


CS162
Operating Systems and
Systems Programming
Lecture 1

What is an Operating System?

January 22nd, 2019
Prof. John Kubiatowicz
<http://cs162.eecs.Berkeley.edu>

Greatest Artifact of Human Civilization...

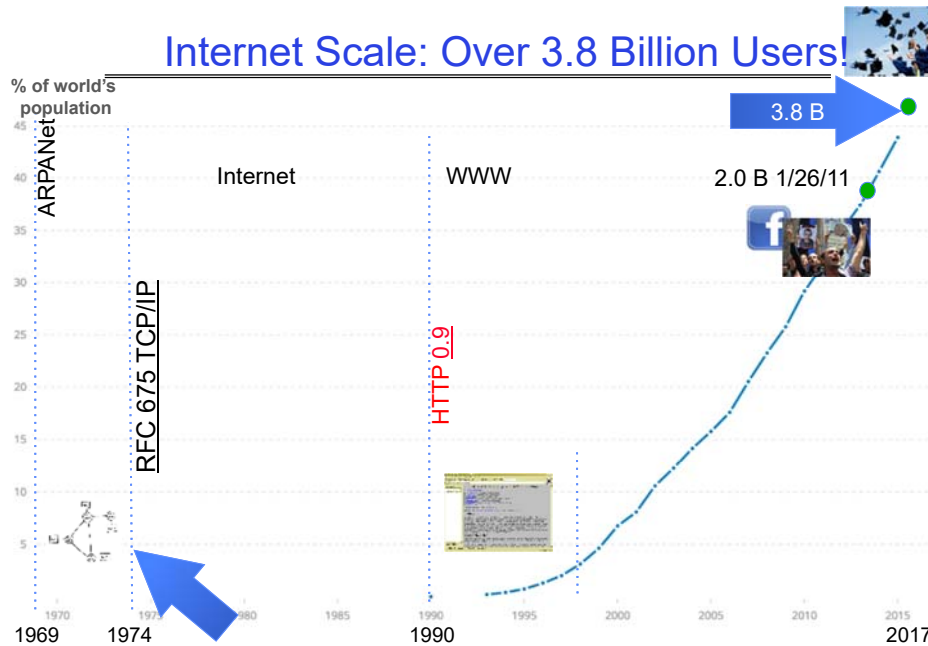


1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.2

Internet Scale: Over 3.8 Billion Users!



1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.3

Operating Systems are at the Heart of it All!

- Provide abstractions to apps
 - File systems
 - Processes, threads
 - VM, containers
 - Naming system
 - ...
- Manage resources:
 - Memory, CPU, storage, ...
- Achieves the above by implementing specific algorithms and techniques:
 - Scheduling
 - Concurrency
 - Transactions
 - Security
 -




1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.4

Example: What's in a Search Query?



- Complex interaction of multiple components in multiple administrative domains
 - Systems, services, protocols, ...

1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.5

Why take CS162?

- Some of you will actually design and build operating systems or components of them.
 - Perhaps more now than ever
- Many of you will create systems that utilize the core concepts in operating systems.
 - Whether you build software or hardware
 - The concepts and design patterns appear at many levels
- All of you will build applications, etc. that utilize operating systems
 - The better you understand their design and implementation, the better use you'll make of them.

1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.6

Goals for Today

- What is an Operating System?
 - And – what is it not?
- What makes Operating Systems so exciting?
- Oh, and “How does this class operate?”

Interactive is important!

Ask Questions!

Slides courtesy of David Culler, Anthony D. Joseph, John Kubiatowicz, AJ Shankar, George Necula, Alex Aiken, Eric Brewer, Ras Bodik, Ion Stoica, Doug Tygar, and David Wagner.


1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.7

What is an operating system?

- Special layer of software that provides application software access to hardware resources
 - Convenient abstraction of complex hardware devices
 - Protected access to shared resources
 - Security and authentication
 - Communication amongst logical entities



