



Operating Systems Orientation

CS 241


January 25, 2012

[Objectives]

- Explain the main purpose of operating systems and describe milestones of OS evolution
- Explain fundamental machine concepts
 - Instruction processing
 - Memory hierarchy
 - Interrupts
 - I/O
- Explain fundamental OS concepts
 - System calls
 - Processes
 - Synchronization
 - Files
- Explain the POSIX standard (UNIX specification)

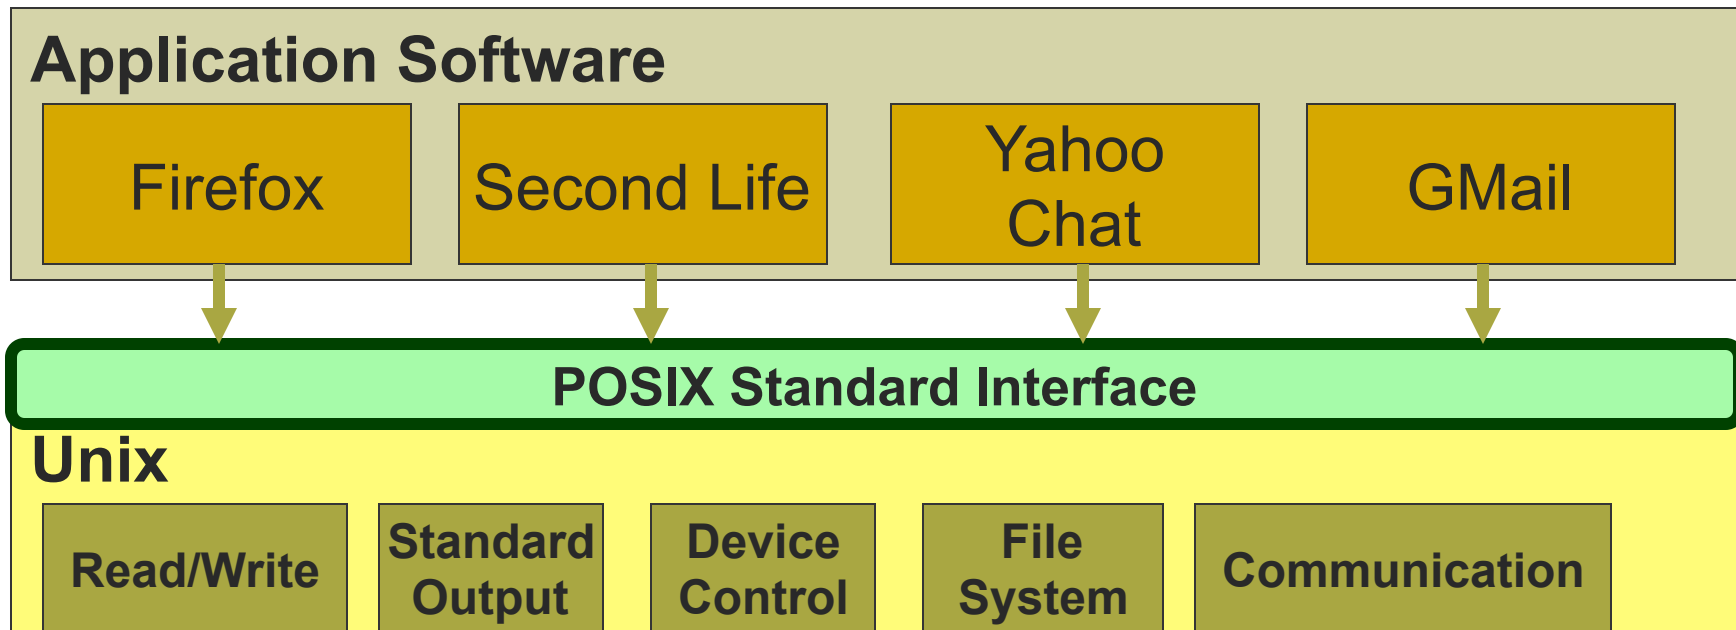


[OS Structure]



POSIX

The UNIX Interface Standard




[What is an Operating System?]


- It is an *extended machine*
 - Hides the messy details that must be performed
 - Presents user with a virtualized and simplified abstraction of the machine, easier to use
- It is a *resource manager*
 - Each program gets time with the resource
 - Each program gets space on the resource



