

CS107 Winter 2020, Lecture 1

Welcome to CS107!

reading:

General Information handout
Bryant & O'Hallaron, Ch. 1

Plan For Today

- Introduction
- CS107 Course Policies
- Unix and the Command Line
- Getting Started With C

Plan For Today

- **Introduction**
- CS107 Course Policies
- Unix and the Command Line
- Getting Started With C

What is CS107?

- The CS106 series:
 - Taught you how to solve problems as a programmer
 - Many times, CS106 instructors had to say, “just don’t worry about that,” or “it probably doesn’t make sense why that happens, but ignore it for now,” or “just type this to fix it”
- CS107 finally takes you **behind the scenes**:
 - Not quite down to hardware or physics/electromagnetism (that’s for later...)
 - It’s how things work **inside C++/Python/Java**, and how your programs map onto the components of computer systems
 - Who doesn’t love being in on a secret? **It just feels good to know!**

What is CS107?



Computer Organization and Systems

- How languages like C++/Java/Python **represent data** under the hood
- How programming structures are encoded in **bits and bytes**
- How to efficiently **manipulate and manage memory**
- How computers **compile** programs
- Uses the **C** programming language
- Programming **style** and software development practices

CS107 Learning Goals

The goals for CS107 are for students to gain **mastery** of

- writing C programs with complex use of memory and pointers
- an accurate model of the address space and compile/runtime behavior of C programs

to achieve **competence** in

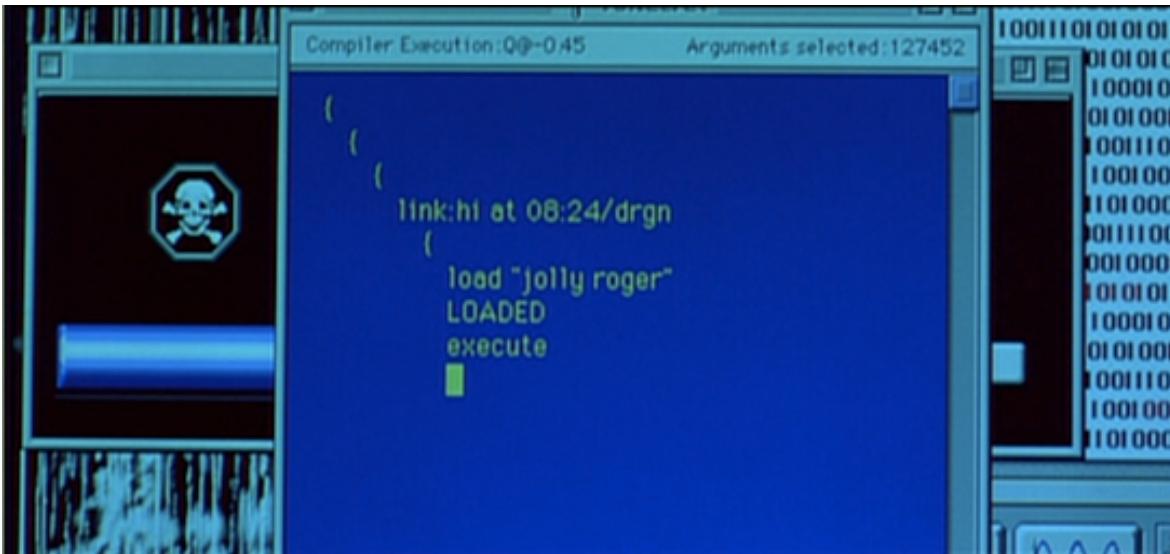
- translating C to/from assembly
- writing programs that respect the limitations of computer arithmetic
- identifying bottlenecks and improving runtime performance
- working effectively in a Unix development environment

and have **exposure** to

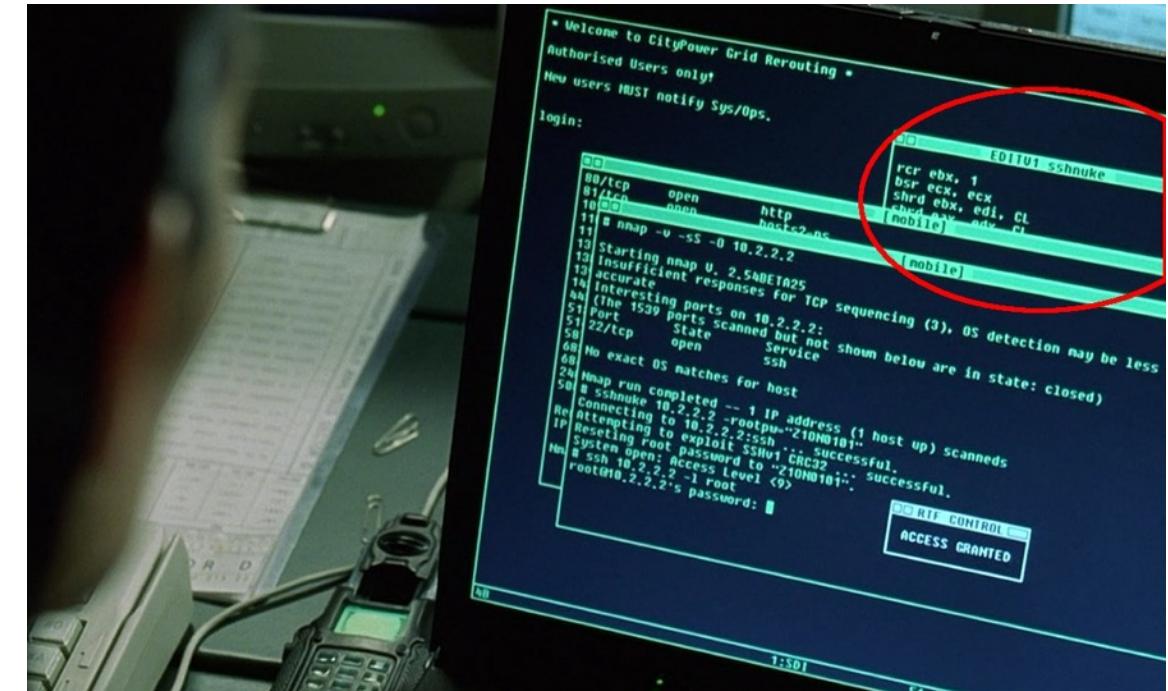
- a working understanding of the basics of computer architecture

