

2. Computer-System Structures

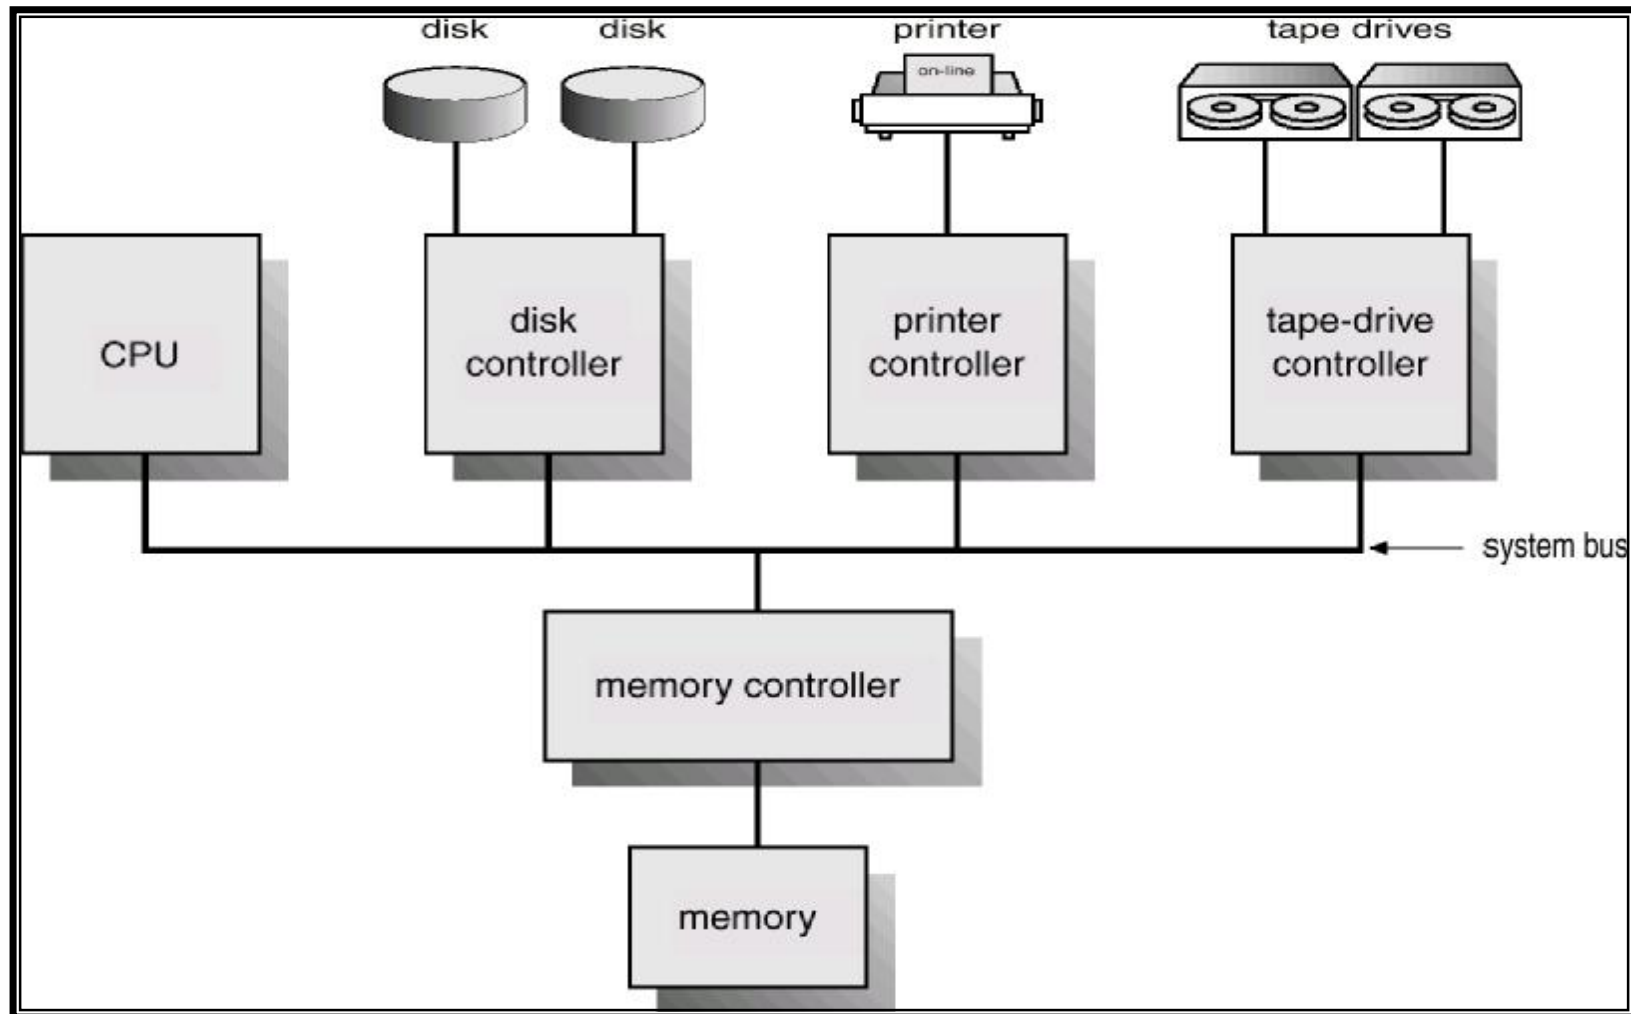
Sungyoung Lee

*College of Engineering
KyungHee University*

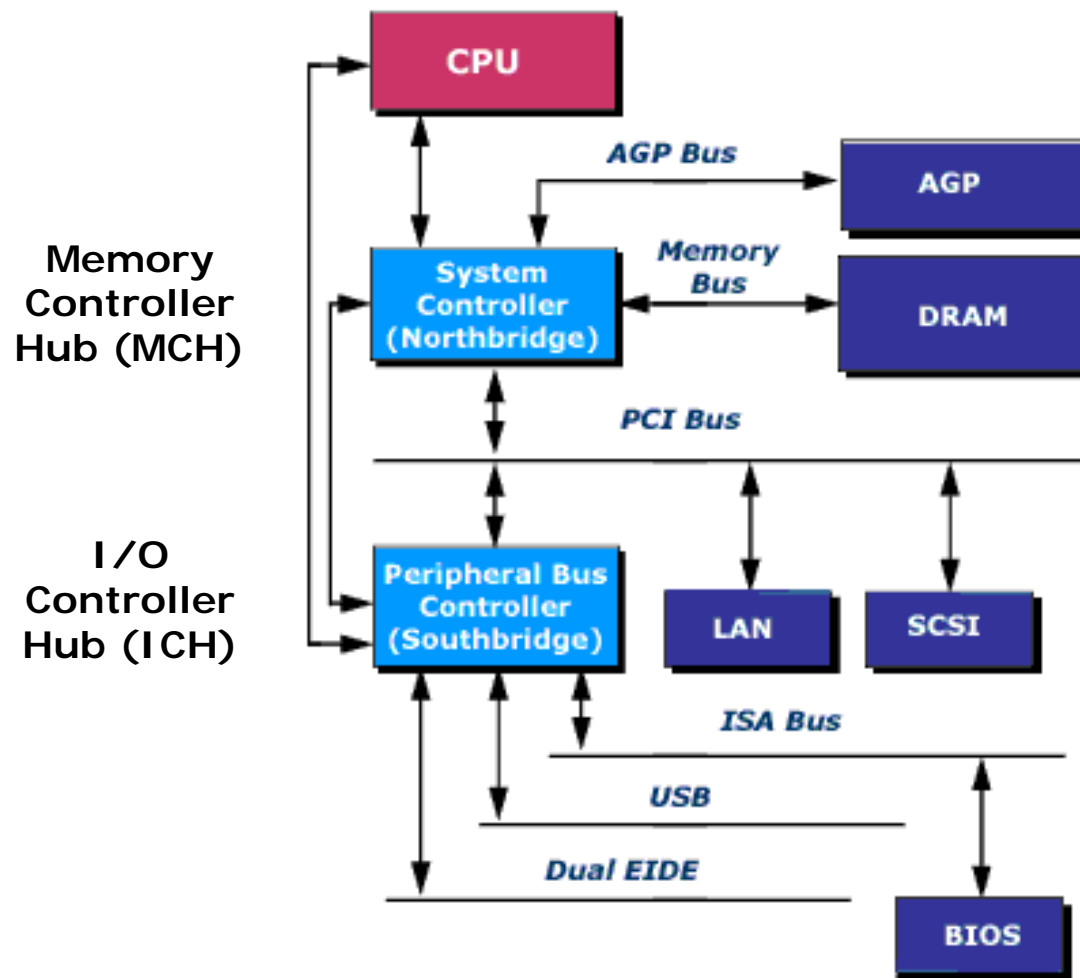
Contents

- n *Computer System Operation*
- n *I/O Structure*
- n *Storage Structure*
- n *Storage Hierarchy*
- n *Hardware Protection*
- n *General System Architecture*

Computer-System Architecture



Modern PC Architecture



CPU

n Instruction Set Architecture

- ü RISC vs. CISC
- ü Intel, SPARC, MIPS, PowerPC, ARM, Alpha, ...

n Pipelining

- ü Fetch, Decode, Execute, Write Back, etc.

n Registers

- ü General-purpose registers
- ü Program Counter
- ü PSW (Program Status Word)

n Instruction-Level Parallelism(ILP)

- ü Superscalar vs. VLIW
- ü Simultaneous Multithreading

OS and Architecture

n Mutual interaction

- ü The functionality of an OS is limited by architectural features
 - § Multiprocessing on DOS/8086?
- ü The structure of an OS can be simplified by architectural support
 - § Interrupt, DMA, etc.
- ü Most proprietary OS's were developed with the certain architecture in mind

Computer-System Operation

- n I/O devices and the CPU can execute concurrently
- n Each device controller is in charge of a particular device type
- n Each device controller has a local buffer
- n CPU moves data from/to main memory to/from local buffers
- n I/O is from the device to local buffer of controller
- n Device controller informs CPU that it has finished its operation by causing an *interrupt*

Common Functions of Interrupts

- n Interrupt transfers control to the interrupt service routine generally, through the *interrupt vector*, which contains the addresses of all the service routines
- n Interrupt architecture must save the address of the interrupted instruction
- n Incoming interrupts are *disabled* while another interrupt is being processed to prevent a *lost interrupt*
- n A *trap (or exception)* is a software-generated interrupt caused either by an error or a user request
- n An operating system is *interrupt driven*