1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.8

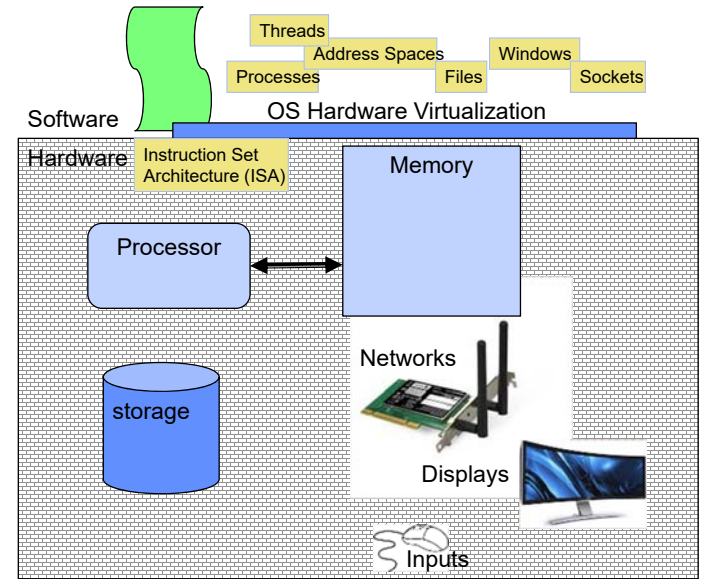
Operator ...




