
CS 422/522 Design & Implementation
of Operating Systems

Lecture 1: Introduction

Zhong Shao
Dept. of Computer Science
Yale University

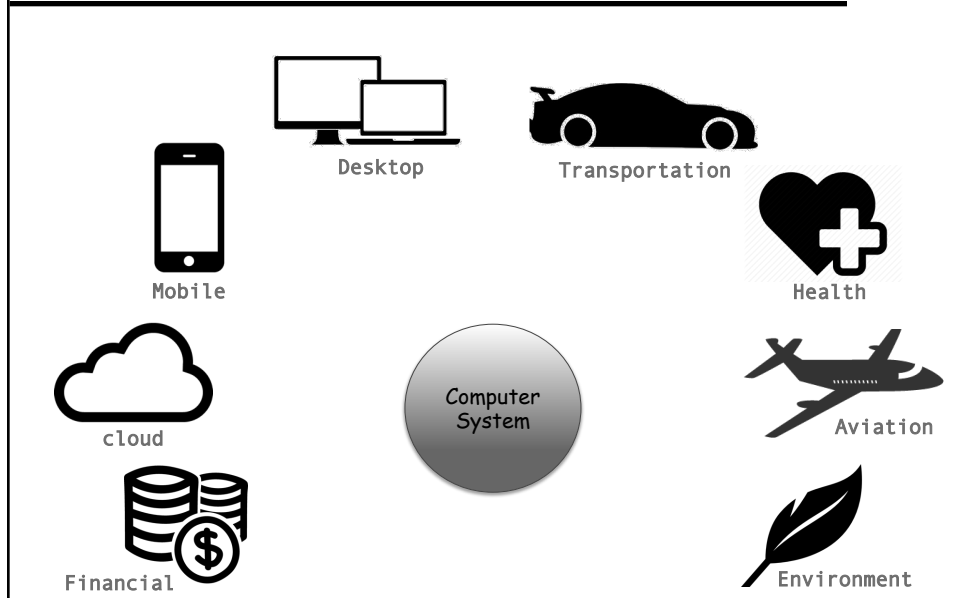
1

Today's lecture

- ◆ Why study operating systems ?
- ◆ What is an OS? What does an OS do?
- ◆ History of operating systems
- ◆ Principles of operating system design
- ◆ Course overview
 - course information
 - schedule, assignments, grading and policy
 - other organization issues
 - see web pages for more information

