size 100,100. This is helpful in case we want to store the elements off-by-one or similar.
- In C++, array sizes must be constant expressions, so we use the keyword const.

```
const int N = 100100;
int a[N];
```

Global Variables

Getting Started

- Running Til C++
- Debugging

- Global variables are unsafe in larger projects, but we never use more than one file so they're fine for our purposes.
- This saves us some typing, since we don't have to explicitly pass them in function calls.
 - The alternative is to pass a local variable *by reference*, using & just as you would in C.
- Global arrays are stored on the heap.
 - Creating a large array (say 10⁷ integers for a prime sieve) on the heap is fine.
 - However, it might not be possible to allocate this much contiguous memory on the stack.
- A convenience: global variables are initialised to the default value (not garbage)
 - For example, a global array of integers will be 0-initialised.

Integer Types

Getting Started

- int is a 32-bit integer. Be wary of overflow.
- Since C++11, long long is a 64-bit integer. On a 64-bit processor, the performance difference is usually negligible.
- Never use long. The language standard does not enforce whether it is 32 or 64 bits.
- You won't need 128-bit or arbitrary-precision integers in this course.

Floating-Point Types

Getting Started

- Never use float. The extra precision from double is necessary in most problems, and the performance difference is again negligible.
- long double can be 64-bit or 128-bit.
 - If you are getting WA due to precision and not too close to time limit, it's worth a try.
- Problems which require you to produce floating-point output will typically allow answers within some relative and/or absolute error. This will always be specified in the output format.

String Types

Getting Started

Running Tir C++ Debugging • C-style character arrays still work, and the <cstring> header corresponds to the C library <string.h>.

• We will usually use the string type instead, which integrates directly with C++-style I/O and various extra functions provided in the <string> header.

• To access the character array underlying string s, we use s.c_str(). This is useful for C-style formatted printing.

Converting Between Types

- You can convert between types using explicit or implicit casting, as in C.
- C++ also introduces the stringstream, which you can both insert to (<<) and extract from (>>).
- This is very useful for converting between strings and other types.
 - For example, you can read a string using cin, insert it to a stringstream, and then extract integers from there.
 - In other situations, stoi() might suffice.

Pairs and Tuples

Getting Started

- Running Time C++ Debugging
- The <utility> header defines the class template pair<T1,T2>, which allows you to couple together two elements as a single unit.
- Since C++11, this has been generalised to tuple<...>, which supports any fixed-size collection of elements.
- The elements of a pair or tuple can be of any type, and they *do not* have to all be of the same type.
- This is extremely useful. For example, it lets you sort items while keeping track of their original indices, by making pairs of the form (a[i],i).

Structures

Getting Started

- C++ introduces classes, but we won't use them.
- It also supports C-style structs, which we will sometimes use.
- You can define a function as a member of a struct, which we will occasionally use in this course.
- You can also define operators such as operator== and operator< for your structs, e.g. so that you can sort a collection of them.

Algorithms

Getting Started

- The <algorithm> header provides many useful functions for working on ranges of elements.
- Many of these (e.g. fill(), search(), count(), max(), min()) are just conveniences that replace two or three lines of very simple code.
- Even so, the more code you write yourself, the more bugs you can introduce.
- Swapping values is notoriously inconvenient in C, but in C++ we can just call swap(x,y).



Getting Started

- The <algorithm> header includes a sort() function, which takes iterators (think generalised pointer) to the range to be sorted (left-inclusive, right-exclusive).
- Since C++11, this sorting algorithm is guaranteed to run in $O(n \log n)$ time in the worst case.
- This function takes an optional third argument, in which you can provide a custom comparison to be used in place of the default (operator<).
 - The default comparison for pairs is to compare the first entries, with ties broken by comparing the second entries.
- For example, to sort the first n entries of an integer array a in *descending* order, we can call sort(a,a+n,greater<int>()).

Example problem: Sodor Census

Getting Started

Running Tim C++ Debugging • **Problem statement** Topham has *n* engines, numbered from 1 to *n*. Topham is fond of his engines, so he has given each a distinct name.

If the engines were to arrange themselves in lexicographic order of their names, what sequence of numbers could be read from them?

- Input An integer n (1 ≤ n ≤ 100,000), followed by n strings representing the names of engines 1, 2, ..., n.
- **Output** A permutation of $\{1, \ldots, n\}$.
- Sample Input

4 Thomas Edward Henry Gordon

- - -

Example problem: Sodor Census

Getting Started

• Algorithm Make an array of pairs (*name*, *number*), sort, and read off the second coordinates.

