

Amortized complexity

Different kinds of complexities

Average Case complexity

= average complexity over all *possible inputs/situations*
(we need to know the likelihood of each of the input!)

Worst Case complexity

= the worst complexity over all *possible inputs/situations*

Best Case complexity (mostly useless)

= the best complexity over all *possible inputs/situations*

Amortized complexity

= average time taken over a sequence of *consecutive* operations

We will often see algorithms where the worst case complexity is much more than the average case complexity. It depends on the usage, if you're processing a lot of data, occasional bad performance might be OK. However, if we are serving a customer and (in the worst case) she has to wait a long time, that's not good for the company's reputation.

In average-, worst- or best-case complexities, we are concerned in the performance of **one** (independent) operation. The amortized complexity is different: we are thinking about the average time complexity among a number of **successive operations**.

The idea is that, sometimes, you deliberately put extra effort in some operations, in order to speed up subsequent operations. For example, you might spend some time cleaning up or reorganizing of your data structure in order to improve the speed of the coming operations.

Whenever you are reorganizing your data structure, it might slow you down now but you'll benefit from this later (and hopefully many times).

Example: Linear search

What is the time complexity of linear search, where the searched value is stored in x ? Assume that the length of the array is n .

- **Worst Case:** x is at the end of the array
 \implies we need to traverse n elements
- **Best Case:** x is at the beginning of the array
 \implies we only compare with the first one

Example: Linear search

What is the time complexity of linear search, where the searched value is stored in x ? Assume that the length of the array is n .

- **Worst Case:** x is at the end of the array
 \implies we need to traverse n elements
- **Best Case:** x is at the beginning of the array
 \implies we only compare with the first one
- **Average Case (assuming that x is in the array)** – we consider two scenarios:
(1) The likelihood of x being on any position is uniform. In other words, the chance that x is on position 0 is the same as for 1 or any other position.

Then, the average number of traversed elements is equal to

$$\frac{1 + 2 + 3 + \dots + n}{n} = \frac{\frac{n(n+1)}{2}}{n} = \frac{n+1}{2}$$

We assume that x is in the array and it is there exactly once! Otherwise, we would have to compute the average complexity slightly differently.

Recall that we have a formula for *triangular numbers*:

$$1 + 2 + \dots + n = \frac{n(n+1)}{2}$$

The *arithmetic progression* $1 + 2 + \dots + n$ can be also written as $\sum_{i=1}^n i$.

Another example of a progression for which we have an exact formula is the following *geometric progression*:

$$1 + a + a^2 + a^3 + \dots + a^{k-1} = \frac{a^k - 1}{a - 1}.$$

We will see an example of this on slide 8 for $a = 2$.

Example: Linear search

- **Average Case (assuming that x is in the array)** – we consider two scenarios:

(2) If the likelihood is not uniform, then the average number of traversed elements is computed as

$$1 \times P(1) + 2 \times P(2) + 3 \times P(3) + \dots + n \times P(n)$$

where $P(i)$ denotes the likelihood (or probability) that x is stored on the i th position. This means that $0 \leq P(i) \leq 1$ and the sum of all likelihoods is equal to 1, i.e.

$$P(1) + P(2) + \dots + P(n) = 1.$$

(In the uniform case $P(i) = \frac{1}{n}$ for every position.)

Remark: The sum $P(1) + P(2) + \dots + P(n)$ can be also written as $\sum_{i=1}^n P(i)$.

Example: Amortized car cost

Oil consumption: 8 litres per 100 miles

Price of 1 litre is £1.20

⇒ £9.6 per 100 miles

Extra expenses

- new tyres £320 every 70 000 miles
- new breaks £250 every 30 000 miles
- fix gearbox £300 every 130 000 miles
- fix clutch £406 every 100 000 miles
- labour £200 per hour to fix these

⇒ extra £4.3 per 100 miles

⇒ £13.9 is **amortized cost** per 100 miles

(★ those numbers are somewhat made up!)

Amortized complexity: Dynamic array (first attempt)

Naive approach:

1. initially allocate an array of 1000 entries
2. whenever the array becomes full, increase its size by 100

To insert n entries, starting from empty, how long does it take?

For simplicity assume that $n = 1000 + 100k$ (for some k).

1000 insertions +
1000 copies + 100 insertions +
1100 copies + 100 insertions +
1200 copies + 100 insertions +
1300 copies + 100 insertions +
...

In total:

- insertions: $1000 + 100k$
- copies:
 $1000(k - 1)$
 $+ 100 \times (1 + 2 + \dots + (k - 1))$
 $= 1000(k - 1) + 50k(k - 1)$

At the beginning we have

```
MAXSIZE = 1000;  
arr = new int[MAXSIZE];  
stored = 0;
```

We add elements to it by storing them at the end and increasing

`stored` by one. Then, anytime `stored == MAXSIZE` (i.e. `arr` becomes full), we have to allocate a new array of size

`MAXSIZE = MAXSIZE + 100` and copy all elements from `arr` into it.

Recall that $1 + 2 + \dots + n = \frac{n(n+1)}{2}$ therefore

$$1 + 2 + \dots + (k - 1) = \frac{(k - 1)k}{2}$$

The following analysis, however, does not depend on the exact choice of the parameters. If we started with an array of length 10 and increased its size by 5 every time it becomes full, for example, the resulting amortized complexity would still be the same.