CS107 Learning Goals

(also learn to identify legitimate programmer scenes in old movies)



Jeff Goldblum's character saving the world by uploading a virus to the alien mothership
Independence Day, 1996



Trinity saving the world by hacking into the power grid using Nmap Network Scanning
The Matrix Reloaded, 2003

Course Topic Overview

1. **Bits and Bytes** - *How can a computer represent integer numbers?*
2. **Chars and C-Strings** - *How can a computer represent and manipulate more complex data like text?*
3. **Pointers, Stack and Heap** – *How can we effectively manage all types of memory in our programs?*
4. **Generics** - *How can we use our knowledge of memory and data representation to write code that works with any data type?*
5. **Floats** - *How can a computer represent floating point numbers in addition to integer numbers?*
6. **Assembly** - *How does a computer interpret and execute C programs?*
7. **Heap Allocators** - *How do core memory-allocation operations like malloc and free work?*

Companion Class: CS107A

- **CS107A** is an extra 1-unit “Pathfinders” or “ACE” section with additional course support, practice and instruction.
- Meets for an additional weekly section and has additional review sessions
- Entry by application – see the FAQ on the course website for details



Jennie Yang

Plan For Today

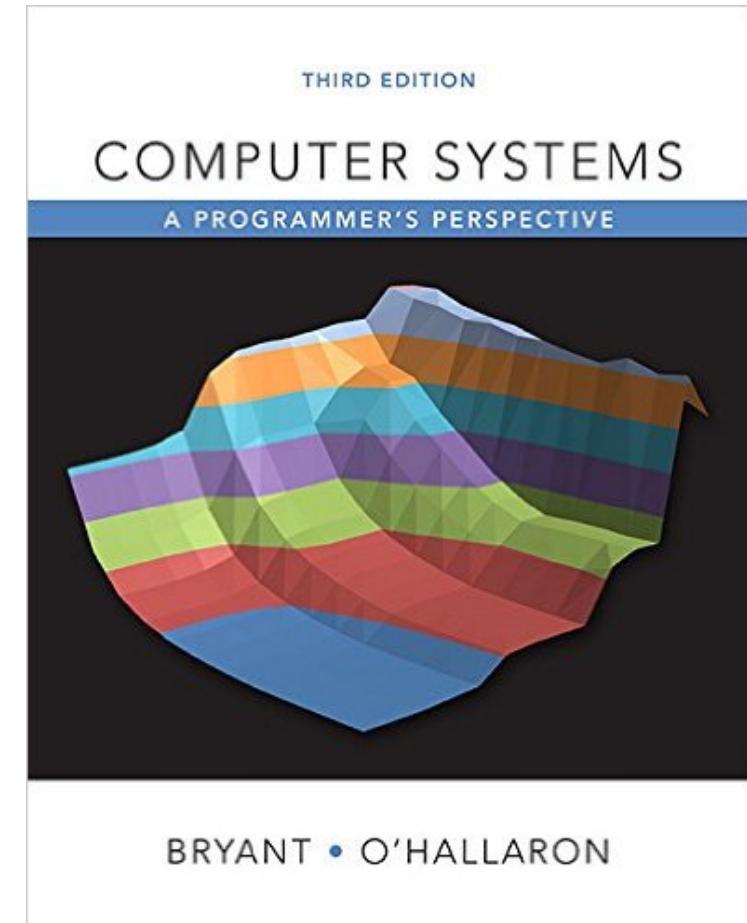
- Introduction
- **CS107 Course Policies**
- Unix and the Command Line
- Getting Started With C

Course Website

cs107.stanford.edu

Textbooks

- *Computer Systems: A Programmer's Perspective*
by Bryant & O'Hallaron, **3rd Edition**
 - **3rd edition matters** – important updates to course materials
- A C programming reference of your choice
 - *The C Programming Language* by Kernighan and Ritchie (free link on course website Resources page)
 - Other C programming books, websites, or reference sheets



Website vs Textbook

- **Course website:**

- Lecture code, slides
- Assignment writeups
- How-to guides
- Walkthrough videos

✓ Great reference guide

✓ Review details from lecture

⚠ Can be overwhelming if you
you don't know where to start

- **Textbook (B&O, K&R):**

- Words written by someone other than CS107 instructors
- Practice problems (conceptual, coding)
- Lots of diagrams, new examples

✓ Tells a comprehensive story of
the course material

✓ Has lots of visuals to
understand concepts better

⚠ Goes more in-depth into material
than what is expected for this course

The textbook (and C programming references) are **very** good resources in this course,
especially post-midterm.

