

Algorithmic Complexity

Amir Kirsh & Adam Segoli Schubert



© All rights reserved

Why this talk?

Why this talk?

Performance is the name of the game

Algorithmic Complexity @ CppCon 2021

Why this talk?

Performance is the name of the game

You all (hopefully) know that O(n) is better than $O(n^2)$

Why this talk?

Performance is the name of the game

You all (hopefully) know that O(n) is better than $O(n^2)$

But there is still important stuff that might be overlooked

Algorithmic Complexity @ CppCon 2021

Why this talk?

Performance is the name of the game

You all (hopefully) know that O(n) is better than $O(n^2)$

But there is still important stuff that might be overlooked

And...

Why this talk?

Performance is the name of the game

You all (hopefully) know that O(n) is better than $O(n^2)$

But there is still important stuff that might be overlooked

And... the academic answer isn't always the practical answer

Algorithmic Complexity @ CppCon 2021





Adam Segoli Schubert

Software consultant and programming instructor.

Part of Dev Advocate Office at

Collaborates and passionate about projects which focus on decentralization, parallelization, with the objective of advancing transparency, freely available distributed knowledge, and autonomy.



Algorithmic Complexity @ CppCon 2021

Algorithmic Co	omplexity
----------------	-----------

Performance is the name of the game

Algorithmic Complexity @ CppCon 2021

Algorithmic Complexity

It's actually about something BIGGER than just performance

What's BIGGER than just performance?

Algorithmic Complexity @ CppCon 2021

Algorithmic Complexity

What's BIGGER than just performance?

Scalability

14

Computational Complexity

Computational Complexity or simply **Complexity** of an algorithm is the *amount of resources* required to run it.

- from Computational complexity in wikipedia

Resources:

The amount of time, storage, or any other resource.

Algorithmic Complexity @ CppCon 2021

Computational Complexity

 $n \rightarrow f(n)$

n - is size of input

f(n) - is the amount of resources required to run the algorithm

Time - number of required elementary operations Often denoted by T(n) or t(n)

Space - Amount of memory required Often denoted by S(n) or s(n)

16



Big O Notation

We say that:

$$f(n) \in O(g(n))$$
 $[f(n) = O(g(n)) \text{ as } n \rightarrow$

 ∞

iff:

$$\exists k > 0 \exists n_0 \forall n > n_0: f(n) \leq k \cdot g(n)$$

Algorithmic Complexity @ CppCon 2021

Only the dominating factor counts

In Big O, we care about asymptotic analysis, when $n \rightarrow \infty$

Thus, for example:

 $t(n) = k_1 \cdot n \cdot \log(n) + k_2 \cdot n$

Algorithmic Complexity @ CppCon 2021

Only the dominating factor counts

In Big O, we care about asymptotic analysis, when $n \rightarrow \infty$

Thus, for example:

 $\mathbf{t}(\mathbf{n}) = \mathbf{k}_1 \cdot \mathbf{n} \cdot \mathbf{log}(\mathbf{n}) + \mathbf{k}_2 \cdot \mathbf{n}$

Algorithmic Complexity @ CppCon 2021

Only the dominating factor counts

In Big O, we care about asymptotic analysis, when $n \rightarrow \infty$

Thus, for example:

 $\mathfrak{t}(n) = k_1 \cdot \mathbf{n} \cdot \log(\mathbf{n}) + k_2 \cdot n \ \in \ O(\mathbf{n} \cdot \log(\mathbf{n}))$

Algorithmic Complexity @ CppCon 2021

Not in this talk...

Notation	Name	Short Explanation
$\Theta(g(n))$	Theta	f(n) is bounded both above and below by g(n) asymptotically
$\Omega(g(n))$	Omega	f(n) is bounded below by g(n) asymptotically
o(g(n))	Small o	f(n) is dominated by g(n) asymptotically
$\omega(g(n))$	Small omega	f(n) dominates g(n) asymptotically
$\tilde{O}(g(n))$	Tilde O	same as big O, but "ignores" logarithmic factors