Switchboard Operator

Computer Operators

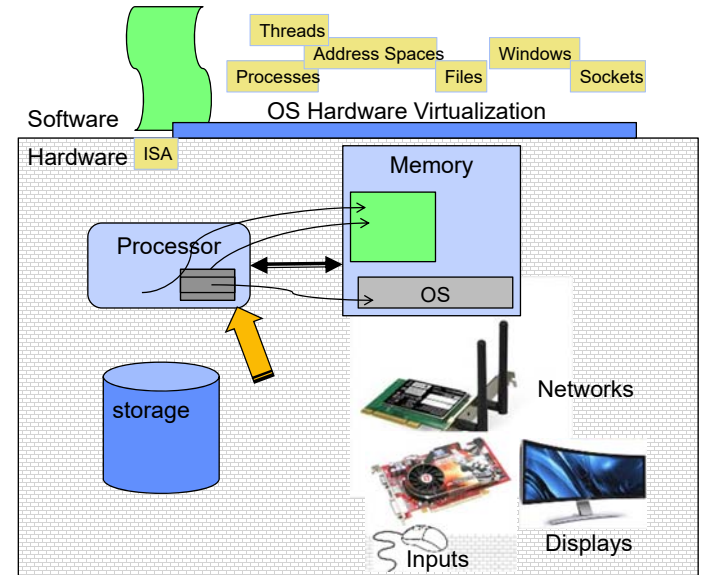
OS Basics: "Virtual Machine" Boundary



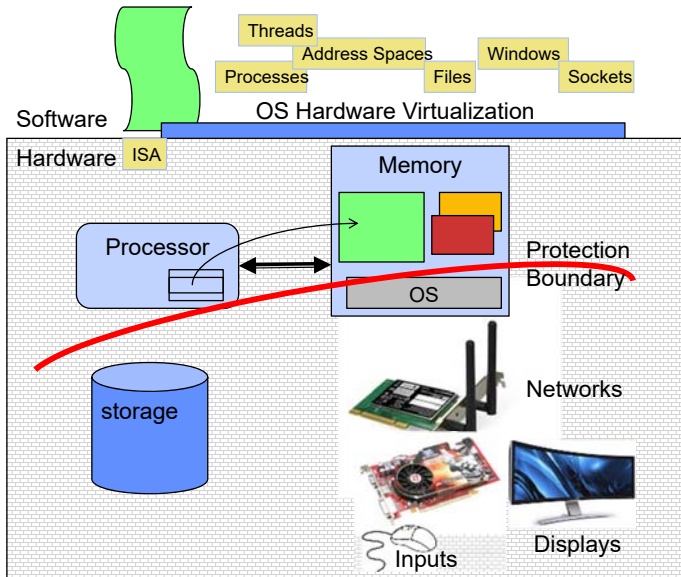
OS Basics: Program ⇒ Process



OS Basics: Context Switch



OS Basics: Scheduling, Protection

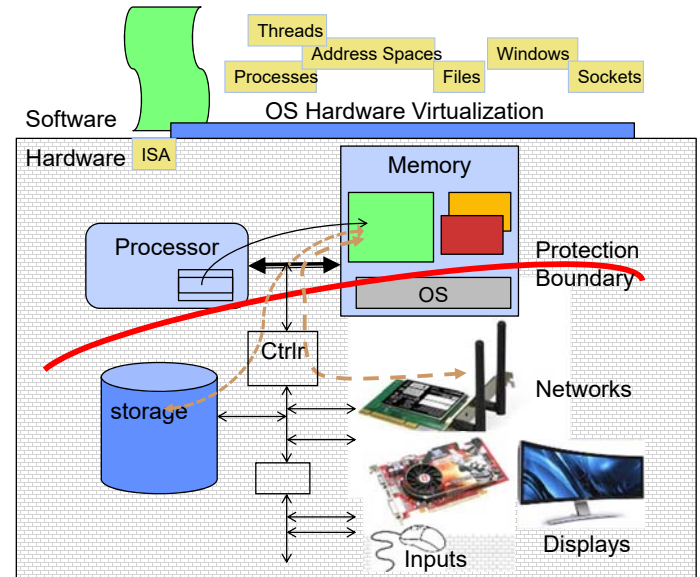


1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.13

OS Basics: I/O

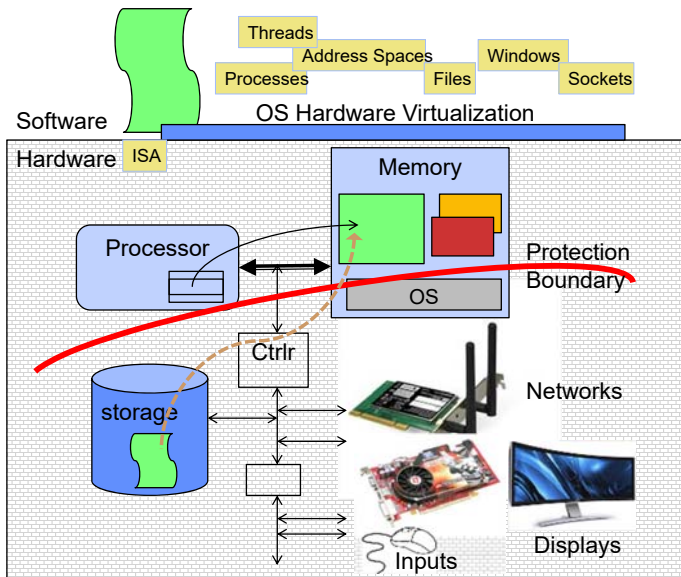


1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.14

OS Basics: Loading



1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.15



What makes Operating Systems Exciting and Challenging?

1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.16

Technology Trends: Moore's Law

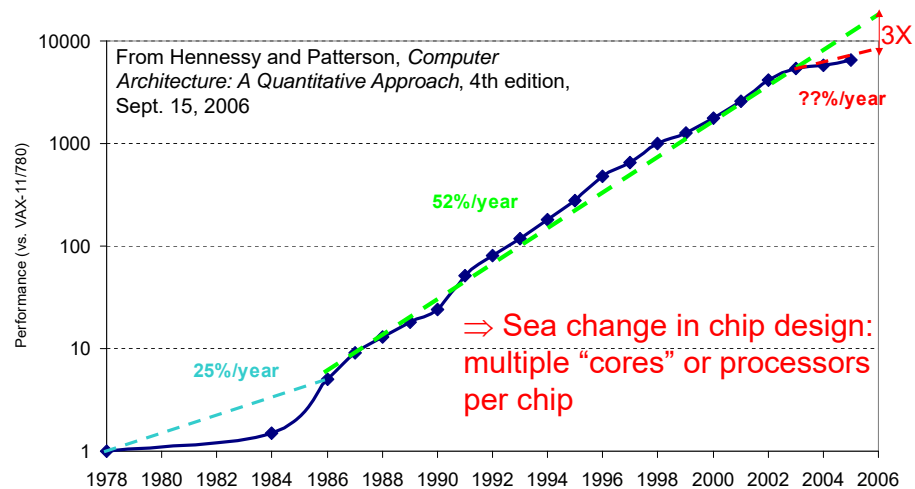


2X transistors/Chip Every 1.5 years
Called "**Moore's Law**"

Microprocessors have become smaller, denser, and more powerful

Gordon Moore (co-founder of Intel) predicted in 1965 that the transistor density of semiconductor chips would double roughly every 18 months

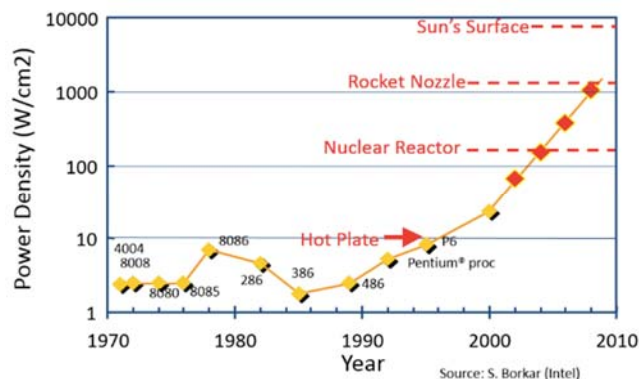
Big Challenge: Slowdown in Joy's law of Performance



