

Operating Systems: Principles and Practice

Tom Anderson

How This Course Fits in the UW CSE Curriculum

- CSE 333: Systems Programming
 - Project experience in C/C++
 - How to use the operating system interface
- CSE 451: Operating Systems
 - How to make a single computer work reliably
 - How an operating system works internally
- CSE 452: Distributed Systems
 - How to make a set of computers work reliably, despite failures of some nodes

New Project: xk

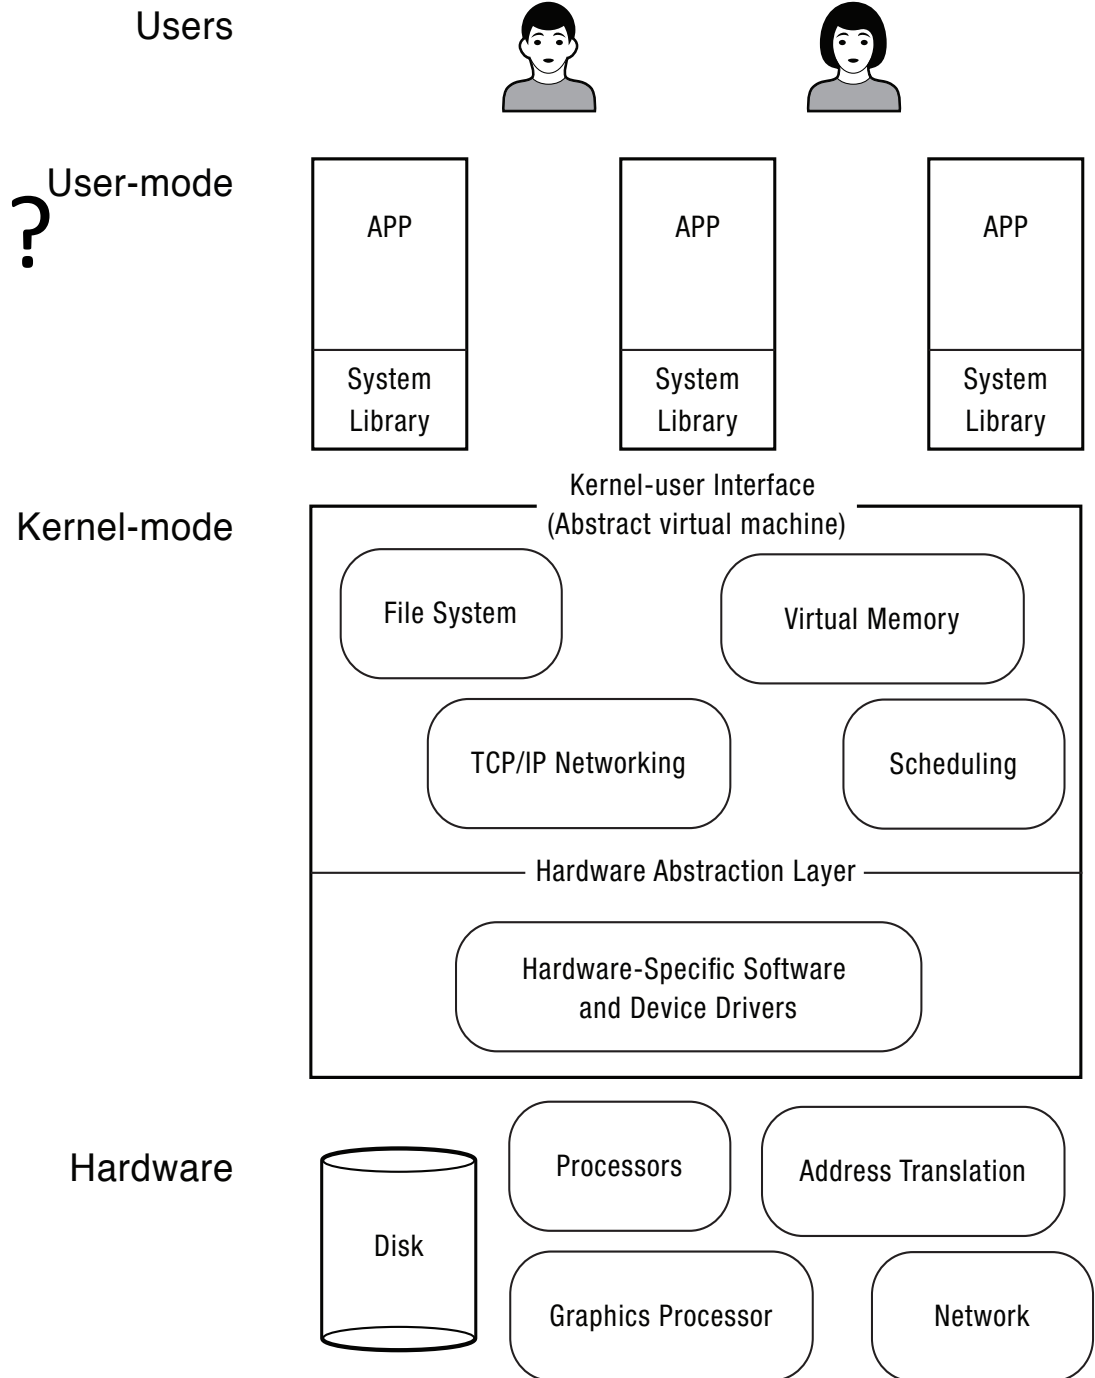
- Build an operating system
 - That can boot on a real system
 - Run multiple processes
 - Page virtual memory
 - Store file data reliably
- We give you some basic starting code
 - Five assignments, that build on each other
 - Work in **groups of 2**
- Instructions on web page
 - Download and browse code before section
 - Bring laptop or smartphone to section