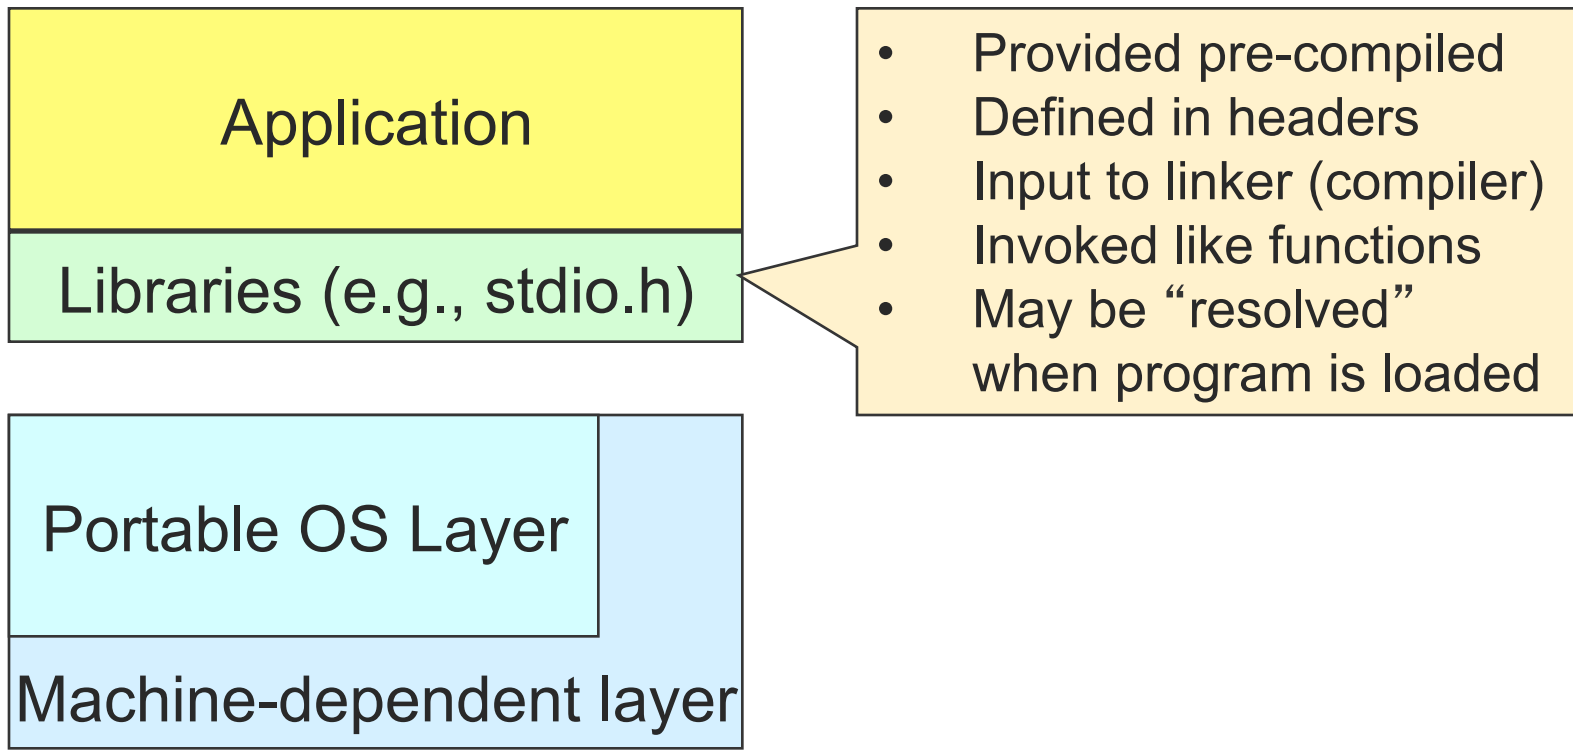
[A Peek into Unix]





[Application]



[Unix: Libraries]



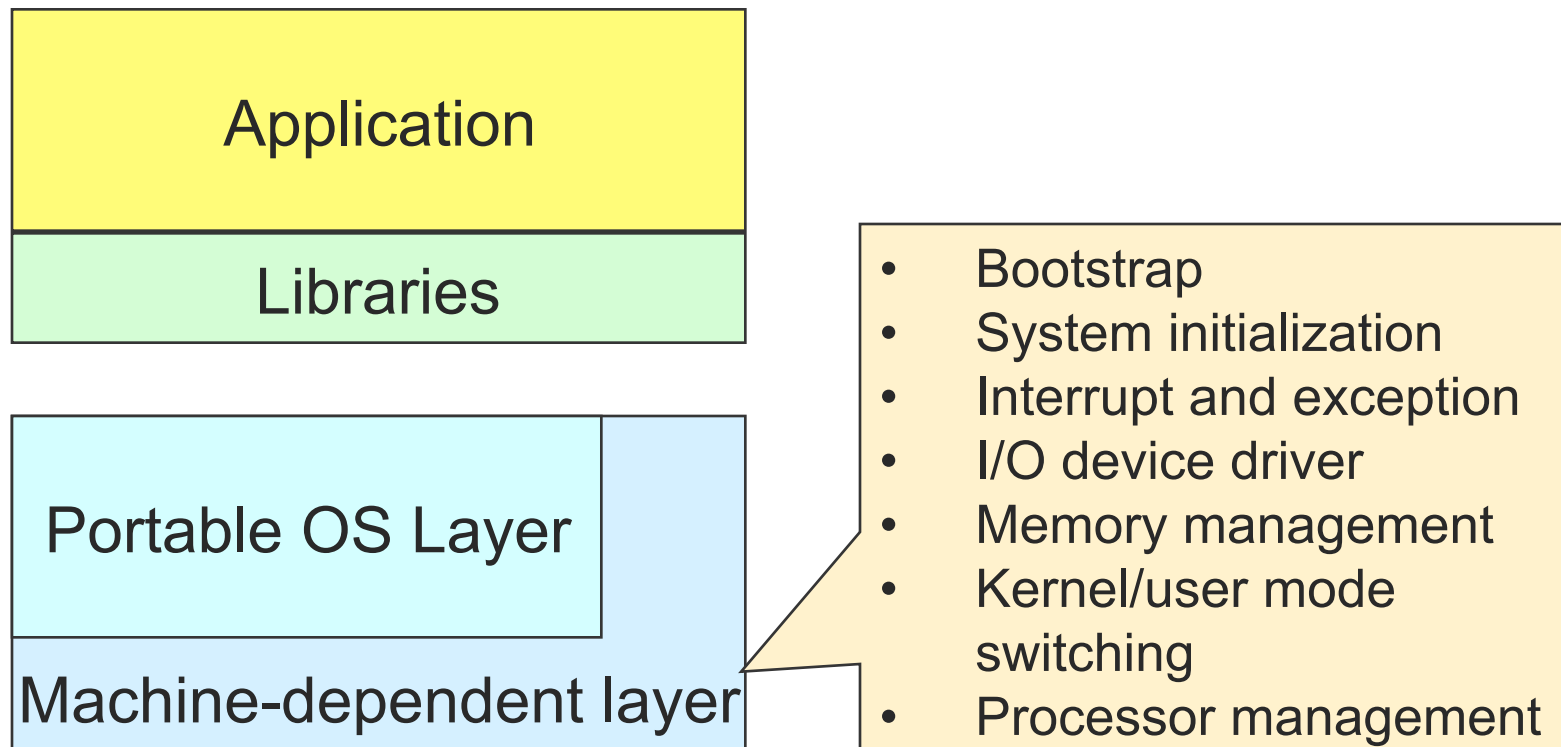
[Typical Unix OS Structure]