- VAX : 25%/year 1978 to 1986
- RISC + x86 : 52%/year 1986 to 2002
- RISC + x86 : ??%/year 2002 to present

⇒ Sea change in chip design: multiple "cores" or processors per chip

Another Challenge: Power Density




- Moore's law extrapolation
 - Potential power density reaching amazing levels!
- Flip side: battery life very important
 - Moore's law yielded more functionality at equivalent (or less) total energy consumption

ManyCore Chips: The future arrived in 2007

- Intel 80-core multicore chip (Feb 2007)
 - 80 simple cores
 - Two FP-engines / core
 - Mesh-like network
 - 100 million transistors
 - 65nm feature size
- Intel Single-Chip Cloud Computer (August 2010)
 - 24 "tiles" with two cores/tile
 - 24-router mesh network
 - 4 DDR3 memory controllers
 - Hardware support for message-passing

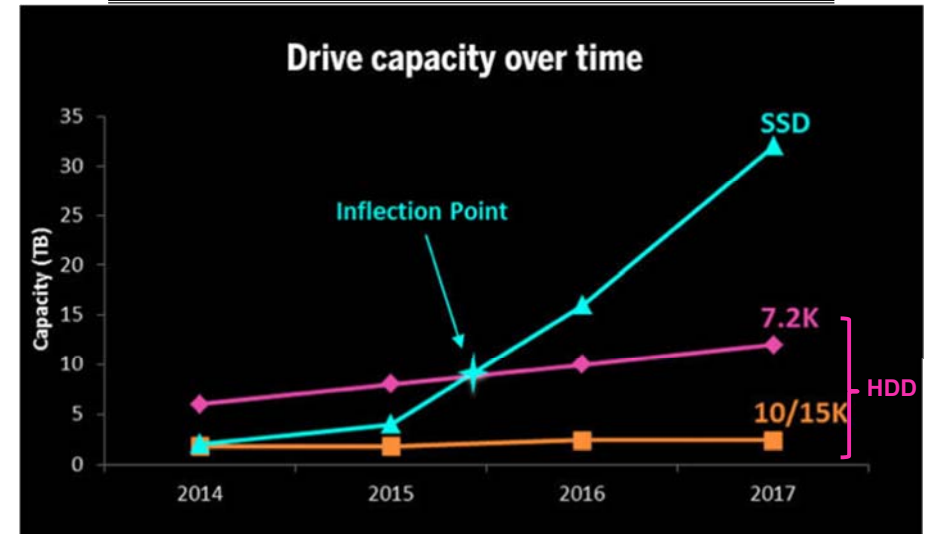
- How to program these?
 - Use 2 CPUs for video/audio
 - Use 1 for word processor, 1 for browser
 - 76 for virus checking???
- **Parallelism must be exploited at all levels**
- Amazon X1 instances (2016)
 - 128 virtual cores, 2 TB RAM

But then Moore's Law Ended...




- Moore's Law has (officially) ended -- Feb 2016
 - No longer getting 2 x transistors/chip every 18 months...
 - or even every 24 months
- May have only 2-3 smallest geometry fabrication plants left:
 - Intel and Samsung and/or TSMC
- Vendors moving to 3D stacked chips
 - More layers in old geometries

Storage Capacity Still Growing




(source: <https://www.networkworld.com/article/3153244/data-center/solid-state-drives-are-now-larger-than-hard-disk-drives-the-impact-for-your-data-center.html>)

Network Capacity Still Increasing



(source: <http://www.ospmag.com/issue/article/Time-Is-Not-Always-On-Our-Side>)

Internet Scale: 1.06 Billion Hosts (Jan 2017)



Source: Internet Systems Consortium (www.isc.org)

Internet Scale: Over 3.8 Billion Users!

WORLD INTERNET USAGE AND POPULATION STATISTICS DEC 31, 2017 - Update						
World Regions	Population (2018 Est.)	Population % of World	Internet Users 31 Dec 2017	Penetration Rate (% Pop.)	Growth 2000-2018	Internet Users %
Africa	1,287,914,329	16.9 %	453,329,534	35.2 %	9,941 %	10.9 %
Asia	4,207,588,157	55.1 %	2,023,630,194	48.1 %	1,670 %	48.7 %
Europe	827,650,849	10.8 %	704,833,752	85.2 %	570 %	17.0 %
Latin America / Caribbean	652,047,996	8.5 %	437,001,277	67.0 %	2,318 %	10.5 %
Middle East	254,438,981	3.3 %	164,037,259	64.5 %	4,893 %	3.9 %
North America	363,844,662	4.8 %	345,660,847	95.0 %	219 %	8.3 %
Oceania / Australia	41,273,454	0.6 %	28,439,277	68.9 %	273 %	0.7 %
WORLD TOTAL	7,634,758,428	100.0 %	4,156,932,140	54.4 %	1,052 %	100.0 %