Main Points (for today)

- Operating system definition
 - Software to manage a computer's resources for its users and applications
- OS challenges
 - Reliability, security, responsiveness, portability, ...
- OS history
 - How did we get here?
- How I/O works

What is an operating system?

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



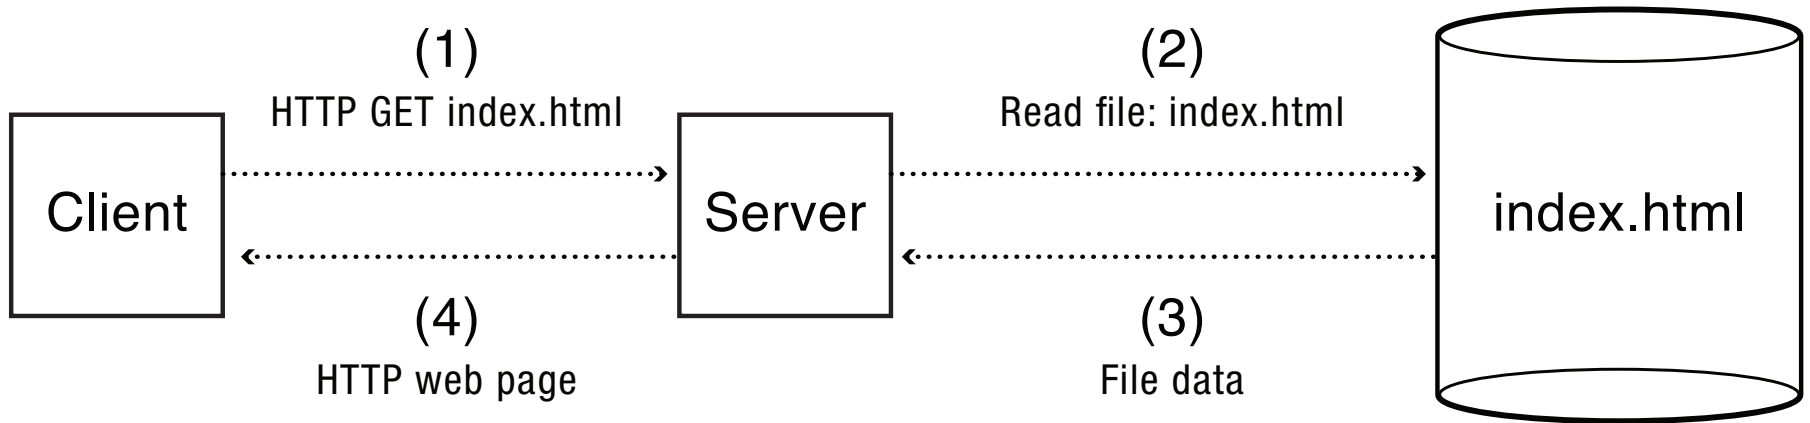
Operating System Roles

- 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, ...