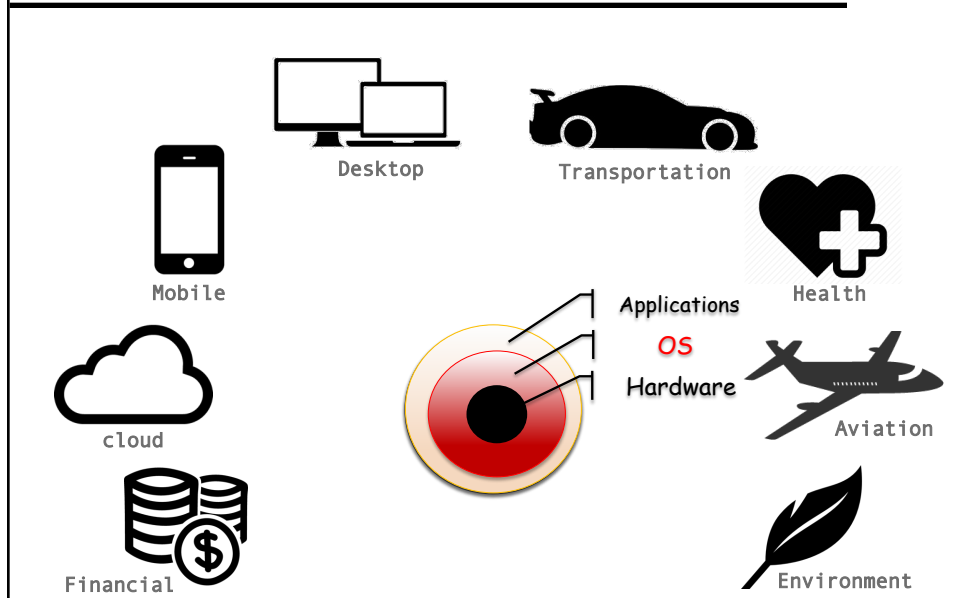
2

OS is pervasive



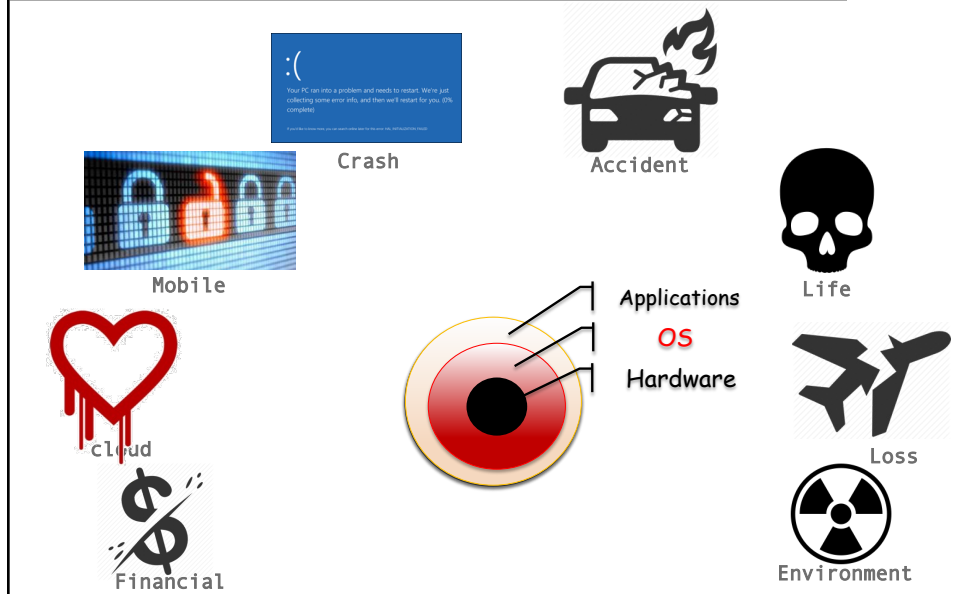
3

OS is pervasive



4

OS is pervasive & extremely important



5

OS is no longer for a single machine



Sean Everett [Follow](#)

Founder & CEO PROME Biologic Intelligence, EIR The Mission (acq. Humanizing Tech). Many startup...
Nov 16, 2016 · 9 min read

The New Battleground: Car Operating Systems

How Apple, Google, Samsung, Tesla, NVIDIA and BlackBerry are already competing behind the scenes

Recommended by Sean Everett (author)



Marius Slavesco
Nov 18, 2016 · 1 min read

Operating System: like BlackBerry's QNX, which Ford is using, but also Apple's CarOS and Android Auto extended to full OS. Currently folks are just running a fork of Ubuntu or the Ro...



The software running on the SDC should be standardized, in my opinion an open source base would be the right approach, similar with what we have in Android ecosystem.

This approach will ensure auto makers cooperation and allow everyone to contribute to the SDC ecosystem, and provide faster advanced, safe and consistent...

6

OS is no longer for a single machine

Harvard
Business
Review

TECHNOLOGY

Smart Cities Are Going to Be a Security Nightmare

by Todd Thibodeaux



In the fictional world of the video game *Watch Dogs*, you can play a hacker who takes over the central operating system of a futuristic, hyper-connected Chicago. With control over the city's security system, you can spy on residents using surveillance cameras, intercept phone calls, and cripple the city's critical infrastructure, unleashing a vicious cyberattack that brings the Windy City to its knees. <https://hbr.org/2017/04/smart-cities-are-going-to-be-a-security-nightmare>

7


OS is no longer for a single machine

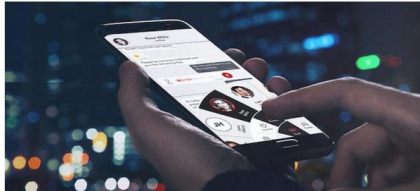
forbes.com/sites/adrianbridgwater/2018/07/03/what-is-a-blockchain-operating-system/#30121abc36ae

marks Customize Links Google Voice Zhong Shao's Cac... Gmail - Inbox Read Later CRASH Wiki FREE Logo Maker...

