

CS162
Operating Systems and
Systems Programming
Lecture 1

What is an Operating System?

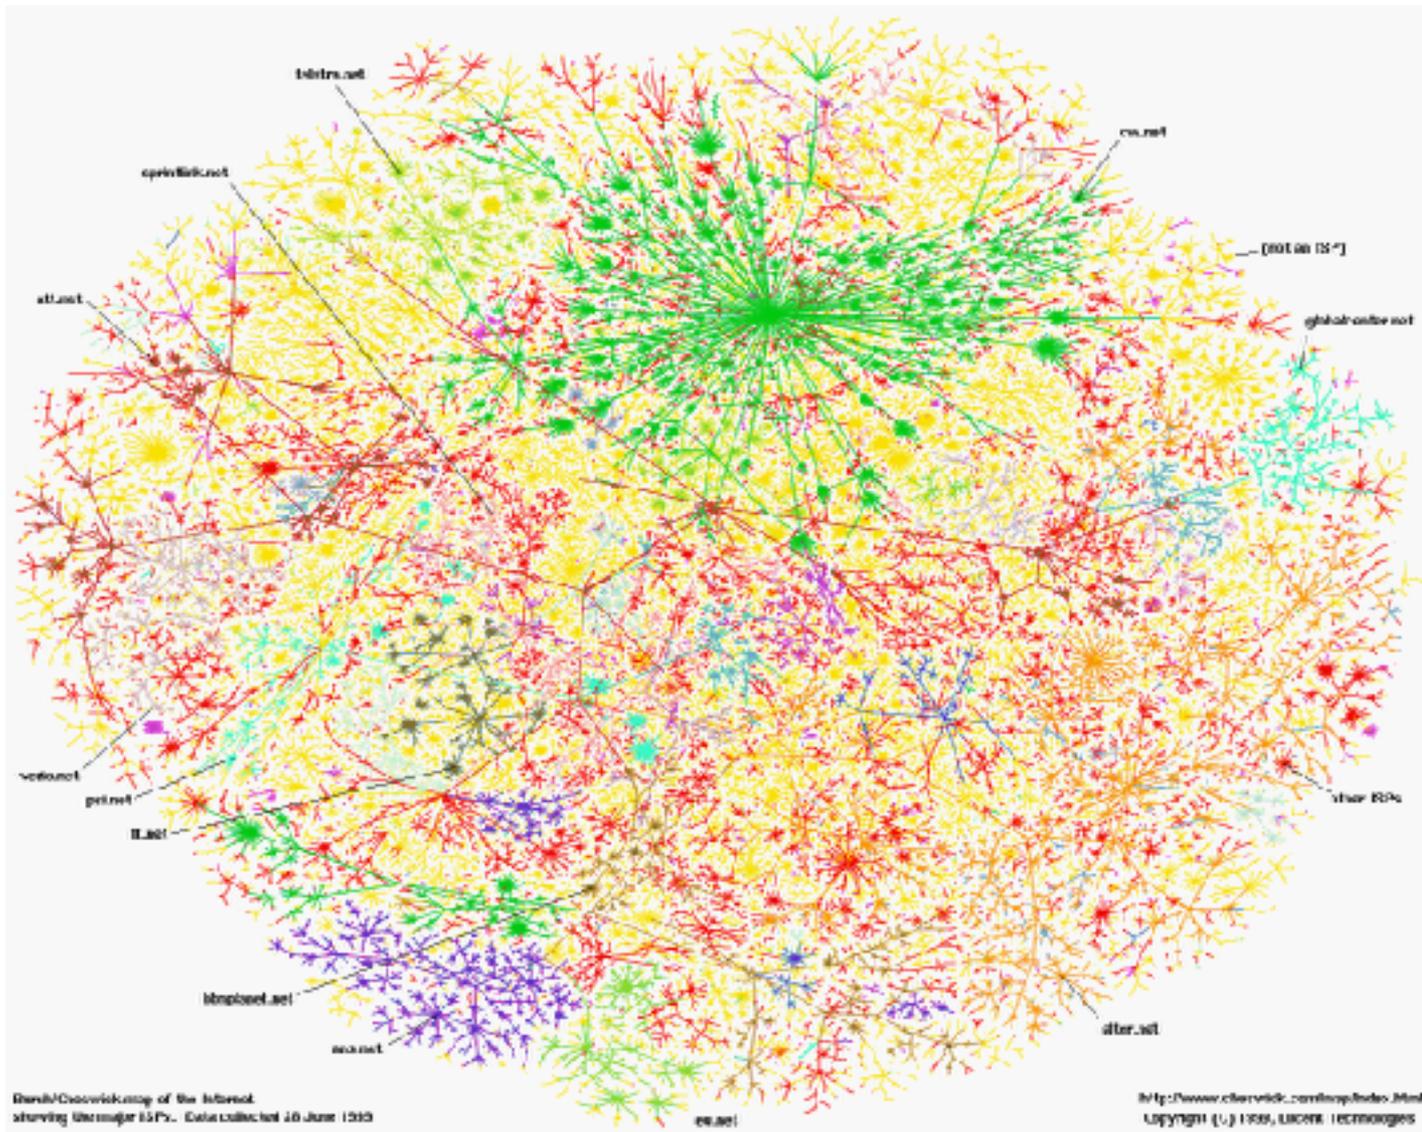
August 26th, 2015

Prof. John Kubiawicz

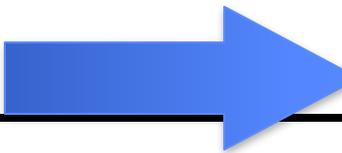
<http://cs162.eecs.Berkeley.edu>

Acknowledgments: Lecture slides are from the Operating Systems course taught by John Kubiawicz at Berkeley, with few minor updates/changes. When slides are obtained from other sources, a reference will be noted on the bottom of that slide, in which case a full list of references is provided on the last slide.

Greatest Artifact of Human Civilization...



3 Billion Internet Users by ...



2.8 B



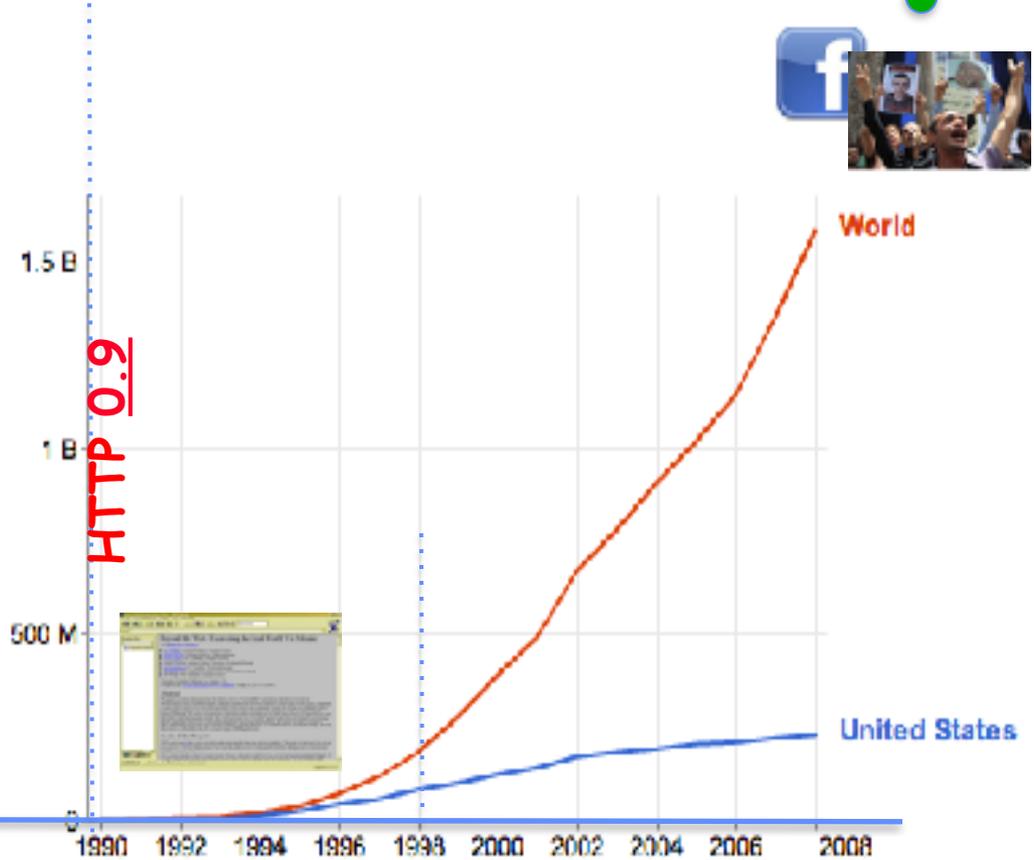
ARPANet

Internet

WWW

2.0 B 1/26/11

RFC 675 TCP/IP



Data source: [World Bank, World Development Indicators](#) - Last updated December 21, 2010

1969 1974

1990

2010

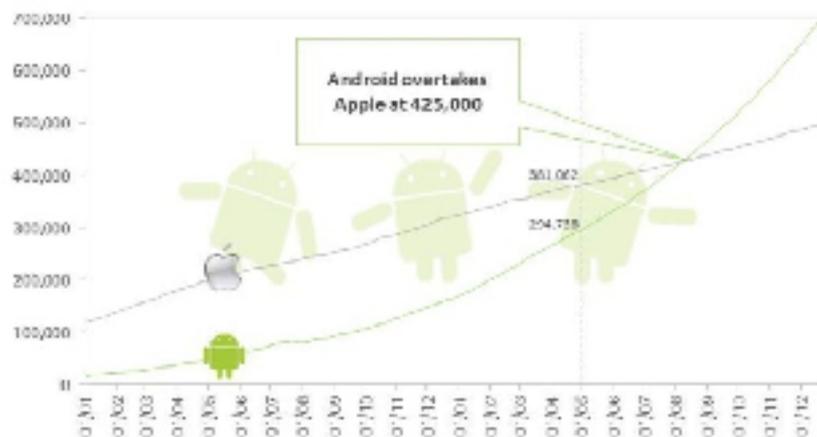
Operating Systems at the heart of it all ...

- Make the incredible advance in the underlying hardware available to a rapid evolving body of applications.
 - Processing, Communications, Storage, Interaction

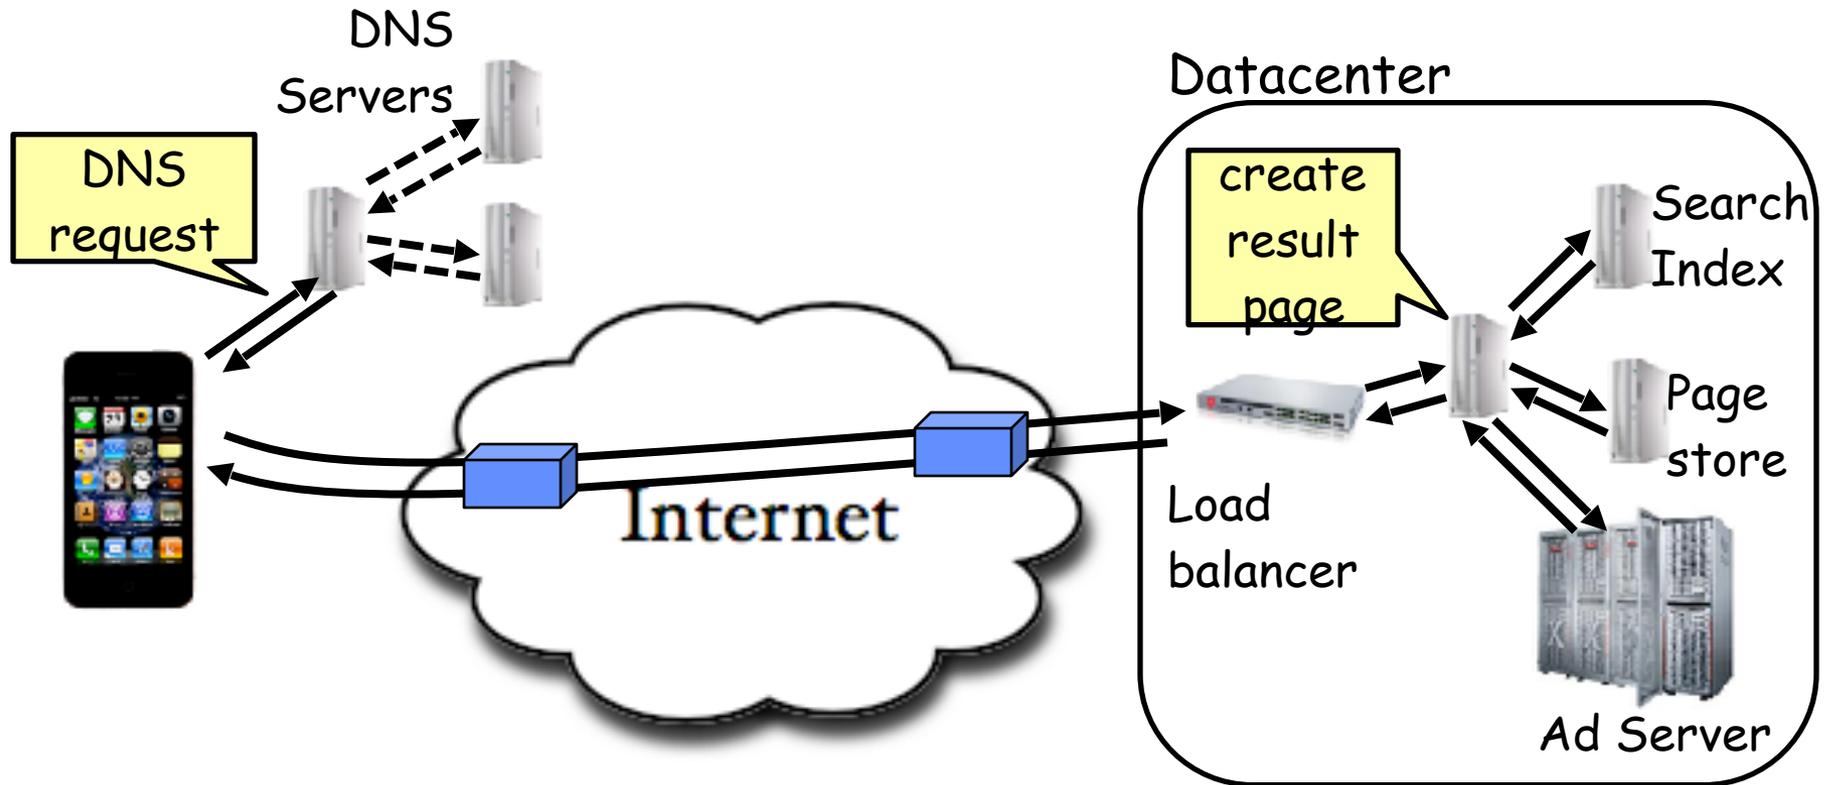
- The key building blocks

- Scheduling
- Concurrency
- Address spaces
- Protection, Isolation, Security
- Networking, distributed systems
- Persistent storage, transactions, consistency, resilience
- Interfaces to all devices

Number of apps in Apple App Store and Android Market [01/2010 – 12/2011E]



Example: What's in a Search Query?



