

# Operating Systems

## Syllabus

### Special Note

Emergency remote instruction continues in Fall 2020. This is not what we would all prefer, but we can work together and succeed. Please bear with us as we adapt to remote teaching. We will support you as you deal with remote learning.

This course was completely redesigned in Winter 2020. There is minimal overlap with previous instances of the course in terms of reading, labs, codebases, and exams. **Because it is a new design, we are very interested in student feedback and will be, to some extent, adaptive to it.**

### Class Resources

- [Course Website https://northwesternos.org](https://northwesternos.org): All course details and schedule.
- [Canvas](#): Grade reports, class recordings, and zoom links. Already enrolled.
- [Piazza](#): Class discussions, questions, and messages. We will enroll you.
- [Github Classroom](#): Lab access and submission. We will enroll you.

### Instructor

Branden Ghena  
branden@northwestern.edu

### Teaching Assistance

Conor Hetland (TA)  
ConorHetland2015@u.northwestern.edu

Calypso McDonnell (PM)  
calypsomcdonnell2021@u.northwestern.edu

Michael Cuevas (PM)  
cuevas@u.northwestern.edu

### Location and Time

Lectures:	Videos available on Canvas.
Q&A Sessions:	Tuesdays/Thursdays 11:20-12:40 on Zoom (also recorded).
Lab Discussions:	TBD on Zoom (also recorded).
Office Hours:	See web page.
Midterm Exam:	TBD, mid-quarter. Details to come.
Final Exam:	Thursday, December 03, 12:00-2:00. Details to come.

## Prerequisites

Required	CS 213 or CE 205 or equivalent
Required	CS 214 or equivalent
Required	Experience with C or C++
Required	Some experience with programming in a Unix environment (e.g., as in CS 211 and CS 213)

Any version of CS 213 or CE 205 is acceptable, but we will expect that you have seen basic concepts such as the existence of exceptional control flow and virtual memory, and the typical Unix system calls for processes, threads, and files+I/O. The syllabus shown in [pdinda.org/ics](http://pdinda.org/ics) is a good starting point.

Any version of CS 214 is acceptable, but we will expect that you have seen basic data structures, algorithms, and their implementation. These include linked lists, balanced search trees, hashing and hash tables, heaps, graphs, sorting, etc.

Experience with C or C++ in part means familiarity with arrays, structs, unions, and, most importantly, pointers and pointer-based data structures. Low-level pointer-based mechanisms are used throughout an OS, and by the underlying hardware.

Experience with programming on Unix means being able to navigate the Unix command line, remote access, use/extend Makefiles, etc.

CS 343 satisfies one of the **Systems Breadth**, Tech Elective, and Project requirements in in the Computer Science curriculum in both McCormick and Weinberg. CS 343 can also be taken for credit within the Computer Engineering curriculum.

## Textbook

Andrew S. Tanenbaum and Herbert Bos, *Modern Operating Systems, 4th Edition*, Pearson, 2014, (ISBN-13: 978-0133591620, ISBN-10: 013359162X)

We also considered several other books for this course, which may be useful as further references:

- Abraham Silberschatz, Peter Galvin, and Greg Gange, *Operating Systems Concepts*, 10th edition, Wiley, 2018.
- Remzi Arpaci-Dusseau, and Andrea Arpaci-Dusseau, *Operating Systems: Three Easy Pieces*, 2018. [freely available]
- Thomas Anderson, and Michael Dahlin, *Operating Systems: Principles and Practice*, Recursive, 2014.
- William Stallings, *Operating Systems: Internals and Design Principles*, 9th edition, Pearson, 2017.

The choice of Tanenbaum as the textbook for this course is a compromise. All of these books have strengths and weaknesses.

It is important to note that your CS 213 textbook (Randal Bryant, and David O'Hallaron, *Computer Systems: A Programmer's Perspective*) has an excellent “what every programmer should know” treatment of some of the topics we will cover, including threads, processes, virtual memory, and the various Linux/Unix system call interfaces.

There are also books on specific operating systems that advanced students might be interested in, particularly on FreeBSD and Linux. Ask if you're curious.

### **Objectives, framework, philosophy, and caveats**

This course introduces you to the basic, foundational concepts and principles of operating systems, many of which generalize to other areas of computer science and engineering. You will learn many of these concepts and principles by applying them in practice on a modern machine through labs that are designed to put you in the shoes of a systems-level developer. OS (and systems more broadly) is very much a learn-by-doing kind of area.

