

CS 106B

Lecture 3: Vectors, Grids, Big O

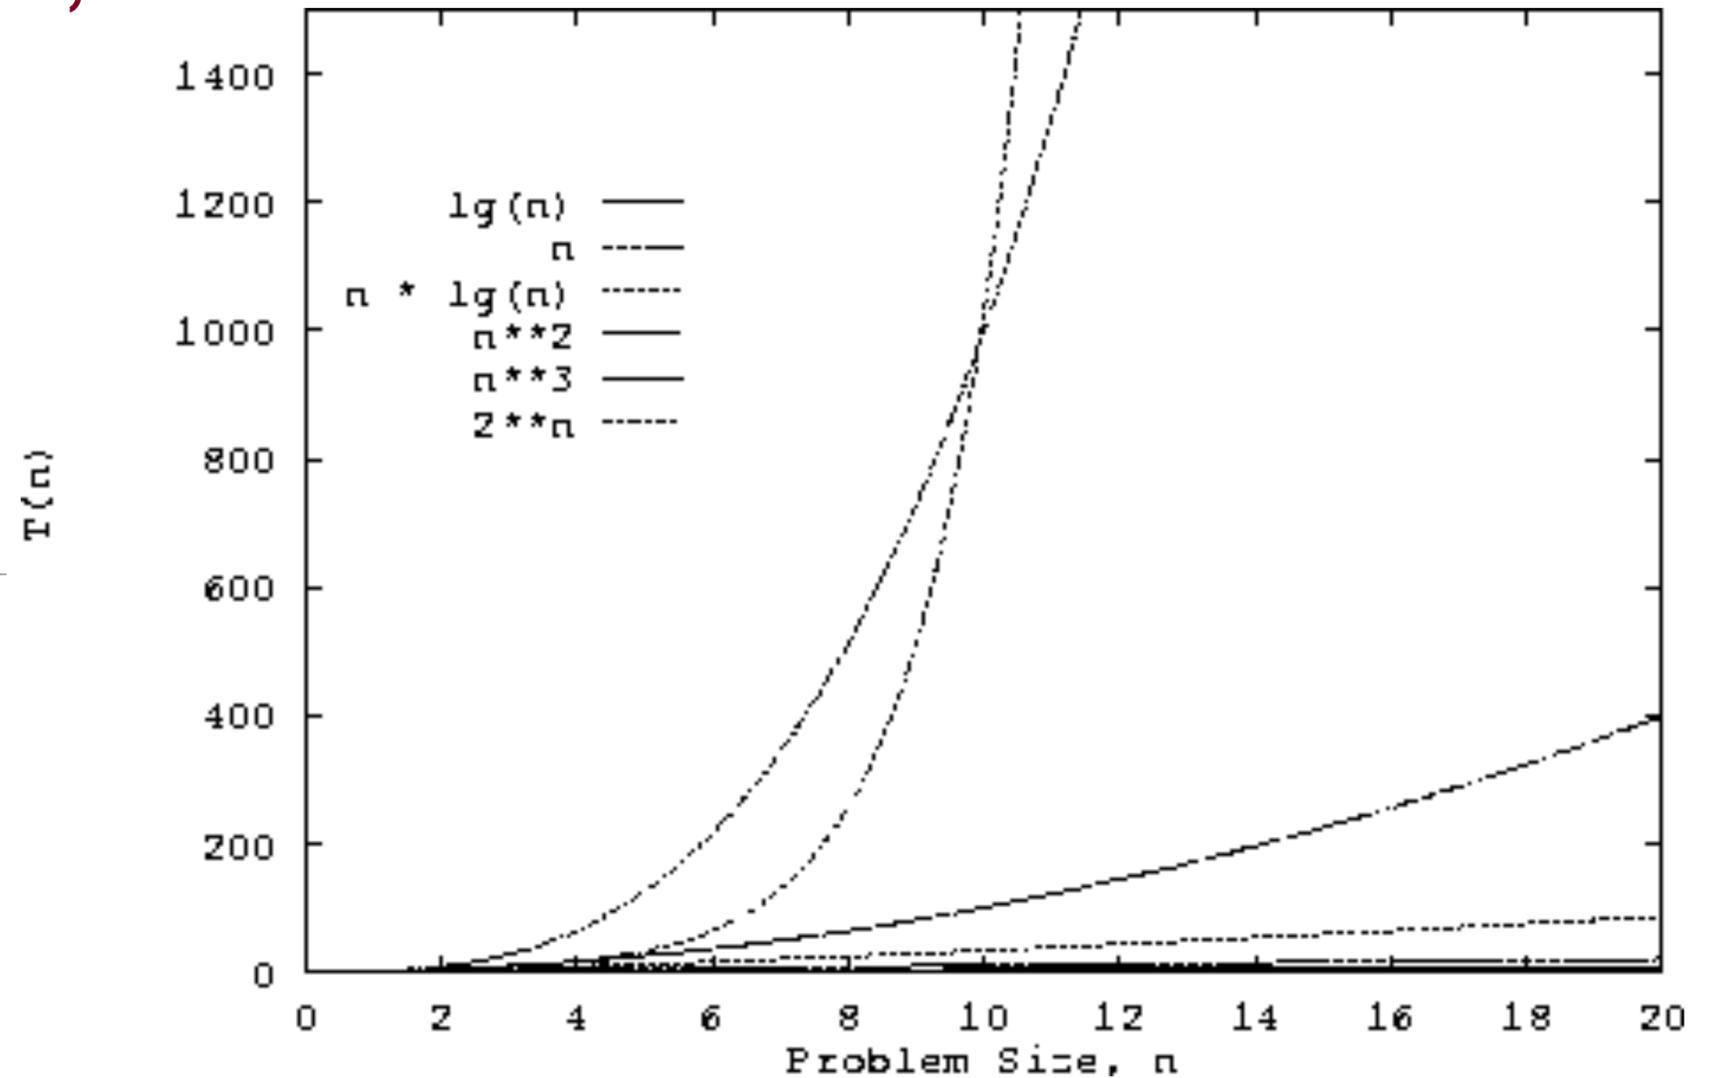
Friday, April 6th, 2018

Programming Abstractions
Spring 2018
Stanford University
Computer Science Department

Lecturer: Chris Gregg

reading:

Programming Abstractions in C++, Chapters
5.1-5.2, Section 10.2



Today's Topics

- Logistics:
 - OAE Letters -- please send to Chris and Nick by **next Wednesday**
 - A note on the honor code
- Review
- Vectors
- Grids
- Introduction to Computational Complexity and "Big O"
- Reading Assignment: Chapter 5.1-5.2, Section 10.2



A Note on the Honor Code

- Honor code handout:

<http://web.stanford.edu/class/cs106b/handouts/3-HonorCode.pdf>



Value semantics

- **value semantics:** In Java and C++, when variables (`int`, `double`) are passed as parameters, their values are copied.
 - Modifying the parameter will not affect the variable passed in.

```
void grow(int age) {
    age = age + 1;
    cout << "grow age is " << age << endl;
}

int main() {
    int age = 20;
    cout << "main age is " << age << endl;
    grow(age);
    cout << "main age is " << age << endl;
    return 0;
}
```

Output:

```
main age is 20
grow age is 21
main age is 20
```



Reference semantics (2.5)

- **reference semantics:** In C++, if you declare a parameter with an `&` after its type, instead of passing a copy of its value, it will link the caller and callee functions to the same variable in memory.
- Modifying the parameter *will* affect the variable passed in.

```
void grow(int &age) {  
    age = age + 1;  
    cout << "grow age is " << age << endl;  
}  
  
int main() {  
    int age = 20;  
    cout << "main age is " << age << endl;  
    grow(age);  
    cout << "main age is " << age << endl;  
    return 0;  
}
```

Output:

```
main age is 20  
grow age is 21  
main age is 21
```



Reference semantics (2.5)

•Notes about references:

- References are super important when dealing with objects that have a lot of elements (Vectors, for instance). Because the reference does not copy the structure, it is fast. You don't want to transfer millions of elements between two functions if you can help it!
- The reference syntax can be confusing, as the "&" (ampersand) character is also used to specify the address of a variable or object. The & is only used as a reference parameter in the function declaration, not when you call the function:

```
void grow(int &age) {  
    age = age + 1;  
    cout << "grow age is "  
        << age << endl;  
}
```

```
int main() {  
    grow(age);  
    grow(&age);  
    return 0;  
}
```



Reference pros/cons

- benefits of reference parameters:
 - a useful way to be able to 'return' more than one value
 - often used with objects, to avoid making bulky copies when passing
- downsides of reference parameters:
 - hard to tell from call whether it is ref; can't tell if it will be changed
foo(a, b, c); // will foo change a, b, or c? :-/
 - (very) slightly slower than value parameters
 - can't pass a literal value to a ref parameter
grow(39); // error

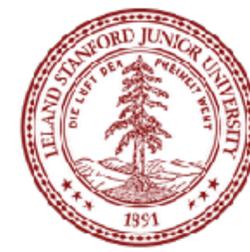


Reference Example

- Without references, you can't write a swap function to swap two integers. This is true about Java. What happens with the following function?

```
/*  
 * Attempts to place a's value into b and vice versa.  
 */  
void swap(int a, int b) {  
    int temp = a;  
    a = b;  
    b = temp;  
}
```

- Answer: the original variables are unchanged, because they are passed as copies (values)!



Reference Example

- *With* references, you *can* write a swap function to swap two integers, because you can access the original variables:

```
/*  
 * Places a's value into b and vice versa.  
 */  
void swap(int &a, int &b) {  
    int temp = a;  
    a = b;  
    b = temp;  
}
```

- Answer: the original variables **are changed**, because they are passed as references !