NOTES: (1) Internet Usage and World Population Statistics estimates in Dec 31, 2017. (2) CLICK on each world region name for detailed regional usage information. (3) Demographic (Population) numbers are based on data from the [United Nations Population Division](#). (4) Internet usage information comes from data published by [Nielsen Online](#), by the [International Telecommunications Union](#), by [GfK](#), by local ICT Regulators and other reliable sources. (5) For definitions, navigation help and disclaimers, please refer to the [Website Surfing Guide](#). (6) The information from this website may be cited, giving the due credit and placing a link back to www.internetworldstats.com. Copyright © 2018, Miniwatts Marketing Group. All rights reserved worldwide.

(source: <http://www.internetworldstats.com/stats.htm>)

Not Only PCs connected to the Internet

- In 2011, smartphone shipments exceeded PC shipments!
- 2011 shipments:
 - 487M smartphones
 - 414M PC clients
 - » 210M notebooks
 - » 112M desktops
 - » 63M tablets
 - 25M smart TVs
- 4 billion phones in the world → smartphones over next few years
- Then...

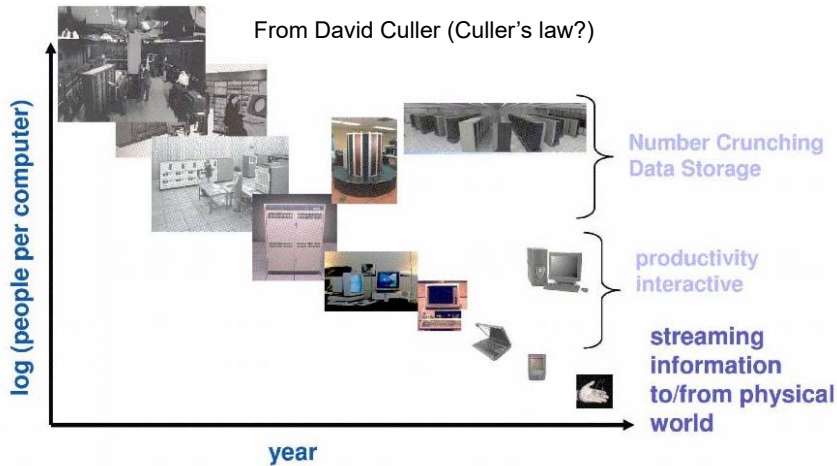
1.53B in 2017

262.5M in 2017

164M in 2017

39.5M in 2017

People-to-Computer Ratio Over Time




- Today: multiple CPUs/person!
 - Approaching 100s?

Societal Scale Information Systems (Or the "Internet of Things"?)

- The world is a large distributed system
 - Microprocessors in everything
 - Vast infrastructure behind them

Who am I?

- Professor John Kubiawicz (Prof “Kubi”)
 - Background in Hardware Design
 - » Alewife project at MIT
 - » Designed CMMU, Modified SPARC processor
 - » Helped to write operating system
 - Background in Operating Systems
 - » Worked for Project Athena (MIT)
 - » OS Developer (device drivers, network file systems)
 - » Worked on Clustered High-Availability systems.
 - Peer-to-Peer
 - » OceanStore project – Store your data for 1000 years
 - » Tapestry and Bamboo – Find you data around globe
 - SwarmLAB/Berkeley Lab for Intelligent Edge
 - » Global Data Plane (GDP)/DataCapsules
 - » Fog Robotics
 - Quantum Computing
 - » Exploring architectures for quantum computers
 - » CAD tool set yields some interesting results



Alewife



Tesselation



OceanStore

1/22/19


Kubiawicz CS162 © UCB Fall 2019

29


CS162 TAs: Sections TBA




William Sheu
(Head TA)



