


Big-O Notation and Algorithmic Analysis

What do you think makes some algorithms "faster"
or "better" than others?




Roadmap

Object-Oriented Programming



Roadmap



Today's question

How can we formalize the
notion of efficiency for
algorithms?

Today's topics

1. Review
2. Big-O Notation
3. Algorithmic Analysis
4. Beyond Algorithmic Analysis

Review

Questions from attendance tickets

- Do structs exist outside of GridLocation?
- Can we use both vectors and stacks at the same time?
- Are grids a thing outside of the Stanford C++ library?
- I don't quite understand why we aren't just learning using regular cpp vectors

Stanford “Vector” vs STL “vector”

What you want to do	Stanford <code>Vector<int></code>	<code>std::vector<int></code>
Create a new, empty vector	<code>Vector<int> vec;</code>	<code>std::vector<int> vec;</code>
Create a vector with <code>n</code> copies of 0	<code>Vector<int> vec(n);</code>	<code>std::vector<int> vec(n);</code>
Create a vector with <code>n</code> copies of a value <code>k</code>	<code>Vector<int> vec(n, k);</code>	<code>std::vector<int> vec(n, k);</code>
Add a value <code>k</code> to the end of a vector	<code>vec.add(k);</code>	<code>vec.push_back(k);</code>
Remove all elements of a vector	<code>vec.clear();</code>	<code>vec.clear();</code>
Get the element at index <code>i</code>	<code>int k = vec[i];</code>	<code>int k = vec[i]; (does not bounds check)</code>
Check size of vector	<code>vec.size();</code>	<code>vec.size();</code>
Loop through vector by index <code>i</code>	<code>for (int i = 0; i < vec.size(); ++i) ...</code>	<code>for (std::size_t i = 0; i < vec.size(); ++i) ...</code>
Replace the element at index <code>i</code>	<code>vec[i] = k;</code>	<code>vec[i] = k; (does not bounds check)</code>

Credit: [CS106L](#)

Ordered ADTs

Elements accessible by indices:

- Vectors (1D)
- Grids (2D)


Elements not accessible by indices:

- Queues (FIFO)
- Stacks (LIFO)

Unordered ADTs


- Sets (elements unique)
- Maps (keys unique)

Nested Data Structures




Map< _____, _____ >

Nested Data Structures



Assignment 2: Fun with Collections




A couple of the ADTs you'll use:

- Grid<bool>
- Set<GridLocation>
- Stack<GridLocation>

Assignment 2: Fun with Collections

A couple of the ADTs you'll use:

- Set<string>
- Map<string, Set<string>>



Nested ADTs Summary

- Powerful
 - Can express highly structured and complex data
 - Used in many real-world systems
- Tricky
 - With increased complexity comes increased cognitive load in differentiating between the levels of information stored at each level of the nesting
 - Specifically in C++, working with nested data structures can be tricky due to the fact that references and copies show up at different points in time. Follow the correct paradigms presented earlier to stay on track!


Note on ADTs (and learning overall)

- We covered six different data structures and their applications over the span of a week, and concluded by implementing BFS using many of these ADTs.
 - This is a lot to take in!
- As Kylie mentioned at the beginning of the quarter, we want to normalize struggle in this class.
 - We cover content very quickly in this class!
 - If you leave lecture feeling you don't understand the algorithm/concept covered that day, don't worry.
 - Lecture is always your first exposure to content – very few people can build deep understanding upon the first exposure
 - The assignments (and section and office hours and LaIR) are your chance to revisit lecture, practice, and really nail down the concepts!
 - Struggling along the way means that you are really *learning*.

How can we formalize the
notion of efficiency for
algorithms?

Why do we care about efficiency?

- Implementing inefficient algorithms may make solving certain tasks impossible, even with unlimited resources



Why do we care about efficiency?