- System calls (read, open..)
- All “high-level” code



[Typical Unix OS Structure]



History of Computer Generations

- Pre-computing generation 1792 - 1871
 - Charles Babbage's "Analytical Engine"
 - Purely mechanical
 - Designed, but never actually built
 - Required high-precision gears/cogs that didn't exist yet
 - A man before his time
 - When this works, we'll need software!
 - First programming language
 - World's first programmer: Ada Lovelace



History of Computer Generations

- Pre-computing generation 1792 – 1871
- First generation 1945 – 1955
 - Vacuum tubes, relays, plug boards
 - Seconds per operation!
 - Focus on numerical calculations
 - No programming language
 - Everything done using pure machine language or wiring electrical circuits!
 - No operating system
 - Sign up for your time slot!
 - Progress: Punch cards!




History of Computer Generations

- Pre-computing generation 1792 – 1871
- First generation 1945 – 1955
- Second generation 1955 - 1965
 - Transistors, mainframes
 - Large human component



History of Operating Systems

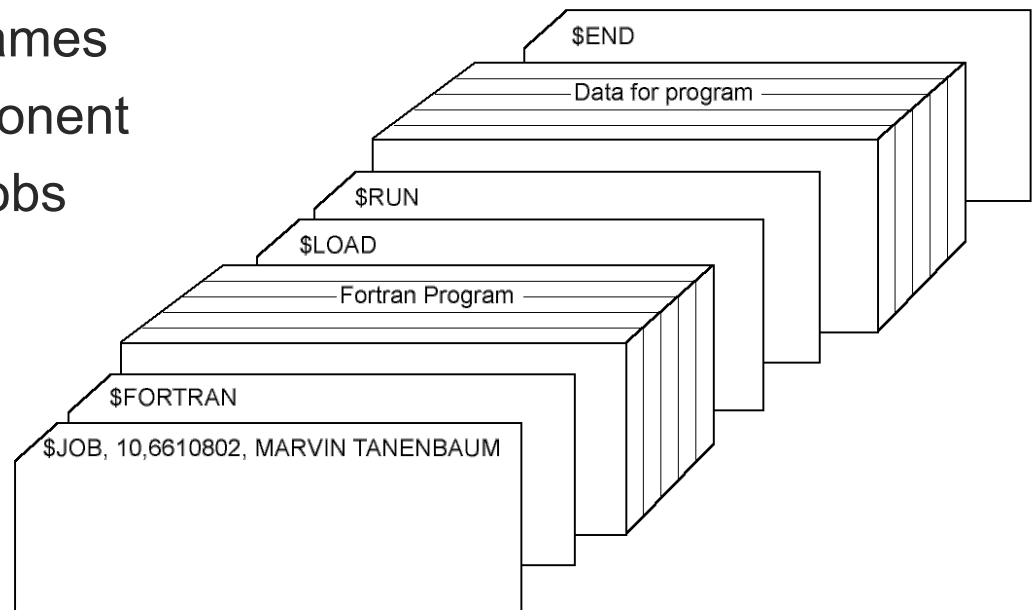


- Problem: a lot of time wasted by operators walking around machine room
- One solution: An early “batch system” by IBM
 - bring cards to 1401
 - read cards to tape
 - put tape on 7094 which does computing
 - put tape on 1401 which prints output



History of Computer Generations

- Pre-computing generation 1792 – 1871
- First generation 1945 – 1955
- Second generation 1955 - 1965
 - Transistors, mainframes
 - Large human component
 - Solution: Batched jobs




History of Computer Generations

- Pre-computing generation 1792 – 1871
- First generation 1945 – 1955
- Second generation 1955 – 1965
- Third generation 1965 – 1980
 - Integrated circuits and multiprogramming
 - IBM's New model: all software and OS must work on all platforms
 - A beast!
 - Progress: Multiprogramming
 - Keep the CPU busy



History of Operating Systems



- Multiprogramming/timesharing system
 - Three jobs in memory – 3rd generation



History of Computer Generations

- Pre-computing generation 1792 – 1871
- First generation 1945 – 1955
- Second generation 1955 – 1965
- Third generation 1965 – 1980
 - Integrated circuits and multiprogramming
 - IBM's New model: all software and OS must work on all platforms
 - Progress: Multiprogramming and timesharing
 - Progress: Spooling
 - Always have something ready to run
 - MULTICS + minicomputers == UNIX!