The following concepts and principles are included:

- **OS Structure:** kernel, device drivers, file systems, network stacks, schedulers, system calls, libraries, toolchains, language virtual machines, user interface/shell, applications, etc.
- **OS Models:** monolithic kernel, microkernel, virtual machine monitor/hypervisor, jail/zone/container, exokernel, unikernel, ...
- **OS Abstractions:** thread, name space, address space, process, IPC, virtual machine, container, file, directory stream, plus abstraction design within the kernel (devices, file systems, ...)
- **Concurrency Sources:** multiprocessors, devices, interrupts, threads, processes, horror stories, ...
- **Concurrency Challenges:** memory system coherence/consistency, race conditions, deadlock, livelock, horror stories, ...
- **Concurrency Control:** interrupt control, atomics, spinlocks, critical sections, blocking vs waiting, mutexes, semaphores, condvars, monitors, barriers, lockfree/waitfree models, plus typical synchronization problems such as producer-consumer, reader-writer, and dining philosophers.
- **Scheduling and Resource Management:** theory, FCFS, GPS, SRPT, dynamic priority (e.g. Unix), lottery, fixed priority, preemptive vs non-preemptive, real-time vs non-real-time, horror stories, ...
- **Virtual Memory:** hardware-software co-design, paging, swapping, segmentation and (possibly) current alternatives.
- **Device Drivers:** interrupts, DMA vs PIO, MMIO vs PMIO, atomics, hardware memory barriers, software memory barriers.

- **Protection and Security:** kernel/user mode, mode/ring transitions, role of encryption, interaction with virtual memory, horror stories.
- **Memory management:** page allocation versus heap allocation, garbage collection, allocation in special contexts (e.g. interrupt context), page replacement, working set.
- **File systems:** issues/interfaces, data structures on block devices, examples (V7, FAT+, ext2+)
- **Principles:** policy versus mechanism, orthogonality, worse-is-better, lazy evaluation, caching, end-to-end argument, mythical man-month, no silver bullet, hardware/software co-design

The hardware environment that we will focus on is Intel/AMD machines running in 64 bit mode ("x64"), which is the commonplace platform for systems ranging from laptops to supercomputers today.<sup>1</sup> Your lab work will be done on Linux in the C programming language.<sup>2</sup> Two of your labs (on concurrency and scheduling) will be done in user-level Linux. The remaining labs will be in the context of the Nautilus kernel framework ("NK"), a research kernel develop at Northwestern and other institutions. The experience you gain in NK will generalize to the Linux kernel, for the most part.<sup>3</sup>

## Lectures / Q&A Sessions / Attendance

Lectures will be recorded and made available on Canvas in advance of the scheduled date for that material. Question and Answer sessions will take place during the normal course time synchronously over zoom. It is important that you review the readings and lecture recording before each scheduled Q&A session so that you can ask questions during the session. **You should prepare at least one question for each Q&A session.**

Q&A sessions will be recorded, but if no one shows up there are not going to be very many answers. The synchronous sessions are our chance to build a class community and gain a deeper understanding together. Please attend them.

---

<sup>1</sup> Most of what you learn about x64 via a vis OS will generalize to the other main platform, ARM, which is the basis for phones and tablets.

<sup>2</sup> Linux is the common OS on everything except laptops and desktops. It is also the OS underlying Android. C is the lingua franca of low-level software development.

<sup>3</sup> In the design of this course, we considered several other options. The most desirable would have been to have you work within the Linux kernel itself. This proved to be intractable from a pedagogical point of view. The complexity we would have to shield you from, particularly in a lab based on paging, would have been overwhelming to manage. We also considered the teaching OS xv6 for IA32 and for RISC-V. IA32 and RISC-V both would require revisiting material students have already learned, for x64, in CS 213, plus xv6 for IA32 would have made a device driver lab particularly challenging to pull off. Another consideration was to use CMU's Pebbles OS specification and have students build Pebbles from scratch as in CMU's course. This was also limited to IA32, and seemed intractable to execute in a single quarter. The intent behind using NK is to give a view inside a modern, x64 codebase with clear internal interfaces that has a development model (e.g., Kbuild, C, etc) that is similar to Linux.

What I am asking of you is: Read. Watch. Attend. Ask. There is no such thing as a dumb question (or too esoteric of a question) - we will try our best to answer or comment on all questions.