- Implementing inefficient algorithms may make solving certain tasks impossible, even with unlimited resources
- Implementing efficient algorithms allows us to solve important problems, often with limited resources available



TIME COST

STRATEGY A

STRATEGY B

ANALYZING WHETHER
STRATEGY A OR B
IS MORE EFFICIENT



THE REASON I AM SO INEFFICIENT

[Source](#)

Questions computer scientists ask:

1. Does it work?
2. Is it fast?

Assignment 1 At A Glance

- In Assignment 1, you implemented three different algorithms for finding perfect numbers

Assignment 1 At A Glance

- In Assignment 1, you implemented three different algorithms for finding perfect numbers
 - Exhaustive Search
 - Runtime predictions to find 5th perfect number: Anywhere from 25-100+ days

Assignment 1 At A Glance

- In Assignment 1, you implemented three different algorithms for finding perfect numbers
 - Exhaustive Search
 - Runtime predictions to find 5th perfect number: Anywhere from 25-100+ days
 - Smarter Search
 - Runtime predictions to find 5th perfect number: Anywhere from a couple minutes to 1 hour

Assignment 1 At A Glance

- In Assignment 1, you implemented three different algorithms for finding perfect numbers
 - Exhaustive Search
 - Runtime predictions to find 5th perfect number: Anywhere from 25-100+ days
 - Smarter Search
 - Runtime predictions to find 5th perfect number: Anywhere from a couple minutes to 1 hour
 - Euclid's Algorithm
 - Actual runtime to find 5th perfect number: Less than a second!

Assignment 1 At A Glance

- In Assignment 1, you implemented three different algorithms for finding perfect numbers
 - Exhaustive Search
 - Runtime predictions to find 5th perfect number: Anywhere from 25-100+ days
 - Smarter Search
 - Runtime predictions to find 5th perfect number: Anywhere from a couple minutes to 1 hour
 - Euclid's Algorithm
 - Actual runtime to find 5th perfect number: Less than a second!
- Core idea: Although each individual experienced dramatically different real runtimes for these three algorithms, there is a clear distinction here between "fast"/"efficient" and "slow"/"inefficient" algorithms

Containers

- In lecture on Tuesday, someone asked, why not use a `Vector<Vector<string>>` instead of a `Grid<string>`?
 - Vector of Vectors would be slower to use
- There are "fast"/"efficient" and "slow"/"inefficient" ways to insert, delete, or manipulate data in containers

Questions computer scientists ask:

1. Does it work?
2. Is it fast?

Why do computer scientists
care about efficiency?

We solve problems at scale.

Google Search

3.8 million searches per minute



Why do we care about efficiency?

- Implementing inefficient algorithms may make solving certain tasks impossible, even with unlimited resources
- Implementing efficient algorithms allows us to solve important problems, often with limited resources available
- If we can quantify the efficiency of an algorithm, we can understand and predict its behavior when we apply it to unseen problems

Estimating Quantities

[polls]

Leveraging Intuition


- Today's activity is going to look a little bit different than usual. There's no code, no pseudocode, and nothing that resembles C++.
- Instead, you're going to be presented with a set of 4 scenarios, where you have two similar items of different magnitudes, one small and one larger. You know the exact magnitude of the smaller item – can you predict what the magnitude of the larger item will be based on the intuitive visual relationship?
- We'll collect your response to all 4 polls first, and then we'll walk through the answers to the exercises. Remember that these are **guesses** based on your intuition – don't try to do any complex calculations!

What is a rate?


 Δy

 Δx

Example 1



Mass: 100kg




These two square plates are made of the same material.


They have the same thickness.

What's your best guess for the mass of the second square?

Example 2



Mass: 60kg




These two square plates are made of the same material.


They have the same thickness.

What's your best guess for the mass of the second square?

Example 3



Mass: 100kg



These two cubes are made of the same material.

What's your best guess for the mass of the second cube?

Example 4



Mass: 1,000kg



These two statues
are made of the
same material.


What's your best
guess for the
mass of the
second statue?

Example 5



All sides of each triangle
are $10m$ long.

Paint required:
90L




All sides of each triangle
are $40m$ long.

How much paint is
needed to paint
the surface of the
larger
icosahedron?


Answers

Example 1



Mass: 100kg

Mass is about 400kg
(4 smaller squares
make up the larger
square)




These two square plates are made of the same material.

They have the same thickness.


What's your best guess for the mass of the second square?

Example 2



Mass: 60kg

Mass is about 240kg
(side length is
doubled, overall area
increases by factor
of 4)




These two square
plates are made
of the same
material.

They have the
same thickness.


What's your best
guess for the
mass of the
second square?

Example 3



Mass: 100kg

Mass is about 800kg
(8 smaller cubes
make up the larger
cube)



These two cubes
are made of the
same material.

What's your best
guess for the
mass of the
second cube?