Jason Chin




Alexander Kozarian



Jonathan Murata



Will Wang



Eric Zhou



Joan Zhu


1/22/19

Kubiawicz CS162 © UCB Fall 2019

Lec 1.30

Infrastructure, Textbook & Readings

- Infrastructure
 - Website: <http://cs162.eecs.berkeley.edu>
 - Piazza: <https://piazza.com/berkeley/spring2019/cs162>
 - Webcast (Through Cal Central):
<https://calcentral.berkeley.edu/academics/teaching-semester/spring-2019/class/compsci-162-2019-B>
- Textbook: Operating Systems: Principles and Practice (2nd Edition) Anderson and Dahlin
- Recommend: Operating Systems Concepts, 9th Edition Silberschatz, Galvin, Gagne
 - Copies in Bechtel
- Online supplements
 - See course website
 - Includes Appendices, sample problems, etc.
 - Networking, Databases, Software Eng, Security
 - **Some Research Papers!**



Lec 1.31

1/22/19

Kubiawicz CS162 © UCB Fall 2019

Syllabus

- OS Concepts: How to Navigate as a Systems Programmer!
 - Process, I/O, Networks and Virtual Machines
- Concurrency
 - Threads, scheduling, locks, deadlock, scalability, fairness
- Address Space
 - Virtual memory, address translation, protection, sharing
- File Systems
 - I/O devices, file objects, storage, naming, caching, performance, paging, transactions, databases
- Distributed Systems
 - Protocols, N-Tiers, RPC, NFS, DHTs, Consistency, Scalability, multicast
- Reliability & Security
 - Fault tolerance, protection, security
- Cloud Infrastructure

1/22/19

Kubiawicz CS162 © UCB Fall 2019

Lec 1.32

Learning by Doing

- Individual Homeworks: Learn Systems Programming
 - 0. Tools, Autograding, recall C, executable
 - 1. Simple Shell
 - 2. Web server
 - 3. Memory allocation
- Three Group Projects (Pintos in C)
 - 1. Threads & Scheduling
 - 2. User-programs
 - 3. File Systems

1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.33

Getting started

- Start homework 0 right away (hopefully Thursday 1/24)
 - Github account
 - Registration survey
 - Vagrant virtualbox – VM environment for the course
 - » Consistent, managed environment on your machine
 - Get familiar with all the cs162 tools
 - Submit to autograder via git
- Sections on Friday – attend any section you want
 - We'll assign permanent sections after forming project groups
- This is an Early Drop Deadline course (February 1st)
 - If you are not serious about taking, please drop early
 - Dept will continue to admit students as other students drop
- On the waitlist ???
 - Unfortunately, we maxed out sections and TA Support
 - If people drop, we can move others off waitlist

1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.34

Group Project Simulates Industrial Environment

- Project teams have 4 members (try really hard to find 4 members – 3 members requires serious justification)
 - Must work in groups in “the real world”
 - Same section much preferred
- Communicate with colleagues (team members)
 - Communication problems are natural
 - What have you done?
 - What answers you need from others?
 - You must document your work!!!
- Communicate with supervisor (TAs)
 - What is the team's plan?
 - What is each member's responsibility?
 - Short progress reports are required
 - **Design Documents: High-level description for a manager!**

1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.35

Grading

- 42% three midterms (14% each). No class day of MT
 - Thursday, 2/28, TBA, tentatively 7-9pm (may ⇒ 8-10pm)
 - Thursday, 4/04, TBA, tentatively 7-9pm
 - Thursday, 5/02, TBA, tentatively 7-9pm
 - Although we have rooms scheduled in 7-9pm time slot, I'm going to try to get rooms scheduled during class time +30 instead (i.e. 5-7pm). Stay tuned!
- 35% projects
- 15% homework
- 8% participation
- *No final exam*
- Projects
 - Initial design document, Design review, Code, Final design
 - Submission via *git push* triggers autograder

1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.36

Personal Integrity

- UCB Academic Honor Code: "As a member of the UC Berkeley community, I act with honesty, integrity, and respect for others."

<http://asuc.org/honorcode/resources/HC%20Guide%20for%20Syllabi.pdf>


CS 162 Collaboration Policy

- Explaining a concept to someone in another group
- Discussing algorithms/testing strategies with other groups
- ✓ Helping debug someone else's code (in another group)
- Searching online for generic algorithms (e.g., hash table)

- ✗ **Sharing code or test cases with another group**
- ✗ **Copying OR reading another group's code or test cases**
- ✗ **Copying OR reading online code or test cases from prior years**

We compare all project submissions against prior year submissions and online solutions and will take actions (described on the course overview page) against offenders

Typical Lecture Format



- 1-Minute Review
- 20-Minute Lecture
- 5- Minute Administrative Matters
- 25-Minute Lecture
- 5-Minute Break (water, stretch)
- 25-Minute Lecture
- Instructor will come to class early & stay after to answer questions

Lecture Goal

Interactive!!!