Grading

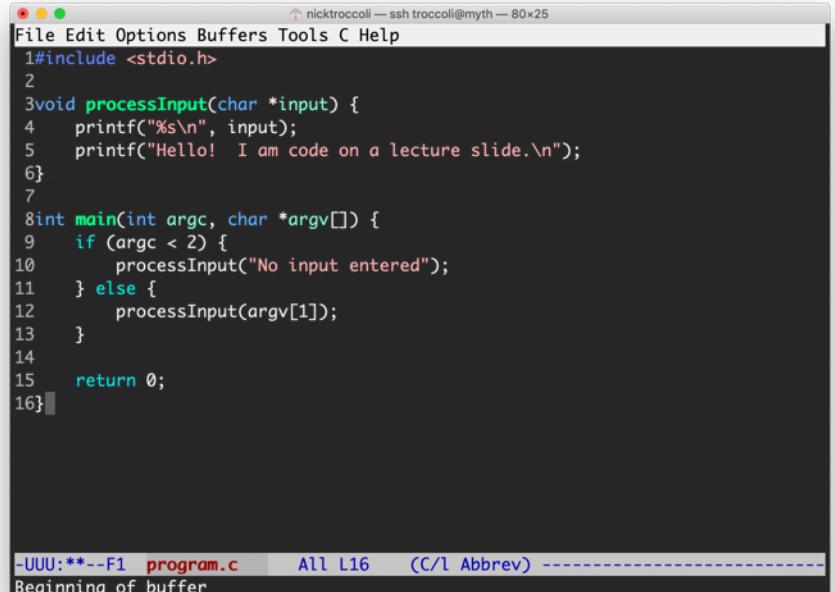
****	35%	Assignments
**	15%	Lab Participation
**	17%	Midterm Exam
****	33%	Final Exam

Grading

***	35%	Assignments
**	15%	Lab Participation
**	17%	Midterm Exam
****	33%	Final Exam

Assignments

- 8 programming assignments completed individually using **Unix command line tools**
 - Free software, pre-installed on Myth machines / available on course website
 - We will give out starter projects for each assignment
- Graded on **functionality** (behavior) and **style** (elegance)
 - Functionality graded using *automated tools*, given as point score
 - Style graded via *automated tests* and TA code review, given as bucket score
 - Grades returned via course website



A screenshot of a terminal window titled 'nicktroccoli — ssh troccoli@myth — 80x25'. The window shows a C program named 'program.c' with the following code:

```
File Edit Options Buffers Tools C Help
1#include <stdio.h>
2
3void processInput(char *input) {
4    printf("%s\n", input);
5    printf("Hello! I am code on a lecture slide.\n");
6}
7
8int main(int argc, char *argv[]) {
9    if (argc < 2) {
10        processInput("No input entered");
11    } else {
12        processInput(argv[1]);
13    }
14
15    return 0;
16}
```

The terminal window also displays the status bar with the text: '-UUU:***--F1 program.c All L16 (C/l Abbrev) ----- Beginning of buffer'.

The Style Bucket System

+	An outstanding job; could be used as course example code for good style.
ok	A good job; solid effort, but also opportunities for improvement.
-	Shows some effort and understanding but has larger problems that should be focused on.
--	Little effort; incomplete or mostly non-functional.
0	No work submitted, or barely any changes from the starter assignment.

Late Policy

- Submitting by the assignment deadline typically earns an **extra-credit on-time bonus, usually ~5%**.
- If you miss the deadline, there may be a grace period for late submissions, typically 48 hours, but submitting during the grace period **does not earn any on-time bonus**.
- "Pre-granted grace period" – additional extensions granted only in *very special* circumstances. **Instructor** must approve extensions.

Grading

**** 35% Assignments

** 15% Lab Participation

** 17% Midterm Exam

**** 33% Final Exam

Lab Sections

- Weekly 1 hour 50-minute labs led by a CA, starting *next* week, offered on Tuesdays, Wednesdays and Thursdays.
- Hands-on practice in pairs at computers with lecture material and course concepts.
- Graded on attendance + participation (verified by submitting work at the end)
- Lab signups open **Friday 01/10 at 10:30AM** and are **first-come-first-served**. Sign up on the labs page of the course website.

Grading

**** 35% Assignments

** 15% Lab Participation

** 17% Midterm Exam

**** 33% Final Exam

Exams

- **Midterm** – Friday, February 14th, 12:30PM-2:20PM (during *full class period*)
- **Final** – Friday, March 20th, 3:30PM-6:30PM
- You **MUST** be able to take both exams at the scheduled time, except for university athletics or OAE accommodations.
- Both exams are closed-book, closed-notes, but you may bring in one double-sided page of notes. You will also be provided with a syntax reference sheet.
- Exams are administered electronically via BlueBook.

Grading

****	35%	Assignments
**	15%	Lab Participation
**	17%	Midterm Exam
****	33%	Final Exam

Getting Help

- Post on **Piazza**
 - Online discussion forum for students; post questions, answer other students' questions
 - Best for course material discussions, course policy questions or general assignment questions (DON'T POST ASSIGNMENT CODE!)
- Visit us at **helper hours**
 - Scheduled throughout the week; schedule be posted to course website later this week
 - Best for **coding/debugging questions**, or **longer course material discussions**
- **Email** the Course Staff
 - cs107@cs.stanford.edu – please do not email CAs individually
 - Best for **private matters** (e.g. grading questions, OAE accommodations).