Example 4

Mass is about 27000kg (statue dimensions increased by factor of 3, and volume increases by factor of 27)



Mass: 1,000kg




These two statues are made of the same material.

What's your best guess for the mass of the second statue?


Example 5

Paint Required is
about 1440L (side
length grows by
factor of 4, area
increases by factor
of $4^2 = 16$)



All sides of each triangle
are 10m long.

Paint required:
90L



All sides of each triangle
are 40m long.

How much paint is
needed to paint
the surface of the
larger
icosahedron?

Key Takeaway

*Knowing the rate at which
some quantity scales allows
you to predict its value in the
future, even if you don't have
an exact formula.*

Big-O Notation

(or, review of rates)

Big-O Notation

- **Big-O notation** is a way of quantifying the rate at which some quantity grows.

Big-O Notation

- **Big-O notation** is a way of quantifying the rate at which some quantity grows.
- Example:
 - A square of side length r has area $O(r^2)$.

Big-O Notation


- **Big-O notation** is a way of quantifying the rate at which some quantity grows.
- Example:
 - A square of side length r has area $O(r^2)$.



The "O" stands for "on the order of", which is a growth prediction, not an exact formula


Big-O Notation

- **Big-O notation** is a way of quantifying the rate at which some quantity grows.
- Example:
 - A square of side length r has area $O(r^2)$.




Big-O Notation

- **Big-O notation** is a way of quantifying the rate at which some quantity grows.
- Example:
 - A square of side length r has area $O(r^2)$.



Big-O Notation

- **Big-O notation** is a way of quantifying the rate at which some quantity grows.
- Example:
 - A square of side length r has area $O(r^2)$.




Doubling r increases area 4x

Tripling r increases area 9x

Big-O Notation

- **Big-O notation** is a way of quantifying the rate at which some quantity grows.
- Example:
 - A square of side length r has area $O(r^2)$.
 - A circle of radius r has area $O(r^2)$.




Doubling r increases area 4x

Tripling r increases area 9x

Big-O Notation

- **Big-O notation** is a way of quantifying the rate at which some quantity grows.
- Example:
 - A square of side length r has area $O(r^2)$.
 - A circle of radius r has area $O(r^2)$.




Doubling r increases area 4x

Tripling r increases area 9x


Big-O Notation

- **Big-O notation** is a way of quantifying the rate at which some quantity grows.
- Example:
 - A square of side length r has area $O(r^2)$.
 - A circle of radius r has area $O(r^2)$.




Doubling r increases area 4x

Tripling r increases area 9x




Big-O Notation

- **Big-O notation** is a way of quantifying the rate at which some quantity grows.
- Example:
 - A square of side length r has area $O(r^2)$.
 - A circle of radius r has area $O(r^2)$.




Doubling r increases area 4x

Tripling r increases area 9x




Big-O Notation

- **Big-O notation** is a way of quantifying the rate at which some quantity grows.
- Example:
 - A square of side length r has area $O(r^2)$.
 - A circle of radius r has area $O(r^2)$.



Doubling r increases area 4x
Tripling r increases area 9x




Doubling r increases area 4x
Tripling r increases area 9x

Big-O Notation


- **Big-O notation** is a way of quantifying the rate at which some quantity grows.
- Example:
 - A square of side length r has area $O(r^2)$.
 - A circle of radius r has area $O(r^2)$.

This just says that these quantities grow at the same relative rates. It does not say that they're equal!



Doubling r increases area 4x

Tripling r increases area 9x



Doubling r increases area 4x

Tripling r increases area 9x

Big-O in the Real World

Big-O Example: Network Value

- Metcalfe's Law
 - The value of a communications network with n users is $O(n^2)$.

Big-O Example: Network Value

- Metcalfe's Law
 - The value of a communications network with n users is $O(n^2)$.
- Imagine a social network has 10,000,000 users and is worth \$10,000,000. Estimate how many users it needs to have to be worth \$1,000,000,000.
- **Reasonable guess:** The network needs to grow its value 100×. Since value grows quadratically with size, it needs to grow its user base 10×, requiring 100,000,000 users.

Big-O Example: Cell Size

- Question: Why are cells tiny?