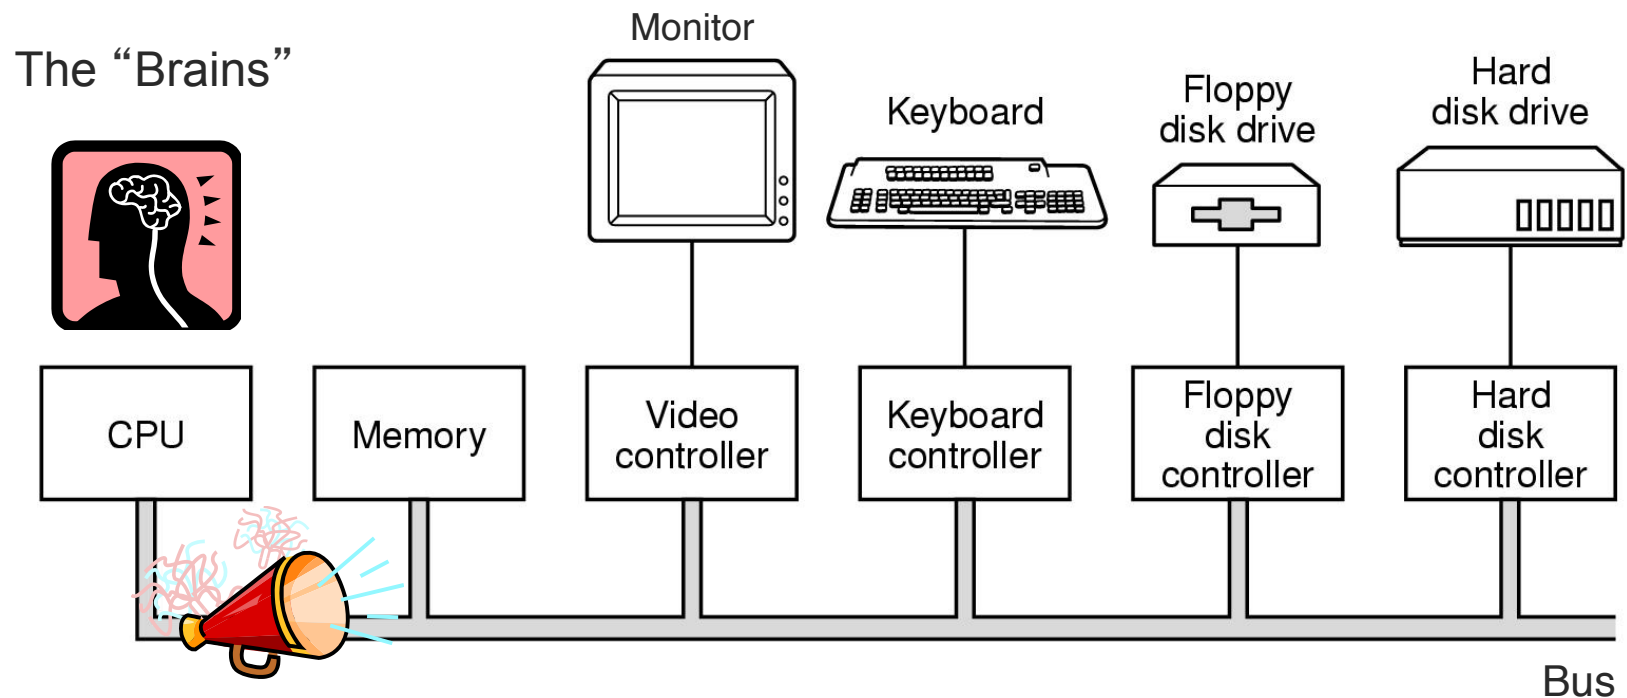


History of Computer Generations

- Pre-computing generation 1792 – 1871
- First generation 1945 – 1955
- Second generation 1955 – 1965
- Third generation 1965 – 1980
- Fourth generation 1980 – present
 - Personal computers
 - Multi-processors
 - Phones
 - ...



