

CSE 451: Operating Systems

Winter 2025

Module 1

Course Introduction

Gary Kimura

Course Staff

Instructors:

Gary Kimura

Ravi Pudipeddi

Teaching Assistants:

Dara Stotland

Tanay Vakharia

Anirudh Kumar

Dawit Faraja Hailu

Aliyan Muhammad

Will Robertson

Course Details

- But first, a word about **Health** and **Safety**
- Information portals
 - **Course webpage, Canvas, Ed, Gradescope, gitlab**
- Class Prerequisites:
 - CSE332: Data Structures and Parallelism
 - CSE333: Systems Programming
 - CSE351: The Hardware/Software Interface
- Lectures and Sections (will be recorded on zoom or panopto, and lecture slides/notes posted)
- Projects – Led by the TAs
 - 4 labs using xk (brush up on programming in C). Due about every two weeks. Check the class website for the actual schedule.
- Grading (subject to change)
 - Projects 70%, Design Docs 10%, Exam 20%
- Occasional short detours or videos on Hard Lessons Learned.

The Projects

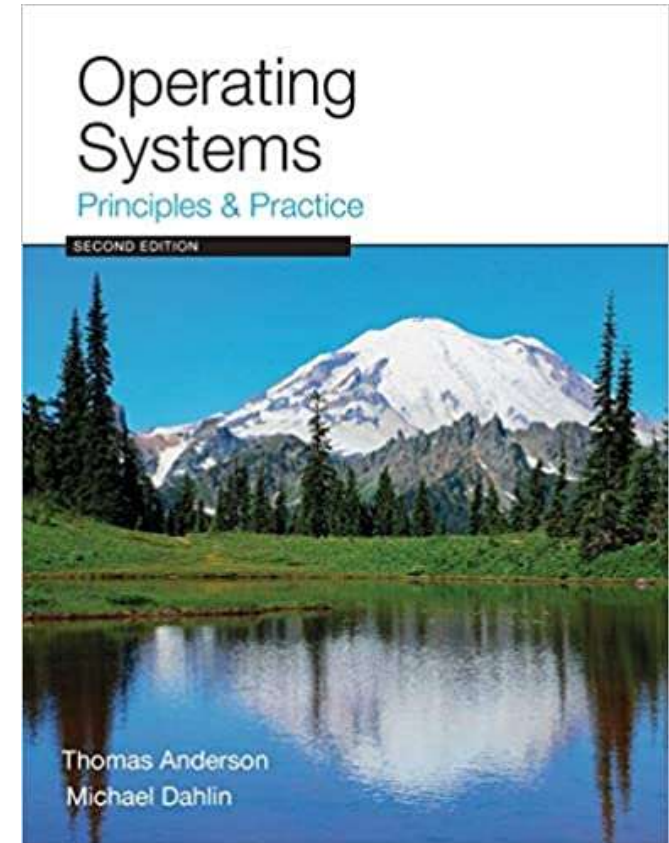
- Start them early
- Four of them
- **Teams of two.** You're likely to be happier if you form a team on your own than if we randomly assign them.
 - Working solo is strongly discouraged, and the TAs are not a replacement for a lab partner
- **Debugging is a skill best learned through experience**
- Do not believe that passing the test cases means your code works...
 - Sorry

Course Objectives

1. (One thing that I hope you learn from this class) How an OS is designed and built. To better know how to use the OS
2. Debugging large programs, adding new features to an existing (incomplete codebase)
3. Quote in the Cutler lab
“Bugs: if you don’t put them in, you don’t have to take them out.”
4. A lot of material to cover in the first few weeks, before it all makes sense
5. **Textbook**: academic head knowledge.
OSPP will be our main textbook
Additional material is available online for free, see class website
6. **Lectures**: enhance and supplement the textbook and do deep dives into some specific implementation issues.
7. **Projects**: get hand dirty, learn by doing

Course Roadmap (subject to change)

1. Chapter 2 Kernel Abstraction
2. Chapter 3 Programming Interface
3. Chapter 4 Concurrency
4. Chapter 5 and 6 Synchronization
5. Chapter 8, 9, 10 Memory Management
6. Chapter 7 Scheduling
7. Chapter 11, 12, 13, 14 Storage
8. Wrap up loose ends