Algorithmic Complexity @ CppCon 2021

Let the Charts talk



















Algorithmic Complexity @ CppCon 2021

n	O(1)	O(log n)	O(n)	O(n log n)	0(n ²)	O(n ³)	0(2 ⁿ)	O(n!)
1	1 µs	1 µs	1 µs	1 µs	1 µs	1 µs	2 µs	1 µs
10	1 µs	3 µs	10 µs	34 µs	100 µs	1 ms	1 ms	3.6 seconds
100	1 µs	6 µs	100 µs	665 µs	10 ms	1 sec	>400 trillion centuries	>googol centuries
1,000	1 µs	9 µs	1 ms	~10 ms	1 sec	16.67 min		
10,000	1 µs	13 µs	10 ms	~133 ms	1.67 min	~12 days		
100,000	1 µs	16 µs	100 ms	1.67 sec	2.78 hours	~32 years		
1,000,000	1 µs	19 µs	1 sec	~20 sec	~12 days	~32,000 years		
100,000 1,000,000	1 μs 1 μs	16 μs 19 μs	100 ms 1 sec	1.67 sec ~20 sec	2.78 hours ~12 days	~32 years ~32,000 years		

Let's see if we got it right

Are you ready for a short quiz?















Amortized Complexity

Amortized complexity considers the total worst case complexity of a sequence of operations, instead of just one operation.

Example 1: If the *total* for *n* operations is in the worst case O(n) then the *amortized complexity* is O(1)

Example 2: If the *total* for *n* operations is in the worst case $O(n^2)$ then the *amortized complexity* is O(n)

Note:

Amortized complexity is NOT the average complexity over different inputs of size n!

See: Tarjan, Robert Endre (April 1985). Amortized Computational Complexity







C++ Specifications - Complexity Requirements

In the spec (*examples*): <u>containers requirements</u> <u>unordered associative containers</u> + <u>requirements</u> <u>complexity of std::sort algorithm</u> <u>complexity of std::ranges::partition algorithm</u>

Then in CppReference (*examples*): <u>complexity of std::vector::insert</u> <u>complexity of std::list::insert</u> <u>complexity of std::unordered_map::insert</u> <u>complexity of std::search algorithm</u> <u>complexity of std::sort algorithm</u>





Back to our quiz

Algorithmic Complexity @ CppCon 2021

Back to our quiz

Are you ready?

(4) What is the Complexity of:

Sorting a vector using std::sort or std::ranges::sort / a list using list::sort

Algorithmic Complexity @ CppCon 2021

(4) What is the Complexity of:

Sorting a vector using std::sort or std::ranges::sort / a list using list::sort

O(n log(n))

See the spec for std::sort and for list::sort

Algorithmic Complexity @ CppCon 2021

(4) What is the Complexity of:	
Sorting a vector using std::sort or std::ranges::sort / a list using list::sort	
O(n log(n))	
See the spec for std::sort and for list::sort	
Note: it might be more efficient to copy the list into a vector, sort the vector, then copy back Why? See benchmark	
Algorithmic Complexity @ CppCon 2021	53

(5) What is the Complexity of:

Finding the median of *n* items

(5) What is the Complexity of:

Finding the median of *n* items

There is an algorithm, <u>PICK</u>, with O(n) worst case complexity!

However, another algorithm, <u>Quickselect</u>, which is $O(n^2)$ at worst case, is usually faster.

They are both O(n) on average.

See https://cs.stackexchange.com/questions/1914/find-median-of-unsorted-array-in-on-time

Algorithmic Complexity @ CppCon 2021

(5) What is the Complexity of:

Finding the median of *n* items

There is an algorithm, <u>PICK</u>, with O(n) worst case complexity!

However, another algorithm, <u>Quickselect</u>, which is $O(n^2)$ at worst case, is usually faster.

They are both O(n) on average.