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, ...

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 do make updates to the web site so that clients always see a consistent view?

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

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?

OS Challenges

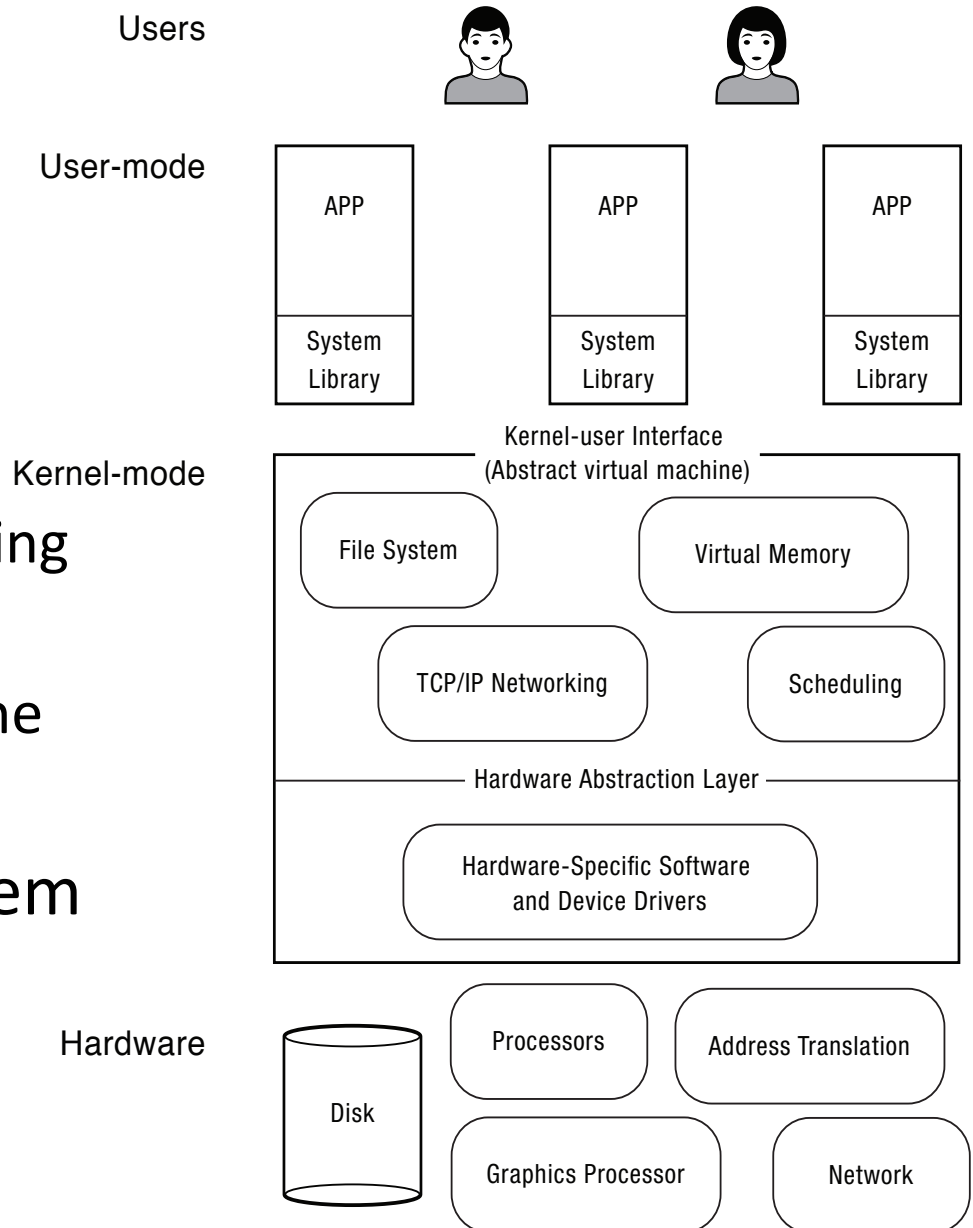
- Portability

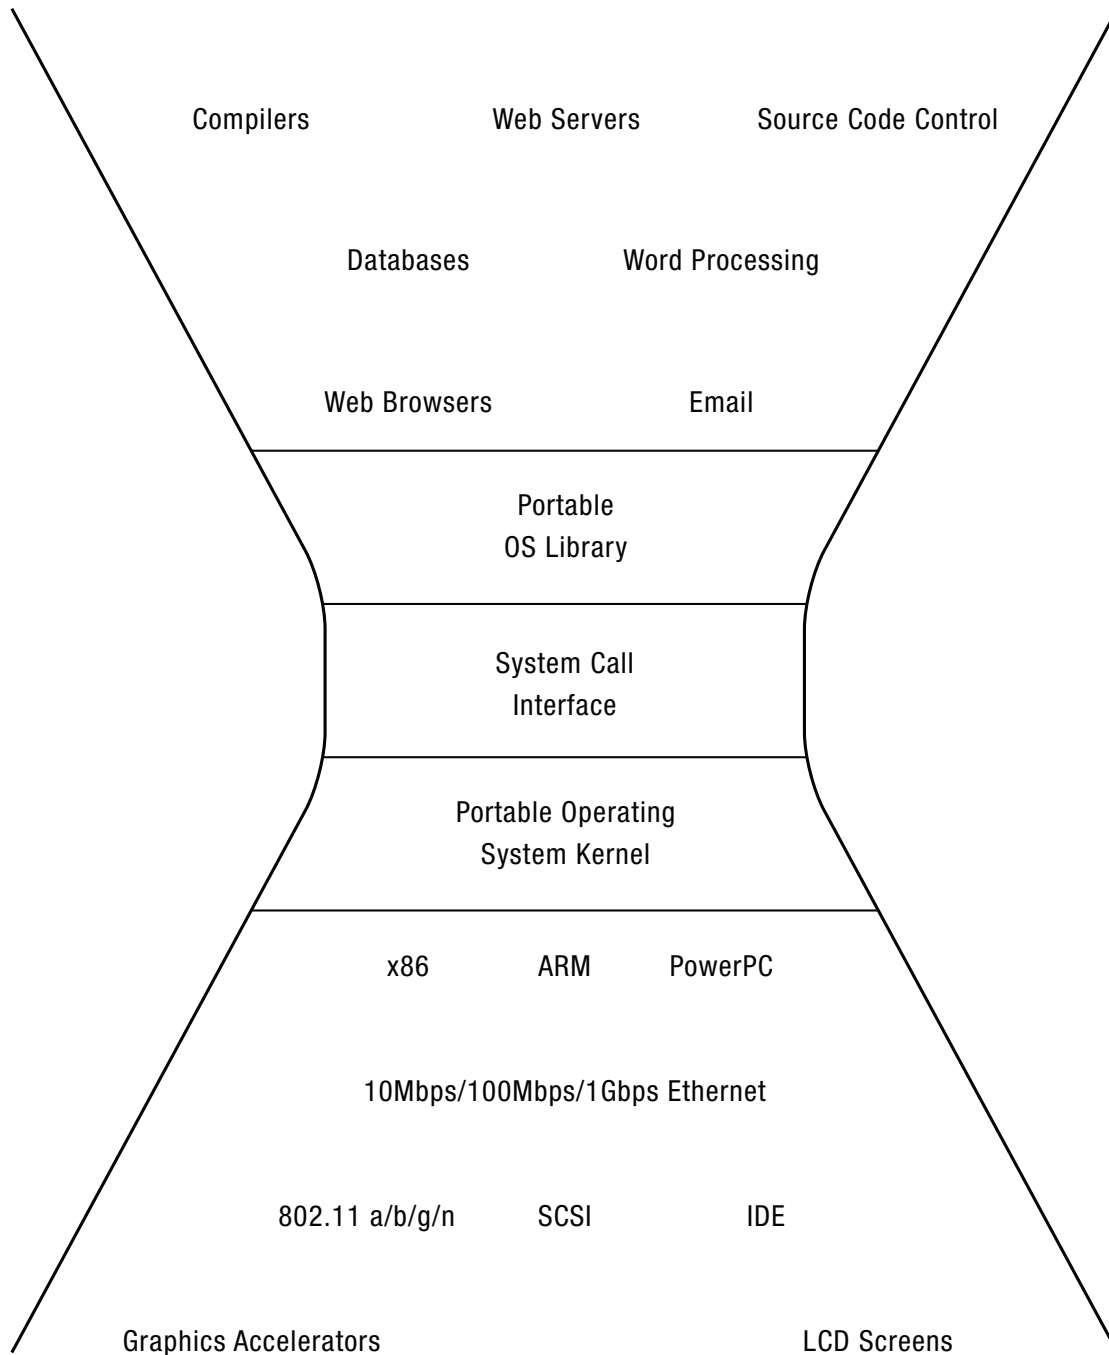
- For programs:

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

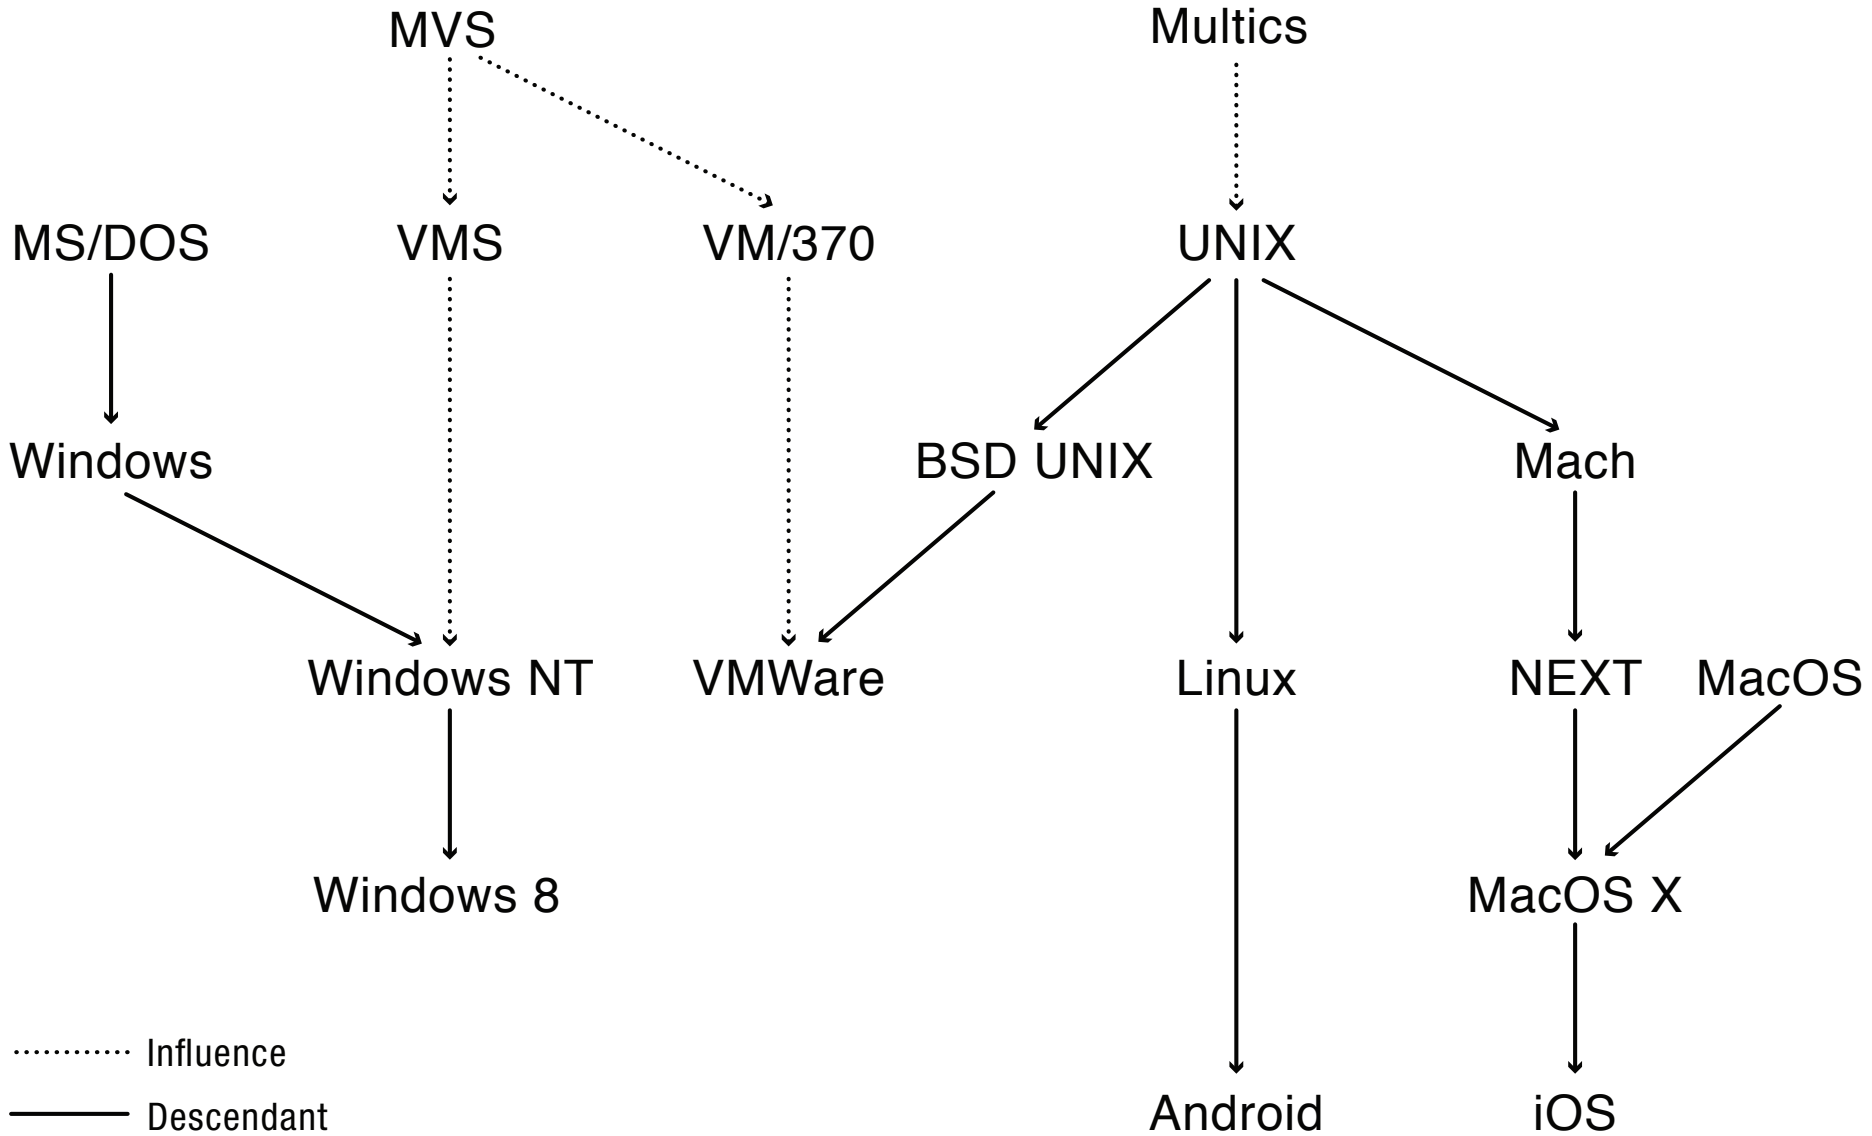
- For the operating system

- Hardware abstraction layer





OS History



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+

Early Operating Systems: Computers Very Expensive

- 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

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

Today's Operating Systems: Computers Cheap

- Smartphones
- Embedded systems
- Laptops
- Tablets
- Virtual machines
- Data center servers