12,551 views | Jul 3, 2018, 06:07am

What Is A Blockchain Operating System?

 **Adrian Bridgwater** Senior Contributor @ Enterprise & Cloud
I track enterprise software application development & data management.



The Nynja blockchain-enabled virtual operating system has a communications layer and a secure payments layer, with a multi-currency wallet. [Nynja](#)

Blockchain, as more and more people are finding out, is a ledger system for keeping records. More specifically, blockchain is an open (i.e. public, not held inside one company, unless it's a private blockchain) distributed (it

8

OS is no longer for a single machine

← → ↻ 🔒 Not Secure | ivoras.net/blog/tree/2009-06-07-the-matrix-as-an-operating-system.html

🔖 Apps 🌟 Bookmarks 📁 Customizable Links 🗣️ Google Voice 📧 Zhong Shao's Cac... 📧 Gmail - Inbox 🕒 Read Later

[The Matrix as an Operating System](#)

Sun 07 June 2009
By [Ivan Voras](#)

In [Misc.](#)

tags: [NonsenseThe Matrix](#)

I wrote this in 2003, and am republishing it here for fun :) Now I know slightly more and would probably rewrite the whole thing from a more microkernel-oriented angle and with better discussion of security properties but even now, it's a complete overkill. Beware of geeking out!

Since one of my fascinations is operating systems' design, implementation and maintenance, ever since I've first seen the Matrix movies I've thought some of the concepts in them can be related to familiar concepts in operating systems:

- The Matrix world: a running operating system, with userland (the "common" world in which people apparently live), and the kernel (the "Matrix" proper). Apparently it's a pretty buggy OS...
- People: processes, both kernel processes and user processes. There's a big distinction between normal, "unprivileged" people, and daemons with root privileges - "agents". Root daemons can open privileged ports, kill random processes, manage memory, etc.
- Matrix: the kernel. It looks like a message passing kernel, not necessarily a microkernel (though they are some microkernel aspects, such as the abundance of kernel processes, strict separation of duty between them, and the already mentioned message passing). Kernel manages all processes, and performs operations on their behalf (such as keeping them alive, servicing them and recycling them). But there's an apparent security defect: some userland processes can (because of a bug) transfer and execute parts of their programs in the kernel space. Only certain syscalls are affected (the "phones"), and this kind of privilege escalation garbles the userland process' return stack, such that if the process receives a signal, it segfaults and is garbage collected (i.e. if you're killed in the Matrix, you're dead for real).
- Oracle: the process (task) scheduler. Has all the numbers from process monitoring (resource usage) and knows in advance (broadly) how to schedule them to run to their optimum.
- Agents: system monitoring / intrusion detection / prevention system (IDS / IPS) with heuristic operation. Most of them have a kernel part (kernel module) but are basically daemons run with superuser privileges in the userland. They are tasked to find and kill processes which attempt to violate system security.
- The trainman: kernel-userland gateway / message passing queue. You've got to go through him if you want to validly pass data between userland and kernel.

9

OS in media / movies & on reddit

- ◆ OS1 in HER <https://www.youtube.com/watch?v=GV01B5kVsC0>
- ◆ Hollywood OS <http://wiki.c2.com/?HollywoodOs>
- ◆ Social Network classroom scenes
https://www.youtube.com/watch?v=-3Rt2_9d7Jg
- ◆ Should every CS major take operating systems?
https://www.reddit.com/r/compsci/comments/50q019/should_every_cs_maj_or_take_operating_systems/
- ◆ How a course in operating systems changed me?
https://www.reddit.com/r/programming/comments/2n0nw5/how_a_course_in_operating_systems_changed_me/

10

Why study operating systems ?