Stanford Honor Code

- The **Honor Code** is an undertaking of the students, individually and collectively:
 - that they will not give or receive aid in examinations; that they will not give or receive unpermitted aid in class work, in the preparation of reports, or in any other work that is to be used by the instructor as the basis of grading;
 - that they will do their share and take an active part in seeing to it that others as well as themselves uphold the spirit and letter of the Honor Code.
- The faculty on its part manifests its confidence in the honor of its students by refraining from proctoring examinations and from taking unusual and unreasonable precautions to prevent the forms of dishonesty mentioned above. The faculty will also avoid, as far as practicable, academic procedures that create temptations to violate the Honor Code.
- While the faculty alone has the right and obligation to set academic requirements, the students and faculty will work together to establish optimal conditions for honorable academic work.

see also: <http://honorcode.stanford.edu/>

It is your responsibility to ensure you have read and are familiar with the honor code guidelines posted on the main page of the CS107 course website. Please read them and come talk to us if you have any questions or concerns.

Honor Code and CS107

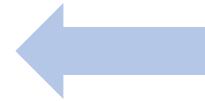
- Please help us ensure academic integrity:
 - Indicate any assistance received on HW (books, friends, etc.).
 - Do not look at other people's solution code or answers
 - Do not give your solutions to others or post them on the web or our Piazza forum.
 - Report any inappropriate activity you see performed by others.
- Assignments are checked regularly for similarity with help of software tools.
- If you realize that you have made a mistake, you may retract your submission to any assignment at any time, no questions asked.
- If you need help, please contact us and we will help you.
 - We do not want you to feel any pressure to violate the Honor Code in order to succeed in this course.

OAE Accommodations

- Please email the course staff as soon as possible with any accommodations you may need for the course.
- We are eager to do everything we can to support you and make you successful in CS107!

Questions?

Course structure

- Lectures: understand concepts, see live lecture demos
- Labs: learn tools, study code, discuss with peers  **Great preview of homework!**
- Assignments: build programming skills, synthesize lecture/lab content

	Monday	Tues-Thurs	Friday
Week N			Lecture: part A
Week N+1	Lecture: part B assignN released	Lab	
Week N+2	assignN due		

- **assign0**: out today, due next Monday (covers today's lecture)

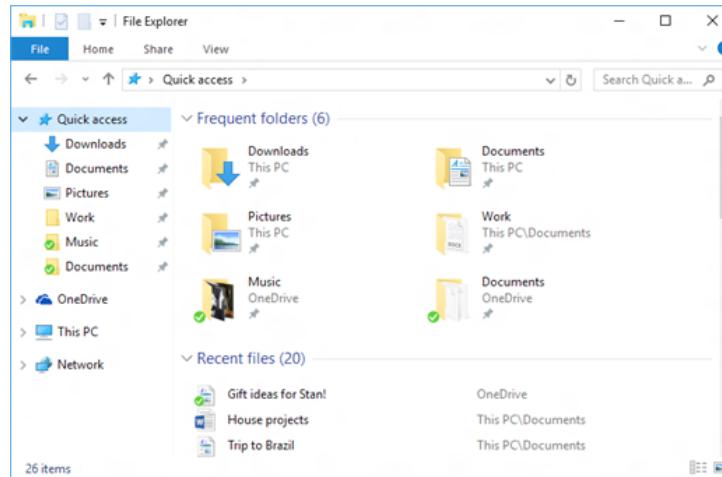
Plan For Today

- Introduction
- CS107 Course Policies
- Unix and the Command Line
- Getting Started With C

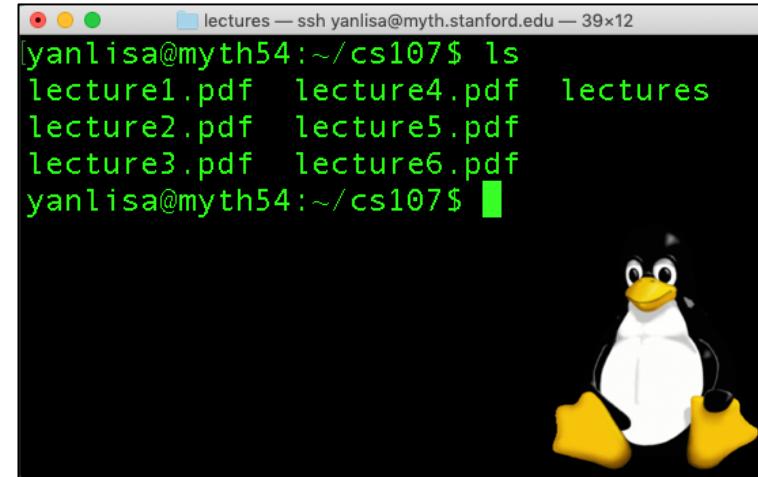
What is Unix?