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

Why take CE424?

- **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.

Goals for Today

- What is an Operating System?
 - And - what is it not?
- Examples of Operating Systems design
- What makes Operating Systems So Exciting?
- Oh, and “How does this class operate?”

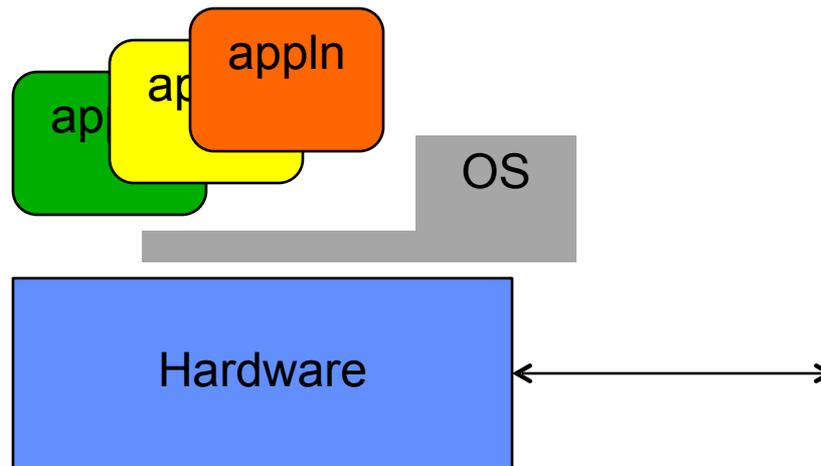
Interactive is important!

Ask Questions!

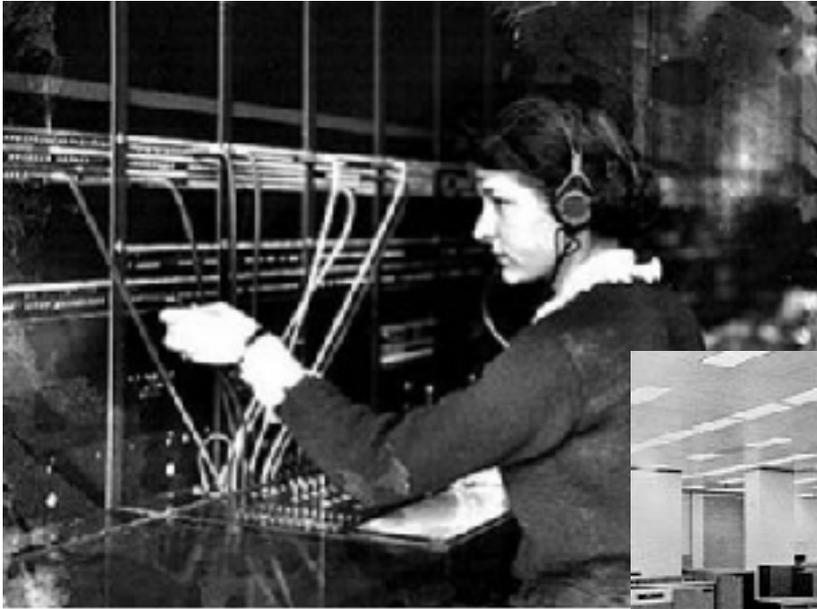
Slides courtesy of David Culler, John Kubiatoicz, AJ Shankar, George Necula, Alex Aiken, Eric Brewer, Ras Bodik, Ion Stoica, Doug Tygar, and David Wagner.

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



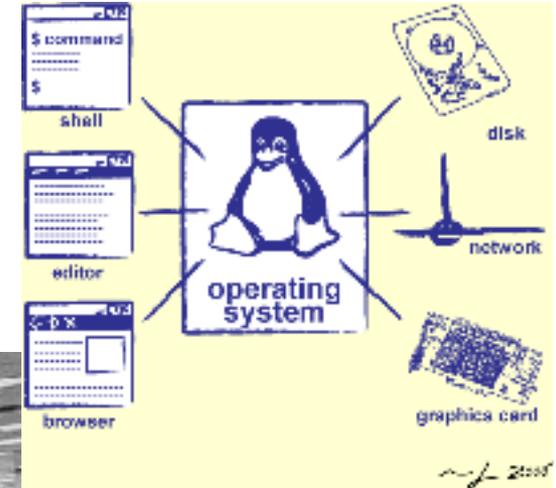
Operator ...