Implementation

```
#include <algorithm>
#include <iostream>
using namespace std:
const int N = 100100;
pair<string, int> engines[N]; // (name, original index)
int main (void) {
    int n:
    cin >> n:
    for (int i = 0; i < n; i++) {
        cin >> engines[i].first;
        engines[i].second = i+1; // 1-based
    3
    sort(engines. engines + n): // sorts by increasing name
                                // and if necessary increasing index
    for (int i = 0; i < n; i++) {
        cout << engines[i].second << (i == n-1 ? '\n' : ' '):
    3
```

Shuffling and Searching



Running Time

C++

• Sometimes, sorted input is the worst case! Use shuffle().

• <algorithm> also helps you avoid writing some binary searches from scratch. We'll discuss binary_search() and related functions in the Paradigms lecture.

Next Permutation

Getting Started

- next_permutation() rearranges the elements in a range into the lexicographically next greater permutation.
- For example, it would transform the character array "permutation" to "permutatnio".
- It returns true if such a permutation exists, or false if the original permutation was already lexicographically greatest.
- An individual call to next_permutation() could take linear time, but the complexity is *amortized constant*!
- Why do we care about this peculiar function?

Next Permutation

Getting Started

Running Time C++ • next_permutation() helps us run exhaustive search (i.e. brute force).

• The following snippet iterates through all bit sequences of length *n* in which exactly *k* bits are set.

```
int selected[n];
fill(selected,selected+n-k,0);
fill(selected+n-k,selected+n,1);
do {
    // ...
} while (next_permutation(selected,selected+n));
```

- Each of these sequences then corresponds to a different unordered selection of *k* items from a collection of *n* items.
- You can designate the selected items as those corresponding to set bits.

	Table of Contents	56
Getting Started		
Running Time C++ Debugging	1 Running Time	
	2 C++	
	3 Debugging	

Debugging

Getting Started

Running Time C++ Debugging • Debugging is often the most difficult, time-consuming and frustrating part of competitive programming.

- The best way to debug is not to make bugs in the first place.
 - Take your time (especially in the weekly problem sets) and read what you're typing
 - In team contests, pair programming helps

How to Debug (pre-submission)

Getting Started

- Running Tim C++
- Debugging

- Suppose you have solved the problem conceptually and implemented your algorithm, but your program fails the sample test cases. What should you do?
 - *Run unit tests* to narrow down what part of your program is malfunctioning.
 - The most primitive approach to debugging is to print heaps of logs as you go.
 - Print out the intermediate states of anything that might be important using cerr (or equivalently fprintf(stderr,...)).
 - You may also want to look into more sophisticated debugging tools.
 - GDB
 - Some IDEs have an inbuilt debugger

How to Debug (post-submission)

Getting Started

Running Time

Debugging

- Suppose your program now passes the sample test cases, but your submission was unsuccessful. What should you do?
- What verdict did you get?
 - TIME-LIMIT
 - check your complexity analysis
 - look for infinite loops
 - RUN-ERROR
 - estimate how much memory you are using
 - check for array out of bounds
 - check for accessing into or deleting from empty data structures

How to Debug (post-submission)

Getting Started

- Running Time C++ Debugging
- For WRONG-ANSWER and some RUN-ERROR verdicts, it's often difficult to even identify what the bug is.
 - Write test cases, including edge cases.
 - Force yourself to read your code carefully line-by-line, e.g. by opening it in a different editor or even printing out and annotating a hard copy.
- Write a slow (potentially brute force) algorithm that certainly produces the correct answer. You can then run it locally on medium to large cases, ignoring the time limit, and compare its answers to your other program's output.
- In the problem sets, ask for help you can discuss this with other students and your tutor.
- If necessary, rethink your algorithm. Are there any hidden assumptions that you didn't examine carefully enough?



your bugs will be minor errors, perhaps even single character fixes.

General Tips

Getting Started

- Running Time C++ Debugging
- Over the course of the term, there will probably be a couple of problems which you come very close to solving.
- This is a normal part of competitive programming, and it does average out in the long run.
- No single problem is worth many marks, so it's perfectly OK to give up on a problem sometimes.
- The problem diary gives you a chance to reflect on your efforts.
- In the weekly problem sets, make sure to also prioritise your other tasks and responsibilities.
- That said, if you can afford to sink days into a problem, there are few feelings quite as gratifying as finally receiving a CORRECT verdict on your 70th submission!