Device I/O

- OS kernel needs to communicate with physical devices
 - Network, disk, video, USB, keyboard, mouse, ...
- Devices operate asynchronously from the CPU
 - Most have their own microprocessor
 - Example: Apple Watch OS runs laptop keyboard

Device I/O

- How does the OS communicate with the device?
 - I/O devices assigned a range of memory addresses
 - Separate from main DRAM memory
 - To issue commands/read results:
 - Special instructions (e.g., inb/outb)
 - Read/write memory locations

Synchronous I/O

- Polling
 - I/O operations take time (physical limits)
 - OS pokes I/O memory on device to issue request
 - While device is working, kernel polls I/O memory to wait until I/O is done
 - Device completes, stores data in its buffers
 - Kernel copies data from device into memory

Faster I/O: Interrupts

- Interrupts
 - OS pokes I/O memory on device to issue request
 - CPU goes back to work on some other task
 - Device completes, stores data in its buffers
 - Triggers CPU interrupt to signal I/O completion
 - Device specific handler code runs
 - When done, resume previous work

Faster I/O: DMA

- Programmed I/O
 - I/O results stored in the device
 - CPU reads and writes to device memory
 - Each CPU instruction is an uncached read/write (over the I/O bus)
- Direct memory access (DMA)
 - I/O device reads/writes the computer's memory
 - After I/O interrupt, CPU can access results in memory