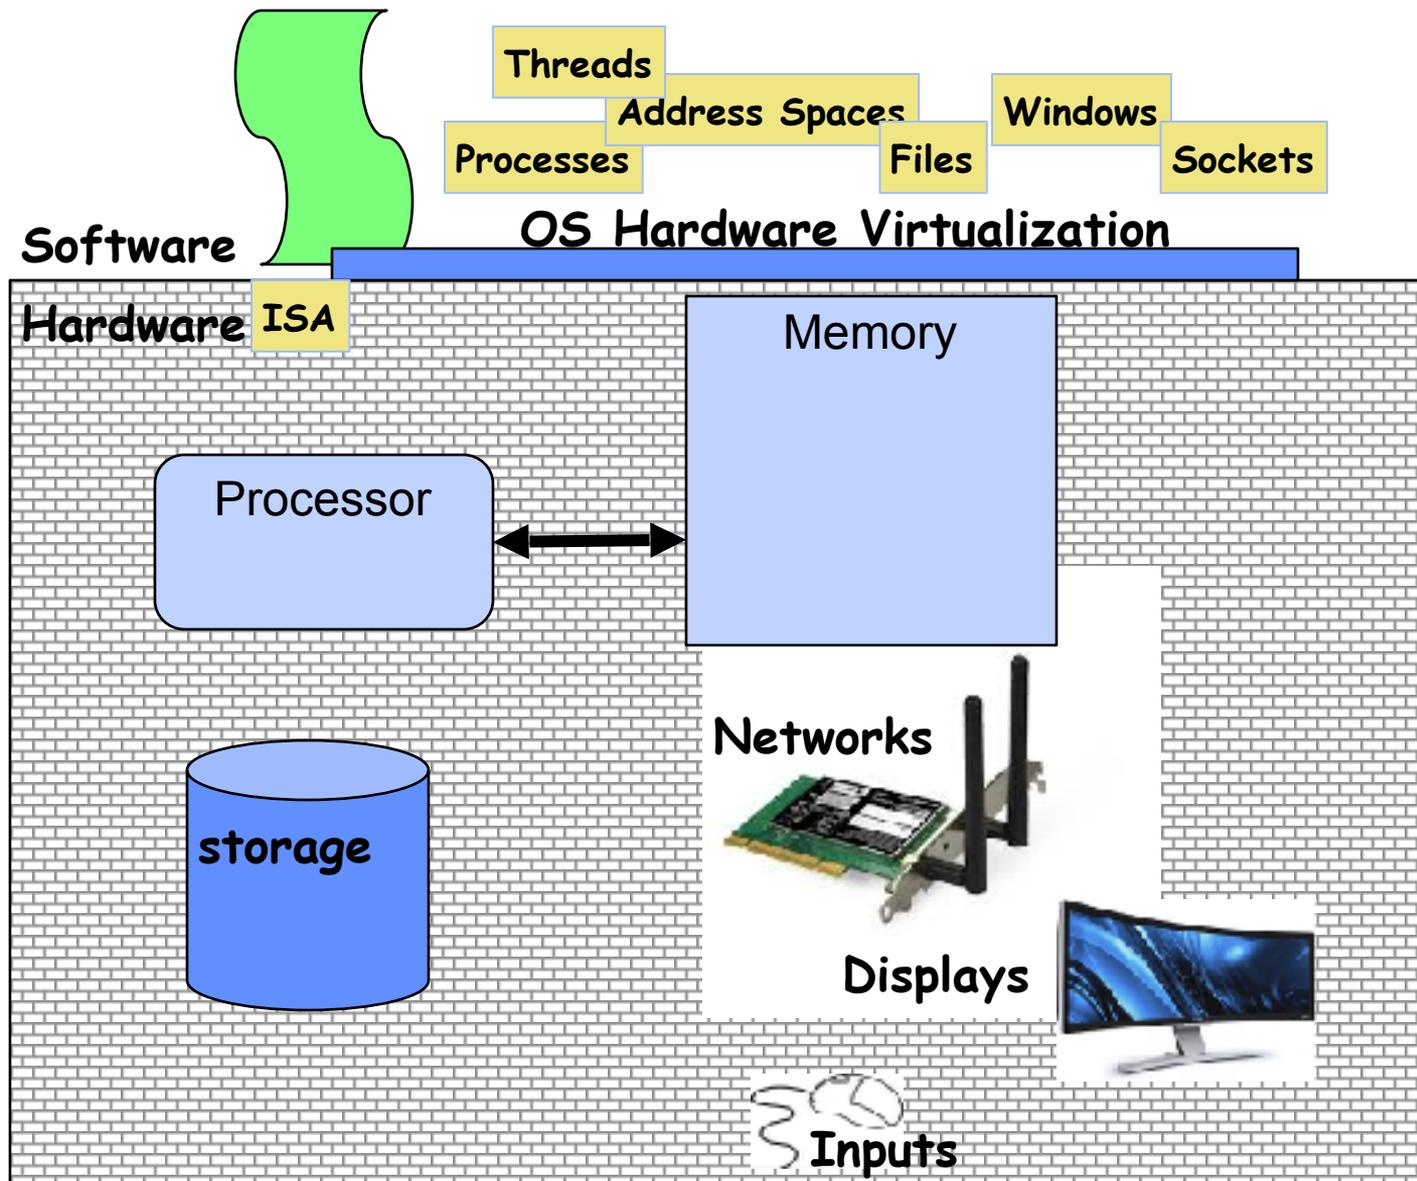
Switchboard Operator



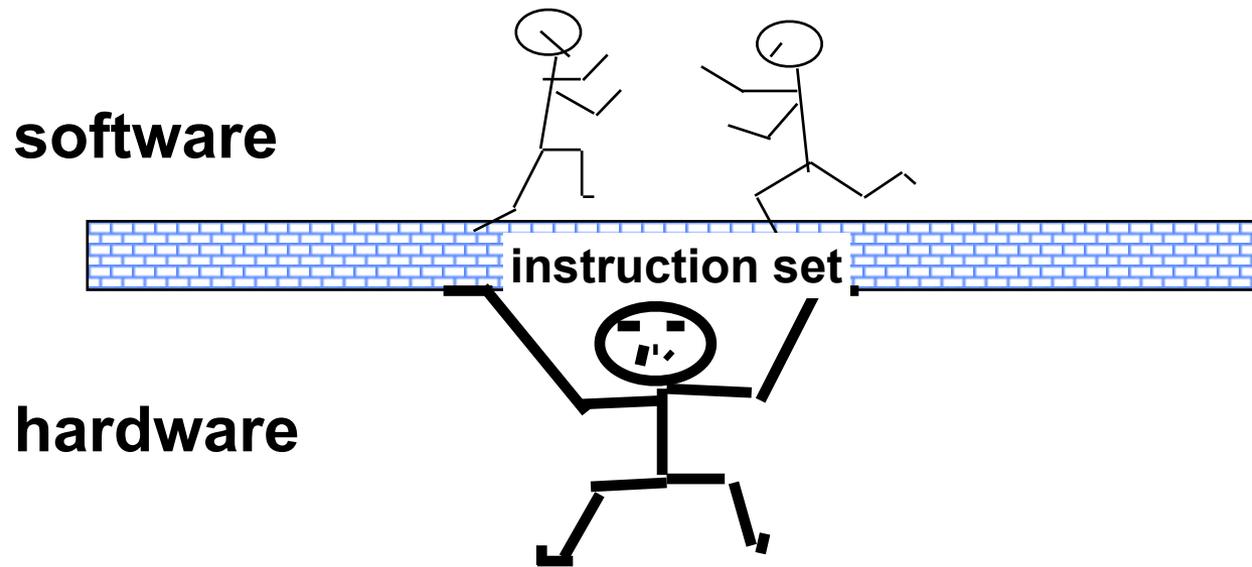
Computer Operators



OS Basics: "Virtual Machine" Boundary

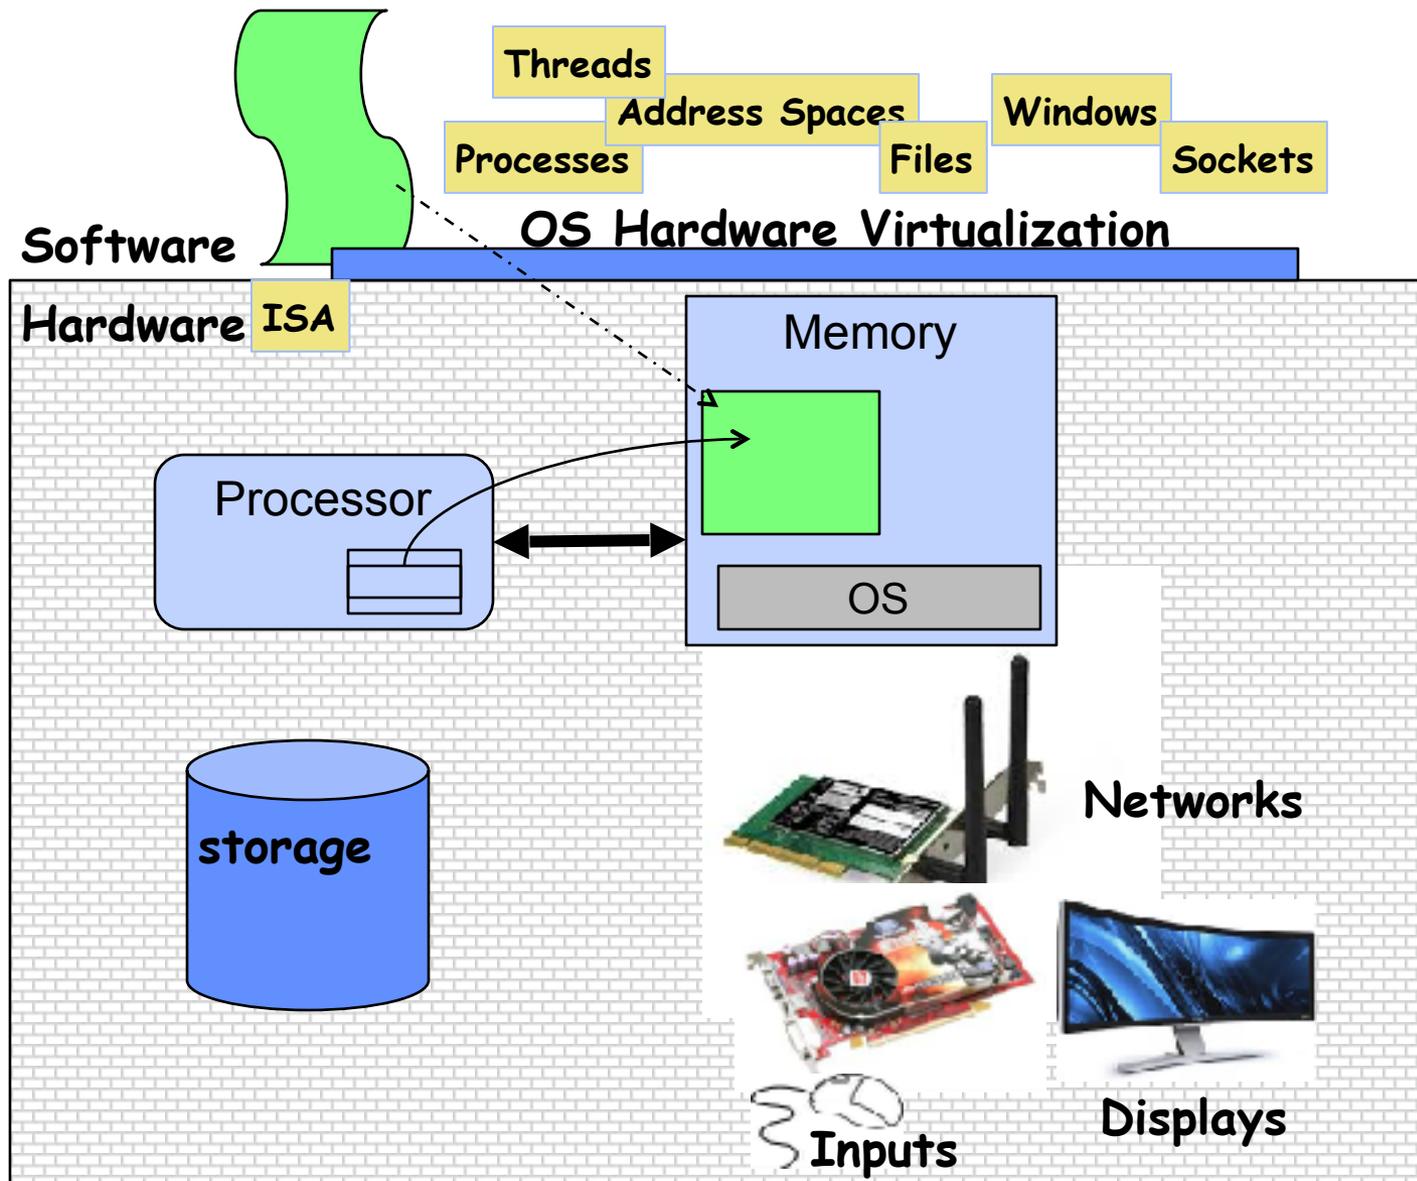


Interfaces Provide Essential Boundaries

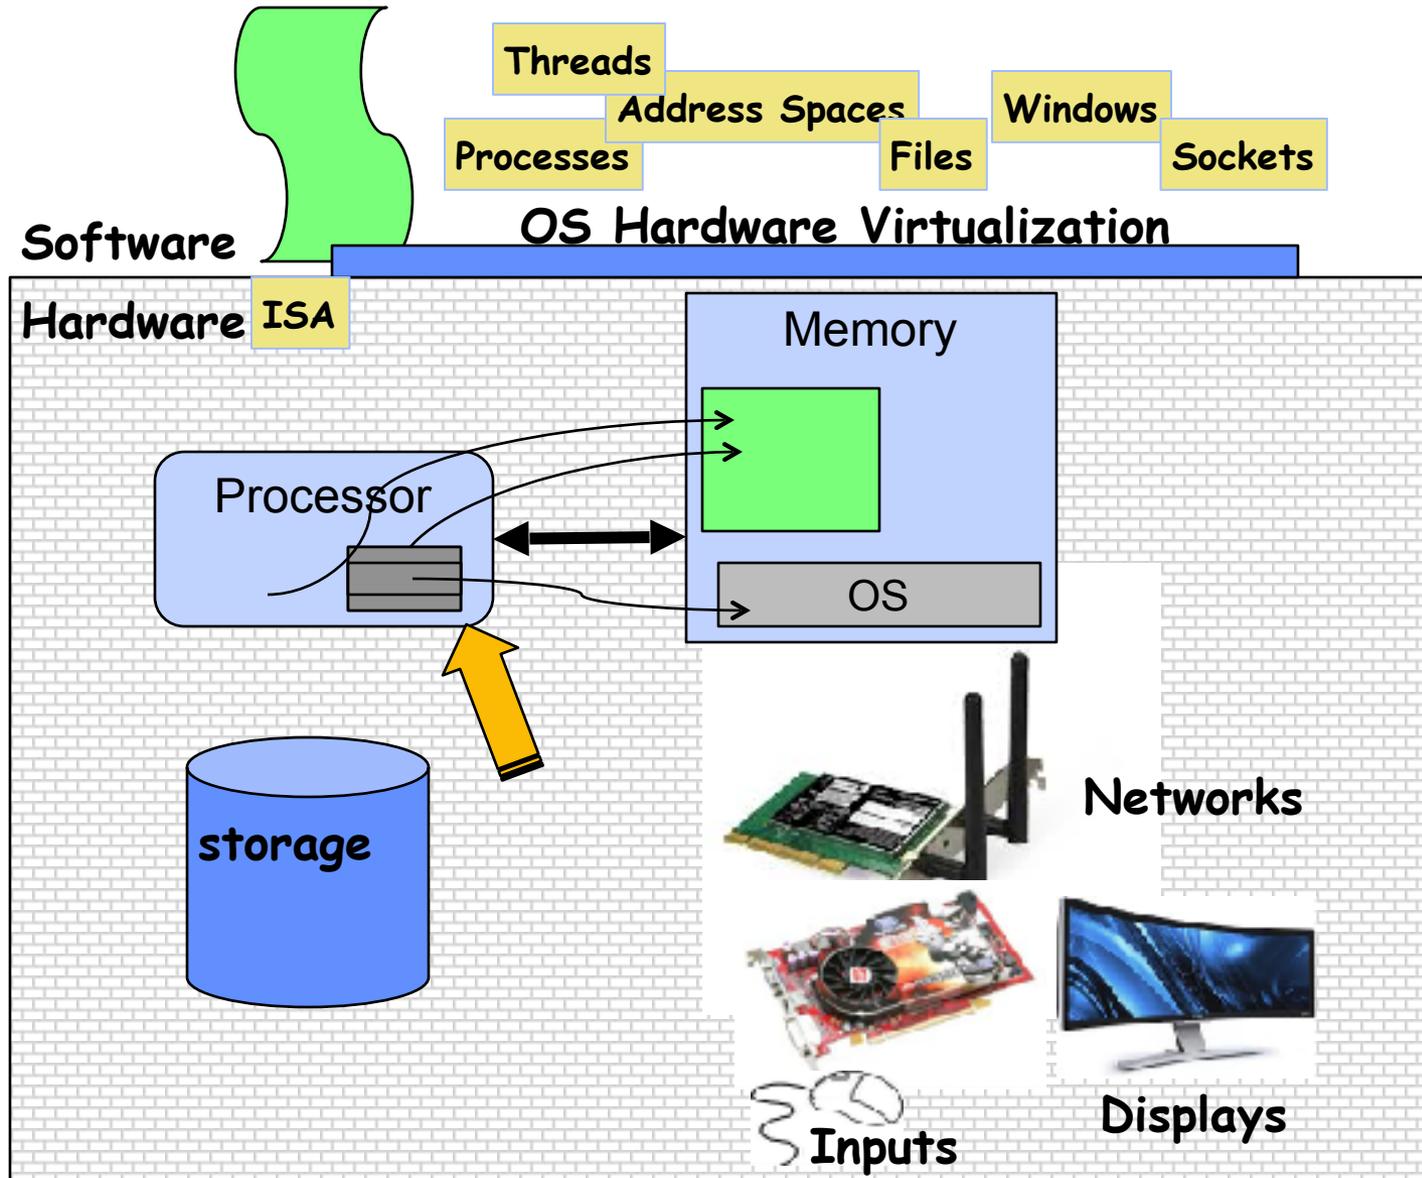


- Why do interfaces look the way that they do?
 - History, Functionality, Stupidity, Bugs, Management
 - CE323 \Rightarrow Machine interface
 - CE227 \Rightarrow Human interface
 - CE418 \Rightarrow Software engineering/management

