

CMSC 412: Operating Systems

Neil Spring

Fall 2017

Instructor Neil Spring

E-mail nspring@cs (include "412" in your subject line)

Class TuTh 11:00-12:15 CSIC 1121

Office hours M 10-12 AVW 4133

Grad TA Chris Maxey

Forum On piazza. <http://piazza.com/>

Web <http://www.cs.umd.edu/class/fall2017/cmssc412/>

Textbook Anderson and Dahlin *Operating Systems: Principles and Practice* ; If you have a copy of Silberschatz (typical 412 textbook), that should be adequate. The online <http://pages.cs.wisc.edu/~remzi/OSTEP> is also good. I do not assign exercises out of books; get a book if you learn best by reading.

1 Goals of this course

At the end of this class, you should be able to write a device driver, modify an operating system, understand how operating systems help or interfere with applications, write concurrent programs without deadlock, and have a strong foundation for programming embedded devices that lack an operating system.

2 Summary

The course will cover the following core topics:

Processes What makes a process, how are they run concurrently, how to create them and communicate between them.

Threads What makes a thread, libraries.

Scheduling How to keep interactive applications responsive and background applications make forward progress.

Synchronization and Deadlock Locks on shared data, and preventing cooperative processes from getting stuck.

Memory and Virtual Memory Swapping, paging, segmentation, allocating memory, copy-on-write, etc.

File System Interface and Implementation the function calls, mounting file systems, organizing blocks on disk, allocation, recovery.

Disk and Storage Systems disk scheduling, RAID, tape hierarchies.

I/O Systems programmed and interrupt-driven I/O.

And given time, the following additional topics:

Protection capabilities, defining access control.

Distributed Coordination Events, atomicity, deadlock in distributed systems where messages can be lost.

Linux how each of the features we learned about are implemented in Linux.

Security Basic crypto, authentication.

Distributed Systems Distributed communication primitives.

Distributed File Systems Global naming of files.

3 Textbook

Anderson and Dahlin, *Operating Systems: Principles and Practice*, remains my chosen textbook, though it has limitations. It costs less than half as much as the typical 412 textbook (Silbershatz, Galvin, and Gagne). It's pretty well written, very well typeset (not much distracting font silliness), and written by people I respect. I've used Dahlin's notes for some topics before.

I know there are a lot of copies of Silbershatz floating around. Having *some* textbook is *required*; which one is not so important to me. I will assign readings out of Anderson (see the schedule), and try to bridge the vocabulary to the extent possible. (Anderson uses "address translation" where Silbershatz uses "virtual memory"; I find Silbershatz's definition of virtual memory counterintuitive.)

I like the textbook by Arpaci-Dusseau, but don't follow its order either. <http://pages.cs.wisc.edu/~remzi/OSTEP>.

4 Prerequisites

CMSC330 – Programming Languages, as well as one of CMSC414, CMSC417, CMSC420, CMSC430, CMSC433.

Experience in CMSC417 (networks) may help you.

You must know what a function pointer is and how it is used. Find a book on *C today* if you do not.

You should understand basic issues of concurrency. That includes the interactions between non-blocking sockets, user-level and kernel-level threads, locking, etc. Too many students seem to think that forking a thread will solve a simple problem without creating many more.

5 Style

I don't use lecture slides; I generally type into text notes in an emacs buffer or use the whiteboard. I expect to be interrupted. I will assume you know more than you do; it is your job to pay attention, and make me clarify when I've left you behind.

Some students like this scheme a lot. Others can't keep up. Students who sit in the back may have the most trouble following a discussion started by student questions.

6 Grading

You may see your scores for individual assignments on <https://grades.cs.umd.edu/>. There, you will also find your linuxlab account, which we will distribute just before the first programming assignment is presented in section.

Many students incorrectly interpret their progress on grades.cs relative to other students ("I'm above average, so I must pass") or relative to an absolute ("90% is an A") scale. Understand that "average" scores are often held back by students who may have abandoned the class and that I do not update the grades or their weighting in real time. In

other words, the information available to you on grades will not be sufficient for you to predict your grade. Ask if you have concerns.

I view my job as to give you the most challenging but fair questions and assignments possible; whether getting 70% of the points represents adequate understanding is not something I worry about.

6.1 Forum / Piazza / Class participation: 2%

In a class so large, I can't expect each of you to speak; participation here is a negative grade, if I think you're doing poorly and it's your own fault for not being engaged with the material, you won't get the participation bump.

Participation is required. That means writing. If I don't see your name on the board asking good questions or answering questions well, and don't know you from in class questions, no points here.

6.2 Quizzes: 12%

There will be 7 quizzes in discussion section on Mondays, each worth 2%. Your lowest score will be dropped: you may miss this quiz physically or mentally.