Tricky Reference Mystery Example

What is the output of this code? Talk to your neighbor!

```
void mystery(int& b, int c, int& a) {  
    a++;  
    b--;  
    c += a;  
}  
  
int main() {  
    int a = 5;  
    int b = 2;  
    int c = 8;  
    mystery(c, a, b);  
    cout << a << " " << b << " " << c << endl;  
    return 0;  
}
```

```
// A. 5 2 8  
// B. 5 3 7  
// C. 6 1 8  
// D. 61 13  
// E. other
```

Note: please don't obfuscate your code like this! :(

See the International Obfuscated C Contest for much, much worse examples



Tricky Reference Mystery Example

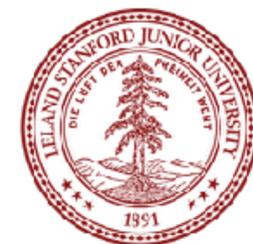
What is the output of this code?

```
void mystery(int& b, int c, int& a) {  
    a++;  
    b--;  
    c += a;  
}  
  
int main() {  
    int a = 5;  
    int b = 2;  
    int c = 8;  
    mystery(c, a, b);  
    cout << a << " " << b << " " << c << endl;  
    return 0;  
}
```

```
// A. 5 2 8  
// B. 5 3 7  
// C. 6 1 8  
// D. 61 13  
// E. other
```

Note: please don't obfuscate your code like this! :(

See the International Obfuscated C Contest for much, much worse examples



Quadratic Exercise -- how do you return multiple things?

- A quadratic equation for variable x is one of the form:

$$\mathbf{ax^2 + bx + c = 0}, \text{ for some numbers } \mathbf{a}, \mathbf{b}, \text{ and } \mathbf{c}.$$

- The two roots of a quadratic equation can be found using the quadratic formula at right.

$$\frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

- Example: The roots of $\mathbf{x^2 - 3x - 4 = 0}$ are $\mathbf{x = 4}$ and $\mathbf{x = -1}$
- How would we write a function named quadratic to solve quadratic equations?
 - What parameters should it accept?
 - Which parameters should be passed by value, and which by reference?
 - What, if anything, should it return?
- We have choices!



Quadratic Exercise -- how do you return multiple things?

```
/*  
 * Solves a quadratic equation  $ax^2 + bx + c = 0$ ,  
 * storing the results in output parameters root1 and root2.  
 * Assumes that the given equation has two real roots.  
 */  
void quadratic(double a, double b, double c,  
              double& root1, double& root2) {  
    double d = sqrt(b * b - 4 * a * c);  
    root1 = (-b + d) / (2 * a);  
    root2 = (-b - d) / (2 * a);  
}
```

$$\frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

- How are we "returning" the results? Answer: by reference
- What other choices could we have made? Talk to your neighbor!



Quadratic Exercise -- how do you return multiple things?

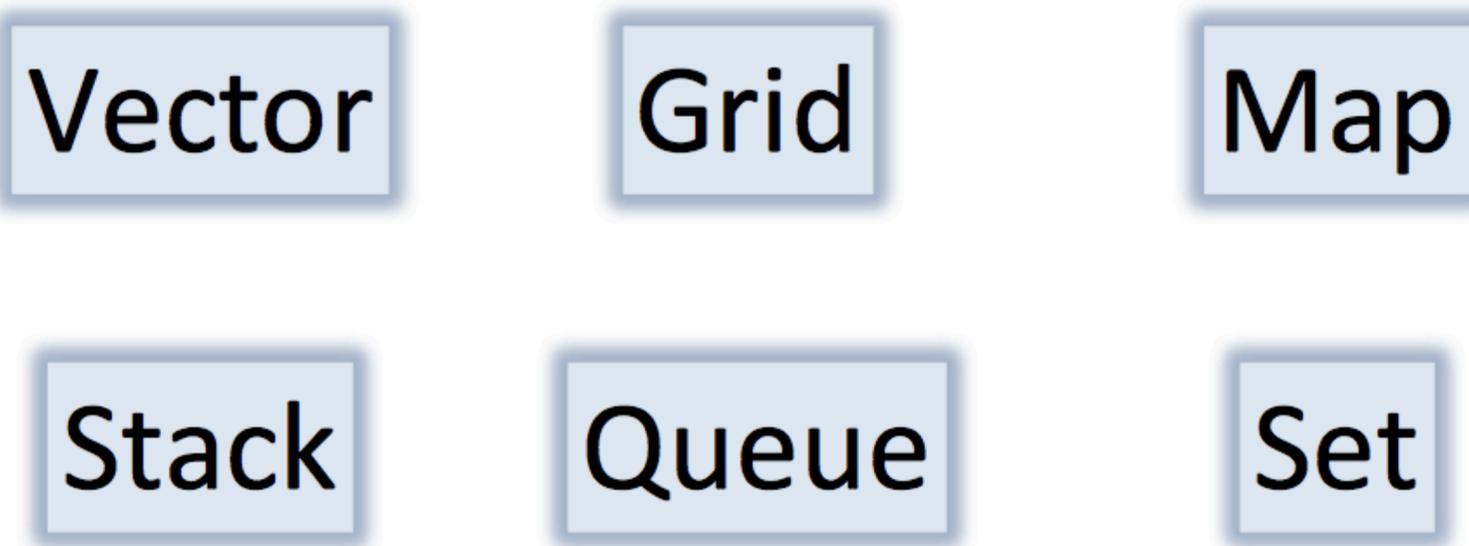
- Possible choices:
 - We could have returned a boolean if the roots were imaginary
 - We could have added extra parameters to support some form of imaginary numbers
 - We could have called an error function inside this function (but that is not always a good idea -- functions like this should generally have an interface through the parameters and/or return value, and should gracefully fail)
 - We could have re-written the function as two functions that return either the positive or negative root, without using references.
 - We could have returned a `Vector<double>` object (tricky syntax!)

$$\frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$



Vectors and Grids

- One of the most powerful aspects of C++ is the ability to have a "collection":



- We will talk about all of these as we go through CS 106B, but you will need to use the Vector and Grid classes for Game of Life.



Vector

- What is it?
 - Like a Java `ArrayList<type>`
 - A list of elements that can grow and shrink.
 - Each element has a place (or index) in the list.
 - Advanced array.
- Important Details
 - Constructor creates an empty list.
 - Bounds checks.
 - Knows its size.
 - Include “vector.h”
- Why not use arrays?



Creating a Vector

```
Vector<int> vec;
```

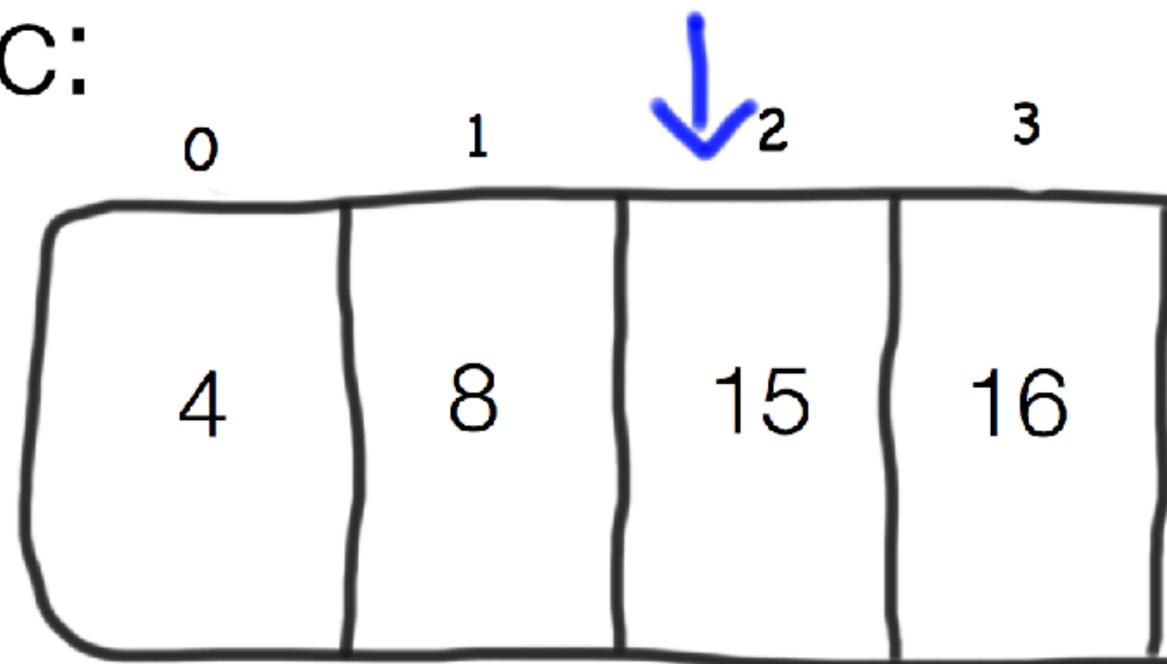
You must specify the type of your vector.
When a vector is created it is initially empty.



Vectors are just arrays under the hood!

```
Vector<int> magic;  
magic.add(4);  
magic.add(8);  
magic.add(15);  
magic.add(16);  
cout << magic[2] << endl;
```

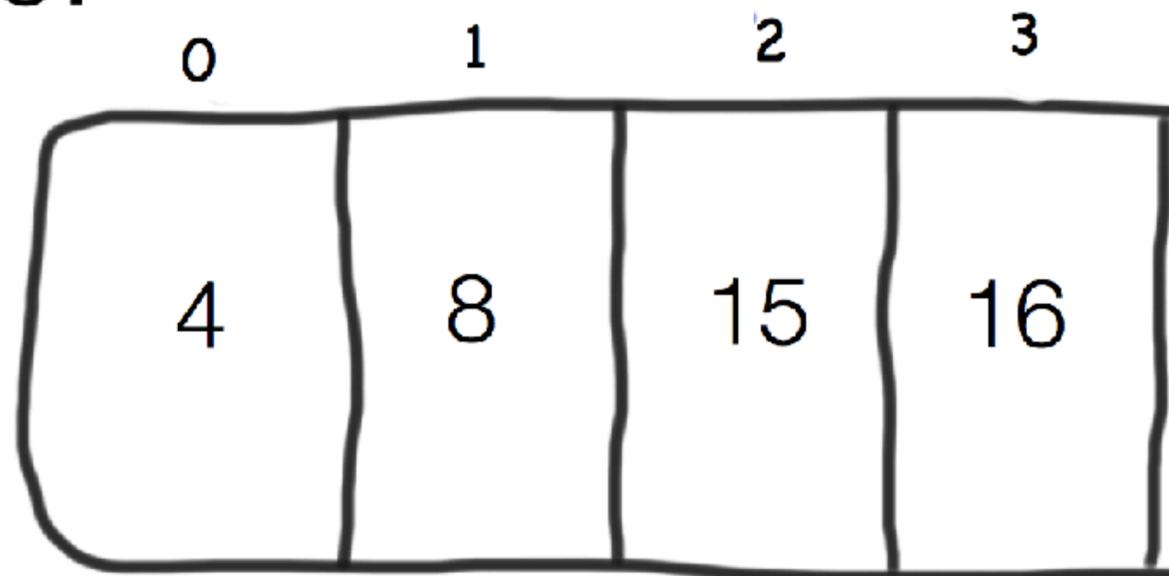
magic:



But they have useful functions, like size()

```
for(int i = 0; i < magic.size(); i++) {  
    cout << magic[i];  
}
```

magic:



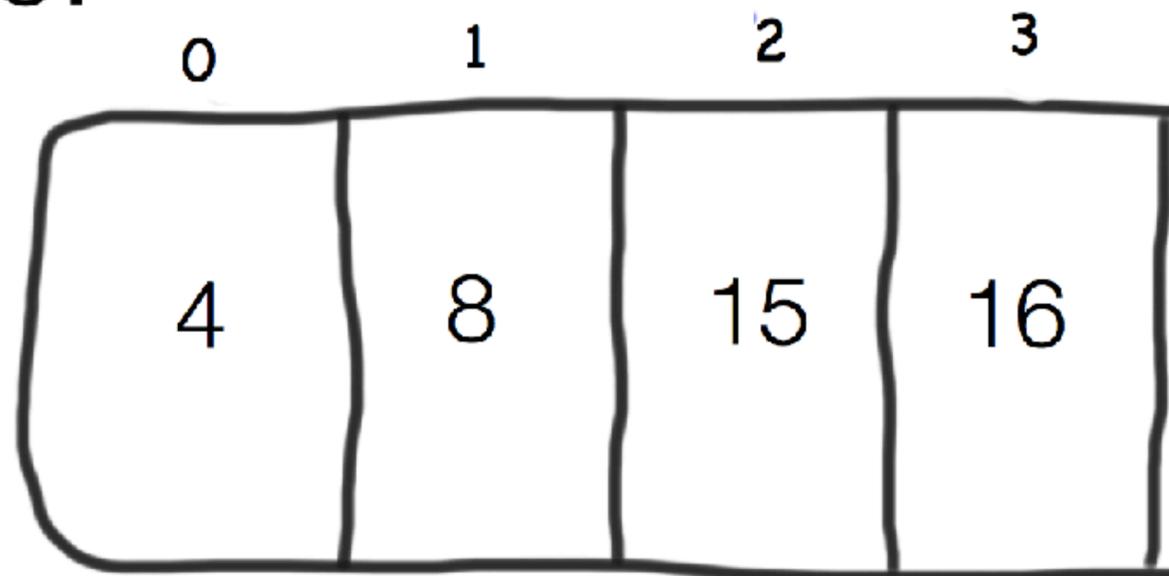
Output: 4
8
15
16



You can use "for each" loops on containers:

```
for(int value : magic) {  
    cout << value << endl;  
}
```

magic:



Output: 4
8
15
16



Vector Methods

`vec.size()`

Returns the number of elements in the vector.

`vec.isEmpty()`

Returns `true` if the vector is empty.

`vec[i]`

Selects the i^{th} element of the vector.

`vec.add(value)`

Adds a new element to the end of the vector.

`vec.insert(index, value)`

Inserts the value before the specified index position.

`vec.remove(index)`

Removes the element at the specified index.

`vec.clear()`

Removes all elements from the vector.

For the exhaustive list check out:

<http://stanford.edu/~stepp/cppdoc/Vector-class.html>



The Grid Container



Grid<type>



WELCOME TO
THE MATRIX!!!!!!

Grid

- What is it?
 - Advanced **2D** array.
 - Think spread sheets, game boards
- Important Details
 - Default constructor makes a grid of size 0
 - Doesn't support “ragged right”.
 - Bounds checks
 - Knows its size.
- *We could* use a combination of Vectors to simulate a 2D matrix, but a Grid is easier!



Creating a Grid

```
Grid<int> matrix(2,2);  
matrix[0][0] = 42;  
matrix[0][1] = 6;  
matrix[1][0] = matrix[0][1];  
cout << matrix.numRows() << endl;  
cout << matrix[0][1] << endl;  
cout << matrix[1][1] << endl;  
cout << matrix[2][3] << endl;
```



Creating a Grid

```
Grid<int> matrix(2,2);  
matrix[0][0] = 42;  
matrix[0][1] = 6;  
matrix[1][0] = matrix[0][1];  
cout << matrix.numRows() << endl;  
cout << matrix[0][1] << endl;  
cout << matrix[1][1] << endl;  
cout << matrix[2][3] << endl;
```

	0	1
0	42	6
1	6	0

*** STANFORD C++ LIBRARY

*** An Exception occurred during program execution:

*** Grid::operator [][]: (3, 2) is outside of valid range [(0, 0)..(2, 1)]



Grid Methods

`grid.numRows ()`

Returns the number of rows in the grid.

`grid.numCols ()`

Returns the number of columns in the grid.

`grid[i][j]`

Selects the element in the i^{th} row and j^{th} column.

`grid.resize (rows, cols)`

Changes the dimensions of the grid and clears any previous contents.

`grid.inBounds (row, col)`

Returns `true` if the specified row , column position is within the grid.

For the exhaustive list check out:

<http://stanford.edu/~stepp/cppdoc/Grid-class.html>



Grid Example: Traversing a Grid

```
void printGrid(Grid<char> &grid) {  
    for(int r = 0; r < grid.numRows(); r++) {  
        for(int c = 0; c < grid.numCols(); c++) {  
            cout << grid[r][c];  
        }  
        cout << endl;  
    }  
}
```

	0	1
0	a	b
1	c	d
2	e	f

	0	1
0	a	b
1	c	d
2	e	f



Collections

1. Defined as Classes

This means they have constructors and member functions

2. Templated

They have a mechanism for collecting different variable types

3. Deep copy assignment

Often pass them by reference!



Collection Common Pitfalls

1. **Vector numbers;**

Needs a type! Should be: **Vector<int> numbers;**

2. **void myFunction(Grid<bool> gridParam);**

Two issues: (a) if you want the original **gridParam** to be changed in the calling function, you're out of luck. (b) inefficient because you have to make a copy of **gridParam**.



Collection Common Pitfalls

```
3. void cout(Grid<bool> & grid) {  
    for(int i = 0; i < grid.numRows(); i++) {  
        for(int j = 0; j < grid.numCols(); j++) {  
            cout << grid[j][i];  
        }  
    }  
}
```

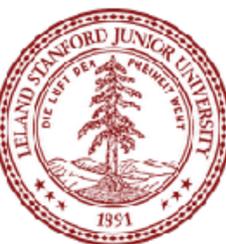
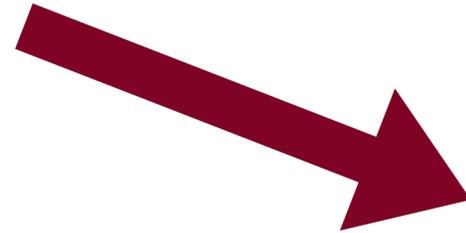
Watch your variable ordering! Better to use r for rows, c for columns.



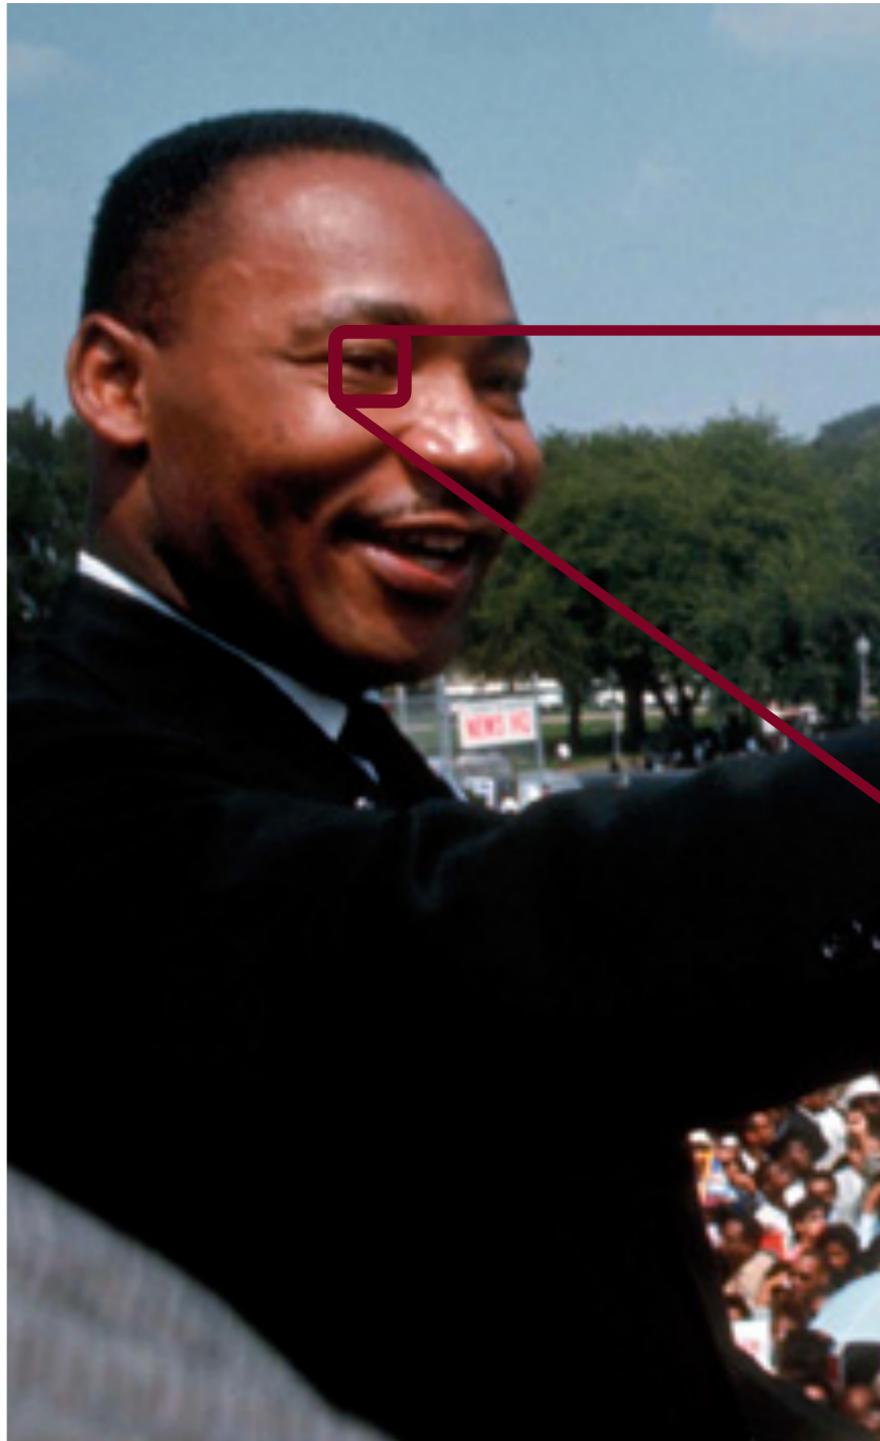
Let's Code Instagram!