Faster I/O: Buffer Descriptors

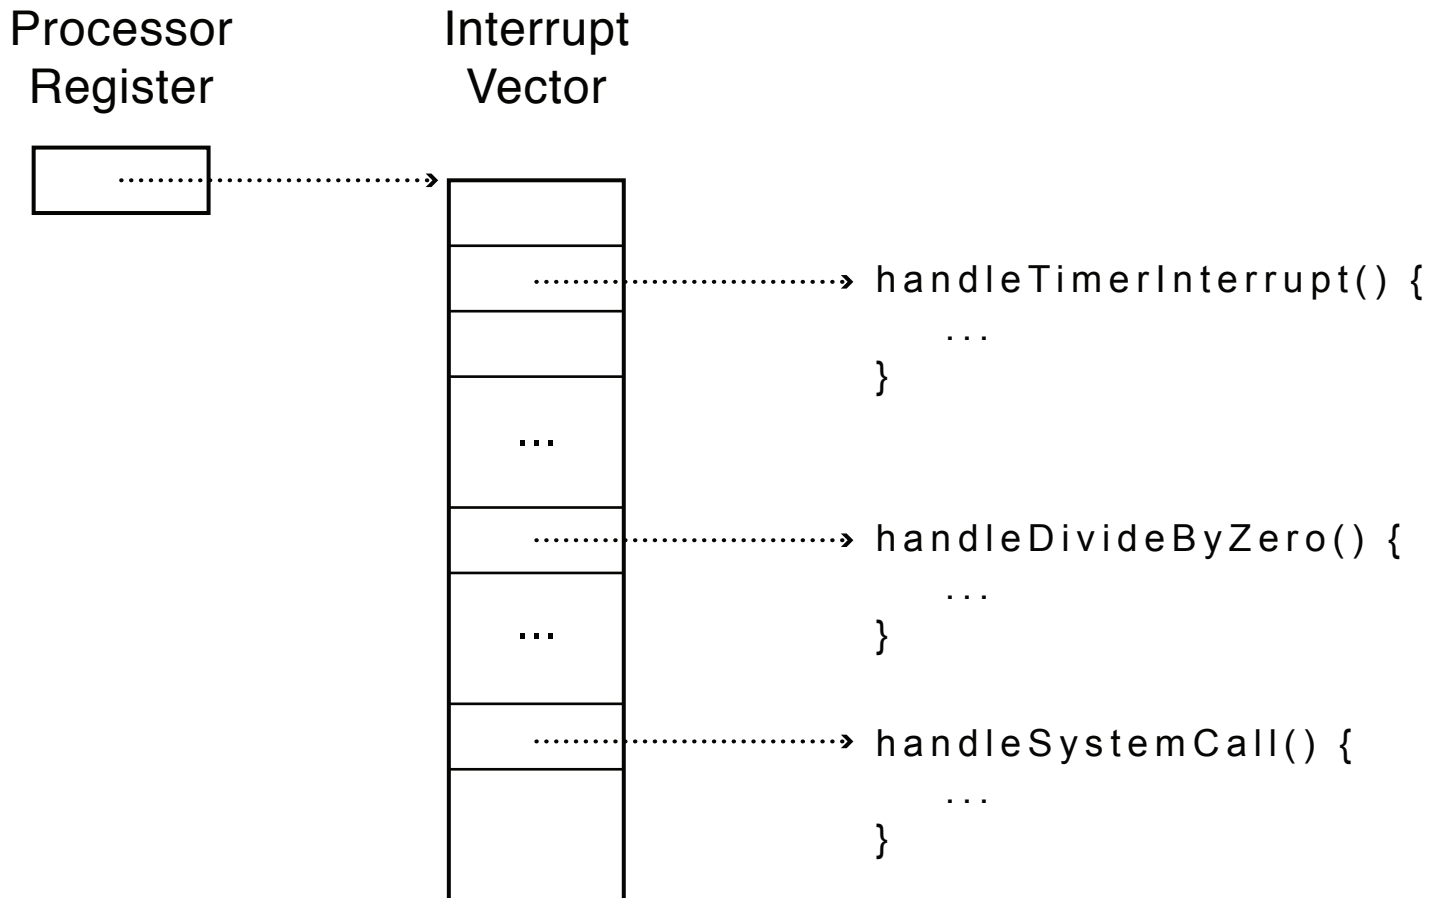
- Buffer descriptor: data structure to specify where to find the I/O request
 - E.g., packet header and packet body
 - Buffer descriptor itself is DMA'ed!
- CPU and device I/O share a queue of buffer descriptors
 - I/O device reads from front
 - CPU fills at tail
- Interrupt only if buffer empties/fills

Device Interrupts

- How do device interrupts work?
 - Where does the CPU run after an interrupt?
 - What is the interrupt handler written in? C? Java?
 - What stack does it use?
 - Is the work the CPU had been doing before the interrupt lost forever?
 - If not, how does the CPU know how to resume that work?

Interrupt Vector

- Table set up by OS kernel; pointers to code to run on different events (in xk, vectors.pl)



Interrupt Masking

- Interrupt handler runs with interrupts off
 - Re-enabled when interrupt completes
- OS kernel can also turn interrupts off
 - Eg., when determining the next process/thread to run
 - On x86
 - CLI: disable interrupts
 - STI: enable interrupts
 - Only applies to the current CPU (on a multicore)
- We'll need this to implement synchronization in chapter 5

Challenge: Saving/Restoring State

- We need to be able to interrupt and transparently resume execution
 - I/O device signals I/O completion
 - Periodic hardware timer to check if app is hung
 - Multiplexing multiple apps on a single CPU
 - Code unaware it has been interrupted!
- Not just the program counter
 - Condition codes, registers used by interrupt handler, ...

Question

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

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?

Textbook

- Lazowska, Spring 2012: “The text is quite sophisticated. You won't get it all on the first pass. The right approach is to [read each chapter before class and] re-read each chapter once we've covered the corresponding material... more of it will make sense then. *Don't save this re-reading until right before the mid-term or final – keep up.*”

Tomorrow's Operating Systems

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