Big-O Example: Cell Size

- Question: Why are cells tiny?
- Assumption: Cells are spheres

Big-O Example: Cell Size


- Question: Why are cells tiny?
- Assumption: Cells are spheres
- A cell absorbs nutrients from its environment through its surface area.
 - Surface area of the cell: $O(r^2)$

Big-O Example: Cell Size

- Question: Why are cells tiny?
- Assumption: Cells are spheres
- A cell absorbs nutrients from its environment through its surface area.
 - Surface area of the cell: $O(r^2)$
- A cell needs to provide nutrients all throughout its volume
 - Volume of the cell: $O(r^3)$

Big-O Example: Cell Size

- Question: Why are cells tiny?
- Assumption: Cells are spheres
- A cell absorbs nutrients from its environment through its surface area.
 - Surface area of the cell: $O(r^2)$
- A cell needs to provide nutrients all throughout its volume
 - Volume of the cell: $O(r^3)$
- As a cell gets bigger, its resource *intake* grows slower than its resource *consumption*, so each part of the cell gets less energy.



Big-O Example: Manufacturing

- You're working at a company producing cat toys. It costs you some amount of money to produce a cat toy, and there was some one-time cost to set up the factory.
- What data would you need to gather to estimate the cost of producing ten million cat toys?

Big-O Example: Manufacturing

- You're working at a company producing cat toys. It costs you some amount of money to produce a cat toy, and there was some one-time cost to set up the factory.
- What data would you need to gather to estimate the cost of producing ten million cat toys?

Cost(n) = n × costPerToy + startupCost

Big-O Example: Manufacturing

- You're working at a company producing cat toys. It costs you some amount of money to produce a cat toy, and there was some one-time cost to set up the factory.
- What data would you need to gather to estimate the cost of producing ten million cat toys?

*This term grows as a
function of n*

Cost(n) = $n \times \text{costPerToy} + \text{startupCost}$

Big-O Example: Manufacturing

- You're working at a company producing cat toys. It costs you some amount of money to produce a cat toy, and there was some one-time cost to set up the factory.
- What data would you need to gather to estimate the cost of producing ten million cat toys?

*This term grows as a
function of n*

*This term does not
grow*

$$\text{Cost}(n) = n \times \text{costPerToy} + \text{startupCost}$$

Big-O Example: Manufacturing

- You're working at a company producing cat toys. It costs you some amount of money to produce a cat toy, and there was some one-time cost to set up the factory.
- What data would you need to gather to estimate the cost of producing ten million cat toys?

*This term grows as a
function of n*

*This term does not
grow*

$$\begin{aligned}\text{Cost}(n) &= n \times \text{costPerToy} + \text{startupCost} \\ &= O(n)\end{aligned}$$

Trick to calculating Big-O

Throw out all the leading coefficients and lower-order terms (including constants).

$$\text{Cost}(n) = \$2 \times n + \$500$$

$$\cancel{\$2 \times n} + \cancel{\$500}$$

$$\text{Cost}(n) = O(n)$$

Nuances of Big-O

- Big-O notation is designed to capture **the rate at which a quantity grows**. It does not capture information about
 - leading coefficients: the area of a square and a circle are both $O(r^2)$.
 - lower-order terms: there may be other factors contributing to growth that get glossed over.
- However, it's still a **very powerful tool for predicting behavior**.

Attendance ticket:

<https://tinyurl.com/30juncs106b>

Please don't send this link to students who are not here. It's on your honor!

Announcements

Announcements

- Assignment 1 due **Friday, July 1, at 11:59 pm** (grace period **Saturday, July 2, 11:59 pm**)
- Assignment 2 will be released today. It will be due on **Thursday, July 9 at 11:59pm PDT.**
 - **YEAH hours TBD**
 - This assignment is a step-up in complexity compared to A1 – make sure to get started ASAP.
- If you haven't already, come visit Kylie and me at our office hours! Huang 19.

Example: One Loop

Problem: does array A contain the integer t?

given A (array of length n)
and t (an integer)

```
for i = 1 to n
    if A[i] == t return TRUE
return FALSE
```

Question: what is the running time?

- (A) $O(1)$
- (B) $O(\log n)$

- (C) $O(n)$
- (D) $O(n^2)$



Example: Two Loops

given A, B (arrays of length n)
and + (an integer) {does A or B
contain +?}

```
for i = 1 to n  
    if A[i] == + return TRUE
```

```
for i = 1 to n  
    if B[i] == + return TRUE
```

```
return FALSE
```

Question: running time?