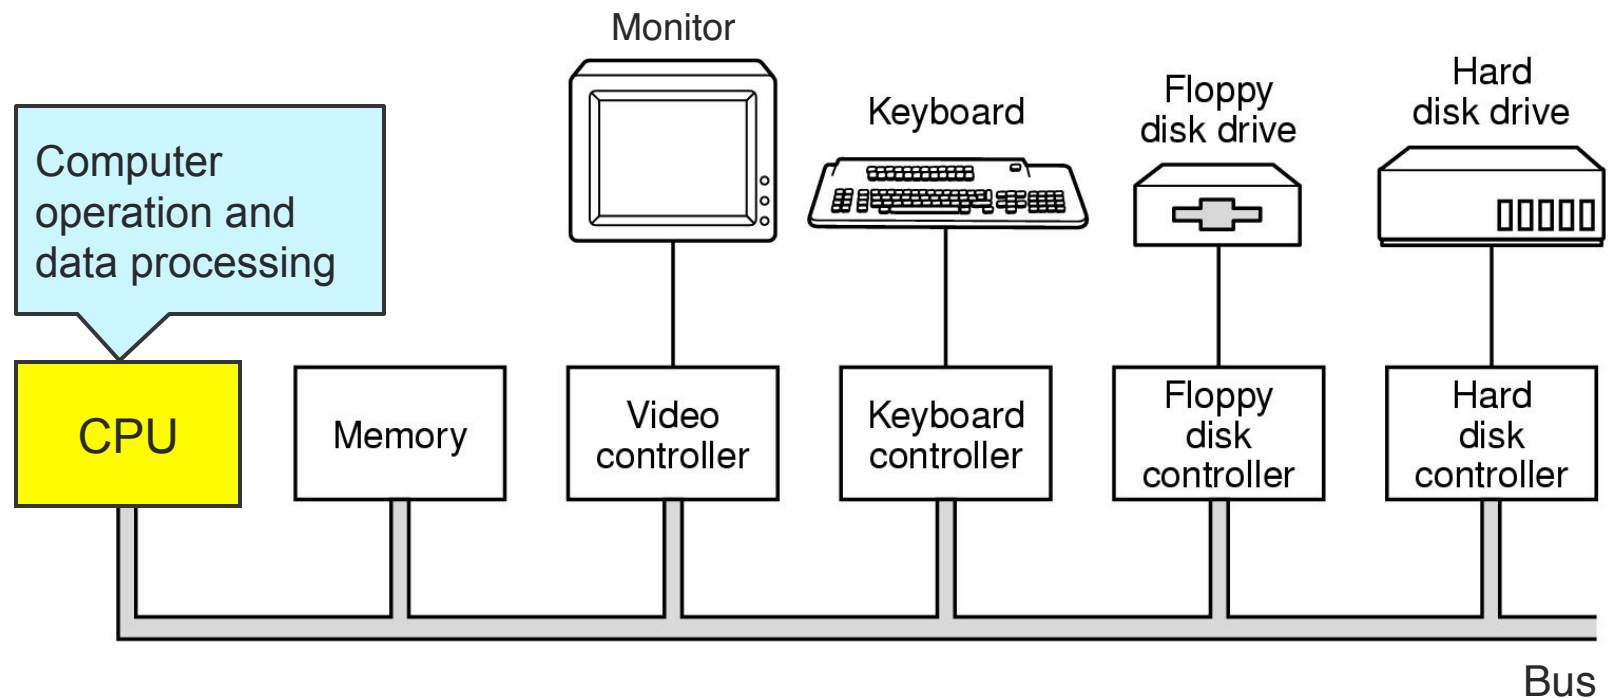
Computer Hardware Review



- Components of a simple personal computer



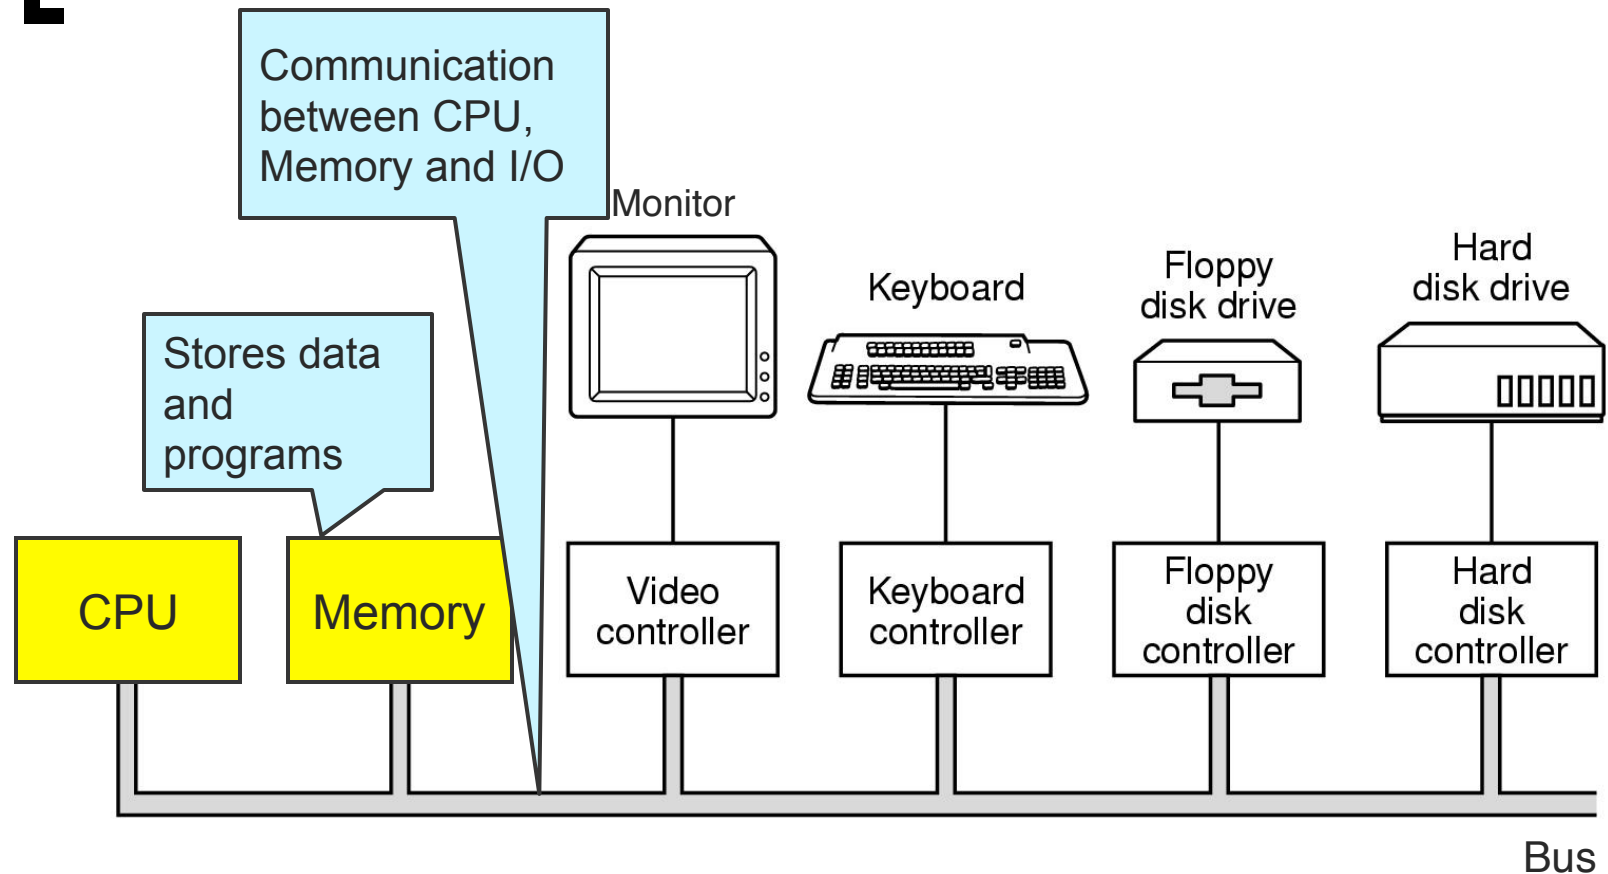
Computer Hardware Review



- Components of a simple personal computer



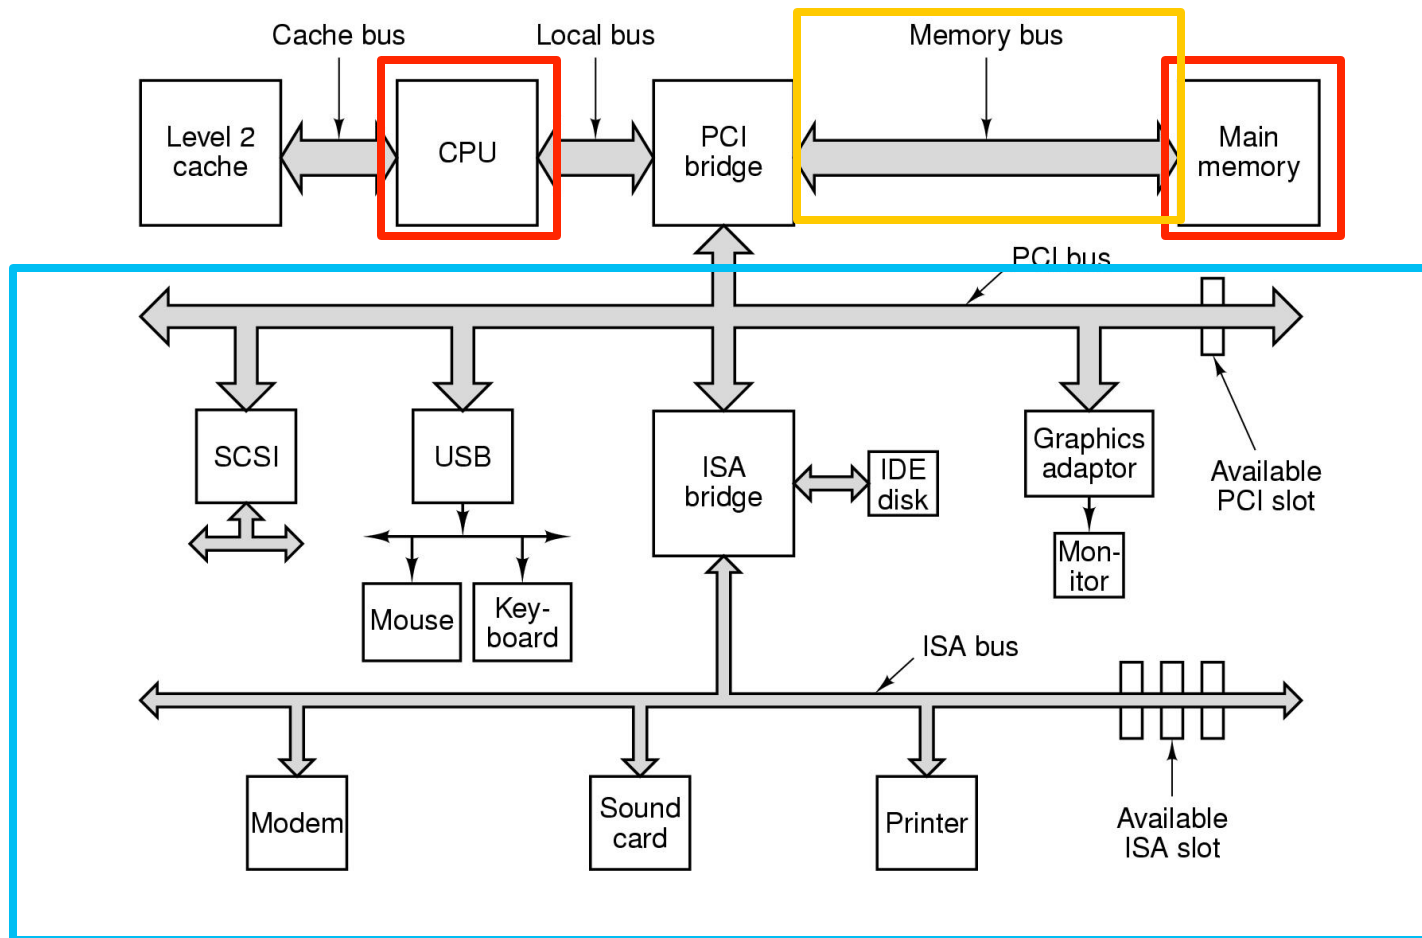
Computer Hardware Review



- Components of a simple personal computer



Early Pentium system



[CPU, From CS231]

- Fetch instruction from code memory
 - Fetch operands from data memory
 - Perform operation (and store result)
 - (Check interrupt line)
 - Go to next instruction
-
- 'Conventional CPU'
(Ignore pipeline, optimization complexities)



[CPU Registers]

- Fetch instruction from code memory
 - Fetch operands from data memory
 - Perform operation (and store result)
 - Go to next instruction
-
- Note: CPU must maintain certain state
 - Current instructions to fetch (program counter)
 - Location of code memory segment
 - Location of data memory segment



[CPU Register Examples]

- Hold instruction operands
- Point to start of
 - Code segment (executable instructions)
 - Data segment (static/global variables)
 - Stack segment (execution stack data)
- Point to current position of
 - Instruction pointer
 - Stack pointer




[CPU Register Examples]

- Hold instruction operands
- Point to start of
 - Code segment
 - Data segment
 - Stack segment
- Point to current position of
 - Instruction pointer
 - Stack pointer
 - Why stack?