- ◆ Understand how “computers” work under the hood
 - Magic for “infinite” CPUs, memory devices, network computing
 - Tradeoffs btw. performance & functionality, division of labor btw. HW & SW
 - Combine language, hardware, data structures, and algorithms
- ◆ Become a much better “programmer & architect” with a deeper level of “computational thinking” skills
- ◆ Help you make informed decisions
 - What “computers / gadgets” to buy? should I upgrade the HW or the OS?
 - “Can do” attitude: what’s going on with new computers, robots, and cars?
- ◆ Give you experience in hacking systems software
 - “this system is so slow; can I do anything about it?”

11

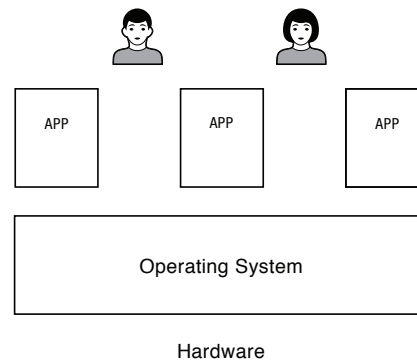
What’s interesting?

- ◆ OS is a key part of a computer system
 - it makes our life better (or worse)
 - it is “magical” and we want to understand how
 - it has “power” and we want to have the power
- ◆ OS is complex
 - how many procedures does a key stroke invoke?
 - real OS is huge and insanely expensive to build
 - * Windows 8: many years, thousands of people. Still doesn’t work well
- ◆ How to deal with complexity?
 - decomposition into many layers of abstraction
 - fail early, fail fast, and learn how to make things work

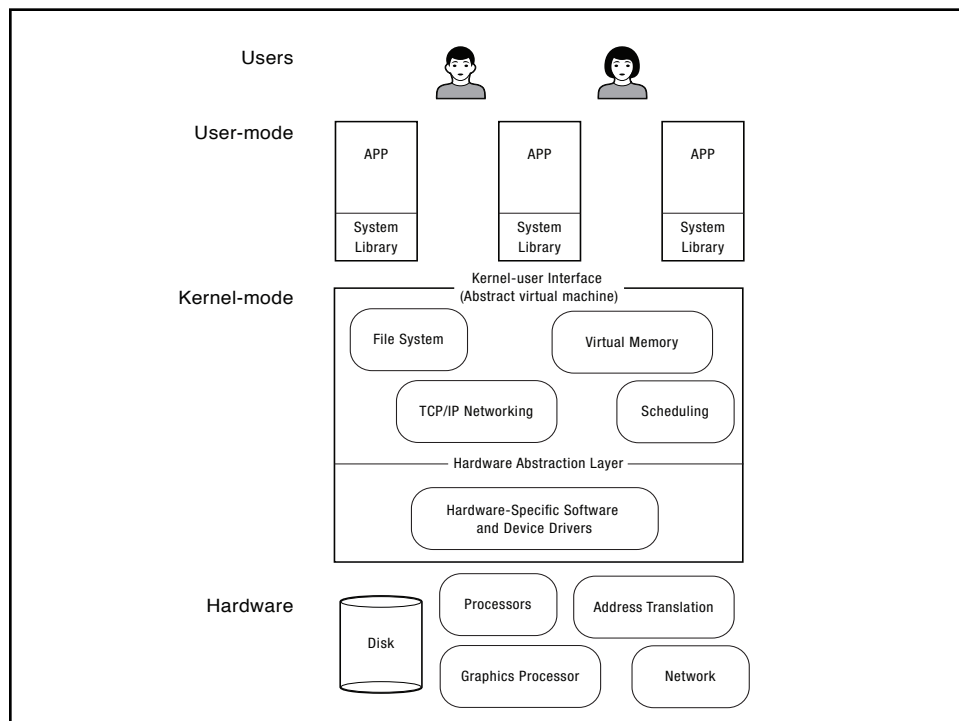
12

What is an OS?