- Which of the following is a **Unix** system? (select all that apply)

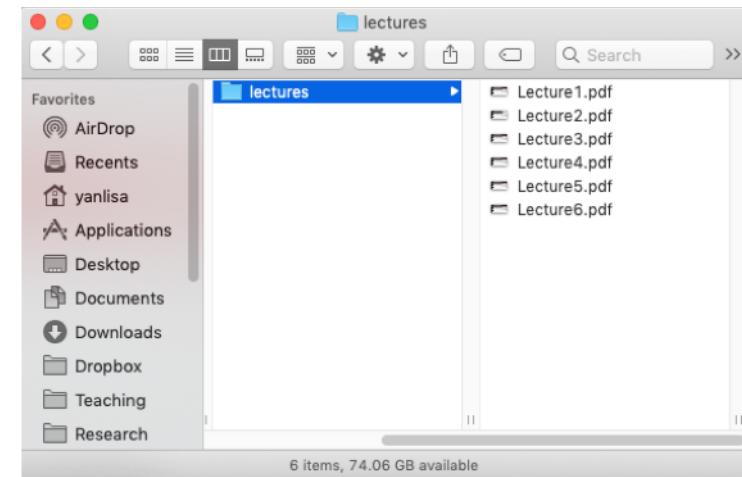
A. Windows



B. Linux



C. Mac OS X

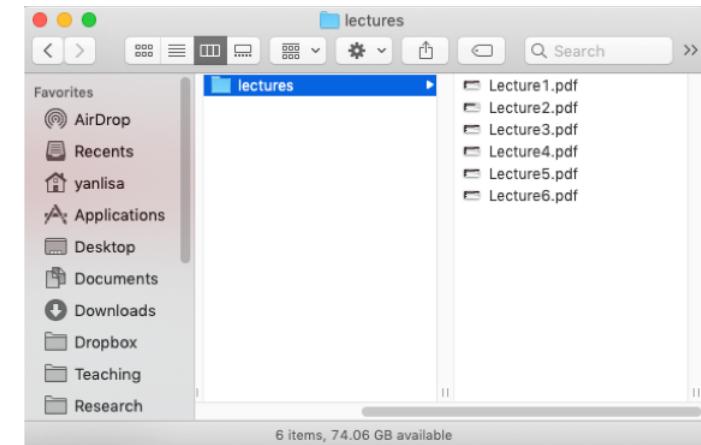


- **Unix**: a set of standards and tools commonly used in software development.
 - **Mac OS X** and **Linux** are operating systems built on top of Unix
- You can navigate a Unix system using the **command line** ("terminal")
- Every Unix system works with the same tools and commands

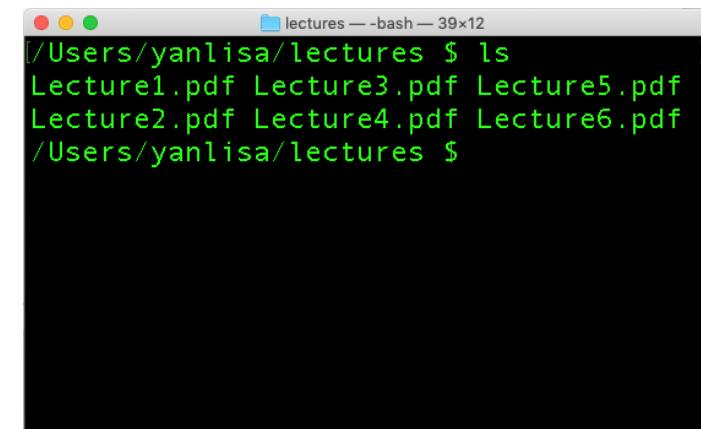


What is the Command Line?

- The **command-line** is a text-based interface (i.e., **terminal** interface) to navigate a computer, instead of a Graphical User Interface (GUI).
- Just like a GUI file explorer interface, a terminal interface can:
 - show you a **specific place** on your computer at any given time.
 - let you go **into folders** and **out of folders**.
 - let you **create new files** and **edit** them.
 - let you **execute programs**.



Graphical User Interface

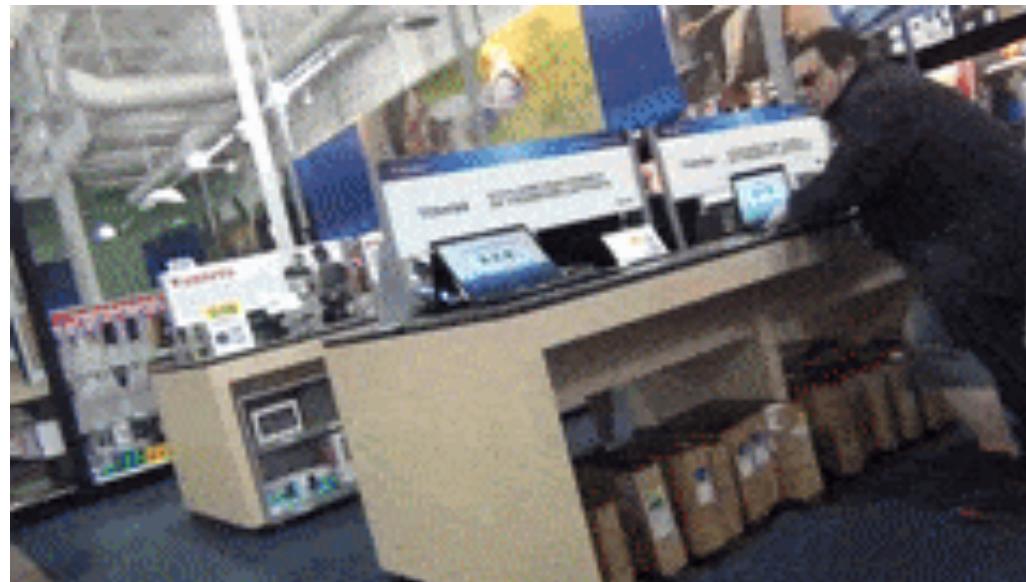


Text-based interface

Why Use Unix / the Command Line?