An important underlying concept

Policy vs. Mechanism

- **Policy:** is what you are trying to achieve
 - All the programs running on a computer get equal access to the CPU
- **Mechanism:** is how you achieve it
 - Use timers and context switching to share the CPU
- The lectures and textbook is mostly about policy, the projects deal with mechanism

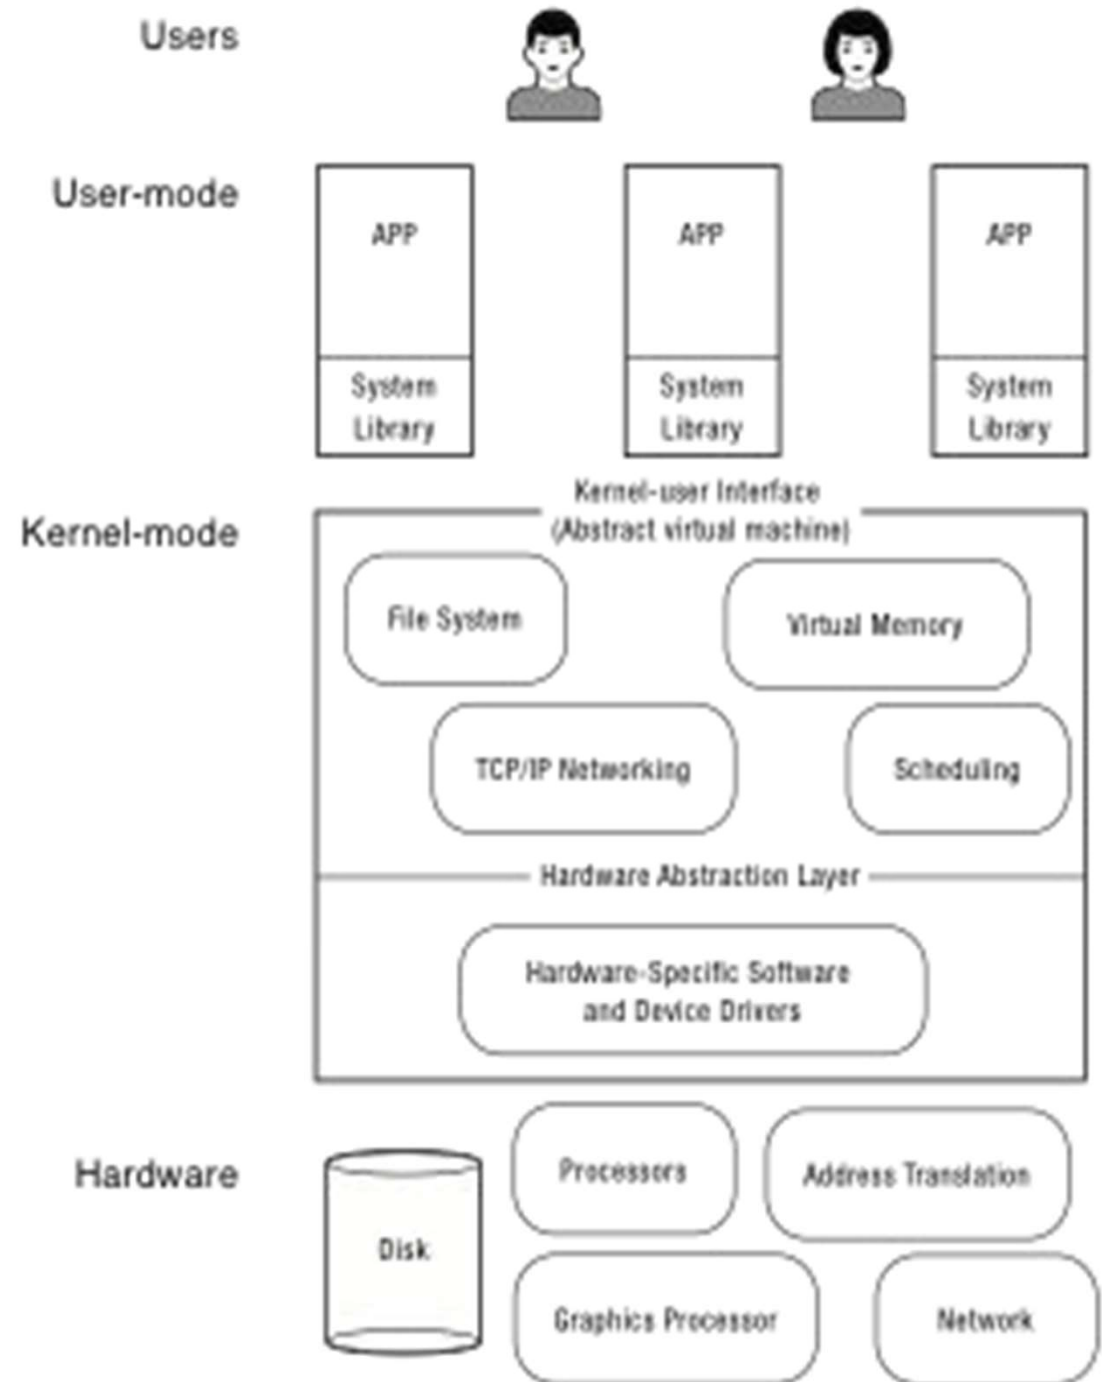
Operating System Mission Statement

- Class exercise, write an OS mission statement (e.g., The mission statement for the UW CSE School might be to prepare students in the art of computer usage for the coming century.)

One possible mission of the OS: provide a means for programs to effortlessly utilize the capability of the computer system

What is an operating system?

- Software to manage a computer's resources for its users and applications



Operating System Roles

- **Referee:**
 - Ensures that everyone plays by the rules
 - Resource allocation among users, applications
 - Isolation of different users, applications from each other
 - Communication between users, applications
- **Illusionist:**
 - Each application appears to have the entire machine to itself
 - Infinite number of processors, (near) infinite amount of memory, reliable storage, reliable network transport
- **Glue:**
 - Libraries, user interface widgets, ...

Also viewed as Virtualization, Concurrency, and Persistence

Example: File Systems

- **Referee:**
 - Prevent users from accessing each other's files without permission
 - Even after a file is deleting and its space re-used
- **Illusionist:**
 - Files can grow (nearly) arbitrarily large
 - Files persist even when the machine crashes in the middle of a save
- **Glue:**
 - Named directories, printf, ...

