

Welcome!

COMS 4118

Operating Systems I

Spring 2022

<http://www.cs.columbia.edu/~jae/4118/?asof=20220119>

Teaching staff

- 7 Teaching Assistants (TAs)
 - Hans Jr Montero hjm2133@columbia.edu – Head TA
 - Kent John Hall kjh2166@columbia.edu – Head TA
 - Tal Zussman tz2294@columbia.edu
 - Xijiao Li xl2950@columbia.edu
 - Kaiwen Xue kx2154@columbia.edu
 - Claire Liu cl3944@columbia.edu
 - Eilam S Lehrman esl2160@columbia.edu
- TA email & office hours
 - Emails to cucs4118-tas@googlegroups.com go to all teaching staff
 - TA office hour calendar: <http://bit.ly/4118-cal>
- Instructor: Jae Woo Lee
 - Email: jae@cs.columbia.edu
 - Office: 715 CEPSR
 - Jae's office hour calendar: <http://bit.ly/jae-cal>
 - Home page: <http://www.cs.columbia.edu/~jae/>

Who am I?

- Jae Woo Lee
 - Senior Lecturer in Computer Science
 - Teaching first, research second
 - Just call me Jae (pronounced ‘Jay’)
 - Note that this is NOT a general rule – address instructors as Professors unless told otherwise
- My background
 - Undergrad in Columbia College
 - Many years of professional experience
 - Designing and coding large-scale software systems
 - Running a start-up company
 - Came back to Columbia for Ph.D.
 - More info at <http://www.cs.columbia.edu/~jae/>

Prerequisites

1. PLEASE PLEASE DO NOT TAKE THIS COURSE IF YOU DON'T KNOW C COLD
2. UNIX
 - Must be comfortable at command line
 - Don't take the course if you never worked on UNIX
3. Computer architecture
 - Basic knowledge of computer hardware: register, cache, bus, etc.
 - Should be able to read simple assembly code: load, store, add, jmp, etc.
4. Data structures
 - Nothing fancy, but must be solid on the basics: list, tree, stack & queue

Columbia courses:

For 1 & 2:

W3157
Advanced
Programming

For 3:

W3827
Fundamentals
of Computer
Systems

For 4:

W3134,
W3136, or
W3137
Data
Structures

Lectures

- In-person lectures
 - May switch to online as needed (like these first two weeks)
 - Auditors are welcome to lectures & listserv, but no GitHub repos, no lab/exam submissions, no TA access
- Zoom lectures are recorded
- Recording in-person lectures are being worked out – no promises

Exam dates

- **Synchronous** and **in-person** exams for all sections
 - Exam 1: **Wednesday Mar 9, 4:20pm**
 - Exam 2: **Wednesday Apr 27, 4:20pm**
- May switch to online format if necessary
- All students must take the exams at those times; there are **no make-up and no alternate exams**
 - Please take OS next semester if you can't make these times
 - If you receive extended time accommodation, you cannot have a class after this class
- Your overall grade
 - HW 40%, Midterm 30%, Final 30%
 - Grading policy may change later

Topics

- General OS theory
 - Throughout the whole semester
- Advanced UNIX programming
 - First 1/3 of the semester
 - UNIX from outside
 - Processes, threads, networking, concurrency, signals, non-blocking & async I/O
- Linux kernel implementation
 - Later 2/3 of the semester
 - UNIX from inside
 - Syscalls, wait queues, scheduler, file systems, virtual memory

28.8 million lines of code

- Linux 5.12: 28.8 million lines of code across more than 60k files in 2021
 - Learn to navigate a large code base
 - Learn to read code rather than documentations that are often vague, out-of-date, or flat-out wrong
- You are likely to encounter a large existing code base at work

Homework

- 7 assignments (not including hw0)
 - Some are individual, some are group assignments
 - Some are short & light, some are long & heavy
 - Assignments carry different weights
- Some assignments may not be graded
 - But you won't know until after the deadline
 - HWs picked for grading will be 40% of your grade
- Late policy
 - 20% penalty after deadline up to 24 hours; zero afterwards

Zero tolerance on cheating

- **REQUIRED READING:**
<http://www.cs.columbia.edu/~jae/honesty.html>
- You are cheating if you:
 - Take code from friends, or search for code on the Internet
 - Look at solutions that your friend has from previous semester
 - Upload any class materials (including your own code) to public repository (ex. GitHub) during or after this semester
- We can tell
 - We compare you submissions to **CURRENT AND PREVIOUS** submissions
 - You submit work history – **minimum 5 commits required**

Class ListServ

- Communication between all of us, including official announcements
 - Do:
 - Ask & answer questions – 1st place to go for non-personal questions
 - Provide helpful tips & links for your classmates
 - Be considerate & friendly
 - Don't:
 - Ask questions without first trying to solve it on your own
 - Post code or critical info that leads directly to solution
 - Be impatient & rude
 - Learn to manage high volume
 - [ANN] in email subject for announcements – set up Gmail filter
- Please use class listserv rather than TA mailing list
 - General questions to the TAs may be redirected to class listserv with your ID removed
 - Never send a same question individually to multiple TAs
- There will be an ongoing anonymous feedback form

Textbooks

1. Operating Systems: Three Easy Pieces

- **Version 1.00, 2018** – by Remzi H. Arpaci-Dusseau, Andrea C. Arpaci-Dusseau
- Free in PDF form: <http://pages.cs.wisc.edu/~remzi/OSTEP/>

2. Linux Kernel Development

- **3rd Edition, 2010**, Addison-Wesley – by Robert Love

3. Advanced Programming in the UNIX Environment

- **3rd Edition, 2013**, Addison-Wesley – by W. Richard Stevens, Stephen A. Rago

4. And a few other online materials that will be assigned

Get them wherever you usually get your textbooks from...

Equipment

- You need a computer that has:
 - 64-bit CPU with multiple cores
 - All computers manufactured within the last five years should have this
 - At least 8GB RAM
- You must run one of the following platforms:
 - Windows on x86 CPU (i.e. Intel or AMD)
 - MacOS on x86 CPU
 - MacOS on Apple M1 chip
 - Linux on x86 CPU
- You will receive VMware for your platform
 - VMware for Apple M1 is in beta

Let's get to work

1. **Subscribe to 4118 ListServ TODAY**

- <https://lists.cs.columbia.edu/mailman/listinfo/cs4118>
- In the textbox “Your name (optional)” put **Your Full Name (UNI)**
 - For example: Jae Woo Lee (jwl3)
- **You must reply to the confirm email (which might be in your spam folder)**
- Then receive “Welcome to the “Cs4118” mailing list”
 - This email contains your password for accessing archives of past postings

2. Read the following two documents:

- <http://www.cs.columbia.edu/education/honesty>
- <http://www.cs.columbia.edu/~jae/honesty.html>

3. **HW0 (50 points) – due Fri, 1/21, 11:59pm**

4. **HW1 (100 points) – due Sun, 1/30, 11:59pm**

5. Reading assignments

- See course home page for HW0, HW1, and reading assignments:
<http://www.cs.columbia.edu/~jae/4118/?asof=20220119>

6. Start forming groups of 3 – feel free to advertise on listserv