## Other Ways of Getting Help

Your TAs and peer mentors will run an optional weekly discussion focused on labs, which we will schedule, with your input, during the first week. The goal of the optional weekly discussion is to provide a place to learn more and to get help in a more structured way than office hours.

Your instructor, TAs, and peer mentors will also have regularly scheduled office hours and be available by appointment if these do not work. We will schedule office hours in the first week to maximize opportunities to attend.

We will use an online discussion group on Piazza as well. We will enroll you. The link is on the course web page. The intent is to have multiple venues for discussion with different styles so that all students feel comfortable participating. If you have a question, answer, or comment, please put it forward. We will try our best to answer.

Labs will be done using GitHub Classroom. One goal here is to make it straightforward for us to see the current state of your lab work, so that we do not have to spend a lot of time reconstructing setups during office hours, etc. Push early and often!

## Computing Resources

You will have Linux accounts on the Wilkinson machines, and it should be possible to do some of your work on them, or other 64-bit Linux machines. You will also have access to a newly purchased high-end server which has a range of software set up for use by this course. This is the easiest option and is also where we will grade labs. The very first lab is intended to get you familiar with this environment by having you build and run a kernel on it.

It is also possible to work on your own machine. Generally speaking, using Linux will be easiest. I often do development with Ubuntu installed in a VirtualBox VM on my Mac or PC. We will provide instructions in Piazza for those who would like to set up their own environment.

## Grading

- 60% Programming labs (breakdown as below)
- 20 % Midterm (covers first half of the course)
- 20 % Final (covers second half of the course)

There is extra credit in many of the programming labs.

Your score in the course is the weighted average of your scores on each of the components. You can view all currently graded material, and your score, at any time on Canvas. Final grades are based on the course score (the weighted average), with the basic model being that the 90s are A territory, 80s are B territory, and so on. This model will be adapted toward lower thresholds if necessary based on overall class performance. That is, this is NOT a curved class.

The instructor ultimately assigns scores and grades in consultation with the TAs and peer mentors. If you have a problem with a score on an assignment/exam or your grade, you are welcome to bring it up with him or the TAs, but only the instructor is empowered to change grades.

## Labs

We will have five programming labs. Except for the first lab, labs should be done in groups of up to three. **Start looking for a partner on day one.** You cannot change groups after they are formed.

In the current design, there are five labs. 60% of the grade in the class will be based on lab work, with a breakdown as follows:

- 5% Getting Started Lab (done individually, no slip-days!)
- 10% Producer-Consumer Lab
- 10% Queueing/Scheduling Lab
- 20% Device Driver Lab
- 15% Paging Lab

We will use GitHub Classroom for disseminating and handing in labs. It is important that you and your partners make sure that your repositories are private. Only your group and the course staff should be able to see your repos.

The Producer-Consumer Lab and Queueing/Scheduling Lab are user-level Linux labs. The others are all done within a research kernel developed at Northwestern. All hardware is x64. All code is in C.

## Lab Late Policy

For each calendar day or portion thereof after the due date for a lab, 20% is lost. After 1 day, the maximum score is 80%, after 2 days, 60%, etc, for a maximum of 5 days.

## Lab Slip Days

To help you handle any issues that arise from remote instruction, we give you **two slip-days**, which allow you to submit a lab late without penalty. Slip-days may only be applied towards labs, and not any other assignments. Slip days are in used in units of entire “days” meaning a lab submitted one minute after the deadline consumes one entire slip day. Please plan accordingly as there is no grace period.

We will track the total number of late days for your submissions and automatically apply slip days to optimize their usage. Slip-days will not be assessed against labs you did not submit. No extra credit is awarded for avoiding the use of slip-days. However, it is in your best interest to avoid turning labs in late. Usually, a new lab will be released very shortly after the current lab is due.

Example slip-day usage:

- Use two slip-days to receive no penalty on a lab submitted two days late
- Use two slip-days to receive no penalty for two separate labs each submitted one day late
- Use two slip-days to receive just a one-day late penalty on a lab submitted three days late

## Exams

There will be a midterm exam and a final exam. The final exam will not be cumulative. We will schedule midterm and final exam review sessions. Details on how exams will be handled remotely will come soon.

## Cheating and Inadvertent Disclosures

Since cheaters are mostly hurting themselves, we do not have the time or energy to hunt them down. We much prefer that you act collegially and help each other to learn the material and to solve problems than to have you live in fear of our wrath and not talk to each other. Nonetheless, if we detect blatant cheating, we will deal with the cheaters as per Northwestern guidelines.