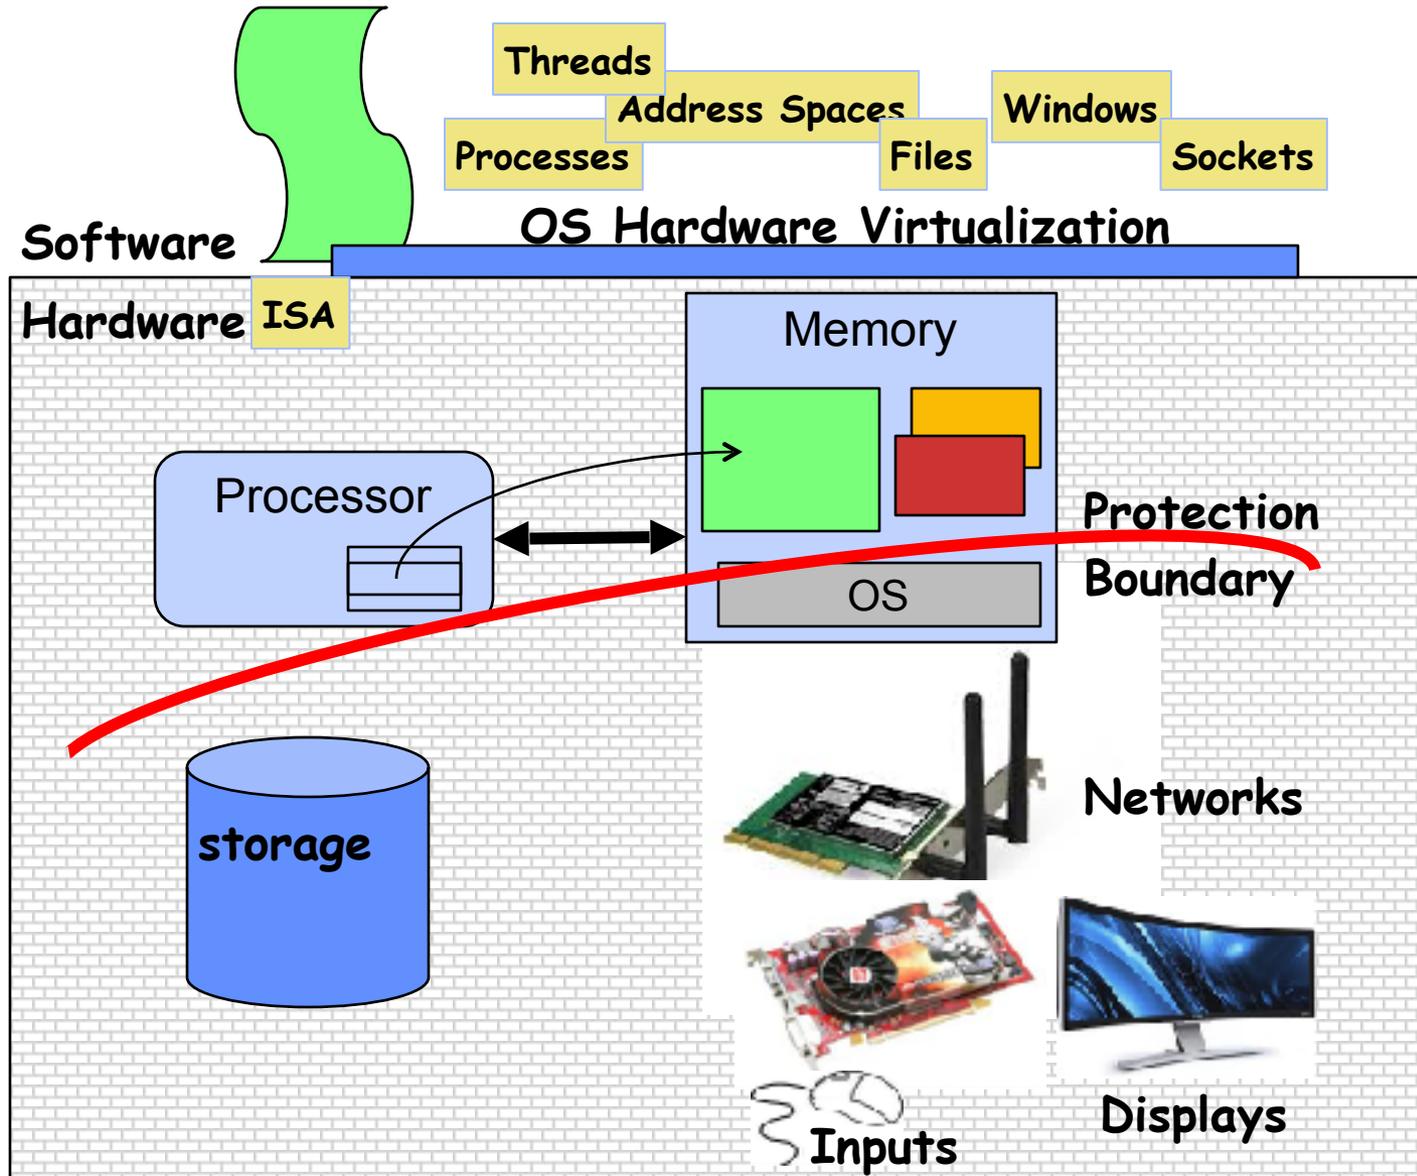
OS Basics: Program => Process



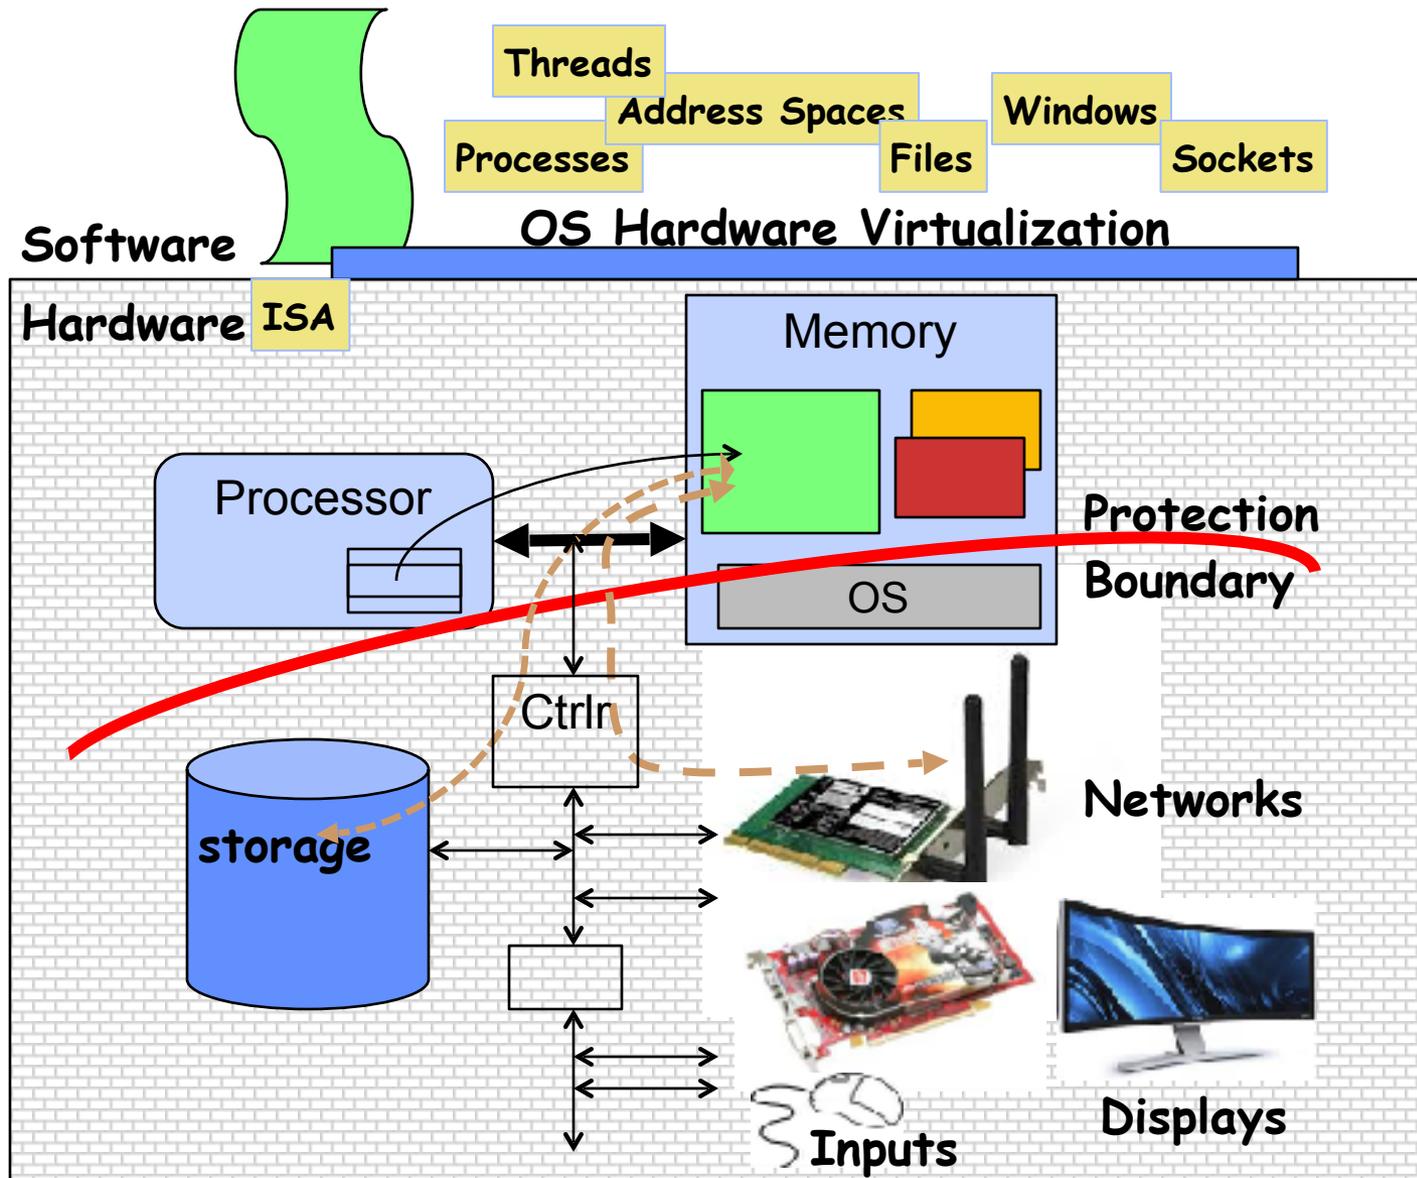
OS Basics: Context Switch



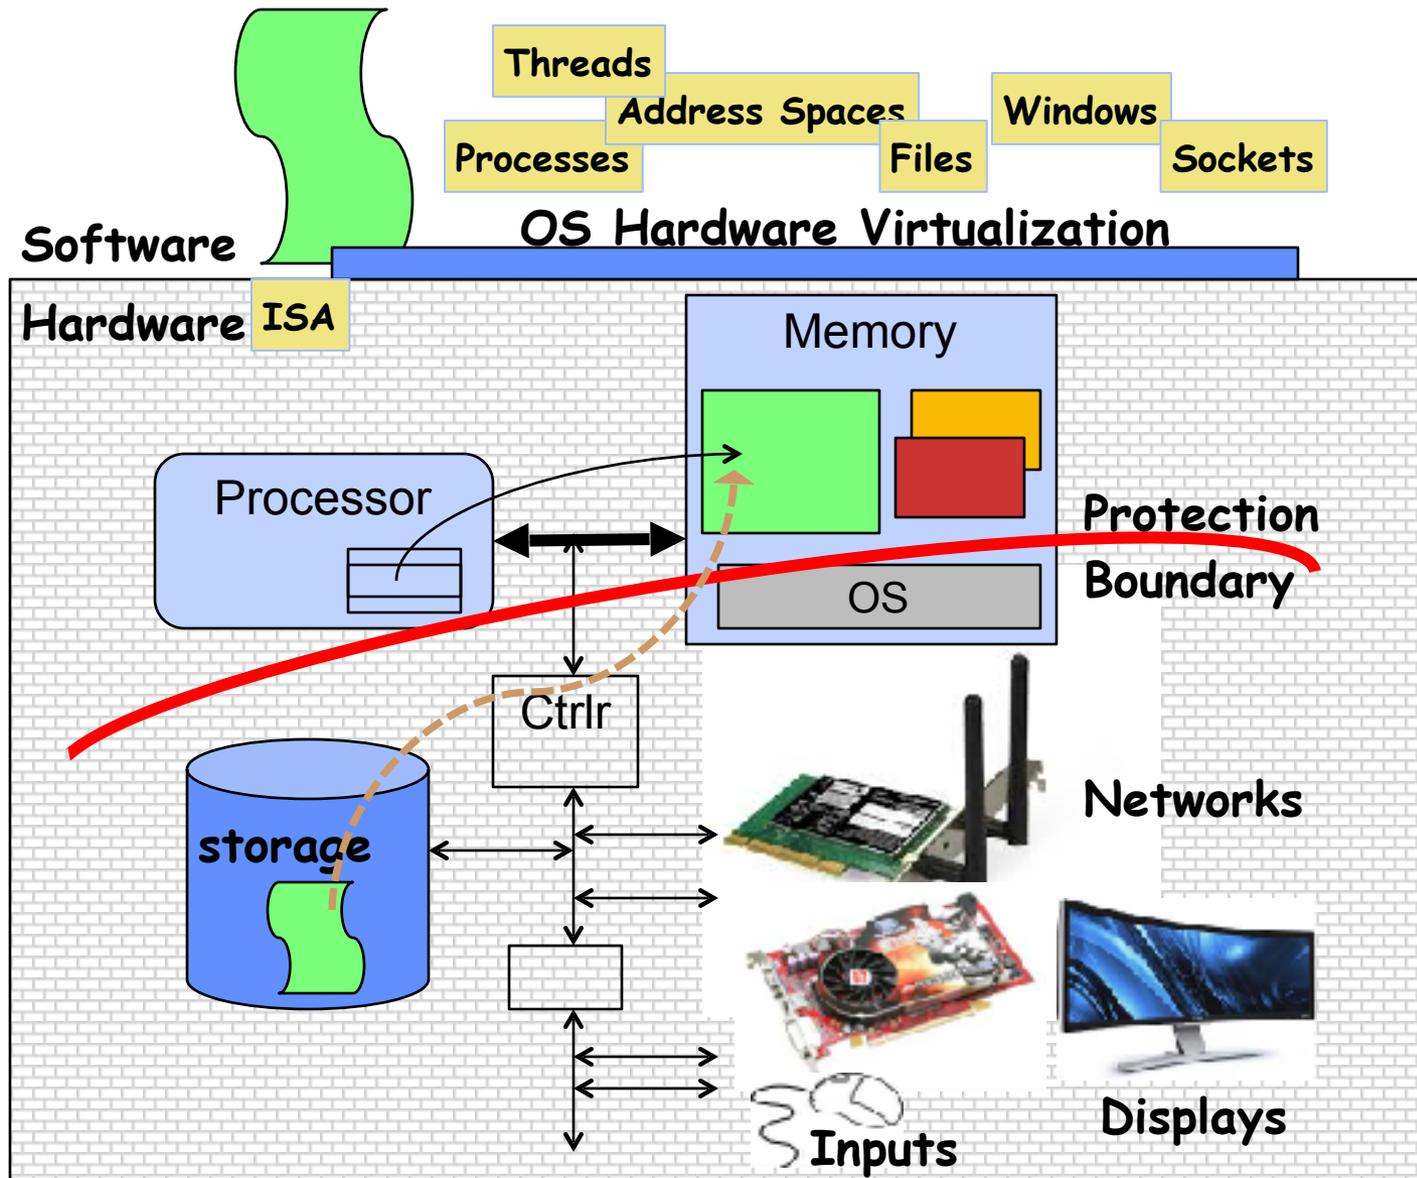
OS Basics: Scheduling, Protection



OS Basics: I/O

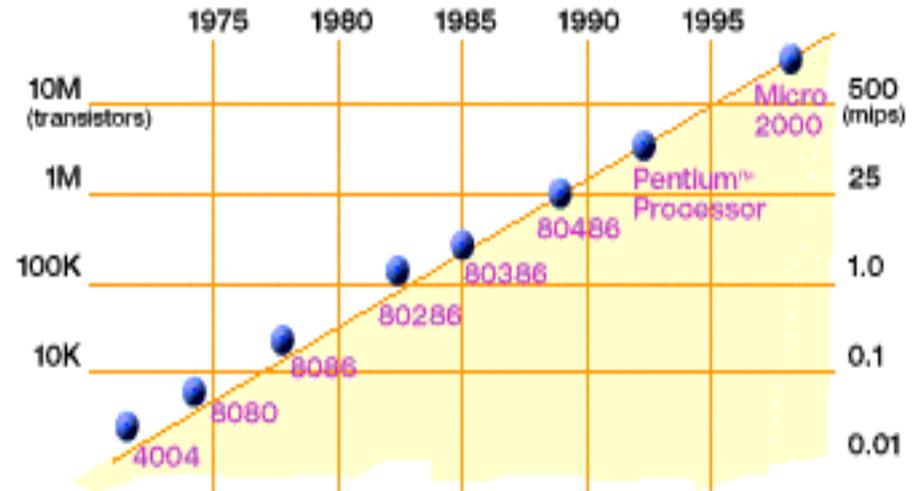
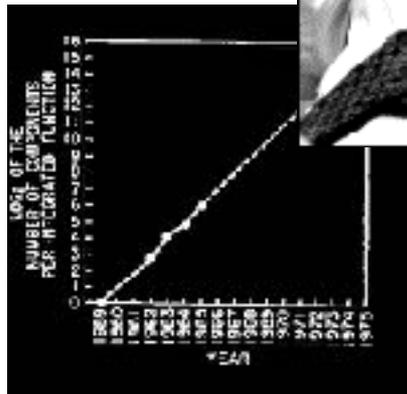
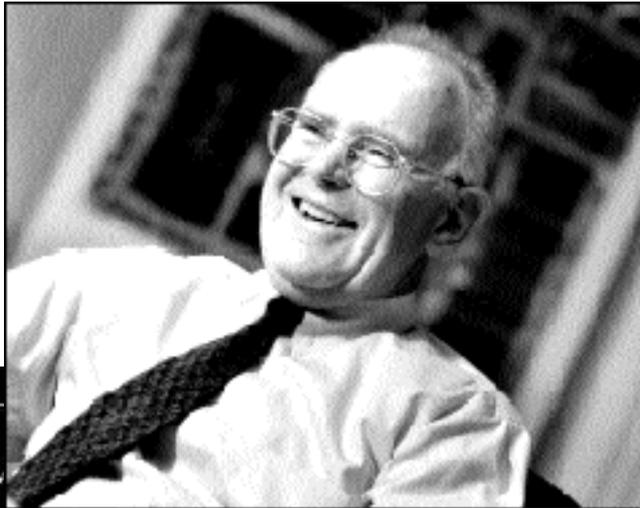