- (A) $O(1)$
- (B) $O(\log n)$

- (C) $O(n)$
- (D) $O(n^2)$



Example: Two Nested Loops

Problem: do arrays A, B have a number in common?
given arrays A, B of length n

```
for i=1 to n
    for j=1 to n
        if A[i] == B[j] return TRUE
return FALSE
```

Question: running time?

- (A) $O(1)$
- (B) $O(\log n)$

 $O(n)$
 $O(n^2)$



Big-O Example: Network Value

- Metcalfe's Law
 - The value of a communications network with n users is $O(n^2)$.
- Imagine a social network has 10,000,000 users and is worth \$10,000,000. Estimate how many users it needs to have to be worth \$1,000,000,000.
 1. 10,000,000
 2. 50,000,000
 3. 100,000,000
 4. 1,000,000,000

Analyzing Code

Analyzing Code

*How can we apply Big-O to
computer science?*

Answering “is it fast?”

- We could use runtime
 - Runtime is the amount of time it takes for a program to run

Answering “is it fast?”

- What is runtime?
 - Runtime is the amount of time it takes for a program to run

```
[SimpleTest] ---- Tests from main.cpp ----
[SimpleTest] starting (PROVIDED_TEST, line 36) timing vectorMax on 10,00... = Correct
  Line 42 Time vectorMax(v) (size =100000000) completed in 0.268 secs
  Line 43 Time vectorMax(v) (size =100000000) completed in 0.264 secs
  Line 44 Time vectorMax(v) (size =100000000) completed in 0.269 secs
You passed 1 of 1 tests. Keep it up!
```

Old 2012
MacBook

Why runtime isn't enough

- What is runtime?
 - Runtime is the amount of time it takes for a program to run

```
[SimpleTest] ---- Tests from main.cpp ----
[SimpleTest] starting (PROVIDED_TEST, line 36) timing vectorMax on 10,00... = Correct
  Line 42 Time vectorMax(v) (size =100000000) completed in 0.268 secs
  Line 43 Time vectorMax(v) (size =100000000) completed in 0.264 secs
  Line 44 Time vectorMax(v) (size =100000000) completed in 0.269 secs
You passed 1 of 1 tests. Keep it up!
```

Old 2012
MacBook

```
[SimpleTest] ---- Tests from main.cpp ----
[SimpleTest] starting (PROVIDED_TEST, line 36) timing vectorMax on 20,00... = Correct
  Line 42 Time vectorMax(v) (size =100000000) completed in 0.181 secs
  Line 43 Time vectorMax(v) (size =100000000) completed in 0.181 secs
  Line 44 Time vectorMax(v) (size =100000000) completed in 0.183 secs
You passed 1 of 1 tests. Que bien!
```

Ed's powerful
computers

Why runtime isn't enough

- Measuring wall-clock runtime is less than ideal, since
 - It depends on what computer you're using,
 - What else is running on that computer,
 - Whether that computer is conserving power,
 - Etc.

It's very hard to standardize.

Why runtime isn't enough

- Measuring wall-clock runtime is less than ideal, since
 - It depends on what computer you're using,
 - What else is running on that computer,
 - Whether that computer is conserving power,
 - Etc.
- Worse, **individual runtimes can't predict future runtimes.**

Answering “Is it fast?”

- We need a standardized way to think about rate of algorithms
- That doesn't make assumptions about our computer, our circumstances, our inputs, etc.

Answering “Is it fast?”

- We need a standardized way to think about rate of algorithms
- That doesn't make assumptions about our computer, our circumstances, our inputs, etc.

Idea: count the number of executions in an algorithm.

- number of times a single operation is done (access an element, compare two items)
- We can analyze this before we even run the program!

Analyzing Code: **vectorMax()**

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Assume any individual statement takes one unit of time to execute.

*If the input **Vector** has n elements, how many executions (time units) will this code take to run?*

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total time based on # of repetitions

1 time unit

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total time based on # of repetitions

1 time unit

1 time unit

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total time based on # of repetitions

1 time unit

1 time unit

1 time unit

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total time based on # of repetitions

1 time unit

1 time unit

1 time unit

N time units

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total time based on # of repetitions

1 time unit

1 time unit

1 time unit

N time units

$N-1$ time units

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total time based on # of repetitions

1 time unit

1 time unit

1 time unit

N time units

$N-1$ time units

$N-1$ time units

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total time based on # of repetitions

1 time unit

1 time unit

1 time unit

N time units

$N-1$ time units

$N-1$ time units

(up to) $N-1$ time units

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total time based on # of repetitions

1 time unit

1 time unit

1 time unit

N time units

N-1 time units

N-1 time units

(up to) N-1 time units

1 time unit

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total amount of time

$$4N + 1$$

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total amount of time

$$4N + 1$$

Is this useful?