OS Challenges

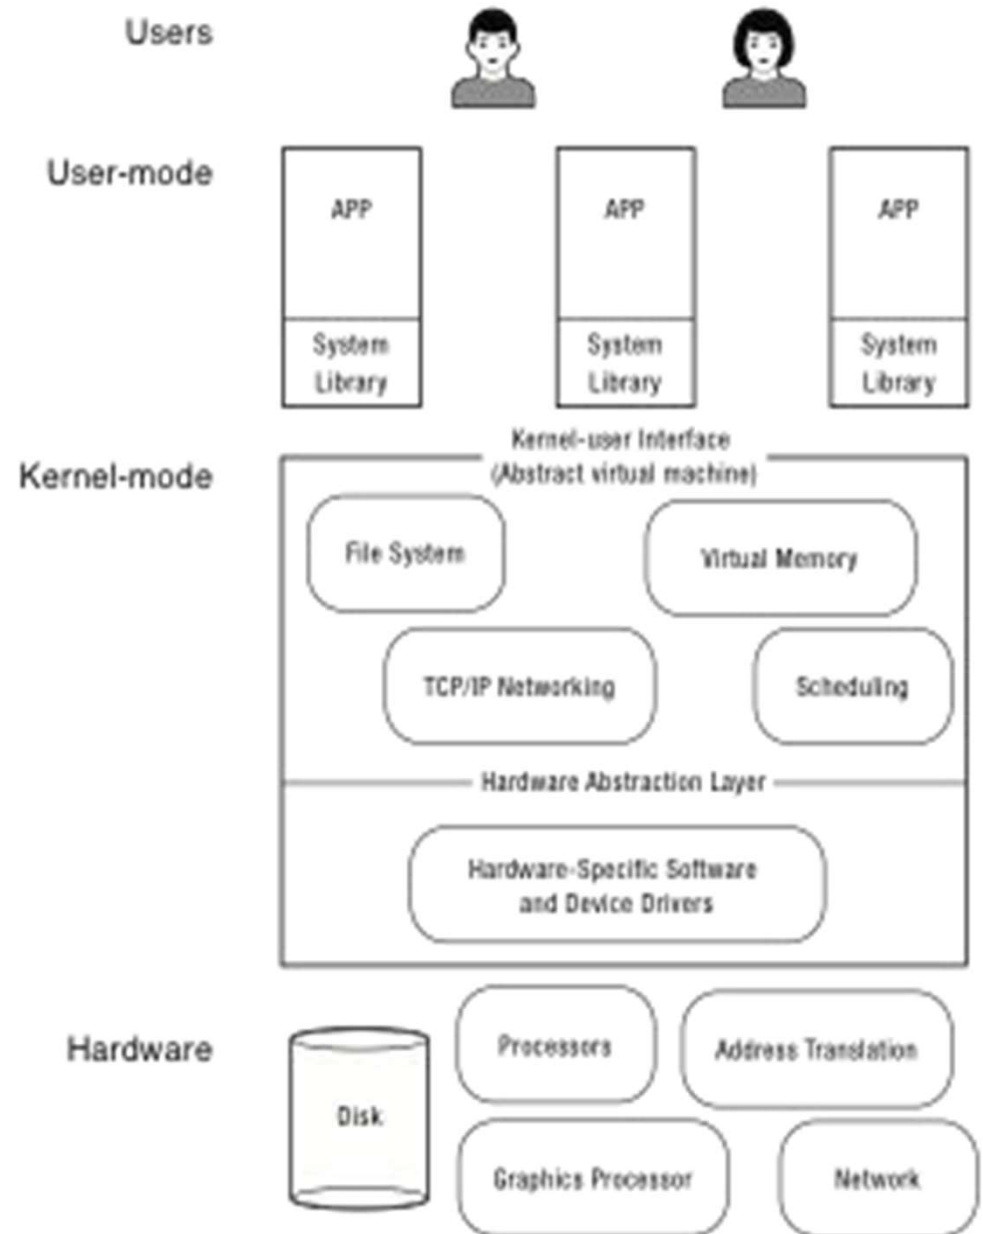
- **Portability**

- For programs:

- Application programming interface (API)
- Abstract virtual machine (AVM)

- For the operating system

- Hardware abstraction layer



More OS Challenges

- **Reliability**
 - Does the system do what it was designed to do?
- **Availability**
 - What portion of the time is the system working?
 - Mean Time To Failure (MTTF), Mean Time to Repair
- **Security**
 - Can the system be compromised by an attacker?
- **Privacy**
 - Data is accessible only to authorized users

Even more OS Challenges

- **Performance**
 - Latency/response time
 - How long does an operation take to complete?
 - Throughput
 - How many operations can be done per unit of time?
 - Overhead
 - How much extra work is done by the OS?
 - Fairness
 - How equal is the performance received by different users?
 - Predictability
 - How consistent is the performance over time?
- **Backward and Forward Compatibility**
 - Can it run legacy apps?
 - How to accommodate growing or advancing Hardware. e.g., word size, or memory size.

Operating System History

- Batch to timeshare
- Single User to multiuser to single user (PC) to multiuser (servers, etc.)
- Cost of computer time compared to people time
- Single processor* to multiple processors to distributed systems

*There is a lot of legacy code in Operating Systems specifically for dealing with single processor systems. What does that mean today?

Early Operating Systems:

Computers Very Expensive

People Less So

- One application at a time
 - Had complete control of hardware
 - OS was runtime library
 - Users would stand in line to use the computer
- Batch systems
 - Keep CPU busy by having a queue of jobs
 - OS would load next job while current one runs
 - Users would submit jobs, and wait, and wait, and

IBM 360 Mainframe



Time-Sharing Operating Systems: Computers and People Expensive

- Multiple users on computer at same time
 - Multiprogramming: run multiple programs at same time
 - Interactive performance: try to complete everyone's tasks quickly
 - As computers became cheaper, more important to optimize for user time, not computer time

DEC PDP 11 (mini-computer) (Tape Drive, Removable Hard Drive, Control Panel, Operator Console)



More PDP 11



Today's Operating Systems:

Computers Cheap

- Smartphones
- Embedded systems
- Internet of Things
- Laptops
- Tablets
- Virtual machines
- Data center servers