OS Basics: Loading



What make Operating Systems exciting and Challenging

Technology Trends: Moore's Law

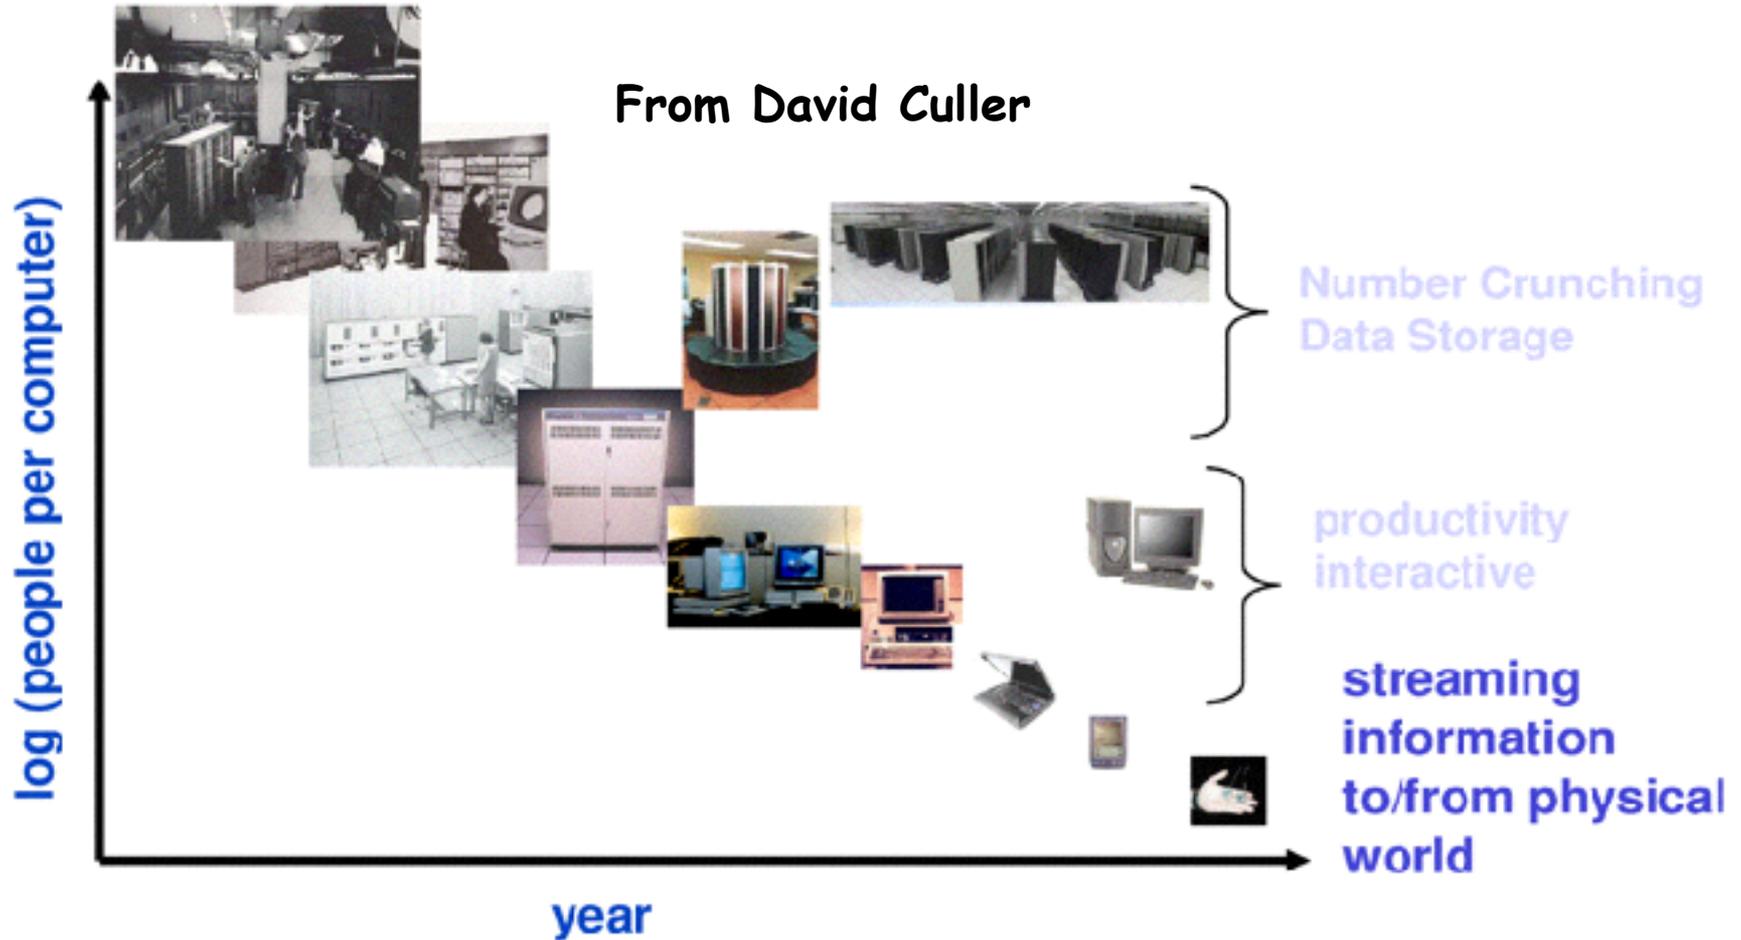


2X transistors/Chip Every 1.5 years
Called "**Moore's Law**"

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

Microprocessors have become smaller, denser, and more powerful.

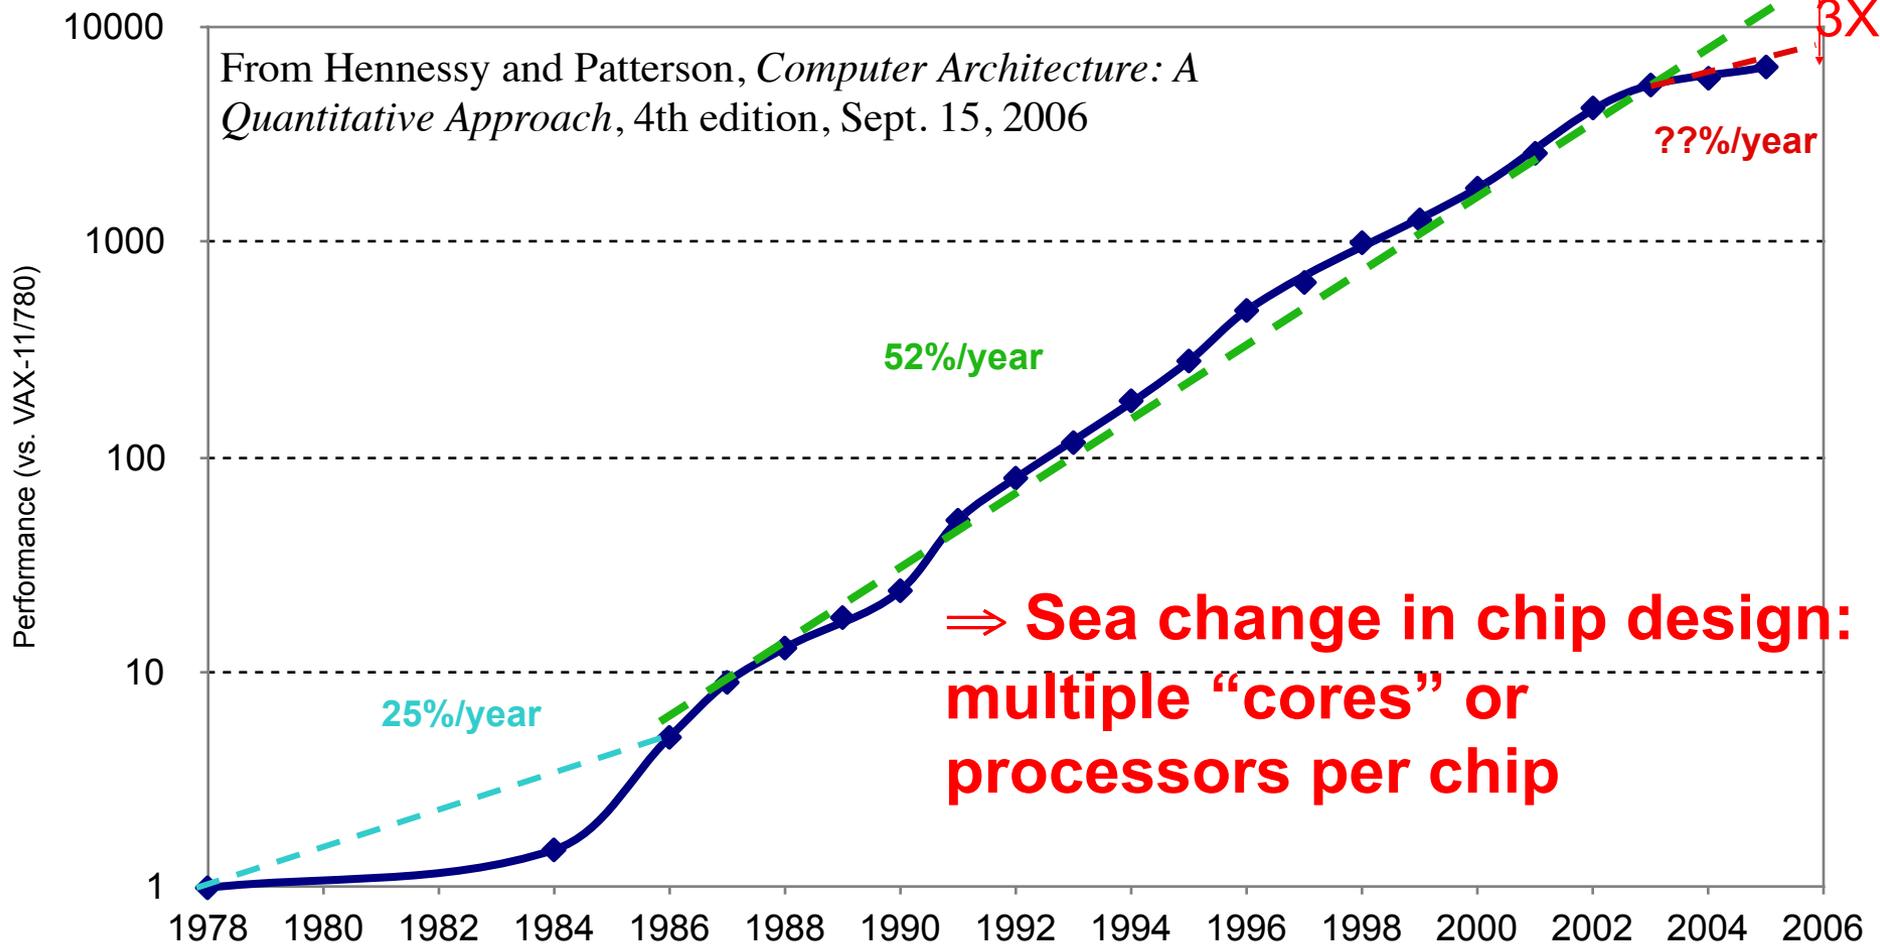
People-to-Computer Ratio Over Time



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

New Challenge: Slowdown in Joy's law of Performance

From Hennessy and Patterson, *Computer Architecture: A Quantitative Approach*, 4th edition, Sept. 15, 2006



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

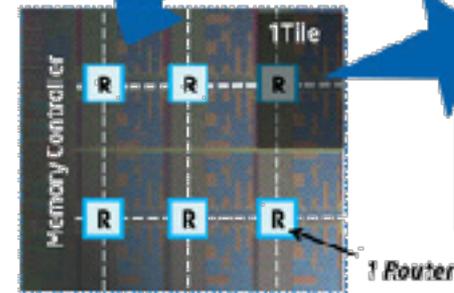
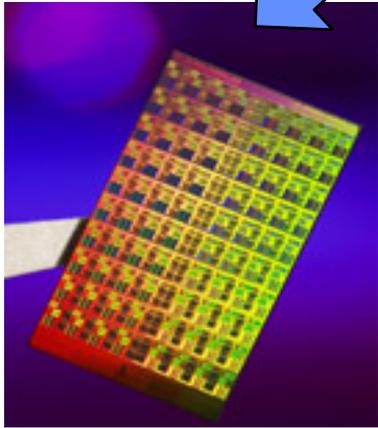
ManyCore Chips: The future is here

- Intel 80-core multicore chip (Feb 2007)