Mike Krieger, Stanford Class of 2008
Founder of Instagram



A Color is an int, and an Image is just a Grid<int>!



A Color is an int, and an Image is just a Grid<int>!

Original



Filtered



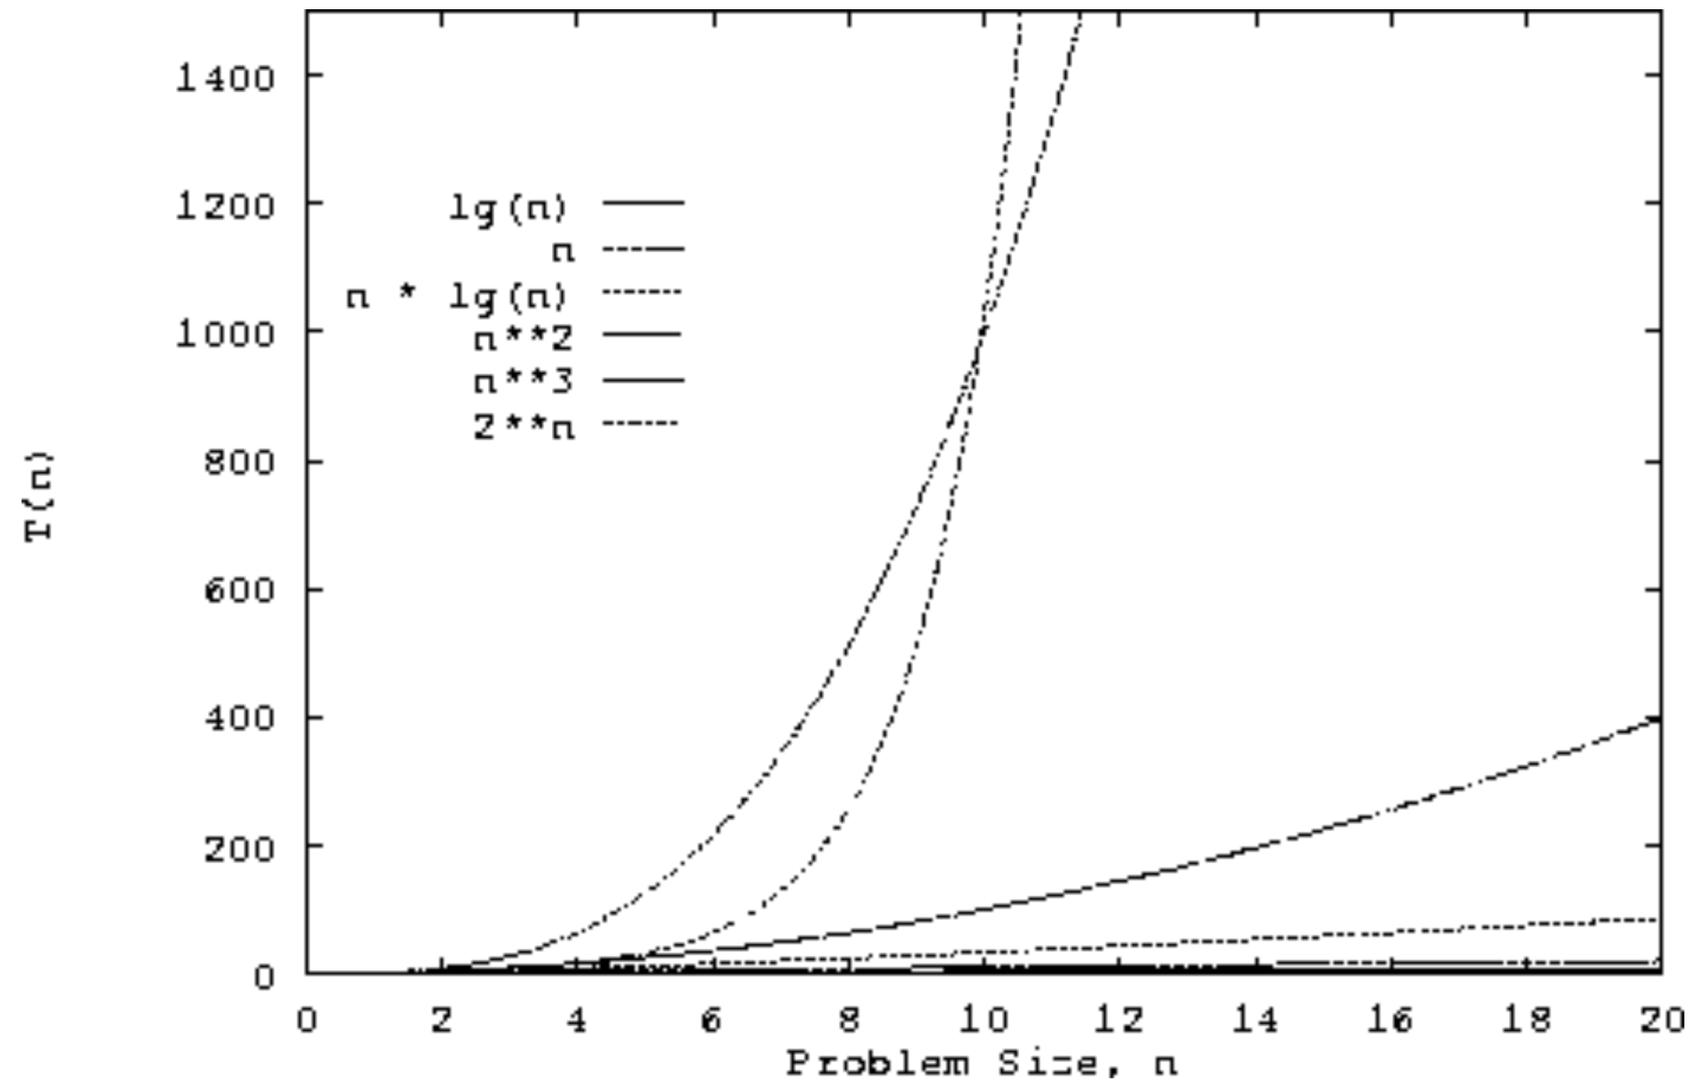
New Palette:



Let's Code!



Computational Complexity



Computational Complexity

How does one go about analyzing programs to compare how the program behaves as it scales? E.g., let's look at a **vectorMax()** function:

```
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;
}
```

What is n ? Why is it important to this function?



Computational Complexity

```
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;
}
```

If we want to see how this algorithm behaves as n changes, we could do the following:

- (1) Code the algorithm in C++
- (2) Determine, for each instruction of the compiled program the time needed to execute that instruction (need assembly language)
- (3) Determine the number of times each instruction is executed when the program is run.
- (4) Sum up all the times we calculated to get a running time.



Computational Complexity

```
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;
}
```

Steps 1-4 on the previous slide...might work, but it is complicated, especially for today's machines that optimize everything "under the hood." (and reading assembly code takes a certain patience).



Assembly Code for vectorMax() function...

```
0x000000010014adf0 <+0>: push    %rbp
0x000000010014adf1 <+1>: mov     %rsp,%rbp
0x000000010014adf4 <+4>: sub     $0x20,%rsp
0x000000010014adf8 <+8>: xor     %esi,%esi
0x000000010014adfa <+10>: mov    %rdi,-0x8(%rbp)
0x000000010014adfe <+14>: mov    -0x8(%rbp),%rdi
0x000000010014ae02 <+18>: callq  0x10014aea0 <std::__1::basic_ostream<char, std::__1::char_traits<char> >::operator<<(long)+32>
0x000000010014ae07 <+23>: mov    (%rax),%esi
0x000000010014ae09 <+25>: mov    %esi,-0xc(%rbp)
0x000000010014ae0c <+28>: mov    -0x8(%rbp),%rdi
0x000000010014ae10 <+32>: callq  0x10014afb0 <std::__1::basic_ostream<char, std::__1::char_traits<char> >::operator<<(long)+304>
0x000000010014ae15 <+37>: mov    %eax,-0x10(%rbp)
0x000000010014ae18 <+40>: movl   $0x1,-0x14(%rbp)
0x000000010014ae1f <+47>: mov    -0x14(%rbp),%eax
0x000000010014ae22 <+50>: cmp    -0x10(%rbp),%eax
0x000000010014ae25 <+53>: jge    0x10014ae6c <vectorMax(Vector<int>&)+124>
0x000000010014ae2b <+59>: mov    -0xc(%rbp),%eax
0x000000010014ae2e <+62>: mov    -0x8(%rbp),%rdi
0x000000010014ae32 <+66>: mov    -0x14(%rbp),%esi
0x000000010014ae35 <+69>: mov    %eax,-0x18(%rbp)
0x000000010014ae38 <+72>: callq  0x10014aea0 <std::__1::basic_ostream<char, std::__1::char_traits<char> >::operator<<(long)+32>
0x000000010014ae3d <+77>: mov    -0x18(%rbp),%esi
0x000000010014ae40 <+80>: cmp    (%rax),%esi
0x000000010014ae42 <+82>: jge    0x10014ae59 <vectorMax(Vector<int>&)+105>
0x000000010014ae48 <+88>: mov    -0x8(%rbp),%rdi
0x000000010014ae4c <+92>: mov    -0x14(%rbp),%esi
0x000000010014ae4f <+95>: callq  0x10014aea0 <std::__1::basic_ostream<char, std::__1::char_traits<char> >::operator<<(long)+32>
0x000000010014ae54 <+100>: mov    (%rax),%esi
0x000000010014ae56 <+102>: mov    %esi,-0xc(%rbp)
0x000000010014ae59 <+105>: jmpq   0x10014ae5e <vectorMax(Vector<int>&)+110>
0x000000010014ae5e <+110>: mov    -0x14(%rbp),%eax
0x000000010014ae61 <+113>: add    $0x1,%eax
0x000000010014ae64 <+116>: mov    %eax,-0x14(%rbp)
0x000000010014ae67 <+119>: jmpq   0x10014ae1f <vectorMax(Vector<int>&)+47>
0x000000010014ae6c <+124>: mov    -0xc(%rbp),%eax
0x000000010014ae6f <+127>: add    $0x20,%rsp
0x000000010014ae73 <+131>: pop    %rbp
0x000000010014ae74 <+132>: retq
```



Algorithm Analysis: Primitive Operations

Instead of those complex steps, we can define *primitive operations* for our C++ code.