6.3 Homework and in-class exercises: 5%

A few homework or in-class exercises will prepare you for the quizzes; each is worth a tiny fraction of your grade.

6.4 Two Midterm Exams: 26%

6.5 Final Exam: 20%

The midterm and final exams will mix multiple choice, simple matching, short answer and long answer questions. The midterms will consume a lecture slot, the final during finals week as scheduled by the university. The exams will be have too many questions to allow all of you to finish the entire exam. You will have to learn and study before the exam.

6.6 Programming Assignments: 35%

The programming assignments in this class will use GeekOS. The assignments are difficult. The assignments will require opening and editing many files, likely using a reasonable programmer's editor to facilitate editing, building, and testing quickly. Time spent early in the semester developing your skills and setting up your environment will pay off during crunch time.

7 Lateness

All programming assignments can be turned in electronically. I will permit one programming assignment to be turned in after the weekend (when due Friday, it can be turned in on Monday). I expect any data loss due to dogs, roommates, lightning strikes or FBI confiscating your machine can be dealt with over a weekend.

Caution: Don't "plan" to use your late deadline; programming assignments have been pretty well tested so you're unlikely to benefit from procrastination. You are, however, likely to underestimate the difficulty you'll have with the project.

8 Administrative Craft

I dislike this section greatly, but codifying each of these policies is important for keeping myself sane and making clear what my expectations are. I'd much prefer a section that said "treat me with respect and I'll do the same for you;" this

section is intended mostly for those who would hope to game the system. Note that I copied verbatim some of these passages; I hope you appreciate irony.

8.1 Excused absences

Students claiming a excused absence must apply in writing and furnish documentary support (such as from a health care professional who treated the student) for any assertion that the absence qualifies as an excused absence. The support should explicitly indicate the dates or times the student was incapacitated due to illness. Self-documentation of illness is not itself sufficient support to excuse the absence. An instructor is not under obligation to offer a substitute assignment or to give a student a make-up assessment unless the failure to perform was due to an excused absence. An excused absence for an individual typically does not translate into an extension for team deliverables on a project.

8.2 Religious observances

I will avoid deadlines October 2-4 and October 11/12. Please inform me in advance of religious observances that will interfere with your ability to complete assignments on time.

8.3 Honor code

The University of Maryland, College Park has a nationally recognized Code of Academic Integrity, administered by the Student Honor Council. This Code sets standards for academic integrity at Maryland for all undergraduate and graduate students. As a student you are responsible for upholding these standards for this course. It is very important for you to be aware of the consequences of cheating, fabrication, facilitation, and plagiarism. For more information on the Code of Academic Integrity or the Student Honor Council, please visit <http://www.studenthonorcouncil.umd.edu/whatis.html>.

8.4 What constitutes cheating?

Copying other assignments, looking over someone's shoulder in the lab, emailing function code, using google to find a code fragment without understanding, looking for code in other people's directories, pulling code printouts off printers, and in any other way attempting to gain a grade without learning.

Consider each programming assignment to be a take-home exam.

Note: leaving someone your code because you want to help is just as bad as borrowing someone else's code. We can tell when code looks and acts too similar to be independent work; we can't (easily) tell which of two implementations was the original.

Restated, it is not even helpful to give away your code, and clearly not permitted.

This policy applies to all course assignments. Explicitly, **it is not permitted to collaborate on homework assignments**. If your answer is not your own, it must be cited (wikipedia, google). If you have questions, post to the forum. If you learned something through a discussion with another student, cite.

Finally, if I find your solutions to any of my programming assignments online at any time, I will refer your case to the office of student conduct for facilitation. Posting project solutions online causes projects to be much harder and less well polished than they should be, permits otherwise unqualified students to pass courses without learning, and has no benefit to the poster's employment prospects, since it exposes the poster to be an inconsiderate moron.

8.5 What constitutes legal collaboration?

Interaction via course discussion forum or discussion of problem and code solutions governed by the Gilligan's Island rule is permitted.¹

Using google where the result is not code is OK. Using wikipedia is encouraged, even during class. If you find a particularly good solution on either, please cite it; there is no penalty for citing sources and I'm more likely to consider

¹You understand the concept only if you can watch one half-hour complete episode of Gilligan's Island and still retain the concept. You may then begin coding with your newfound knowledge safe that it is your own work. Without the thirty minute pause, it is not your work.

answers that disagree with textbook or lecture legitimate. If you find a question far too easy because an answer is present on-line, please let me know.

This policy applies to all course assignments. Explicitly, it is not permitted to collaborate on homework assignments. If your answer is not your own, it must be cited (wikipedia, google). If you have questions, post to the forum. If you learned something through a discussion with another student, cite.