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



13



14

Operating system roles

- ◆ Referee:
 - 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, ...

19

Example: file systems

- ◆ Referee
 - Prevent users from accessing each other's files without permission
 - Even after a file is deleted 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, ...

20

Question

- ◆ What (hardware, software) do you need to be able to run an untrustworthy application?

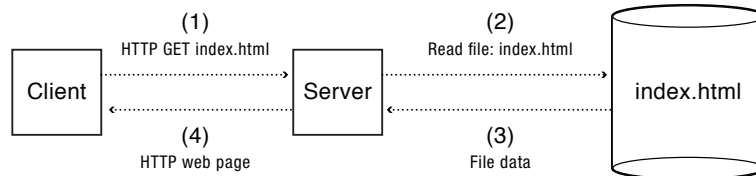
21

Question

- ◆ How should an operating system allocate processing time between competing uses?
 - Give the CPU to the first to arrive?
 - To the one that needs the least resources to complete? To the one that needs the most resources?

22

Example: web service



- ◆ How does the server manage many simultaneous client requests?
- ◆ How do we keep the client safe from spyware embedded in scripts on a web site?
- ◆ How to make updates to the web site so that clients always see a consistent view?

23

What does an OS do ?

- ◆ OS converts bare HW into nicer abstraction
 - **provide coordination:** allow multiple applications/users to work together in efficient and fair way (memory protection, concurrency, file systems, networking)
 - **provide standard libraries and services** (program execution, I/O operations, file system manipulations, communications, resource allocation and accounting)
- ◆ For each OS area, you ask
 - what is the hardware interface --- the physical reality ?
 - what is the application interface (API) --- the nicer abstraction?

24

Example of OS coordination: protection

Goal: isolate bad programs and people (security)

Solutions:

- CPU Preemption
 - * give application something, can always take it away (via clock interrupts)
- Dual mode operation
 - * when in the OS, can do anything (kernel-mode)
 - * when in a user program, restricted to only touching that program's memory (user-mode)
- Interposition
 - * OS between application and "stuff"
 - * track all pieces that application allowed to use (in a table)
 - * on every access, look in table to check that access legal
- Memory protection: address translation

25

Example: address translation

Restrict what a program can do by restricting what it can touch!

◆ Definitions:

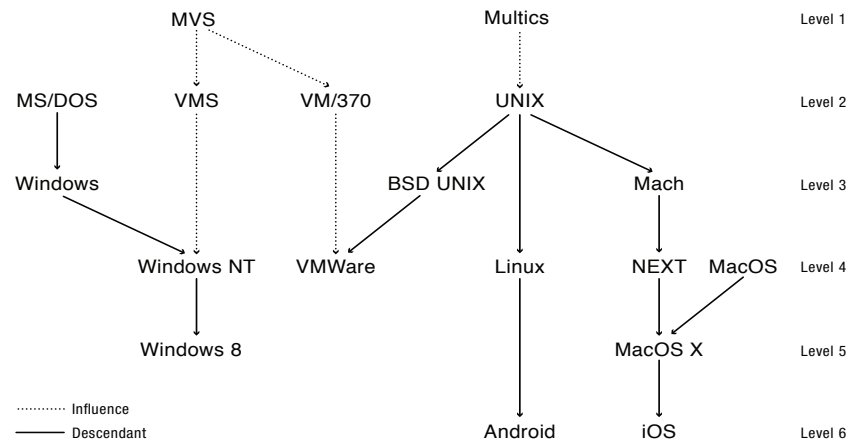
- Address space: all addresses a program can touch
- Virtual address: addresses in process' address space
- Physical address: address of real memory
- Translation: map virtual to physical addresses

◆ Virtual memory

- Translation done using per-process tables (page table)
- done on every load and store, so uses hardware for speed
- protection? If you don't want process to touch a piece of physical memory, don't put translation in table.