- Assigning a value to a variable
- Calling a function
- Arithmetic (e.g., adding two numbers)
- Comparing two numbers
- Indexing into a Vector
- Returning from a function

We assign "1 operation" to each step. We are trying to gather data so we can compare this to other algorithms.



Algorithm Analysis: Primitive Operations

```
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;  
}
```

executed once (2 ops)

executed once (2 ops)

executed $n-1$ times
($2*(n-1)$ ops)

ex. $n-1$ times ($2*(n-1)$ ops)

ex. at most $n-1$ times
($2*(n-1)$ ops), but as few as
zero times

ex. once (1 op)



executed
once (1 op)

ex. n times (n ops)

Algorithm Analysis: Primitive Operations

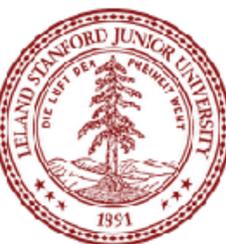
Summary:

Primitive operations for **vectorMax ()** :

at least: $2 + 2 + 1 + n + 4 * (n - 1) + 1 = 5n + 2$

at most: $2 + 2 + 1 + n + 6 * (n - 1) + 1 = 7n$

i.e., if there are n items in the Vector, there are between $5n+2$ operations and $7n$ operations completed in the function.



Algorithm Analysis: Primitive Operations

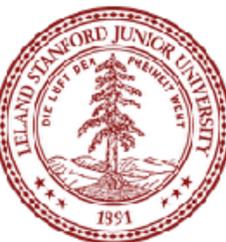
Summary:

Primitive operations for **vectorMax ()** :

best case: $5n + 2$

worst case: $7n$

In other words, we can get a "best case" and "worst case" count



Algorithm Analysis: Simplify!

Do we *really* need this much detail? Nope!

Let's simplify: we want a "big picture" approach.

It is enough to know that `vectorMax()` grows

linearly proportionally to n

In other words, as the number of elements increases, the algorithm has to do proportionally more work, and that relationship is linear. 8x more elements? 8x more work.



Algorithm Analysis: Big-O

Our simplification uses a mathematical construct known as “Big-O” notation — think “O” as in “on the Order of.”

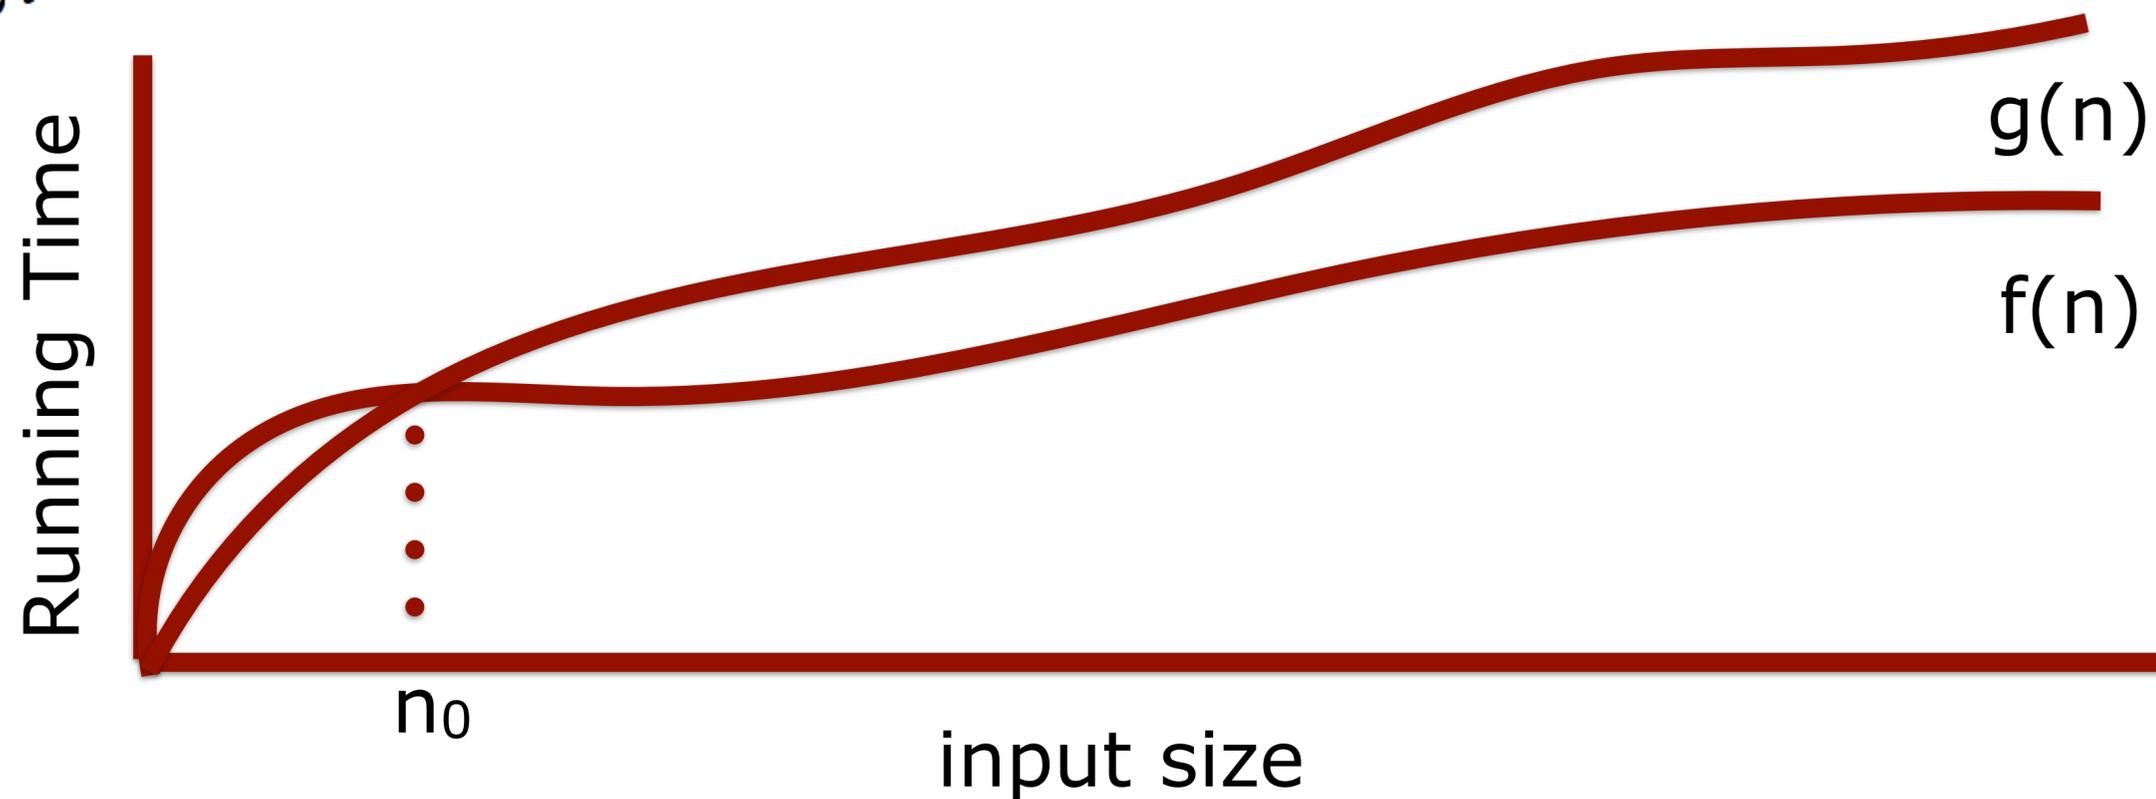
Wikipedia:

“Big-O notation describes the limiting behavior of a function when the argument tends towards a particular value or infinity, usually in terms of simpler functions.”



Algorithm Analysis: Big-O

Let $f(n)$ and $g(n)$ be functions mapping nonnegative integers to real numbers. We say that $f(n)$ is $O(g(n))$ if there is a real constant $c > 0$ and an integer constant $n_0 \geq 1$, such that $f(n) \leq cg(n)$ for every integer $n \geq n_0$. This definition is often referred to as the “big-Oh” notation. We can also say, “ $f(n)$ is *order* $g(n)$.”



Algorithm Analysis: Big-O

Dirty little trick for figuring out Big-O: look at the number of steps you calculated, throw out all the constants, find the “biggest factor” and that’s your answer:

$$5n + 2 \text{ is } O(n)$$

Why? Because constants are not important at this level of understanding.



Algorithm Analysis: Big-O

We will care about the following functions that appear often in data structures:

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

When you are deciding what Big-O is for an algorithm or function, simplify until you reach one of these functions, and you will have your answer.



Algorithm Analysis: Big-O

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

Practice: what is Big-O for this function?

$$20n^3 + 10n \log n + 5$$

Answer: $O(n^3)$

First, strip the constants: $n^3 + n \log n$

Then, find the biggest factor: n^3



Algorithm Analysis: Big-O

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

Practice: what is Big-O for this function?

$$2000 \log n + 7n \log n + 5$$

Answer: $O(n \log n)$

First, strip the constants: $\log n + n \log n$

Then, find the biggest factor: $n \log n$



Algorithm Analysis: Back to 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;
}
```

When you are analyzing an algorithm or code for its *computational complexity* using Big-O notation, you can ignore the primitive operations that would contribute less-important factors to the run-time. Also, you always take the *worst case* behavior for Big-O.



Algorithm Analysis: Back to 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;  
}
```

When you are analyzing an algorithm or code for its *computational complexity* using Big-O notation, you can ignore the primitive operations that would contribute less-important factors to the run-time. Also, you always take the *worst case* behavior for Big-O.