The Matrix (1999)



Why Use Unix / the Command Line?

- You can navigate almost any device using the same tools and commands:
 - Servers
 - Laptops and desktops
 - Embedded devices (Raspberry Pi, etc.)
 - Mobile Devices (Android, etc.)
- Used frequently by software engineers:
 - **Web development**: running servers and web tools on servers
 - **Machine learning**: processing data on servers, running algorithms
 - **Systems**: writing operating systems, networking code and embedded software
 - **Mobile Development**: running tools, managing libraries
 - And more...
- We'll use Unix and the command line to implement and execute our programs.

Demo: Using Unix and the Command Line



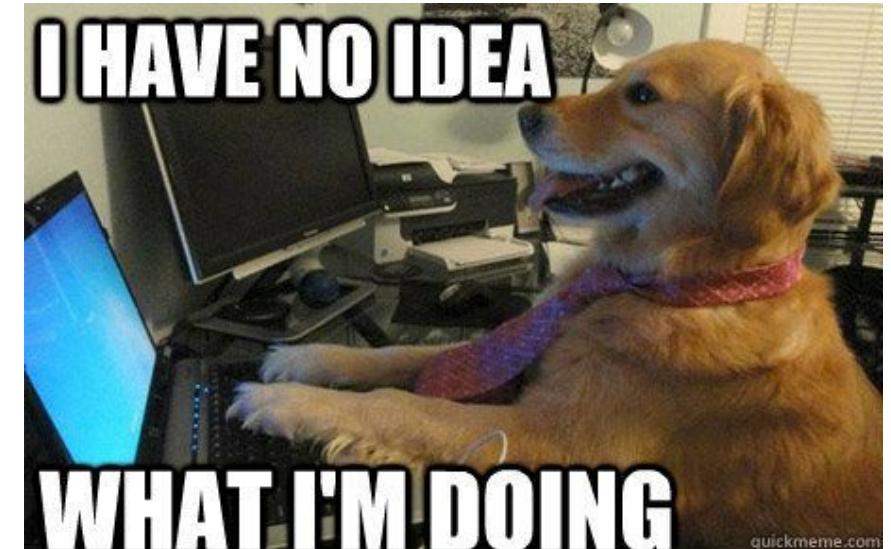
Unix Commands Recap

- **cd** – change directories (..)
- **ls** – list directory contents
- **mkdir** – make directory
- **emacs** – open text editor
- **rm** – remove file or folder
- **man** – view manual pages

See the Resources page of the course website for more commands, and a complete reference.

Practice builds confidence

- Using Unix and the command line can be intimidating at first:
 - It looks really retro!
 - How do I know what to type?
- It's like learning a new language:
 - At the beginning, you might have to constantly look things up (**Resources** page on course website)
 - It's important you spend as much time as you can (during lab and assignments) building muscle memory with the tools.
- We'll continue to work on your Unix while we introduce C.



Plan For Today

- Introduction
- CS107 Course Policies
- Unix and the Command Line
- **Getting Started With C**

The C Language

C was created around 1970 to make writing Unix and Unix tools easier.

- Part of the C/C++/Java family of languages (C++ and Java were created later)
- Design principles:
 - Small, simple abstractions of hardware
 - Minimalist aesthetic
 - Prioritizes efficiency and minimalism over safety and high-level abstractions

C vs. C++ and Java/Python

They all share:

- Syntax
- Basic data types
- Arithmetic, relational, and logical operators

C doesn't have:

- More advanced features like operator overloading, default arguments, pass by reference, classes and objects, ADTs, etc.
- Extensive libraries (no graphics, networking, etc.) – this means not much to learn C!
- many compiler and runtime checks, so extra care must be taken to write code that works and works without security concerns

Programming Language Philosophies

C is procedural: you write functions, rather than define new variable types with classes and call methods on objects. C is small, fast and efficient.

C++ is procedural, with objects: you write functions, and you also define new variable types with classes, and call methods on objects.