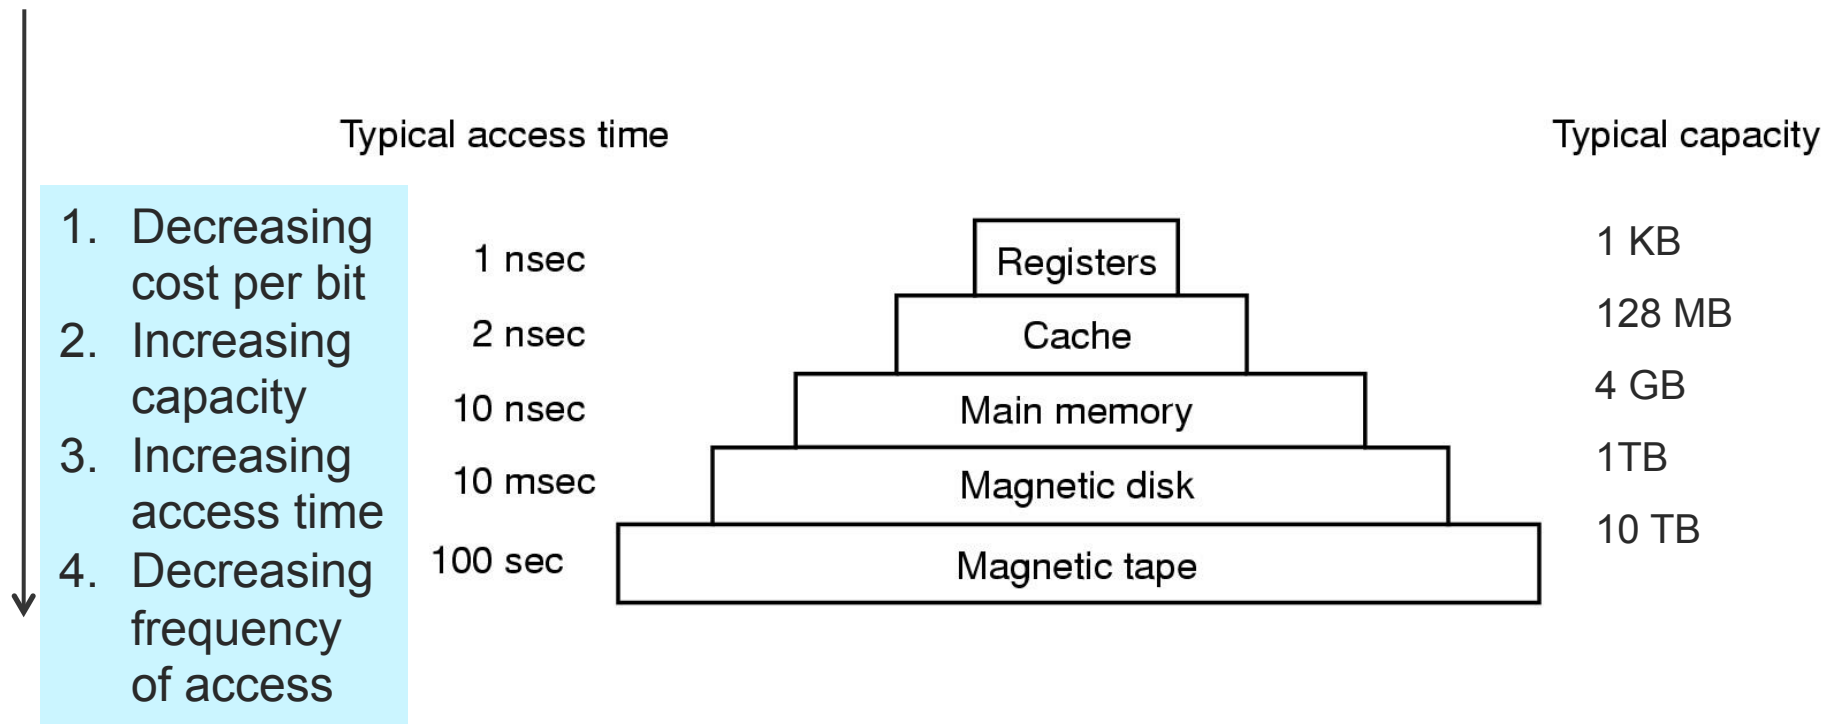


Sample Layout for program image in main memory

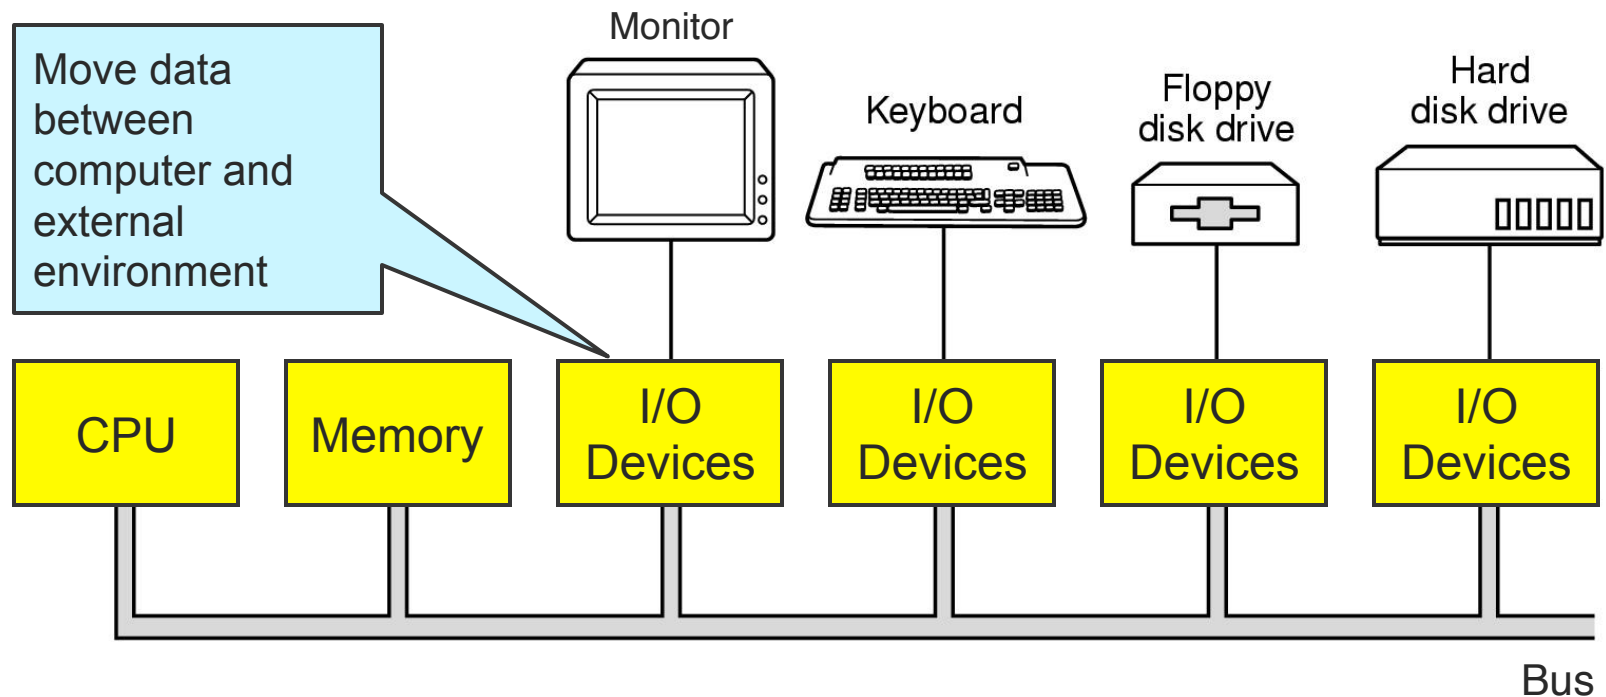


Memory Hierarchy

- Leverage locality of reference



Computer Hardware Review



- Components of a simple personal computer

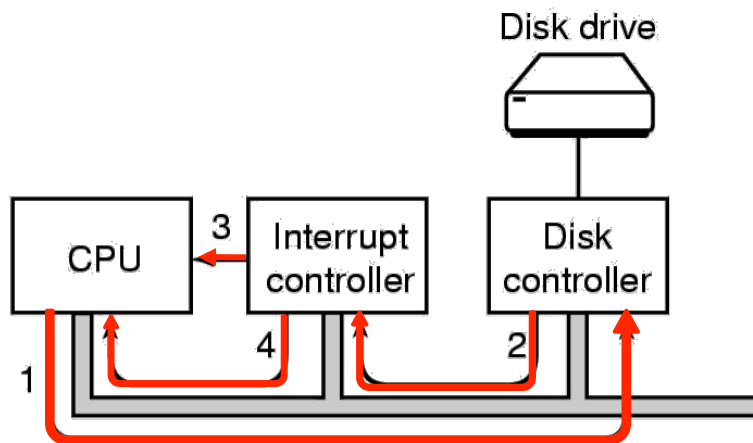


[I/O Device Access]

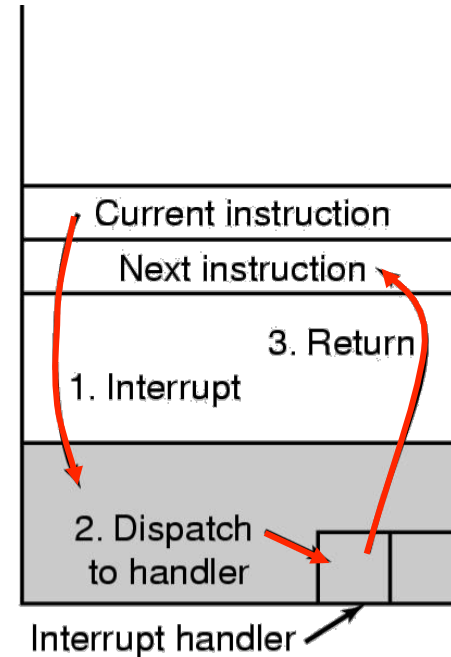
- System Calls
 - Application makes a system call reading from or writing to a device
 - Kernel blocks application
 - Kernel translates request to specific driver
 - Driver starts I/O
 - Polls device for completion
 - Or waits for **interrupt** from device
 - Kernel obtains results, un-blocks application



I/O Interrupt Mechanism



(a)



(b)

1. Application writes into device registers, Controller starts device
2. When done, device controller signals interrupt controller
3. Interrupt controller asserts pin on CPU
4. Interrupt controller puts I/O device number on bus to CPU



[Operating System Concepts]

- Shared resources
 - I have B KB of memory, but need 2B KB
 - I have N processes trying to access the disk at the same time
 - How would you control access to resources?
- Challenges
 - Who gets to use the resources?
 - How do you control fair use of the resources over time?
 - Deadlock




[Operating System Concepts]

- Process

- An executable instance of a program
- Only one process can use a (single-core) CPU at a time

- A process tree

- A created two child processes, B and C
- B created three child processes, D, E, and F