See <u>https://cs.stackexchange.com/questions/1914/find-median-of-unsorted-array-in-on-time</u> See also <u>spec requirement for std::nth_element</u>

(6) What is the Complexity of:

find / insert - unordered_map

Algorithmic Complexity @ CppCon 2021

(6) What is the Complexity of:

find / insert - unordered_map

O(1) average case

O(n) worst case

See <u>the spec for find</u> See <u>the spec for insert</u>

(6) What is the Complexity of:

Performing equality (==) between two unordered_maps of the same type

Algorithmic Complexity @ CppCon 2021

(6) What is the Complexity of:

Performing equality (==) between two unordered_maps of the same type

O(n) average case

O(n²) worst case

See the spec

Algorithmic Complexity @ CppCon 2021

Examples: O(1)

std::list::insert (at any position)

std::list::erase (for a single iterator)

std::vector::pop_back

Examples: O(n)

std::find

std::max

std::min

Algorithmic Complexity @ CppCon 2021

Examples: O(log n)

std::binary_search

std::map::find

std::map::insert

Examples: O(n log n)

std::sort

Algorithmic Complexity @ CppCon 2021

Examples: O(n²)

Bubble sort

Algorithmic Complexity @ CppCon 2021

Examples: O(n²)

Bubble sort

Is there any reason whatsoever for using bubble sort?

Algorithmic Complexity @ CppCon 2021

Examples: O(n²)

Bubble sort

Is there any reason whatsoever for using bubble sort?

- In space complexity, maybe?
- Being stable? (what is stable sorting algorithm?)

68

Examples: O(n²)

Bubble sort

Is there any reason whatsoever for using bubble sort?

- In space complexity, maybe?
- Being stable? (what is stable sorting algorithm?)

Well, no - even though Bubble Sort is O(1) for space complexity and it is stable - there are other sorting algorithms with same attributes and better complexity.

Computing a perfect Strategy for n x n Chess

Print the <u>Power set</u> of a set of size n.

Think about finding a collision in SHA256 / SHA512 ...

Generating the constant c in t(n) = c*n, i.e O(n) What is the complexity of the code below? std::vector<Widget> vec; for(auto& widget: vec) { for(int j=0; j<100; ++j) { // assume that below is 0(1) widget.doSomething(); } }</pre>







unsigned long sum = std::accumulate(vec.begin(), vec.end(), 0); double inner_product = std::inner_product(vec.begin(), vec.end(), vec.begin(), 0.0);

Algorithmic Complexity @ CppCon 2021

Two calls to std algorithms

```
unsigned long sum = std::accumulate(vec.begin(), vec.end(), 0);
double inner_product =
    std::inner_product(vec.begin(), vec.end(), vec.begin(), 0.0);
```

Above calls iterate over vec twice.

Would it be better to perform the two operations inside a single loop?

Algorithmic Complexity @ CppCon 2021

76

```
unsigned long sum = std::accumulate(vec.begin(), vec.end(), 0);
double inner_product =
    std::inner_product(vec.begin(), vec.end(), vec.begin(), 0.0);
```

Above calls iterate over vec twice.

Would it be better to perform the two operations inside a single loop?

Two loops \sim n + n = O(n) Single loop with two operations \sim 2n = O(n)

Algorithmic Complexity @ CppCon 2021

Two calls to std algorithms

```
unsigned long sum = std::accumulate(vec.begin(), vec.end(), 0);
double inner_product =
    std::inner_product(vec.begin(), vec.end(), vec.begin(), 0.0);
```

Above calls iterate over vec twice. Would it be better to perform the two operations inside a single loop?

```
Two loops \sim n + n = O(n)
Single loop with two operations \sim 2n = O(n)
```

So are they the same?

```
unsigned long sum = std::accumulate(vec.begin(), vec.end(), 0);
double inner_product =
```

std::inner_product(vec.begin(), vec.end(), vec.begin(), 0.0);

Above calls iterate over vec twice.

Would it be better to perform the two operations inside a single loop?

Two loops \sim n + n = O(n) Single loop with two operations \sim 2n = O(n)

So are they the same? Complexity-wise yes, practically - not necessarily!

Algorithmic Complexity @ CppCon 2021

Two calls to std algorithms

```
unsigned long sum = std::accumulate(vec.begin(), vec.end(), 0);
double inner_product =
    std::inner_product(vec.begin(), vec.end(), vec.begin(), 0.0);
```

Above calls iterate over vec twice.

Would it be better to perform the two operations inside a single loop?

It might be better due to *data locality* see benchmarks with <u>std::list</u> and <u>std::vector</u> (and see also <u>SO discussion</u> with additional alternatives).