26

OS history



27

Challenges in writing OS

- ◆ Concurrent programming is hard
- ◆ Hard to use high-level programming languages
 - device drivers are inherently low-level
 - real-time requirement (garbage collection? probably not)
 - lack of debugging support (use simulation)
- ◆ Tension between functionality and performance
- ◆ Portability and backward compatibility
 - many APIs are already fixed (e.g., GUI, networking)
 - OS design tradeoffs change as HW changes !

28

Challenges in writing OS (cont'd)

- ◆ 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

29

Main techniques & design principles

- ◆ Keep things simple !
- ◆ Use abstraction (<https://www.youtube.com/watch?v=8Are9dDbW24>)
 - hide implementation complexity behind simple interface
- ◆ Use modularity
 - decompose system into isolated pieces
- ◆ But what about performance
 - find bottlenecks --- the 80-20 rule
 - use prediction and exploits locality (cache)
- ◆ What about security and reliability?

More research is necessary!

30

Course information

Required textbook:

Operating Systems: Principles & Practice (2nd Edition) by T. Anderson and M. Dahlin

information, assignments, & lecture notes are available on-line
we won't use much paper

Official URL: <http://flint.cs.yale.edu/cs422>

for help, go to the [Ed Discussion](#) site (accessible via canvas):

<https://edstem.org/us/courses/13152/discussion/>

31

Course information (cont'd)

- ◆ **13 week lectures on OS fundamentals**
 - class participation is strongly recommended
- ◆ **Course requirements**
 - 70% on assignments (as1 - as6)
 - 30% in-class midterm (Monday, November 7, tentative)
- ◆ **Assignments (as1-as6) and course policies**
 - build a small but real OS kernel, bootable on real PCs.
 - extensive hacking (in C & x86 assembly) but highly rewarding
 - 2 persons / team (one person team is OK too).
 - 6 free late days (3 day late max per assignment).

32

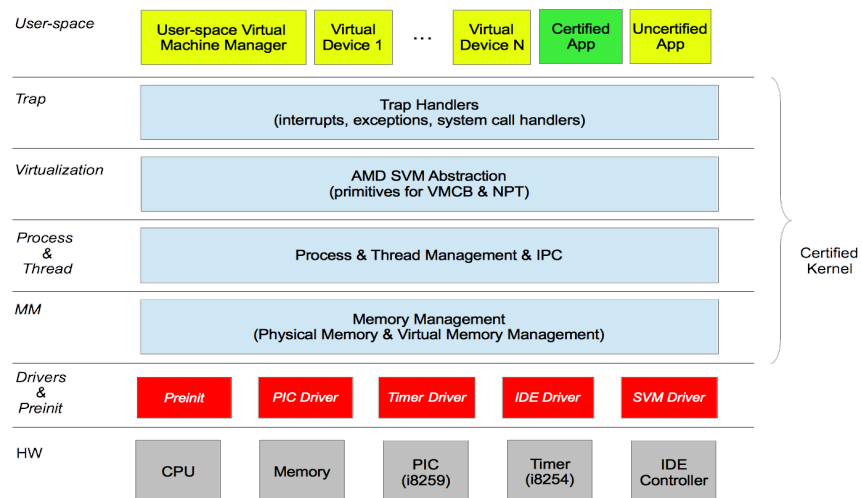
Programming assignments

- ◆ Assignment topics (tentative)
 - Bootloader & physical memory management
 - Container and virtual memory management
 - Process management & trap handling
 - Multicore and preemption
 - File system & other topics
 - IPC, Shell, and other topics
- ◆ How
 - Each assignment takes two weeks
 - Most assignments due on Thursdays 11:59pm ET
- ◆ The Lab
 - Linux cluster in ZOO (3rd Floor of AKW or Room 111 at 17HH)
 - You can setup your own machine to do projects

33

Programming assignments (cont'd)

Based on mCertIKOS (Yale FLINT) & JOS (from MIT)



34

- ◆ Break kernel interdependency by insisting on careful layer decomposition
- ◆ With the right methodology, every CS major should be able to write an OS kernel from scratch