What does this tell us?

Answering “Is it fast?”

- We need a standardized way to think about rate of algorithms
- That doesn't make assumptions about our computer, our circumstances, our inputs, etc.

Idea: count the number of executions in an algorithm.

- Maybe this is still too much detail
- Constant factors might still depend on the system

Answering “Is it fast?”

- We need a standardized way to think about rate of algorithms
- That doesn’t make assumptions about our computer, our circumstances, our inputs, etc.

Better idea: find the **Big-O** of this algorithm.

- General enough to help us compare across computers
- **It's a rate that represents:** As the input size grows, how does the runtime grow?
- A computer-independent metric for efficiency!

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total amount of time

$$4N + 1$$

Is this useful?

What does this tell us?

vectorMax()

```
int vectorMax(Vector<int> &v) {  
    int currentMax = v[0];  
    int n = v.size();  
    for (int i = 1; i < n; i++) {  
        if (currentMax < v[i]) {  
            currentMax = v[i];  
        }  
    }  
    return currentMax;  
}
```

Total amount of time

$O(n)$

More practical: Doubling the size of the input roughly doubles the runtime. Therefore, the input and runtime have a linear ($O(n)$) relationship.

Analyzing Code: `printStars()`

printStars()

```
void printStars(int n) {  
    for (int i = 0; i < n; i++) {  
        for (int j = 0; j < n; j++) {  
            cout << '*' << endl;  
        }  
    }  
}
```

How much time will it take for this code to run, as a function of n ?

Answer using big- O notation.

printStars()

```
void printStars(int n) {  
    for (int i = 0; i < n; i++) {  
        for (int j = 0; j < n; j++) {  
            cout << '*' << endl;  
        }  
    }  
}
```

How much time will it take for this code to run, as a function of n ?

Answer using big- O notation.

printStars()

```
void printStars(int n) {  
    for (int i = 0; i < n; i++) {  
        for (int j = 0; j < n; j++) {  
            // do a fixed amount of work  
        }  
    }  
}
```

How much time will it take for this code to run, as a function of n ?

Answer using big- O notation.

printStars()

```
void printStars(int n) {  
    for (int i = 0; i < n; i++) {  
        for (int j = 0; j < n; j++) {  
            // do a fixed amount of work  
        }  
    }  
}
```

How much time will it take for this code to run, as a function of n ?

Answer using big- O notation.

printStars()

```
void printStars(int n) {  
    for (int i = 0; i < n; i++) {  
  
        // do O(n) time units of work  
  
    }  
}
```

How much time will it take for this code to run, as a function of n ?

Answer using big- O notation.

printStars()

```
void printStars(int n) {  
    for (int i = 0; i < n; i++) {  
        // do O(n) time units of work  
    }  
}
```

How much time will it take for this code to run, as a function of n ?

Answer using big- O notation.

printStars()

```
void printStars(int n) {  
    // do  $O(n^2)$  time units of work  
}
```

How much time will it take for this code to run, as a function of n ?

Answer using big- O notation.

printStars()

```
void printStars(int n) {
    for (int i = 0; i < n; i++) {
        for (int j = 0; j < n; j++) {
            cout << '*' << endl;
        }
    }
}
```

$O(n^2)$

A final analyzing code
example

hmmThatsStrange()

```
void hmmThatsStrange(int n) {  
    cout << "Mirth and Whimsy" << n << endl;  
}
```

The runtime is *completely independent* of the value **n**.

hmmThatsStrange()

```
void hmmThatsStrange(int n) {  
    cout << "Mirth and Whimsy" << n << endl;  
}
```

How much time will it take for this code to run, as a function of n ?
Answer using big- O notation.

hmmThatsStrange()


```
void hmmThatsStrange(int n) {  
    cout << "Mirth and Whimsy" << n << endl;  
}
```

$O(1)$

Applying Big-O to ADTs

Efficiency Categorizations So Far

- Constant Time – $O(1)$
 - Super fast, this is the best we can hope for!
 - Euclid's Algorithm for Perfect Numbers
- Linear Time – $O(n)$
 - This is okay, we can live with this
- Quadratic Time – $O(n^2)$
 - This can start to slow down really quickly
 - Exhaustive Search for Perfect Numbers
- How do all the ADT operations we've seen so far fall into these categories?