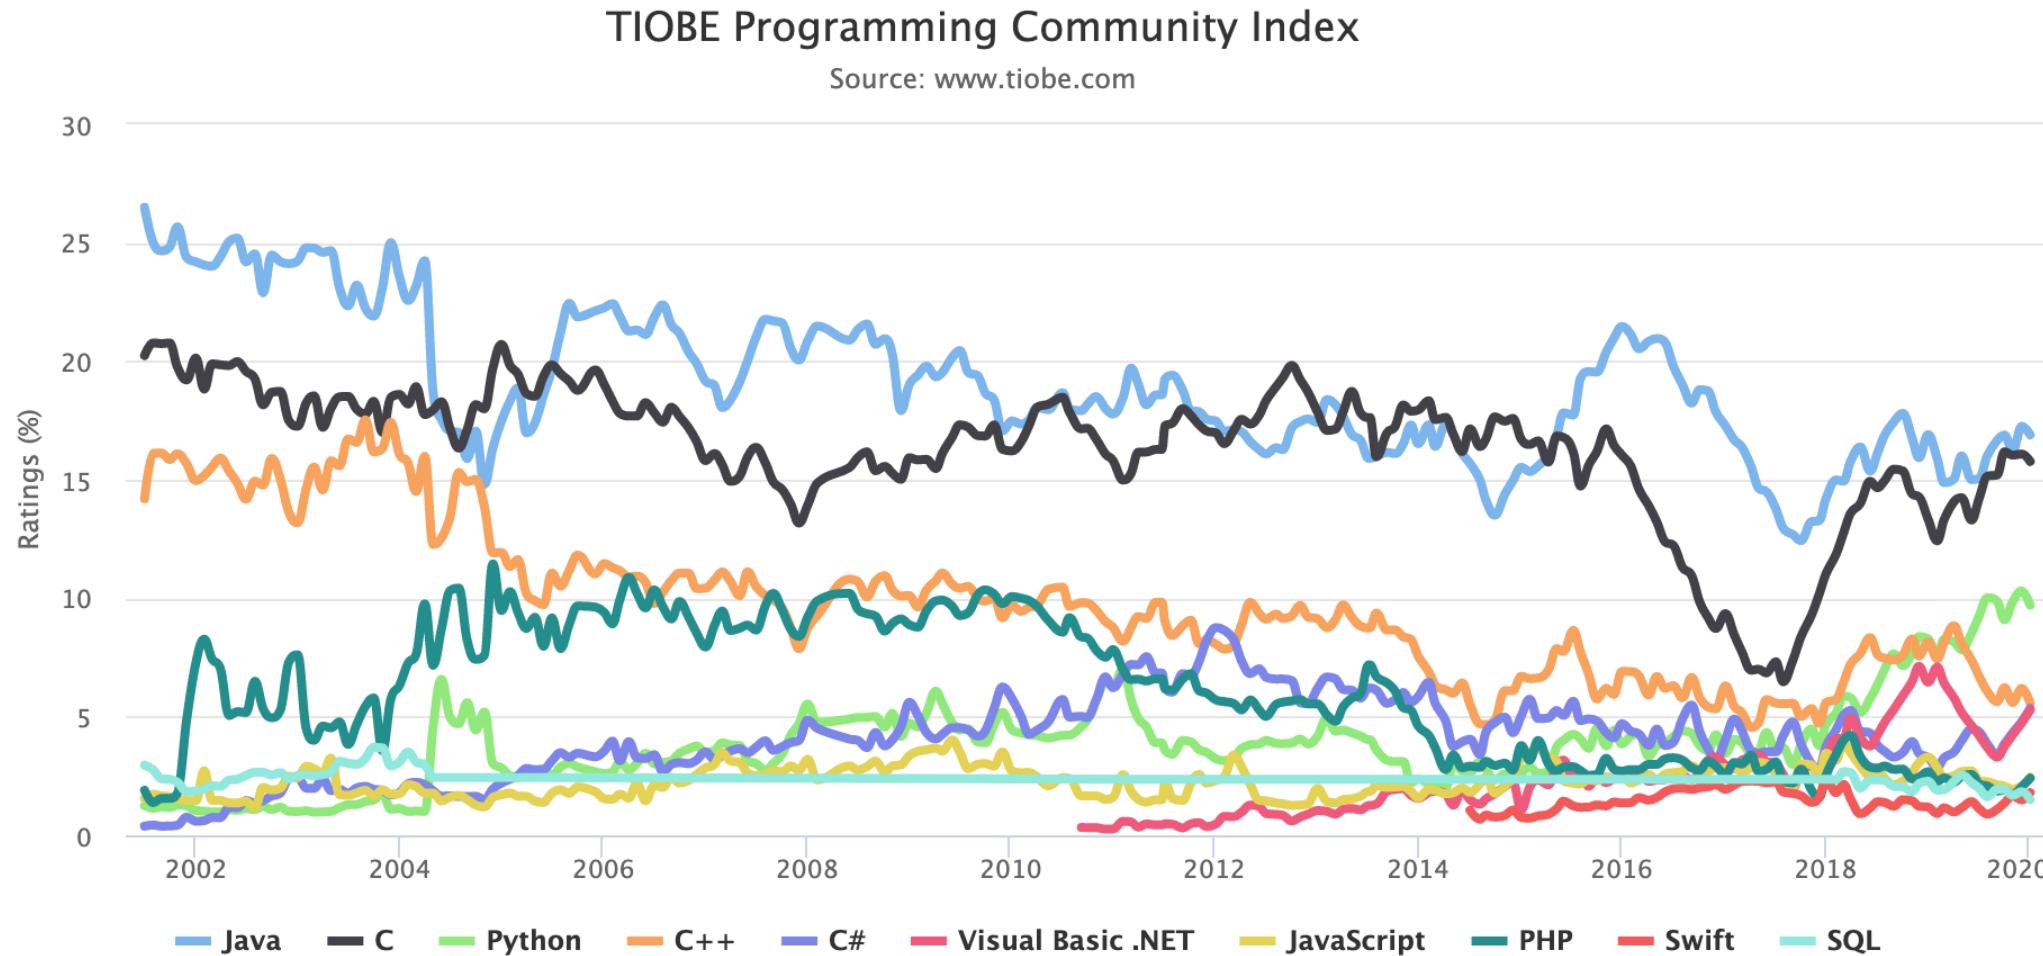
Python is also procedural, but dynamically typed: you still write functions and call methods on objects, but the development process is very different.

Java is object-oriented: virtually everything is an object, and everything you write needs to conform to the OO paradigm

Why C?

- Many tools (and even other languages, like Python!) are implemented in C.
- C is the language of choice for fast, highly efficient programs.
- C is popular for systems programming (operating systems, networking, etc.)
- C lets you work at a lower level to manipulate and understand the underlying system.