So, for vectorMax(): ignore the original two variable initializations, the return statement, the comparison, and the setting of currentMax in the loop.



Algorithm Analysis: Back to 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;  
}
```

So, for vectorMax(): ignore the original two variable initializations, the return statement, the comparison, and the setting of currentMax in the loop.

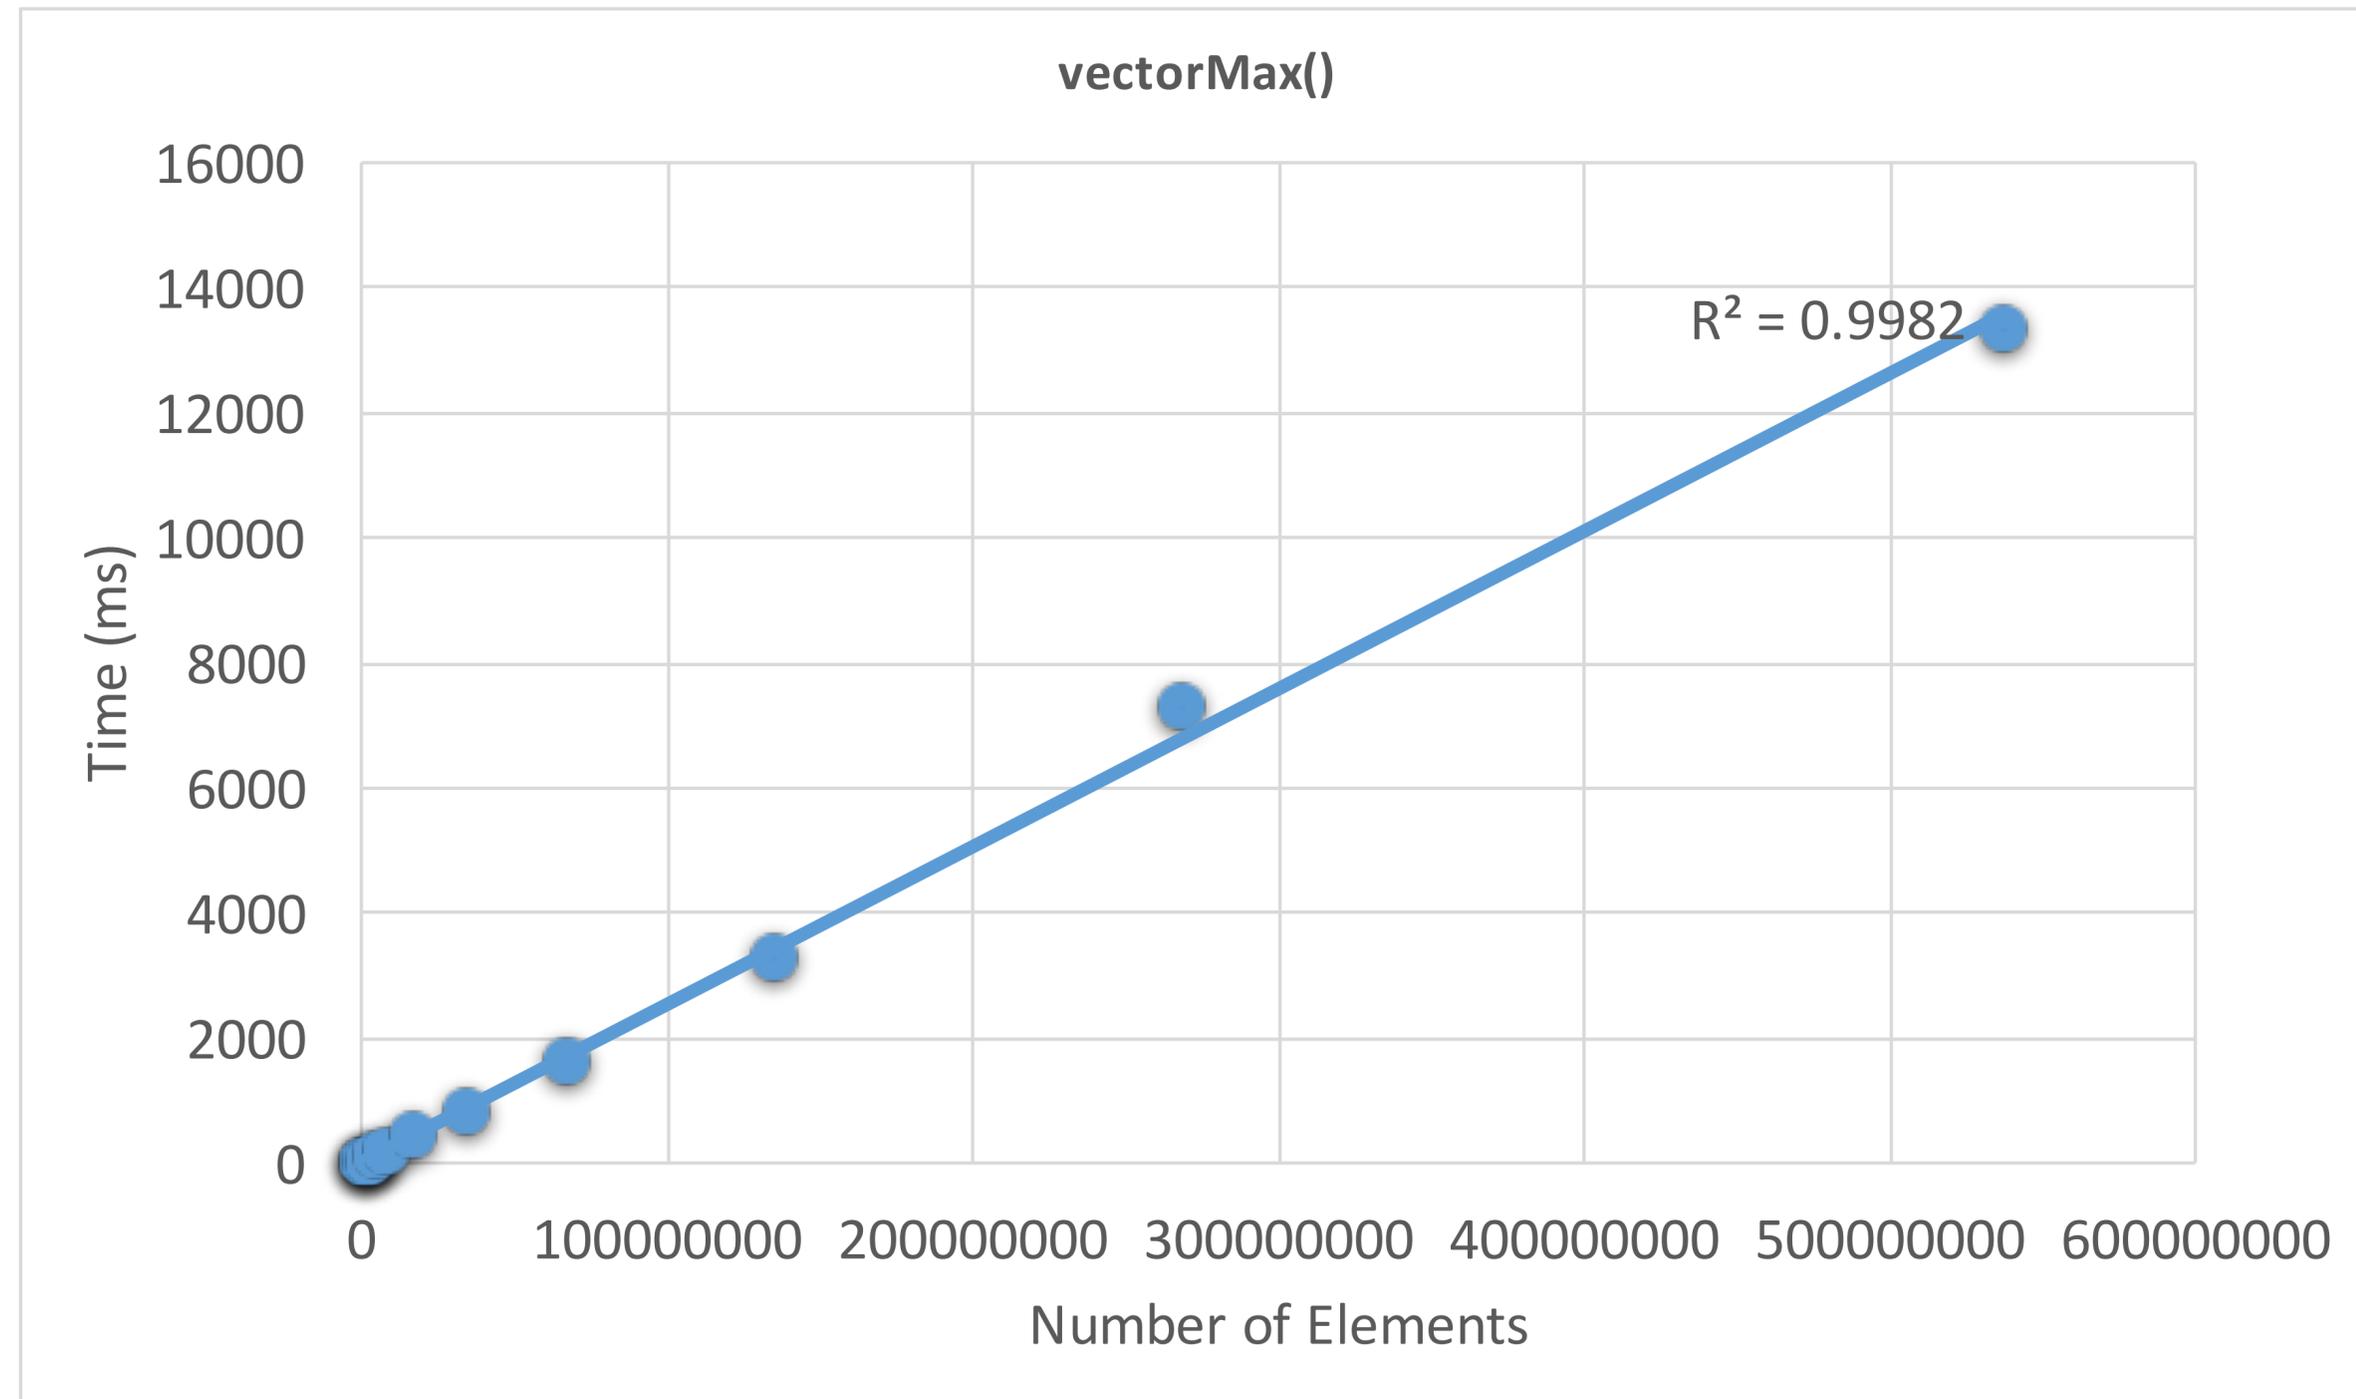
Notice that the important part of the function is the fact that the loop conditions will change with the size of the array: for each extra element, there will be one more iteration. This is a *linear* relationship, and therefore $O(n)$.



Algorithm Analysis: Back to vectorMax()

Data: In the lecture code, you will find a test program for vectorMax(), which runs the function on an increasing (by powers of two) number of vector elements. This is the data I gathered from my computer.

As you can see, it's a linear relationship!



Algorithm Analysis: Nested Loops