What is an Operating System?



- Referee
 - Manage sharing of resources, Protection, Isolation
 - » Resource allocation, isolation, communication



- Illusionist
 - Provide clean, easy to use abstractions of physical resources
 - » Infinite memory, dedicated machine
 - » Higher level objects: files, users, messages
 - » Masking limitations, virtualization



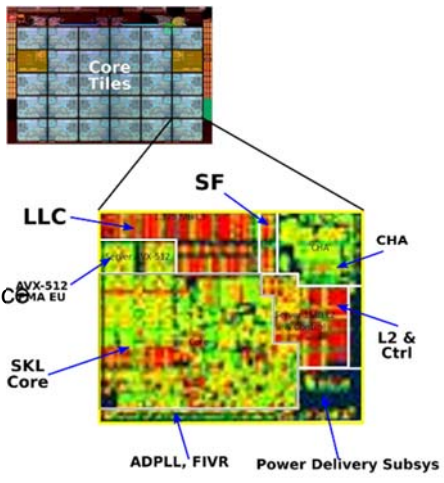
- Glue
 - Common services
 - » Storage, Window system, Networking
 - » Sharing, Authorization
 - » Look and feel

Challenge: Complexity

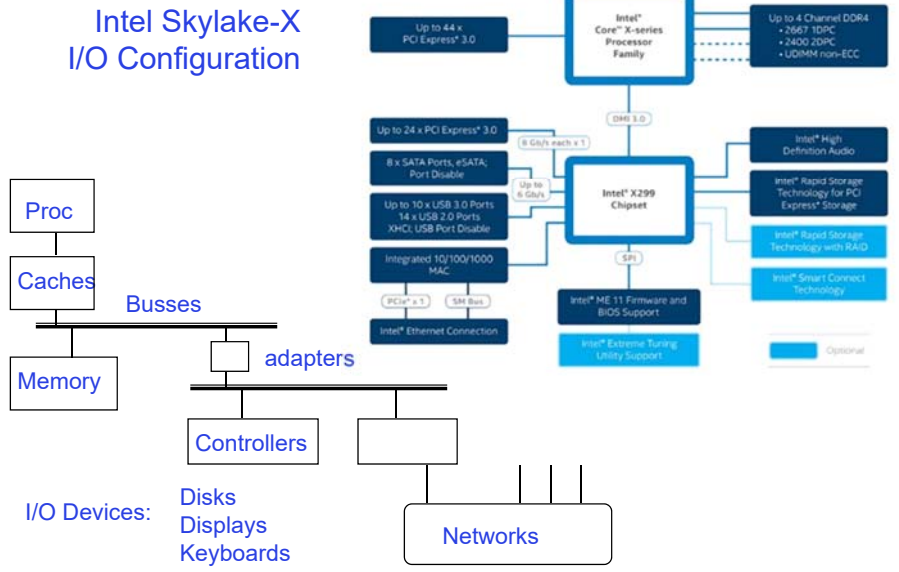
- Applications consisting of...
 - ... a variety of software modules that ...
 - ... run on a variety of devices (machines) that
 - » ... implement different hardware architectures
 - » ... run competing applications
 - » ... fail in unexpected ways
 - » ... can be under a variety of attacks
- Not feasible to test software for all possible environments and combinations of components and devices
 - The question is not whether there are bugs but how serious are the bugs!

The World Is Parallel: Intel SkyLake (2017)


- Up to 28 Cores, 58 Threads
 - 694 mm² die size (estimated)
- Many different instructions
 - Security, Graphics
- Caches on chip:
 - L2: 28 MiB
 - Shared L3: 38.5 MiB (non-inclusive)
 - Directory-based cache coherence
- Network:
 - On-chip Mesh Interconnect
 - Fast off-chip network directly supports 8-chips connected
- DRAM/chips
 - Up to 1.5 TiB



HW Functionality comes with great complexity!



Increasing Software Complexity




Millions of Lines of Code

(source <https://informationisbeautiful.net/visualizations/million-lines-of-code/>)

Example: Some Mars Rover ("Pathfinder") Requirements

- Pathfinder hardware limitations/complexity:
 - 20Mhz processor, 128MB of DRAM, VxWorks OS
 - cameras, scientific instruments, batteries, solar panels, and locomotion equipment
 - Many independent processes work together



- Can't hit reset button very easily!
 - Must reboot itself if necessary
 - Must always be able to receive commands from Earth
- Individual Programs must not interfere
 - Suppose the MUT (Martian Universal Translator Module) buggy
 - Better not crash antenna positioning software!
- Further, all software may crash occasionally
 - Automatic restart with diagnostics sent to Earth
 - Periodic checkpoint of results saved?
- Certain functions time critical:
 - Need to stop before hitting something
 - Must track orbit of Earth for communication
- A lot of similarity with the Internet of Things?
 - Complexity, QoS, Inaccessibility, Power limitations ... ?