Programming Language Popularity



<https://www.tiobe.com/tiobe-index/>

Our First C Program

```
/*
 * hello.c
 * This program prints a welcome message
 * to the user.
 */
#include <stdio.h> // for printf

int main(int argc, char *argv[]) {
    printf("Hello, world!\n");
    return 0;
}
```

Our First C Program

```
/*
 * hello.c
 * This program prints a welcome message
 * to the user.
 */
#include <stdio.h> // for printf

int main(int argc, char *argv[]) {
    printf("Hello, world!\n");
    return 0;
}
```

Program comments

You can write block or inline comments.

Our First C Program

```
/*
 * hello.c
 * This program prints a welcome message
 * to the user.
 */
#include <stdio.h> // for printf

int main(int argc, char *argv[]) {
    printf("Hello, world!\n");
    return 0;
}
```

Import statements

C libraries are written with angle brackets.

Local libraries have quotes:

```
#include "lib.h"
```

Our First C Program

```
/*
 * hello.c
 * This program prints a welcome message
 * to the user.
 */
#include <stdio.h> // for printf

int main(int argc, char *argv[]) {
    printf("Hello, world!\n");
    return 0;
}
```

Main function – entry point for the program
Should always return an integer (0 = success)

Our First C Program

```
/*
 * hello.c
 * This program prints a welcome message
 * to the user.
 */
#include <stdio.h> // for printf

int main(int argc, char *argv[]) {
    printf("Hello, world!\n");
    return 0;
}
```

Main parameters – **main** takes two parameters, both relating to the *command line arguments* used to execute the program.

argc is the *number* of arguments in **argv**

argv is an *array of arguments (char * is C string)*

Our First C Program

```
/*
 * hello.c
 * This program prints a welcome message
 * to the user.
 */
#include <stdio.h> // for printf

int main(int argc, char *argv[]) {
    printf("Hello, world!\n");
    return 0;
}
```

printf – prints output to the screen

Familiar Syntax

```
int x = 42 + 7 * d;                      // variables, types
double pi = 3.14159;
char c = 'Q';                            /* two comment styles */

for (int i = 0; i < 10; i++) {           // for loops
    if (i % 2 == 0) {                   // if statements
        x += i;
    }
}

while (x > 0 && !(c == 'Q' || c == 'q')) { // while loops, logic
    x = x / 2;
    if (x == 42) { return 0; }
}

binky(x, 17, c);                      // function call
```

Boolean Variables

To declare Booleans, (e.g. `bool b = ...`), you must include `stdbool.h`:

```
#include <stdio.h>      // for printf
#include <stdbool.h>    // for bool

int main(int argc, char *argv[]) {
    bool x = argc >= 2 && strcmp(argv[1], "--informal") == 0;
    if (x) {
        printf("Hello, world!\n");
    } else {
        printf("Howdy, world!\n");
    }
    return 0;
}
```

Console Output: printf

```
printf(text, arg1, arg2, arg3);
```

```
// Example
char *classPrefix = "CS";
int classNumber = 107;
printf("You are in %s%d", classPrefix, classNumber); // You are in CS107
```

printf makes it easy to print out the values of variables or expressions.

If you include *placeholders* in your printed text, printf will replace each placeholder *in order* with the values of the parameters passed after the text.

%s (string)

%d (integer)

%f (double)

Writing, Debugging and Compiling

This quarter, we will use:

- the **emacs** or **vim** text editors to write our C programs
- the **make** tool to compile our C programs
- the **gdb** debugger to debug our programs
- the **valgrind** tools to debug memory errors and measure program efficiency

Demo: Compiling And Running A C Program



Working On C Programs Recap

- **ssh** – remotely log in to Myth computers
- **emacs** – text editor to edit C program source files
 - Use the mouse to position cursor, scroll, and highlight text
 - **ctrl-x ctrl-s** to save, **ctrl-x ctrl-c** to quit
 - You can use **vim**, too – see website for documentation on both editors
- **make** – compile program using provided **Makefile**
- **./myprogram** – run executable program (optionally with arguments)
- **make clean** – remove executables and other compiler files
- Lecture code is accessible at **/afs/ir/class/cs107/samples/lect[N]**
 - Make your own copy: **cp -r /afs/ir/class/cs107/samples/lectures/lect[N] lect[N]**
 - See the course Resources page for even more commands, and a complete reference.

assign0

Assignment 0 (Introduction to Unix and C) will be released late today on the course website and is due in one week on **Mon. 1/13 at 11:59PM PST**.

There are **5** parts to the assignment, which is meant to get you comfortable using the command line, and editing/compiling/running C programs:

- Visit the **Resources** page to become familiar with different Unix commands
- **Clone** the assign0 starter project
- **Answer** several questions in `readme.txt`
- **Compile** a provided C program and **modify** it
- **Submit** the assignment

Recap

- CS107 is a programming class in C that teaches you about what goes on under the hood of programming languages and software.
- We'll use Unix and command line tools to write, debug and run our programs.
- Please visit the course website, **cs107.stanford.edu**, where you can read the General Information Handout, information about the Honor Code in CS107, and more about CS107 course policies and logistics.

We're looking forward to an awesome quarter!

Preview: Next Time

- Make sure to reboot Boeing Dreamliners [every 248 days](#)
- Comair/Delta airline had to [cancel thousands of flights](#) days before Christmas
- Many operating systems [may have issues](#) storing timestamp values beginning on Jan 19, 2038
- [Reported vulnerability CVE-2019-3857](#) in libssh2 may allow a hacker to remotely execute code

Next time: *How can a computer represent integer numbers? What are the limitations?*