```
int nestedLoop1(int n) {  
    int result = 0;  
    for (int i=0; i<n; i++) {  
        for (int j=0; j<n; j++) {  
            result++;  
        }  
    }  
    return result;  
}
```

**Also go through the outer loop
n times**

Inner loop complexity: $O(n)$

**Total complexity: $O(n^2)$
(quadratic)**

In general, we don't like $O(n^2)$ behavior! Why?

As an example: let's say an $O(n^2)$ function takes 5 seconds for a container with 100 elements.

How much time would it take if we had 1000 elements?

500 seconds! This is because 10x more elements is (10^2) x more time!



Algorithm Analysis: Nested Loops

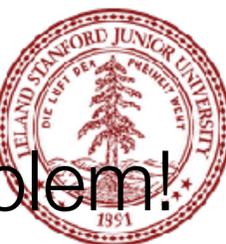
```
int nestedLoop1(int n) {  
    int result = 0;  
    for (int i=0; i<n; i++) {  
        for (int j=0; j<n; j++) {  
            for (int k=0; k<n; k++)  
                result++;  
        }  
    }  
    return result;  
}
```

What would the complexity be of a 3-nested loop?

Answer: n^3 (polynomial)

In real life, this comes up in 3D imaging, video, etc., and it is **slow!**

Graphics cards are built with hundreds or thousands of processors to tackle this problem!



Algorithm Analysis: Linear Search

```
void linearSearchVector(Vector<int> &vec, int numToFind){
    int numCompares = 0;
    bool answer = false;
    int n = vec.size();

    for (int i = 0; i < n; i++) {
        numCompares++;
        if (vec[i]==numToFind) {
            answer = true;
            break;
        }
    }
    cout << "Found? " << (answer ? "True" : "False") << ", "
         << "Number of compares: " << numCompares << endl << endl;
}
```

Best case? $O(1)$

Worst case? $O(n)$

Complexity: $O(n)$ (linear, worst case)

You have to walk through the entire vector one element at a time.



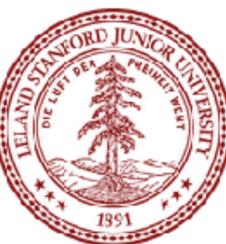
Algorithm Analysis: *Binary* Search

There is another type of search that we can perform on a list that is in order: binary search (as seen in 106A!)

If you have ever played a "guess my number" game before, you will have implemented a binary search, if you played the game efficiently!

The game is played as follows:

- one player thinks of a number between 0 and 100 (or any other maximum).
- the second player guesses a number between 1 and 100
- the first player says "higher" or "lower," and the second player keeps guessing until they guess correctly.



Algorithm Analysis: *Binary* Search

The most efficient guessing algorithm for the number guessing game is simply to choose a number that is between the high and low that you are currently bound to. Example:

bounds: 0, 100

guess: 50 (no, the answer is lower)

new bounds: 0, 49

guess: 25 (no, the answer is higher)

new bounds: 26, 49

guess: 38

etc.

With each guess, the search space is *divided into two*.



Algorithm Analysis: Binary Search

```
void binarySearchVector(Vector<int> &vec, int numToFind) {
    int low=0;
    int high=vec.size()-1;
    int mid;
    int numCompares = 0;
    bool found=false;
    while (low <= high) {
        numCompares++;
        //cout << low << ", " << high << endl;
        mid = low + (high - low) / 2; // to avoid overflow
        if (vec[mid] > numToFind) {
            high = mid - 1;
        }
        else if (vec[mid] < numToFind) {
            low = mid + 1;
        }
        else {
            found = true;
            break;
        }
    }
    cout << "Found? " << (found ? "True" : "False") << ", " <<
    "Number of compares: " << numCompares << endl << endl;
}
```

Best case? $O(1)$

Worst case? $O(\log n)$

Complexity: $O(\log n)$
(logarithmic, worst case)

Technically, this is $O(\log_2 n)$, but we will not worry about the base.

The general rule for determining if something is logarithmic: if the problem is one of "divide and conquer," it is logarithmic. If, at each stage, the problem size is cut in half (or a third, etc.), it is logarithmic.



Algorithm Analysis: Constant Time

When an algorithm's time is *independent* of the number of elements in the container it holds, this is *constant time* complexity, or $O(1)$. We love $O(1)$ algorithms! Examples include (for efficiently designed data structures):

- Adding or removing from the *end* of a Vector.
- Later in the course:
 - Pushing onto a stack or popping off a stack.
 - Enqueuing or dequeuing from a queue.
 - Inserting or searching for a value in a hash table



Algorithm Analysis: Exponential Time

There are a number of algorithms that have *exponential* behavior. If we don't like quadratic or polynomial behavior, we *really* don't like exponential behavior.

Example: what does the following beautiful recursive function do?