A note:

std::ranges allows consecutive algorithm calls to be "lazily attached"

into a single loop

Algorithmic Complexity @ CppCon 2021

Two calls to std algorithms

A note:

std::ranges allows consecutive algorithm calls to be "lazily attached" into a single loop

ranges require its own talk, but if you are interested... <u>Here is a relevant code example</u> (courtesy of Dvir Yitzchaki) You may also want to watch <u>Dvir's CppCon 2019 talk on ranges</u>

82

Best Practices

Algorithmic Complexity @ CppCon 2021

Which loop is more important?

```
for (int i=0; i < n; ++i) {
    Operation 1
    for (int j=0; j < n; ++j) {
        Operation 2
    }
}</pre>
```

Which loop is more important?

```
for (int i=0; i < n; ++i) {
    Operation 1 ← preformed n times
    for (int j=0; j < n; ++j) {
        Operation 2 ← performed n<sup>2</sup> times
    }
}
Algorithmic Complexity @ CppCon 2021
```

Break out of loops

Design your algorithm to break as early as possible from the loop

e.g. Bubble Sort with a flag on whether no swaps made in inner loop.

Break when it takes too long... Design your algorithms to break if it doesn't conclude within reasonable time. Otherwise, you are stucking your entire process, usually when the result is no longer required.

Algorithmic Complexity @ CppCon 2021

Simple calls may hide non-constant complexity

Remember to take into account the inner loops

```
std::vector<Trigger> triggers;
triggers.reserve(matrix.rows());
for (const auto& row: matrix) {
    triggers.push_back(Trigger::create(row));
}
```

Time vs. Space Complexity

max_occurences_item(vec)

Algorithmic Complexity @ CppCon 2021

Time vs. Space Complexity

max_occurences_item(vec)

Option 1 - sort and count:

O(1) in space O(n log n) in time

Option 2 - index and count:

O(n) in space O(n) in time [worst case] [worst and average case]

[worst case] [on average]

Setup Time (1)

Setup time vs. query time: indexing (e.g. previous slide)

Algorithmic Complexity @ CppCon 2021

Setup Time (2)

What is the best practical way to sort a list? It may be: copy to a vector, sort the vector, assign back to a list (already presented above...)

<u>See benchmark</u> <u>See also SO discussion</u>

Setup Time (3)

Setup time to achieve cache locality / branch prediction / other accelerations: This is one of the most famous guestions in SO

On the other hand, benchmarks are quite confusing...

- <u>a benchmark without optimization</u> (not a good way to benchmark)
- <u>a benchmark with -O3</u> (unsorted wins!)
- another benchmark with -O3 (now pre-sort wins!)

Algorithmic Complexity @ CppCon 2021

Picking the right container

std::vector is the best, it's not us who say that, it is the spec:

When choosing a container, remember **vector** is best; leave a comment to explain if you choose from the rest!

std::unordered_map

make sure to provide a good enough hash function for your key, or forget about amortized O(1) operations...

hash function requirements in the spec

Using std algorithms

Don't reinvent the wheel

e.g. don't implement your own sort, you may accidentally implement bubble sort

Algorithmic Complexity @ CppCon 2021

Summary

Summary

Implications of bad algorithms and improper use of data structures are potentially much bigger than other micro-performance improvements

Switching to a better algorithm can decrease runtime dramatically!

Algorithmic Complexity @ CppCon 2021

Summary

Thinking about algorithmic complexity is not pre-optimization

It's an essential element of your design and its ability to scale

Summary

The theoretical worst case Big O shouldn't be your only decision factor:

- In real life, **constants** are important: 2n is better than 4n
- In real life, we might choose an algorithm with better **average performance** but *worse worst case complexity*
- Memory locality is highly important

Summary

Remember the tradeoffs:

- **Prior setup** (e.g. sorting / indexing)
- **Space vs. Time** using space to save runtime (e.g. caching, indexing)

Summary

A final note on Space Complexity

Algorithmic Complexity @ CppCon 2021

Summary

A final note on Space Complexity

the *Conference* ⇔ *Wardrobe* complexity problem

102

Thank you!

```
void conclude(auto greetings) {
    while(still_time() && have_questions()) {
        ask();
    }
    greetings();
}
conclude([]{ std::cout << "Thank you!"; });</pre>
```

Algorithmic Complexity @ CppCon 2021