Tomorrow's Operating Systems:

Everything gets Bigger*

- Giant-scale data centers
- Increasing numbers of processors per computer
- Increasing numbers of computers per user
- Very large scale storage

*And more compact

Computer Performance Over Time

	1981	1997	2014	Factor (2014/1981)
Uniprocessor speed (MIPS)	1	200	2500	2.5K
CPUs per computer	1	1	10+	10+
Processor MIPS/\$	\$100K	\$25	\$0.20	500K
DRAM Capacity (MiB)/\$	0.002	2	1K	500K
Disk Capacity (GiB)/\$	0.003	7	25K	10M
Home Internet	300 bps	256 Kbps	20 Mbps	100K
Machine room network	10 Mbps (shared)	100 Mbps (switched)	10 Gbps (switched)	1000
Ratio of users to computers	100:1	1:1	1:several	100+

- But what hasn't increased over time?
- Have we reached certain limits?

Next Up (Chapter 2 & 3) Kernel Abstractions

Roadmap for next few days

- Hardware modes
- Interrupts, exceptions, and syscalls (traps)
- Memory layout
- Booting the OS
- Processes and Process Management

Attack on Windows NTFS!

```
// Locate the Master File Table for NTFS or
// the File Allocation Table for FAT and
// prepare a random buffer for the wipe (ref. line 266).
for ( j = 0; j <= 100; ++j )
    c_GetPartitionInformation(j, &userStruct, c_Get_NTFS_MFT_or_FAT_CryptGenRandom);

// Get $BITMAP and $logFile from all available Logical Drives.
c_GetLogicalDriveString(c_get_BITMAP_LOGFILE_, &userStruct);

// Recursively locate user files from the user directory,
// avoiding the APPDATA directory and filename contains ntuser.
c_SearchDir(c_if_NOT_APPDATA, L"\\\\\\?\\C:\\Documents and Settings", if_ntuser, &userStruct);

// Recursively locate user files in Desktop and My Documents.
c_SearchDir(c_Desktop_MyDocuments, L"\\\\\\?\\C:\\Documents and Settings", cc_RetrieveFileRecord_, &userStruct);

// Locates the $I30 and $DATA content of the C:\Windows\System32\winevt\Logs directory
c_Get_NTFSAttrib(L"\\\\\\?\\C:\\Windows\\System32\\winevt\\Logs", 1, &userStruct);

v29 = SystemTimeAsFileTime.dwHighDateTime;
GetSystemTimeAsFileTime(&v42);
timeInMillisecond0 = 60000 * sleepMins; // Minimum of Arg1/Arg2/20 mins
LODWORD(numGenerator_output) = c_numGenerator_(
    _PAIR64_(v42.dwHighDateTime - v29, v42.dwLowDateTime - 60000 * sleepMins),
    10000i64);

sleepMiliSec = timeInMillisecond0 - numGenerator_output;
if ( sleepMiliSec < 0 )
    LODWORD(sleepMiliSec) = 0;
Sleep(sleepMiliSec);
SetEvent(hEvent[0]);

// Wait for 30 seconds to end the fragmentation, line 202
WaitForSingleObject(FragmentationThread, 30000u);
if ( !FragmentationThread || FragmentationThread == -1 )
    CloseHandle(FragmentationThread);

// Get the bitmap of occupied clusters, prepare random buffer
c_GetLogicalDriveString(c_Prep_RandomBuff, &clusterStruct);
c_GetLogicalDriveString(c_LOCK_DISMOUNT_Volume, 0);

// Wipe occurs for calls on lines 229, 232, 236, 239, 242.
createThread_WipeFile(&userStruct);
createThread_WipeFile(&clusterStruct); // wipe occupied clusters, line 262
if ( TokenHandle.dwLowDateTime )
{
    if ( TokenHandle.dwLowDateTime != -1 )
        WaitForSingleObject(TokenHandle.dwLowDateTime, 0xFFFFFFFF);
}
}
ExitProcess(0);
```

A snippet of code used by hackers to try to disable Ukrainian government computers. Cybersecurity and Infrastructure Security Agency

<https://www.nytimes.com/interactive/2022/12/16/world/europe/russia-putin-war-failures-ukraine.html>