```
long mysteryFunc(int n) {  
    if (n == 0) {  
        return 0;  
    }  
    if (n == 1) {  
        return 1;  
    }  
    return mysteryFunc(n-1) + mysteryFunc(n-2);  
}
```

This is the *fibonacci sequence*! 0, 1, 1, 2, 3, 5, 8, 13, 21 ...



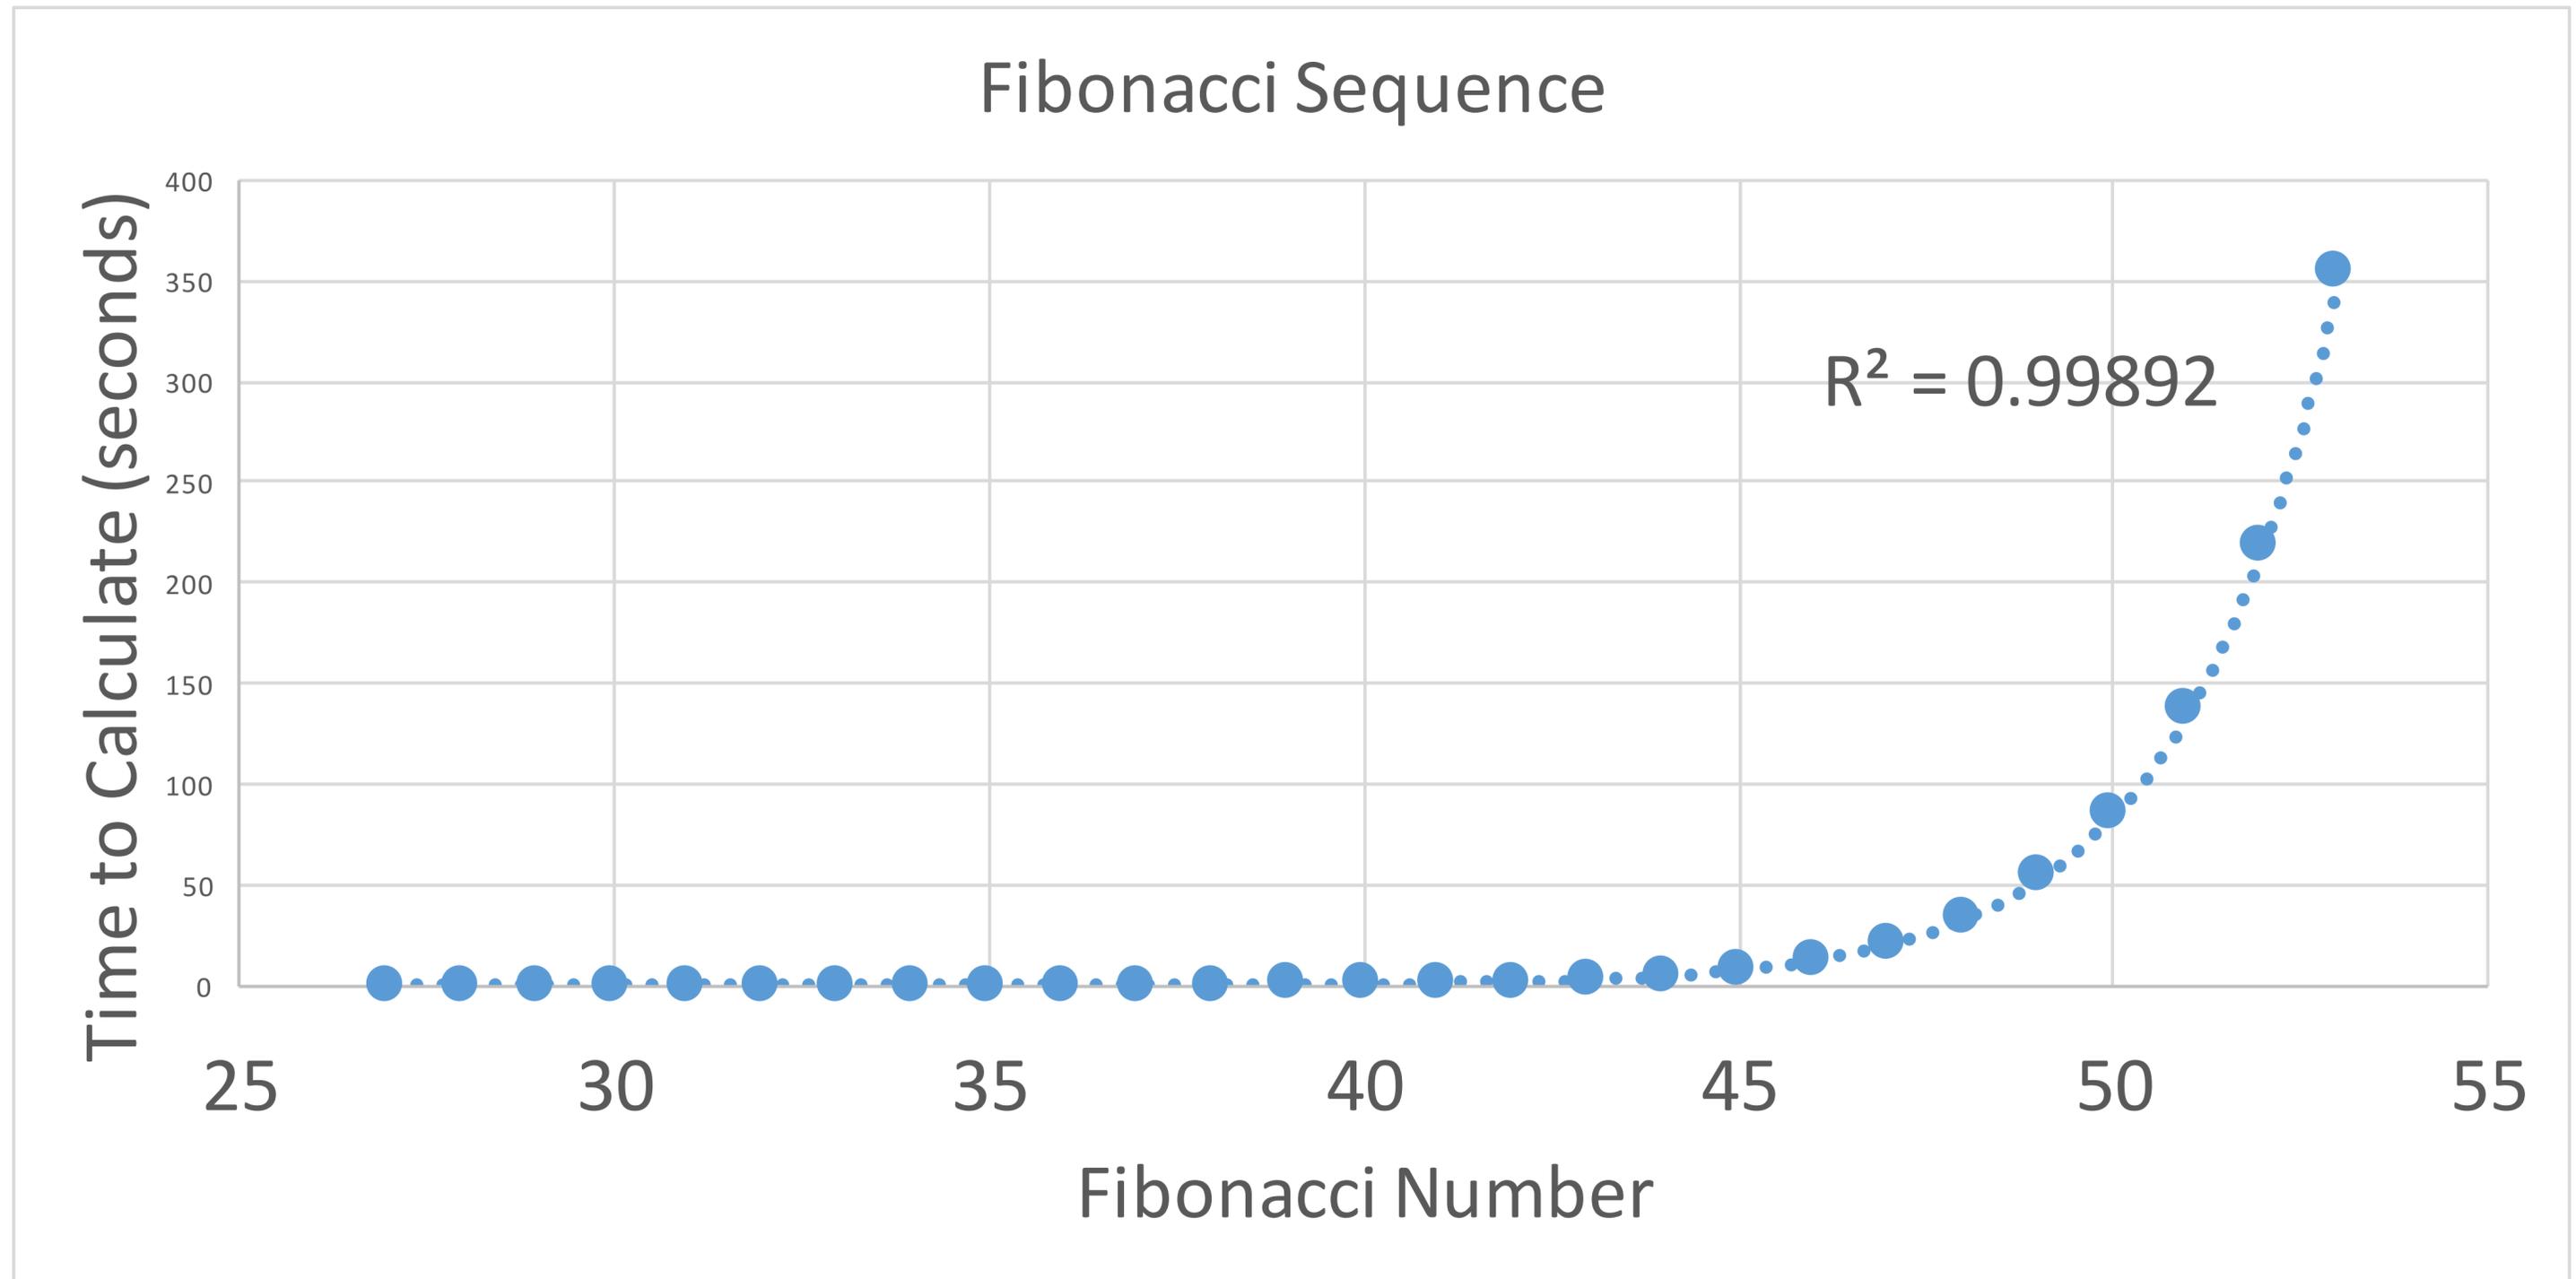
Algorithm Analysis: Exponential Time

```
long fibonacci(int n) {  
    if (n == 0) {  
        return 0;  
    }  
    if (n == 1) {  
        return 1;  
    }  
    return fibonacci(n-1) + fibonacci(n-2);  
}
```

Beautiful, but a flawed algorithm! Yes, it works, but why is it flawed? Let's look at the call tree for fib(6):



Fibonacci Sequence Time to Calculate Recursively



Ramifications of Big-O Differences

Some numbers:

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

<i>constant</i>	<i>logarithmic</i>	<i>linear</i>	<i>n log n</i>	<i>quadratic</i>	<i>polynomial (n³)</i>	<i>exponential (a=2)</i>
$O(1)$	$O(\log n)$	$O(n)$	$O(n \log n)$	$O(n^2)$	$O(n^k)$ ($k \geq 1$)	$O(a^n)$ ($a > 1$)
1ns	10ns	1microsec	10microsec	1millisec	1 sec	10^{292} years



Ramifications of Big-O Differences

Some numbers:

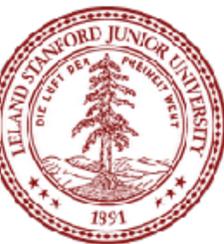
If we have an algorithm that has 1000 elements, and the $O(\log n)$ version runs in 10 milliseconds...

<i>constant</i>	<i>logarithmic</i>	<i>linear</i>	<i>$n \log n$</i>	<i>quadratic</i>	<i>polynomial (n^3)</i>	<i>exponential ($a=2$)</i>
$O(1)$	$O(\log n)$	$O(n)$	$O(n \log n)$	$O(n^2)$	$O(n^k)$ ($k \geq 1$)	$O(a^n)$ ($a > 1$)
1ms	10ms	1sec	10sec	17 minutes	277 hours	heat death of the universe



Recap (Big O)

- Asymptotic Analysis / Big-O / Computational Complexity
 - We want a "big picture" assessment of our algorithms and functions
 - We can ignore constants and factors that will contribute less to the result!
 - We most often care about *worst case* behavior.
 - We love $O(1)$ and $O(\log n)$ behaviors!
- Big-O notation is useful for determining how a particular algorithm behaves, but be careful about making comparisons between algorithms -- sometimes this is helpful, but it can be misleading.
- Algorithmic complexity can determine the difference between running your program over your lunch break, or waiting until the Sun becomes a Red Giant and swallows the Earth before your program finishes -- that's how important it is!



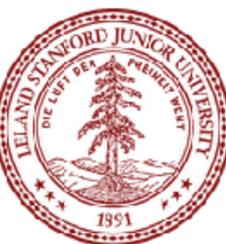
References and Advanced Reading (Big O)

•References:

- Wikipedia on BigO: https://en.wikipedia.org/wiki/Big_O_notation
- Binary Search: https://en.wikipedia.org/wiki/Binary_search_algorithm
- Fibonacci numbers: https://en.wikipedia.org/wiki/Fibonacci_number

•Advanced Reading:

- Big-O Cheat Sheet: <http://bigocheatsheet.com>
- More details on Big-O: http://web.mit.edu/16.070/www/lecture/big_o.pdf
- More details: http://dev.tutorialspoint.com/data_structures_algorithms/asymptotic_analysis.htm
- GPUs and GPU-Accelerated computing: <http://www.nvidia.com/object/what-is-gpu-computing.html>
- Video on Fibonacci sequence: <https://www.youtube.com/watch?v=Nu-IW-Ifyec>
- Fibonacci numbers in nature: <http://www.maths.surrey.ac.uk/hosted-sites/R.Knott/Fibonacci/fibnat.html>



References and Advanced Reading (Vectors and Grids)

- **References:**

- Stanford Vector Class: <http://stanford.edu/~stepp/cppdoc/Vector-class.html>

- Stanford Grid Class: <http://stanford.edu/~stepp/cppdoc/Grid-class.html>

- **Advanced Reading:**

- Standard Template Library vector class (some different functions!): <http://www.cplusplus.com/reference/vector/vector/>

- Adobe Photoshop on Wikipedia: https://en.wikipedia.org/wiki/Adobe_Photoshop