- 80 simple cores
- Two FP-engines / core
- Mesh-like network
- 100 million transistors
- 65nm feature size

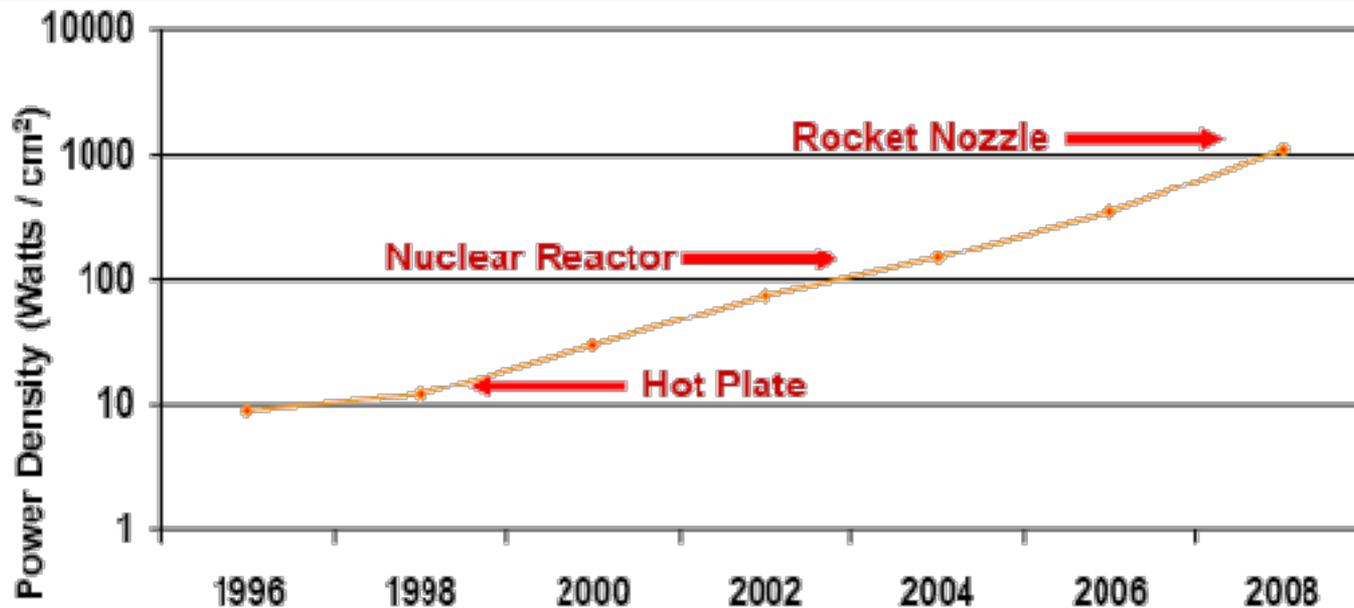
- Intel Single-Chip Cloud Computer (August 2010)

- 24 "tiles" with two cores/tile
- 24-router mesh network
- 4 DDR3 memory controllers
- Hardware support for message-passing



- "ManyCore" refers to many processors/chip
 - 64? 128? Hard to say exact boundary
- 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**

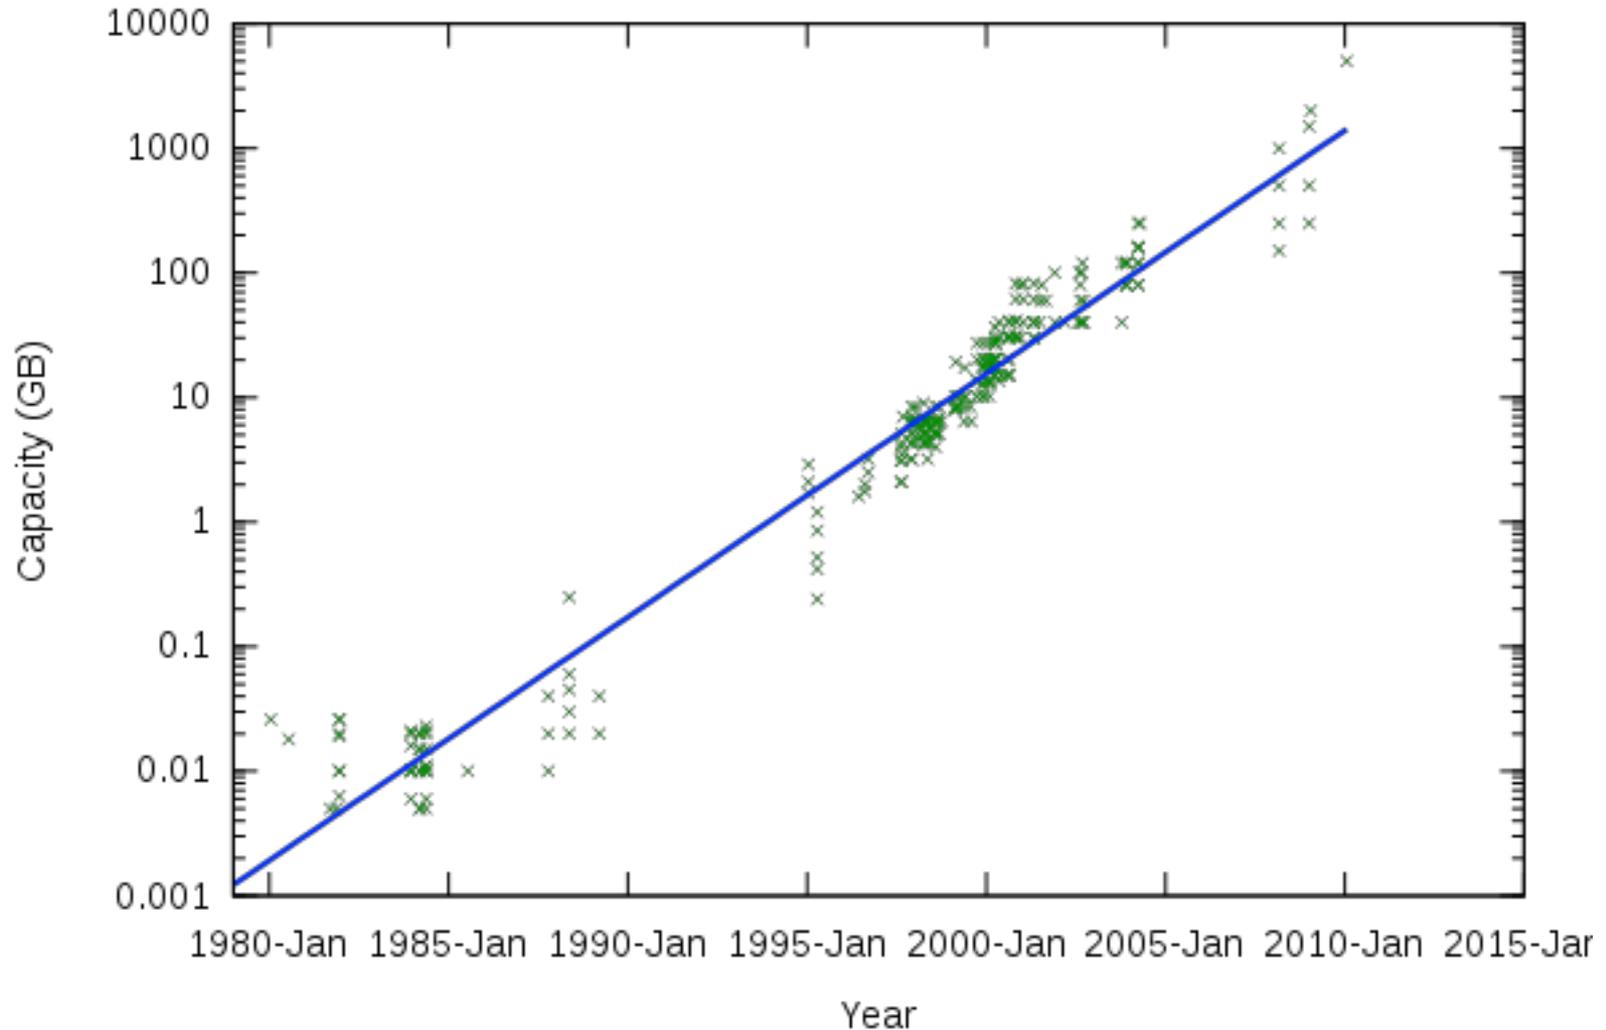
Another Challenge: Power Density



Power Density Becomes Too High to Cool Chips Inexpensively

- **Moore's Law Extrapolation**
 - Potential power density reaching amazing levels!
- **Flip side: Battery life very important**

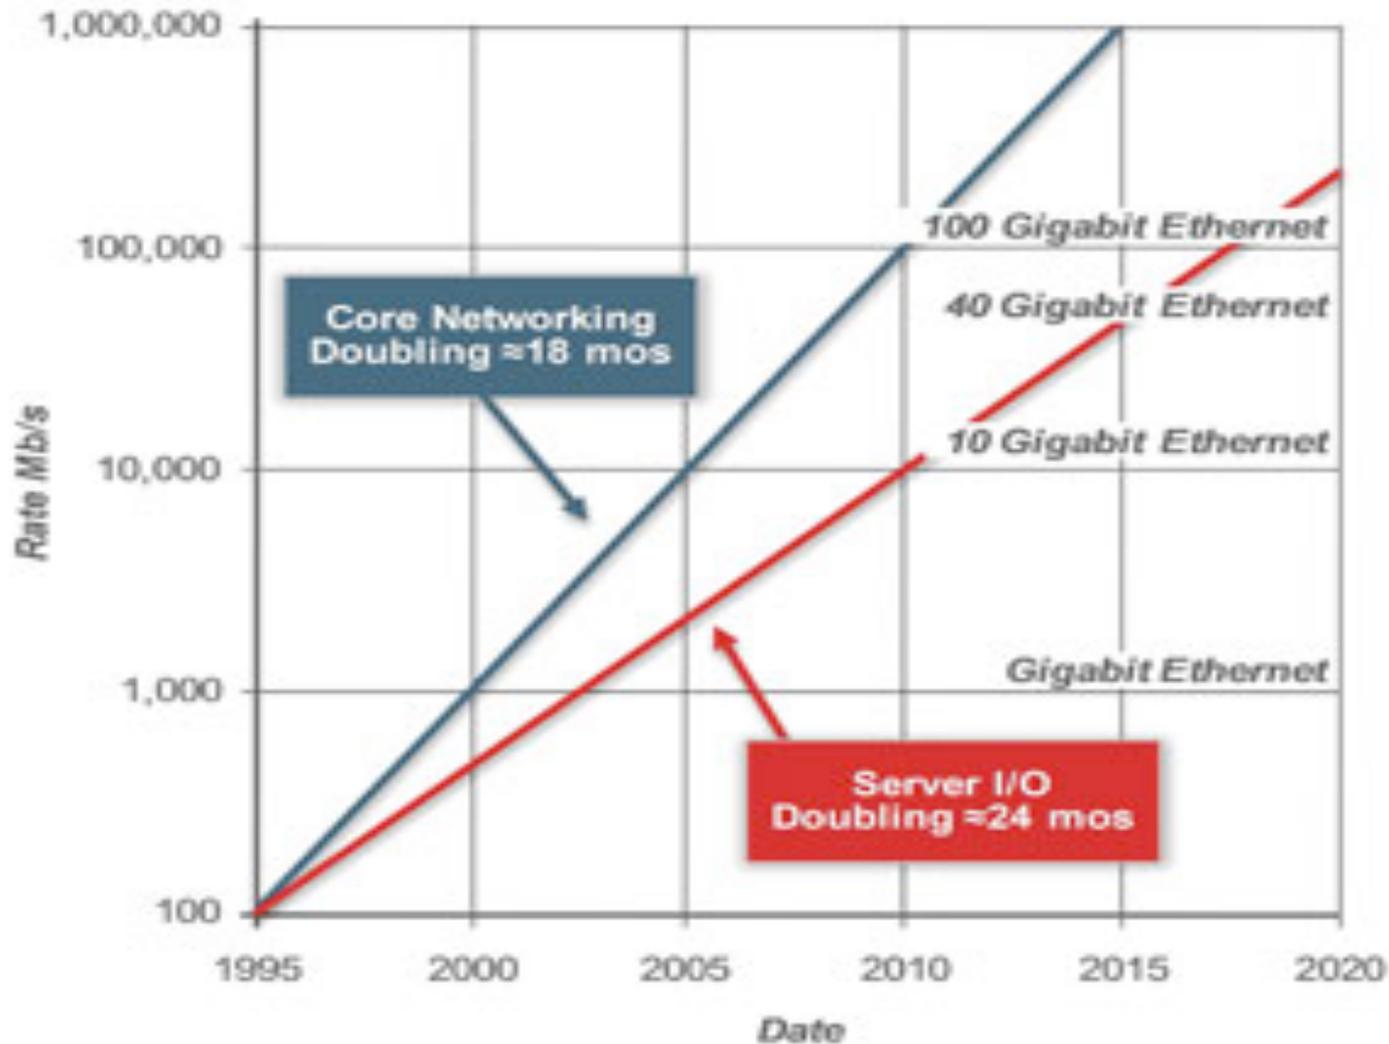
Storage Capacity



- Retail hard disk capacity in GB

(source: <http://www.digitaltonto.com/2011/our-emergent-digital-future/>)