Big-O Terminology

constant	<i>logarithmic</i>	<i>linear</i>	<i>n log n</i>	<i>quadratic</i>	<i>polynomial (other than n^2)</i>	<i>exponential</i>
$O(1)$	$O(\log n)$	$O(n)$	$O(n \log n)$	$O(n^2)$	$O(n^k)$ ($k \geq 1$)	$O(a^n)$ ($a > 1$)

A fast algorithm is when the worst-case run-time grows **SLOWLY** with the input size.

ADT Big-O Matrix

ADT Big-O Matrix

- Vectors
 - `.size()` – $O(1)$
 - `.add()` – $O(1)$
 - `v[i]` – $O(1)$
 - `.insert()` – $O(n)$
 - `.remove()` – $O(n)$
 - `.sublist()` – $O(n)$
 - `traversal` – $O(n)$
- Grids
 - `.numRows()/numCols()` – $O(1)$
 - `g[i][j]` – $O(1)$
 - `.inBounds()` – $O(1)$
 - `traversal` – $O(n^2)$

ADT Big-O Matrix

- Vectors
 - `.size()` – $O(1)$
 - `.add()` – $O(1)$
 - `v[i]` – $O(1)$
 - `.insert()` – $O(n)$
 - `.remove()` – $O(n)$
 - `.sublist()` – $O(n)$
 - `traversal` – $O(n)$
- Grids
 - `. numRows()/. numCols()` – $O(1)$
 - `g[i][j]` – $O(1)$
 - `.inBounds()` – $O(1)$
 - `traversal` – $O(n^2)$
- Queues
 - `.size()` – $O(1)$
 - `.peek()` – $O(1)$
 - `.enqueue()` – $O(1)$
 - `.dequeue()` – $O(1)$
 - `.isEmpty()` – $O(1)$
 - `traversal` – $O(n)$
- Stacks
 - `.size()` – $O(1)$
 - `.peek()` – $O(1)$
 - `.push()` – $O(1)$
 - `.pop()` – $O(1)$
 - `.isEmpty()` – $O(1)$
 - `traversal` – $O(n)$

ADT Big-O Matrix

- Vectors
 - `.size()` – $O(1)$
 - `.add()` – $O(1)$
 - `v[i]` – $O(1)$
 - `.insert()` – $O(n)$
 - `.remove()` – $O(n)$
 - `.sublist()` – $O(n)$
 - traversal – $O(n)$
- Grids
 - `.numRows()/numCols()` – $O(1)$
 - `g[i][j]` – $O(1)$
 - `.inBounds()` – $O(1)$
 - traversal – $O(n^2)$
- Queues
 - `.size()` – $O(1)$
 - `.peek()` – $O(1)$
 - `.enqueue()` – $O(1)$
 - `.dequeue()` – $O(1)$
 - `.isEmpty()` – $O(1)$
 - traversal – $O(n)$
- Stacks
 - `.size()` – $O(1)$
 - `.peek()` – $O(1)$
 - `.push()` – $O(1)$
 - `.pop()` – $O(1)$
 - `.isEmpty()` – $O(1)$
 - traversal – $O(n)$
- Sets
 - `.size()` – $O(1)$
 - `.isEmpty()` – $O(1)$
 - `.add()` – ???
 - `.remove()` – ???
 - `.contains()` – ???
 - traversal – $O(n)$
- Maps
 - `.size()` – $O(1)$
 - `.isEmpty()` – $O(1)$
 - `m[key]` – ???
 - `.contains()` – ???
 - traversal – $O(n)$

ADT Big-O Matrix

- Vectors
 - `.size()` – $O(1)$
 - `.add()` – $O(1)$
 - `v[i]` – $O(1)$
 - `.insert()` – $O(n)$
 - `.remove()` – $O(n)$
 - `.sublist()` – $O(n)$
 - traversal – $O(n)$
- Grids
 - `.numRows()/numCols()`
 - `g[i][j]` – $O(1)$
 - `.inBounds()` – $O(1)$
 - traversal – $O(n^2)$