Copying and inserting together: $1000k + 100k + 50k(k - 1)$

Amortized cost of one insertion:

$$\frac{1100k + 50k(k - 1)}{1000 + 100k}$$

Copying and inserting together: $1000k + 100k + 50k(k - 1)$

Amortized cost of one insertion:

$$\frac{1100k + 50k(k - 1)}{1000 + 100k}$$

The numerator is in $\mathcal{O}(k^2)$ and denominator is in $\mathcal{O}(k)$

\implies the whole fraction is in $\mathcal{O}(k)$.

But $\mathcal{O}(n) = \mathcal{O}(k)$ because $n = 1000 + 100k$

\implies the amortized complexity of insertion is $\mathcal{O}(n)$.

The amortized cost is computed as the average number of operations needed for one insertion. In our case:

$$\frac{\text{number of copying and inserting}}{\text{number of inserting}}$$

We see that copying is the problem. In the following smarter approach we try to suggest a different strategy which makes sure that copying happens less often.

Amortized complexity: Dynamic array (= Java's ArrayList)

Smart approach:

1. initially allocate an array of 1000 entries
2. whenever the array becomes full, double its size

To insert n entries, starting from empty, how long does it take?

For simplicity assume that $n = 1000 \times 2^k$ (for some k).

1000 insertions +
1000 copies + 1000 insertions +
2000 copies + 2000 insertions +
4000 copies + 4000 insertions +
8000 copies + 8000 insertions +
...

In total:

- insertions:

$$1000 + 1000 \times (1 + 2 + 4 + \dots + 2^{k-1})$$

- copies:

$$1000 \times (1 + 2 + 4 + \dots + 2^{k-1})$$

Copying and inserting together:

$$1000 + 2 \times 1000 \times (1 + 2 + 4 + \dots + 2^{k-1})$$

Because $1 + 2 + 4 + \dots + 2^{k-1} = 2^k - 1$, this is equal to

$$1000 + 2 \times 1000 \times (2^k - 1)$$

Amortized cost of one insertion:

$$\frac{1000 + 2 \times 1000 \times 2^k}{n}$$

Copying and inserting together:

$$1000 + 2 \times 1000 \times (1 + 2 + 4 + \dots + 2^{k-1})$$

Because $1 + 2 + 4 + \dots + 2^{k-1} = 2^k - 1$, this is equal to

$$1000 + 2 \times 1000 \times (2^k - 1)$$

Amortized cost of one insertion:

$$\frac{1000 + 2 \times 1000 \times 2^k}{n}$$

Because $n = 1000 \times 2^k$, the numerator is in $\mathcal{O}(n)$ and denominator is in $\mathcal{O}(n)$

\implies the whole fraction is in $\mathcal{O}(1)$

\implies amortized complexity of insertion is $\mathcal{O}(1)$

Comparison

Inserting at the end of an array

	Average Case	Best Case	Worst Case	Amortized
Naive alg.	—	$\mathcal{O}(1)$	$\mathcal{O}(n)$	$\mathcal{O}(n)$
Smart alg.	—	$\mathcal{O}(1)$	$\mathcal{O}(n)$	$\mathcal{O}(1)$

(Average Case complexity doesn't make sense to consider here.)

Search in a sorted array

	Average Case	Best Case	Worst Case	Amortized
Linear srch	$\mathcal{O}(n)$	$\mathcal{O}(1)$	$\mathcal{O}(n)$	$\mathcal{O}(n)$
Binary srch	$\mathcal{O}(\log n)$	$\mathcal{O}(\log n)$	$\mathcal{O}(\log n)$	$\mathcal{O}(\log n)$

(Amortized complexity is the same as Average Case complexity because the previous searches don't have any effect on the next one.)

We see that the best case and worst case complexities of the Naive algorithm and the Smart algorithm are the same. The only difference is the amortized complexity. The reason why the amortized complexity of the naive algorithm is worse is because the worse case happens too often.

Average Case complexity doesn't make sense to consider in the first table. The time complexity does not depend on the "size" of the value that is being added. It only depends on the current number of elements stored in the array.

Trees




Trees

We store values (e.g. integers, Booleans, Students, ...) in the nodes of the trees.

Depending on the number of branches:

- Unary trees
= Linked Lists,
- Binary trees,
- Ternary trees,
- Quad trees, ...




Trees

We store values (e.g. integers, Booleans, Students, ...) in the nodes of the trees.

Depending on the number of branches:

- Unary trees
= Linked Lists,
- Binary trees,
- Ternary trees,
- Quad trees, ...




size = number of nodes = 5
height = number of levels = 3

Terminology: root, nodes, branches, leafs, depth or level of a node, height or size of the tree

Binary Search Tree

Binary Search Tree is a tree which is either **empty** or


1. values in the left subtree are **smaller** than in the root
2. values in the right subtree are **larger** than in the root
3. root's left and right subtrees are also Binary Search Trees



Binary Search Tree

Binary Search Tree is a tree which is either **empty** or

1. values in the left subtree are **smaller** than in the root
2. values in the right subtree are **larger** than in the root
3. root's left and right subtrees are also Binary Search Trees



⇒ values in the Binary Search Tree represent a sorted sequence!


Searching Binary Search Trees

Starting from the root node, how do we determine whether a value x is in the tree?

If the tree is empty, x is not in the tree! Otherwise, compare x and the value stored in the root. There are three possibilities:

- They are equal \implies we found it!
- x is smaller \implies we search the left subtree.
- x is larger \implies we search the right subtree.

Compare complexities:



vs.