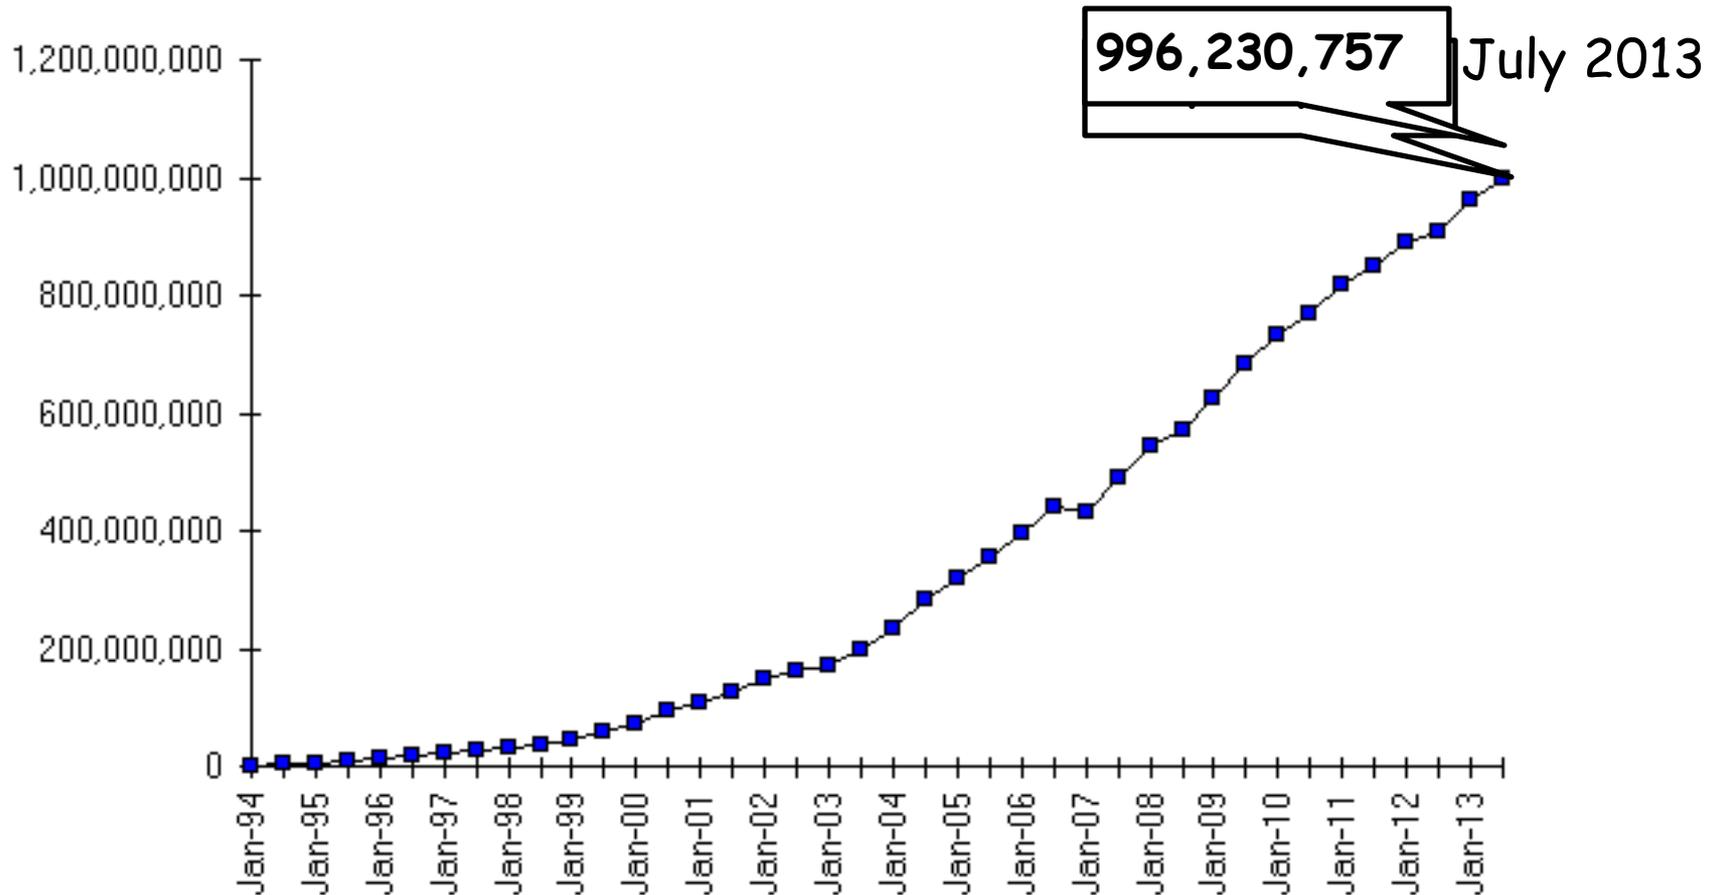
Network Capacity



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

Internet Scale: .96 Billion Hosts

Internet Domain Survey Host Count



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

<https://www.isc.org/solutions/survey>

Internet Scale: 2.8 Billion Users!

WORLD INTERNET USAGE AND POPULATION STATISTICS December 31, 2013

World Regions	Population (2014 Est.)	Internet Users Dec. 31, 2000	Internet Users Latest Data	Penetration (% Population)	Growth 2000-2014	Users % of Table
Africa	1,125,721,038	4,514,400	240,146,482	21.3 %	5,219.6 %	8.6 %
Asia	3,996,408,007	114,304,000	1,265,143,702	31.7 %	1,006.8 %	45.1 %
Europe	825,802,657	105,096,093	566,261,317	68.6 %	438.8 %	20.2 %
Middle East	231,062,860	3,284,800	103,829,614	44.9 %	3,060.9 %	3.7 %
North America	353,860,227	108,096,800	300,287,577	84.9 %	177.8 %	10.7 %
Latin America / Caribbean	612,279,181	18,068,919	302,006,016	49.3 %	1,571.4 %	10.8 %
Oceania / Australia	36,724,649	7,620,430	24,804,226	67.5 %	225.5 %	0.9 %
WORLD TOTAL	7,181,858,619	360,985,492	2,802,478,934	39.0 %	676.3 %	100.0 %

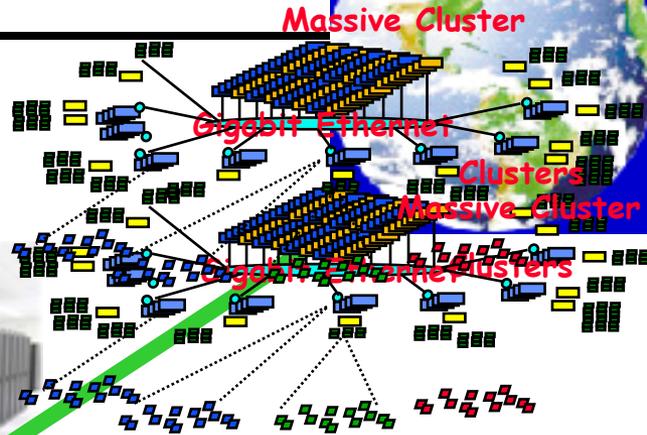
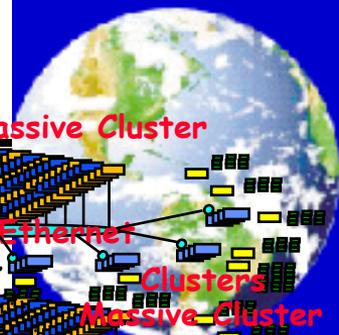
NOTES: (1) Internet Usage and World Population Statistics are for December 31, 2013. (2) CLICK on each world region name for detailed regional usage information. (3) Demographic (Population) numbers are based on data from the [US Census Bureau](#) and local census agencies. (4) Internet usage information comes from data published by [Nielsen Online](#), by the [International Telecommunications Union](#), by [GfK](#), local ICT Regulators and other reliable sources. (5) For definitions, disclaimers, navigation help and methodology, please refer to the [Site Surfing Guide](#). (6) Information in this site may be cited, giving the due credit to www.internetworldstats.com. Copyright © 2001 - 2014, Miriwatts Marketing Group. All rights reserved worldwide.

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

Not Only PCs connected to the Internet

- Smartphone shipments now exceed PC shipments!
- 2011 shipments:
 - 487M smartphones
 - 414M PC clients
 - » 210M notebooks
 - » 112M desktops
 - » 63M tablets
 - 25M smart TVs
- 4 billion phones in the world → smartphone over next decade





• The world is a large distributed system

- Microprocessors in everything
- Vast infrastructure behind them

Internet
Connectivity

Scalable, Reliable,
Secure Services

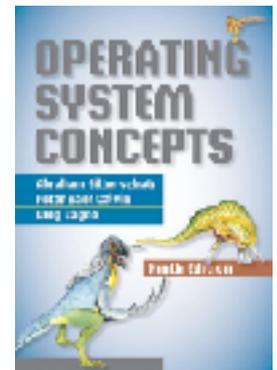
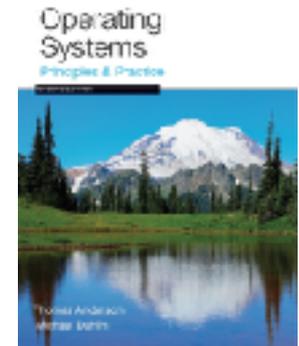


Databases
Information Collection
Remote Storage
Online Games
Commerce
...

MEMS for
Sensor Nets

Infrastructure, Textbook & Readings

- Infrastructure
 - Website: <http://sharif.edu/~kharrazi/courses/40424-961>
 - Mailing list
- Textbook: Operating Systems: Principles and Practice (2nd Edition) Anderson and Dahlin
- Recommend: Operating Systems Concepts, 9th Edition Silberschatz, Galvin, Gagne
- Online supplements
 - See course website
 - Networking, Databases, Software Eng, Security



Syllabus

- **OS Concepts: How to Navigate as a Systems Programmer!**
 - Process, I/O, Networks and VM
- **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 (8)**
 - Protocols, N-Tiers, RPC, NFS, DHTs, Consistency, Scalability, multicast

Learning by Doing

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

Getting started

- Start homework 0 immediately
 - Github account
 - Vagrant virtualbox - VM environment for the course
 - » Consistent, managed environment on your machine
 - Get familiar with all the tools
- Waitlist ???
 - If you are not serious about taking the course, please drop the course now

Group Project Simulates Industrial Environment

- Project teams have 4 members (try really hard to get 4 members - 3 members requires serious justification)
 - Must work in groups in "the real world"
- 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!**

Grading

- 15% Midterms
- 25% Final
- 25% Homework
- 40% Group HWs
- Group HWs grading
 - [10 pts] Initial design
 - [10 pts] Design review
 - [10 pts] Design document
 - [60 pts] Code (3 checkpoints)
 - [10 pts] Final design
- Submission via git push to release branch
- Regular git push so TA sees your progress

Ce 424 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 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

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



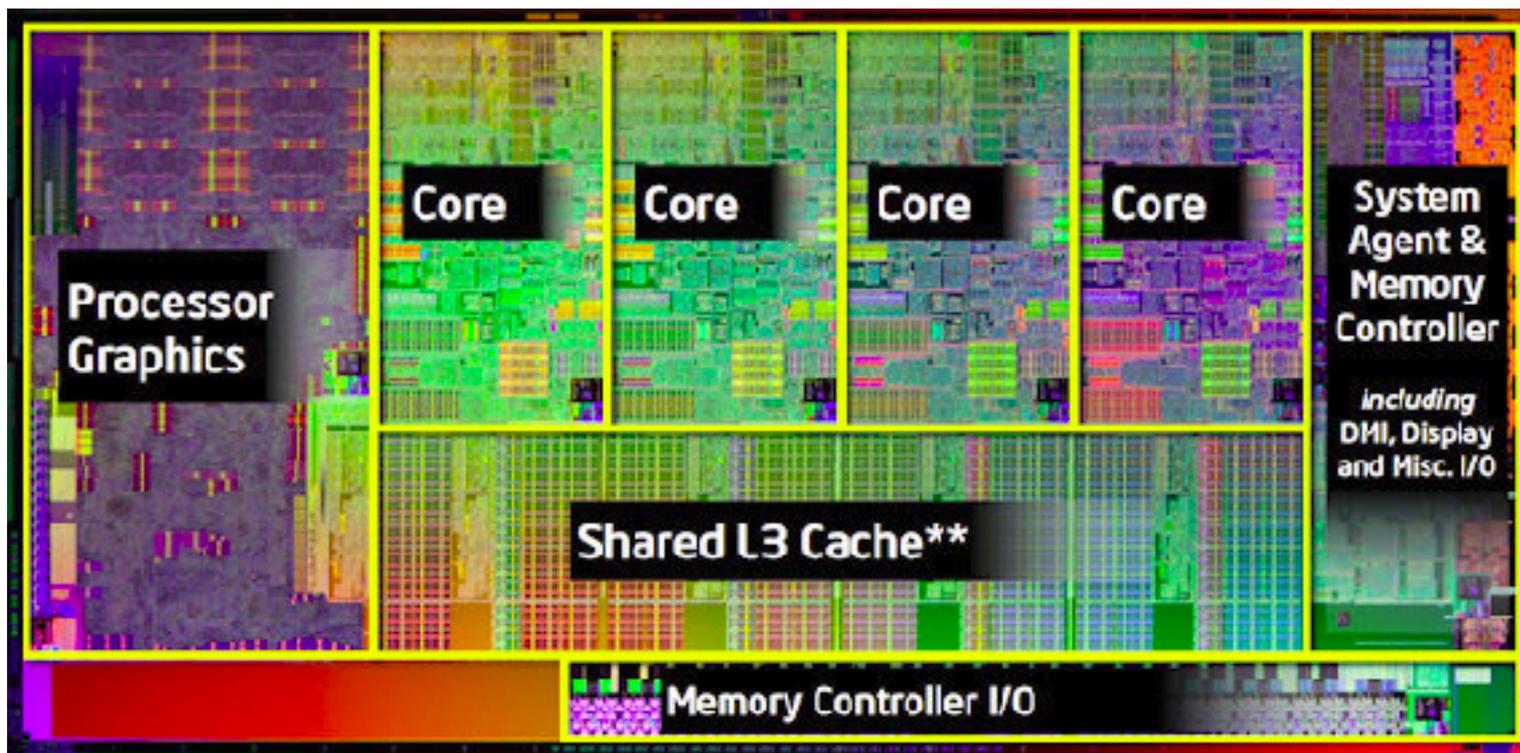
Glue

- 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!