- Queues
 - `.size()` – $O(1)$

- Sets
 - `.size()` – $O(1)$
 - `.empty()` – $O(1)$
 - `l()` – ???
 - `.move()` – ???
 - `tains()` – ???
 - `ersal` – $O(n)$
 - `aps`
 - `():` – $O(1)$
 - `.empty()` – $O(1)$
 - `ey]` – ???
 - `tains()` – ???
- `.isEmpty()` – $O(1)$
- traversal – $O(n)$

How can we achieve faster
than $O(n)$ runtime when
searching/storing n
elements?

Beyond Algorithmic Analysis

(credit: Katie Creel)

Ramifications of Big O Differences

- If we have an algorithm that has 1000 elements, and the $O(\log n)$ version runs in 10 milliseconds...

constant	logarithmic	linear	$n \log n$	quadratic	polynomial (other than n^2 , (n^2))	exponential
1 milliseconds	10 milliseconds	1 second	10 seconds	17 minutes	277 hours	heat death of the universe

Algorithmic complexity analysis can be the difference between a program that runs in a few seconds and one that won't finish before the heat death of the universe. It is often necessary.


Questions computer scientists ask:

1. Does it work?
2. Is it fast?

Is that all we care about?

WHAT ELSE IS IN THE BIG PICTURE VIEW OF AN ALGORITHM?





Case Study:

In 2006, the State of Indiana awarded IBM/ACS a contract for >\$1 billion to “modernize” Indiana’s welfare case management system, including applications for food stamps and Medicaid.

After only 19 months, it was clearly not going as planned.

Case Study

Some "lowlights" of the system's failures:

- "Applicants waited 20 or 30 minutes on hold, only to be denied benefits for "failure to cooperate in establishing eligibility" if they were unable to receive a callback after having burned through their limited cellphone minutes."
- "By February 2008, the number of households receiving food stamps in Delaware County, which includes Muncie, Indiana, dropped more than 7 percent, though requests for food assistance had climbed 4 percent in Indiana overall."

Case Study

In light of these failures, the State of Indiana cancelled its contract with IBM and sued the company for breach of contract.

In court, IBM argued that they were *not responsible* for issues related to wait times, appeals, wrongful denials, lost documents, etc. as the contract only stated that a successful system would *succeed by increasing efficiency and reducing costs and fraud.* IBM's system did reduce costs, but did so by denying people the benefits they needed.

**EFFICIENCY & OPTIMIZATION CAN ONLY BE
AS GOOD AS GOOD AS WHAT IS OPTIMIZED...**

Questions computer scientists ask:

1. Does it work?
2. Is it fast?

Is that all we care about?



When Less Efficient Algorithms Rule

Passwords are often encrypted with a hash.

What prevents a hacker from guessing randomly, perhaps millions of times per minute, until the password is discovered?

~ *Algorithmic Inefficiency* ~

bcrypt and other popular encryption functions are intentionally designed to be slow, memory intensive, or both, making guessing more costly.

Computation requires energy.

BAY AREA

Stanford power outage: University preparing for a restoration that could 'take days'



Annie Vainshtein

June 22, 2022 | Updated: June 22, 2022 6:36 p.m.



BUSINESS • GOOGLE

The Secret Cost of Google's Data Centers: Billions of Gallons of Water to Cool Servers



Bitcoin consumes 'more electricity than Argentina'


By Cristina Criddle
Technology reporter

10 February 2021



Green Computing

Big-O analysis can also be part of a “green computing”: a commitment to decreasing the environmental impact of computing.




Green Computing Includes ...

Selecting sources of electricity with lower greenhouse gas emissions, including by moving computing to the cloud when necessary

Decreasing carbon footprint of datacenters

Recycling and reducing use of raw materials during manufacturing, especially metals that are socially and environmentally damaging to mine

Reducing energy consumption of computation itself, including by *increasing algorithmic efficiency!*



Big data needs Big O
and efficient
algorithms

Improving
efficiency can
reduce energy

But storing big
data has a big
carbon cost




Green Computing's Efficiency Tradeoffs

- When is the answer to increase efficiency?
- When is the answer to choose smaller data?

What's next?

Roadmap

Object-Oriented Programming



Recursion