As we note above, it is important that you control access to your GitHub repos.

Please do not place class materials from on any public site. If it's on the course web site, it's already public and will remain public. If it's from the discussion group or from the handout directory on the course servers, it should not be shared publicly.

## Accessibility

Any student requesting accommodations related to a disability or other condition is required to register with ANU ([accessiblenu@northwestern.edu](mailto:accessiblenu@northwestern.edu); 847-467-5530) and provide professors with an accommodation notification from AccessibleNU, preferably within the first two weeks of class. All information will remain confidential.

Should you need them, additional campus resources are available, including, but not limited to:

- Accessible NU [www.northwestern.edu/accessiblenu/](http://www.northwestern.edu/accessiblenu/)
- CAPS [www.northwestern.edu/counseling/index.html](http://www.northwestern.edu/counseling/index.html)
- Student Enrichment Services [www.northwestern.edu/enrichment/](http://www.northwestern.edu/enrichment/)

I believe in providing reasonable accommodations that allow for full access to learning for all. Please contact me if there is anything that I should be aware of that might have an impact on your participation in this course (documented disability, language challenges, absences for religious observations, etc.)

## Diversity and Inclusion

I consider this classroom to be a place where you will be treated with respect, and I welcome individuals of all ages, backgrounds, beliefs, ethnicities, genders, gender identities, gender expressions, national origins, religious affiliations, sexual orientations, ability—and other visible and nonvisible differences. All members of this class are expected to contribute to a respectful, welcoming, and inclusive environment for every other member of the class.



**Schedule**

Lecture	Date	Topics	Readings	Labs
1	09/17 Th	Introduction, OS Structure, OS Models, HW/SW interface, History	Chapter 1, 8.1.2, 10.1, 10.2	Start lab out
2	09/22 Tu	Concurrency Sources: hw, interrupts, threads, processes, ...	2.1, 2.2, 5.1.5, 8.1.1	
<i>9/22 is the last day for adding courses or changing sections.</i>				
3	09/24 Th	Concurrency Sources: continued	2.1, 2.2, 5.1.5, 8.1.1	Start lab in, PC lab out
4	09/29 Tu	Concurrency Challenges and Control: races, mutual exclusion, critical sections	2.3, Concurrency,	
5	10/01 Th	Concurrency Challenges and Control: blocking, mutexes, spinlocks, semaphores, condvars, barriers, monitors, etc.	2.3 (cont.), 8.1.3, Therac	
6	10/06 Tu	Concurrency Challenges and Control: deadlocks, detection, avoidance, prevention, starvation, lockfree/waitfree data structures	6, 2.5	
7	10/08 Th	Scheduling: classic treatment	2.4, 10.3, 8.1.4	PC lab in, Queue lab out
8	10/13 Tu	Scheduling: workload, queueing, and real-time perspectives	Workload, Queueing, Mars	
9	10/15 Th	<i>Special topic or Slack</i>		
10	10/20 Tu	Devices and drivers: principles	5.1-5.3, 10.5	
<i>Midterm Exam Review: TBD</i>				
<i>Midterm Exam: Around here, time+location TBD</i>				
11	10/22 Th	Devices and drivers: examples	5.4-5.8	Queue Lab in, Driver lab out
12	10/27 Tu	OS design principles	Chapter 12	
13	10/29 Th	Virtual memory with paging and segmentation	3.1, 3.2, 3.3, 3.7	
14	11/03 Tu	Paging and swapping policies and their effects, working set, allocation	3.4-3.6	
15	11/05 Th	Paging on x64 and Linux	10.4	
11/06 is the last day to drop a class.				

16	11/10 Tu	Security and Protection	9.1-9.6, Spectre	Driver lab in, Paging lab out
17	11/12 Th	File systems: principles and issues	4.1-4.4	
18	11/17 Tu	File systems: examples	4.5, 10.6	
19	11/19 Th	Virtualization, containerization, the cloud, etc	Chapter 7	
20	11/24 Tu	<i>Special topic or Slack</i>		Paging lab in
<i>Finals week – Exam is Thursday, December 03, 12-2pm</i>				

Readings are from the textbook, with these exceptions:

Therac	THERAC-25 article
Mars	Mars Pathfinder article
Spectre	Meltdown/Spectre article
Unix	Unix Systems Programming in a Nutshell Handout
Workload	Workload Characterization Handout
Queueing	Queueing Theory Handout
Concurrency	Concurrency Handout