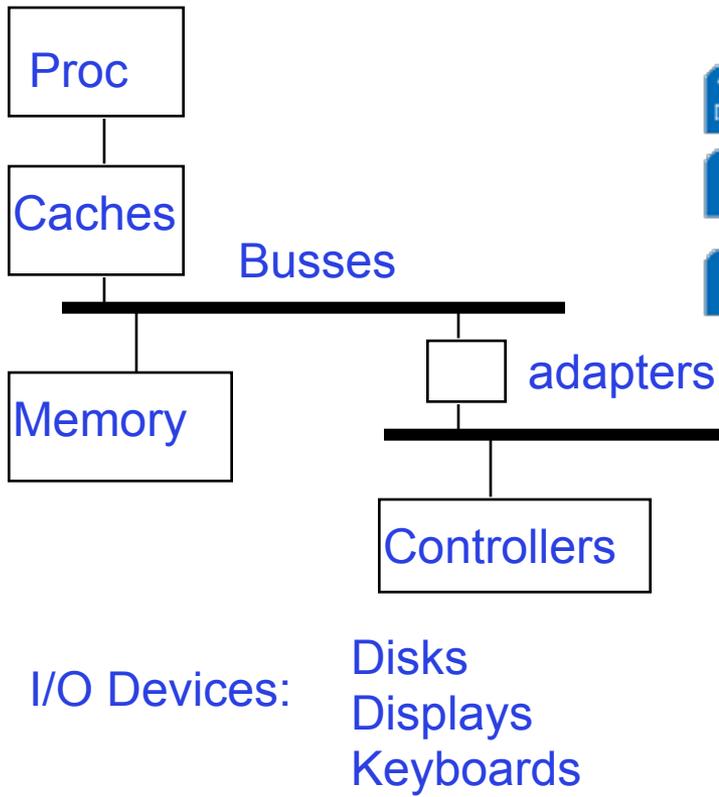
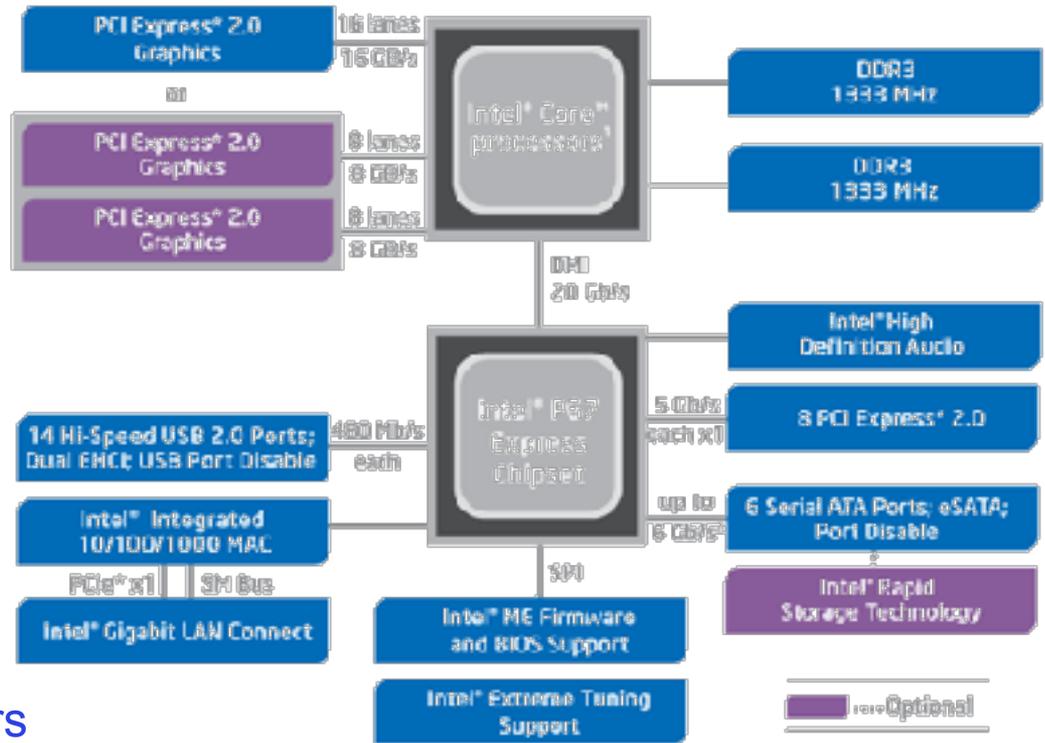
A modern processor: SandyBridge



- **Package: LGA 1155**
 - 1155 pins
 - 95W design envelope
- **Cache:**
 - L1: 32K Inst, 32K Data (3 clock access)
 - L2: 256K (8 clock access)
 - Shared L3: 3MB - 20MB (not out yet)
- **Transistor count:**
 - 504 Million (2 cores, 3MB L3)
 - 2.27 Billion (8 cores, 20MB L3)
- **Note that ring bus is on high metal layers - above the Shared L3 Cache**

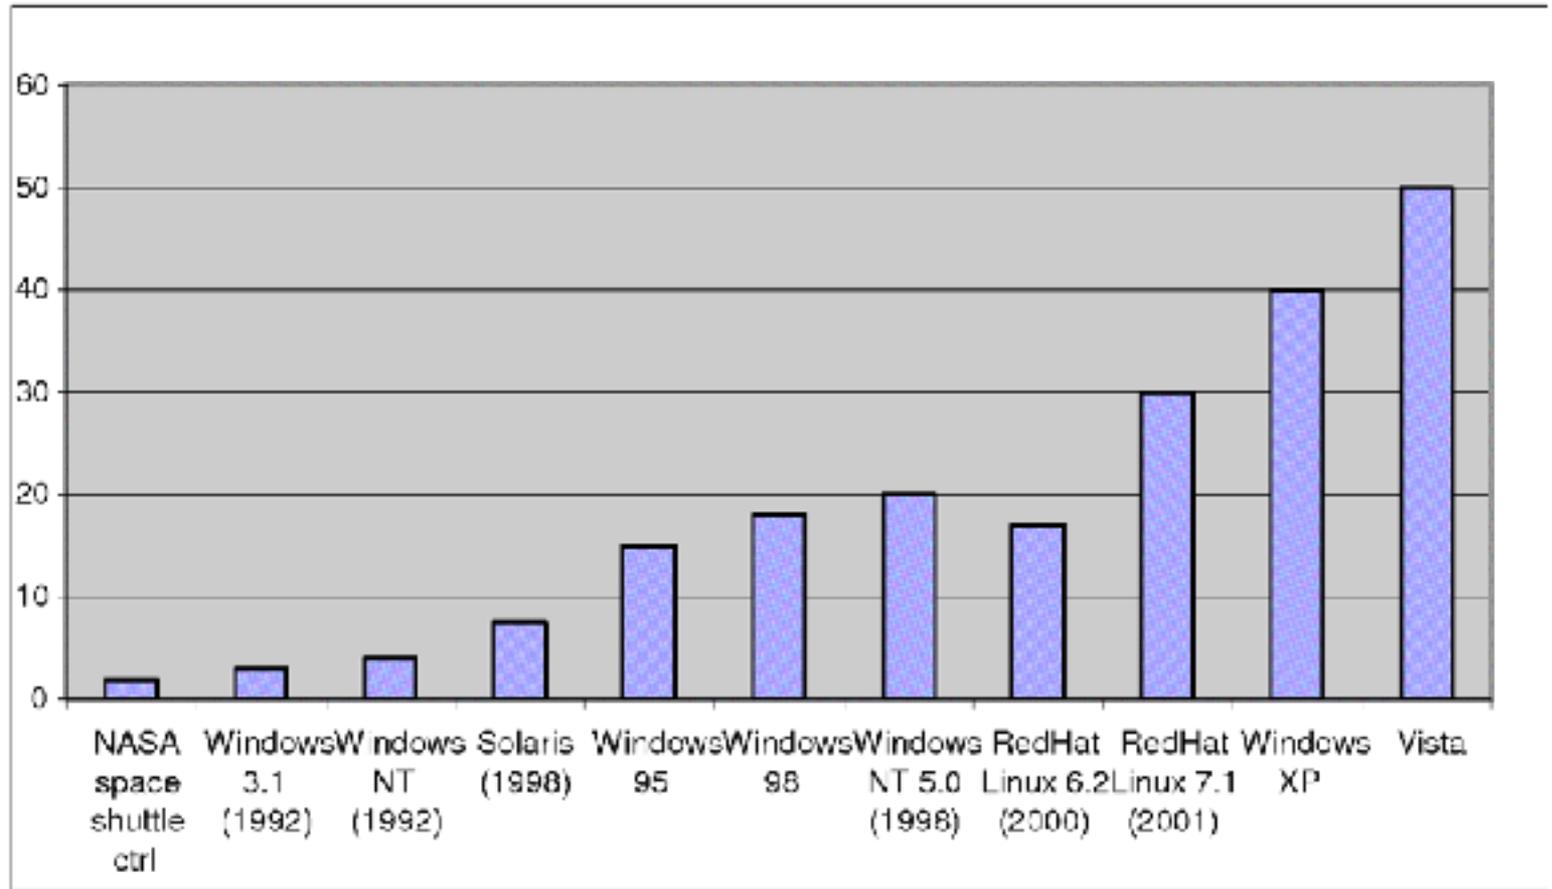
Functionality comes with great complexity!

SandyBridge I/O Configuration



Increasing Software Complexity

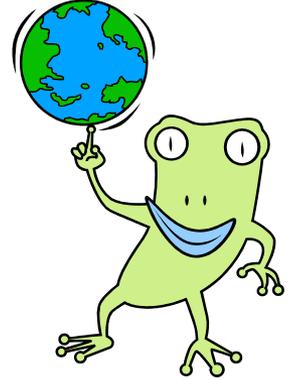
Millions of lines of
source code



From MIT's 6.033 course

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

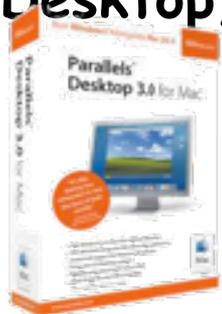
Physical Machine Interface

Hardware

- Software Engineering Problem:
 - Turn hardware/software quirks \Rightarrow 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)

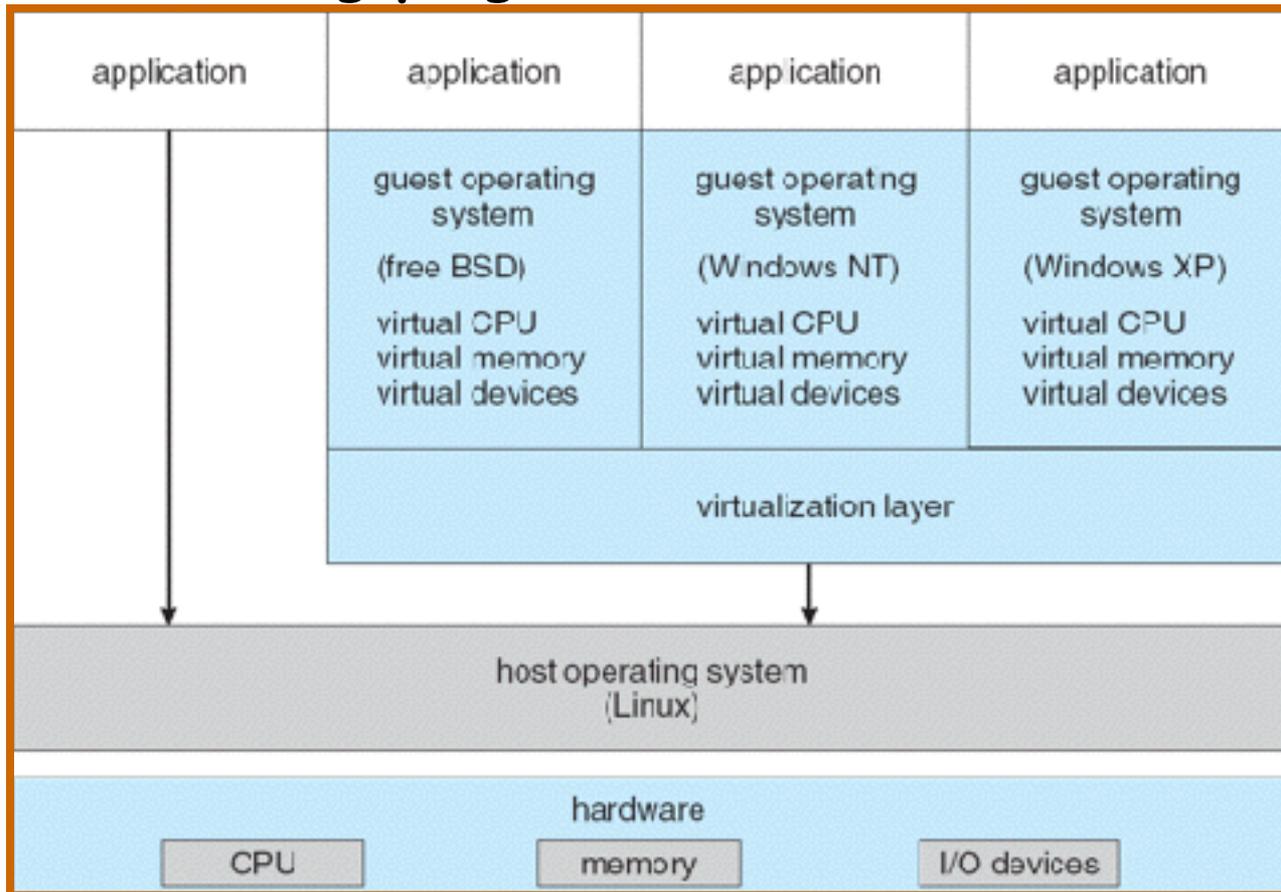


Process VMs

- **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

System Virtual Machines: Layers of OSs

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



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? 😊

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 coordinate resources and protect users from each other
- Operating systems simplify application development by providing standard services
- Operating systems can provide an array of fault containment, fault tolerance, and fault recovery