[Operating System Concepts]

- How would you switch CPU execution from one process to another?
- Solution: Context Switching
 - Store/restore state on CPU, so execution can be resumed from same point later in time
 - Triggers: multitasking, interrupt handling, user/kernel mode switching
 - Involves: Saving/loading registers and other state into a “process control block” (PCB)
 - PCBs stored in kernel memory




[Operating System Concepts]

- Context Switching
 - What are the costs involved?


Item	Time	Scaled Time in Human Terms (2 billion times slower)
Processor cycle	0.5 ns (2 GHz)	1 s
Cache access	1 ns (1 GHz)	2 s
Memory access	15 ns	30 s
Context switch	5,000 ns (5 micros)	167 m
Disk access	7,000,000 ns (7 ms)	162 days
System quanta	100,000,000 (100 ms)	6.3 years



Operating System Concepts



(a) A potential deadlock




(b) An actual deadlock

- Another challenge: Deadlock
 - Set of actions waiting for each other to finish
- Example:
 - Process A has lock on file 1, wants to acquire lock on file 2
 - Process B has a lock on file 2, wants to acquire lock on file 1



[Operating System Concepts]

- Inter-process Communication
 - Now process A needs to exchange information with process B
 - How would you enable communication between processes?



[Next up]

- MP1 released
- Discussion sections tomorrow
- Friday: System calls



[Summary]

- Resource Manager
- Hardware independence
- Virtual Machine Interface
- POSIX
- Concurrency & Deadlock