How do we tame complexity?

- Every piece of computer hardware different
 - Different CPU
 - » Pentium, PowerPC, ColdFire, ARM, MIPS
 - Different amounts of memory, disk, ...
 - Different types of devices
 - » Mice, Keyboards, Sensors, Cameras, Fingerprint readers
 - Different networking environment
 - » Cable, DSL, Wireless, Firewalls,...
- Questions:
 - Does the programmer need to write a single program that performs many independent activities?
 - Does every program have to be altered for every piece of hardware?
 - Does a faulty program crash everything?
 - Does every program have access to all hardware?

OS Tool: Virtual Machine Abstraction

Application

Virtual Machine Interface

Operating System

Physical Machine Interface

Hardware

- Software Engineering Problem:
 - Turn hardware/software quirks ⇒ what programmers want/need
 - Optimize for convenience, utilization, security, reliability, etc...
- For any OS area (e.g. file systems, virtual memory, networking, scheduling):
 - What's the hardware interface? (physical reality)
 - What's the application interface? (nicer abstraction)

Virtual Machines

- Software emulation of an abstract machine
 - Give programs illusion they own the machine
 - Make it look like hardware has features you want
- Two types of “Virtual Machine”s
 - Process VM: supports the execution of a single program; this functionality typically provided by OS
 - System VM: supports the execution of an entire OS and its applications (e.g., VMWare Fusion, Virtual box, Parallels Desktop, Xen)



1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.49

Process VMs

- Programming simplicity
 - Each process thinks it has all memory/CPU time
 - Each process thinks it owns all devices
 - Different devices appear to have same high level interface
 - Device interfaces more powerful than raw hardware
 - » Bitmapped display ⇒ windowing system
 - » Ethernet card ⇒ reliable, ordered, networking (TCP/IP)
- Fault Isolation
 - Processes unable to directly impact other processes
 - Bugs cannot crash whole machine
- Protection and Portability
 - Java interface safe and stable across many platforms

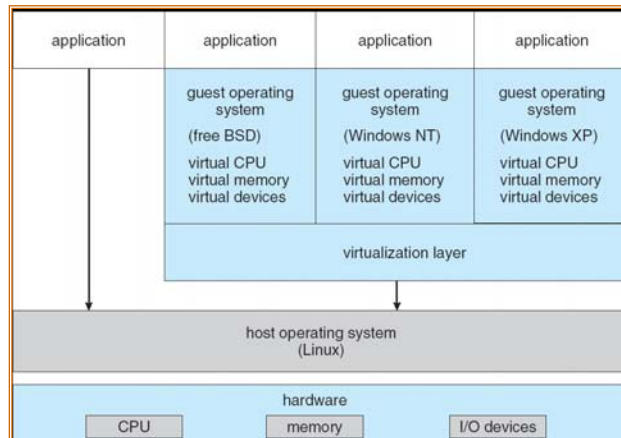
1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.50

System Virtual Machines: Layers of OSs

- Useful for OS development
 - When OS crashes, restricted to one VM
 - Can aid testing programs on other OSs



1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.51

What is an Operating System,... Really?

- Most Likely:
 - Memory Management
 - I/O Management
 - CPU Scheduling
 - Communications? (Does Email belong in OS?)
 - Multitasking/multiprogramming?
- What about?
 - File System?
 - Multimedia Support?
 - User Interface?
 - Internet Browser? ☺
- Is this only interesting to Academics??

1/22/19

Kubiatowicz CS162 © UCB Fall 2019

Lec 1.52

Operating System Definition (Cont.)

- No universally accepted definition
- “Everything a vendor ships when you order an operating system” is good approximation
 - But varies wildly
- “The one program running at all times on the computer” is the **kernel**
 - Everything else is either a system program (ships with the operating system) or an application program

“In conclusion...”

- Operating systems provide a virtual machine abstraction to handle diverse hardware
 - Operating systems simplify application development by providing standard services
- Operating systems coordinate resources and protect users from each other
 - Operating systems can provide an array of fault containment, fault tolerance, and fault recovery
- CS162 combines things from many other areas of computer science:
 - Languages, data structures, hardware, and algorithms