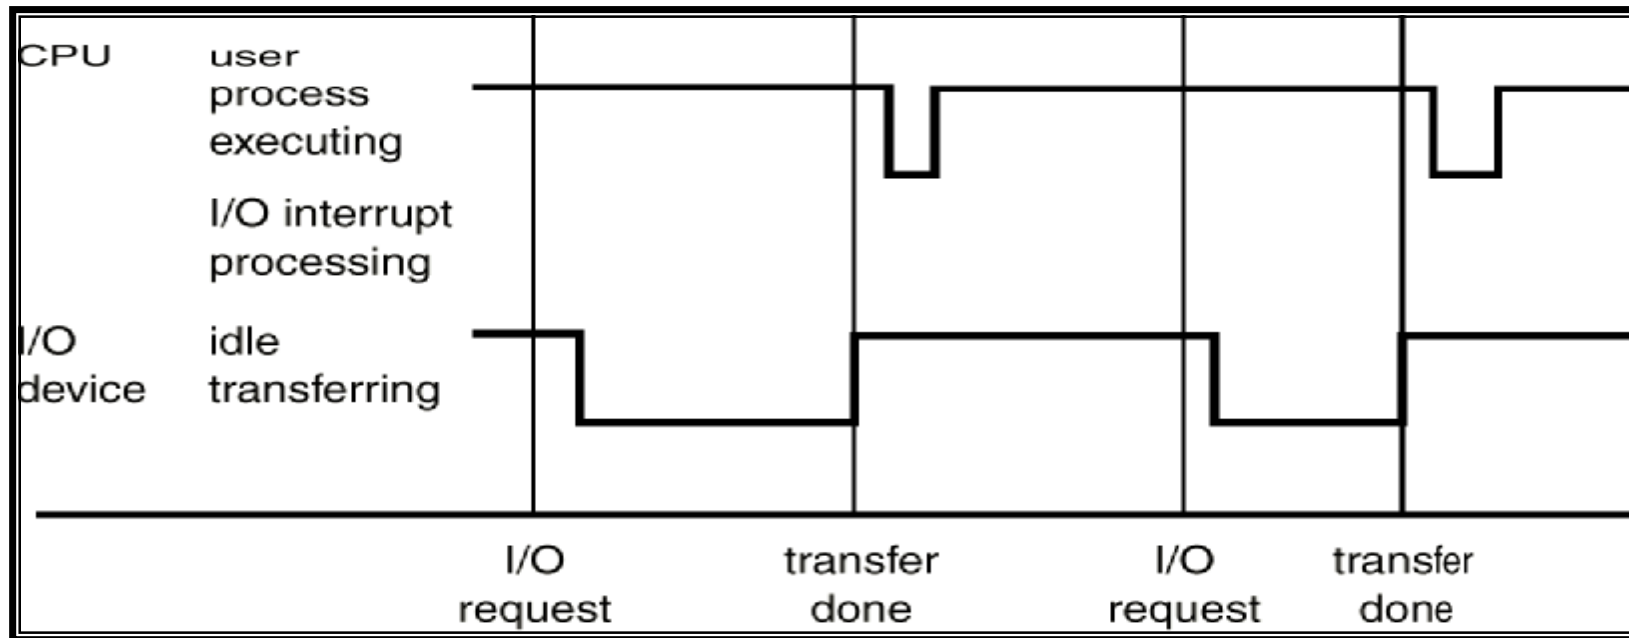
Interrupt Handling

- n The operating system preserves the state of the CPU by storing registers and the program counter

- n Determines which type of interrupt has occurred:
 - ü *polling*
 - ü *vectored* interrupt system

- n Separate segments of code determine what action should be taken for each type of interrupt

Interrupt Time Line For a Single Process Doing Output



Interrupts and Exceptions

n Interrupts

- ü Generated by hardware devices
 - § Triggered by a signal in INTR or NMI pins (Pentium)
- ü Asynchronous

n Exceptions

- ü Generated by software executing instructions
 - § INT instruction in IA32
 - § Page fault, protection fault
- ü Synchronous
- ü Trap (expected) or fault (unexpected)

I/O Structure

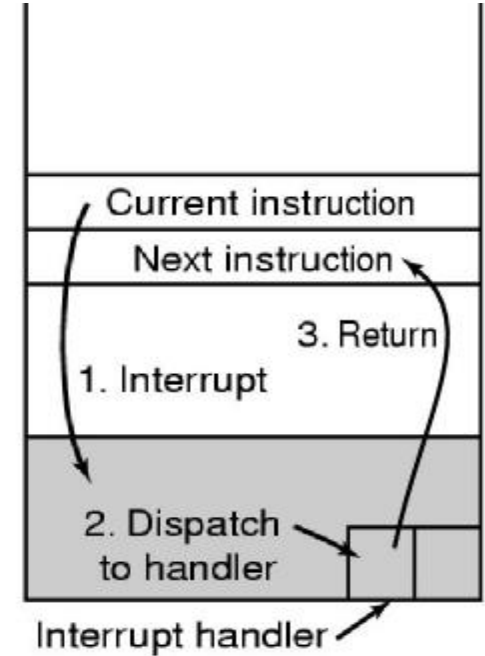

- n After I/O starts, control returns to user program only upon I/O completion
 - ü Wait instruction idles the CPU until the next interrupt
 - ü Wait loop (contention for memory access)
 - ü At most one I/O request is outstanding at a time, no simultaneous I/O processing

- n After I/O starts, control returns to user program without waiting for I/O completion
 - ü *System call* – request to the operating system to allow user to wait for I/O completion
 - ü *Device-status table* contains entry for each I/O device indicating its type, address, and state
 - ü Operating system indexes into I/O device table to determine device status and to modify table entry to include interrupt

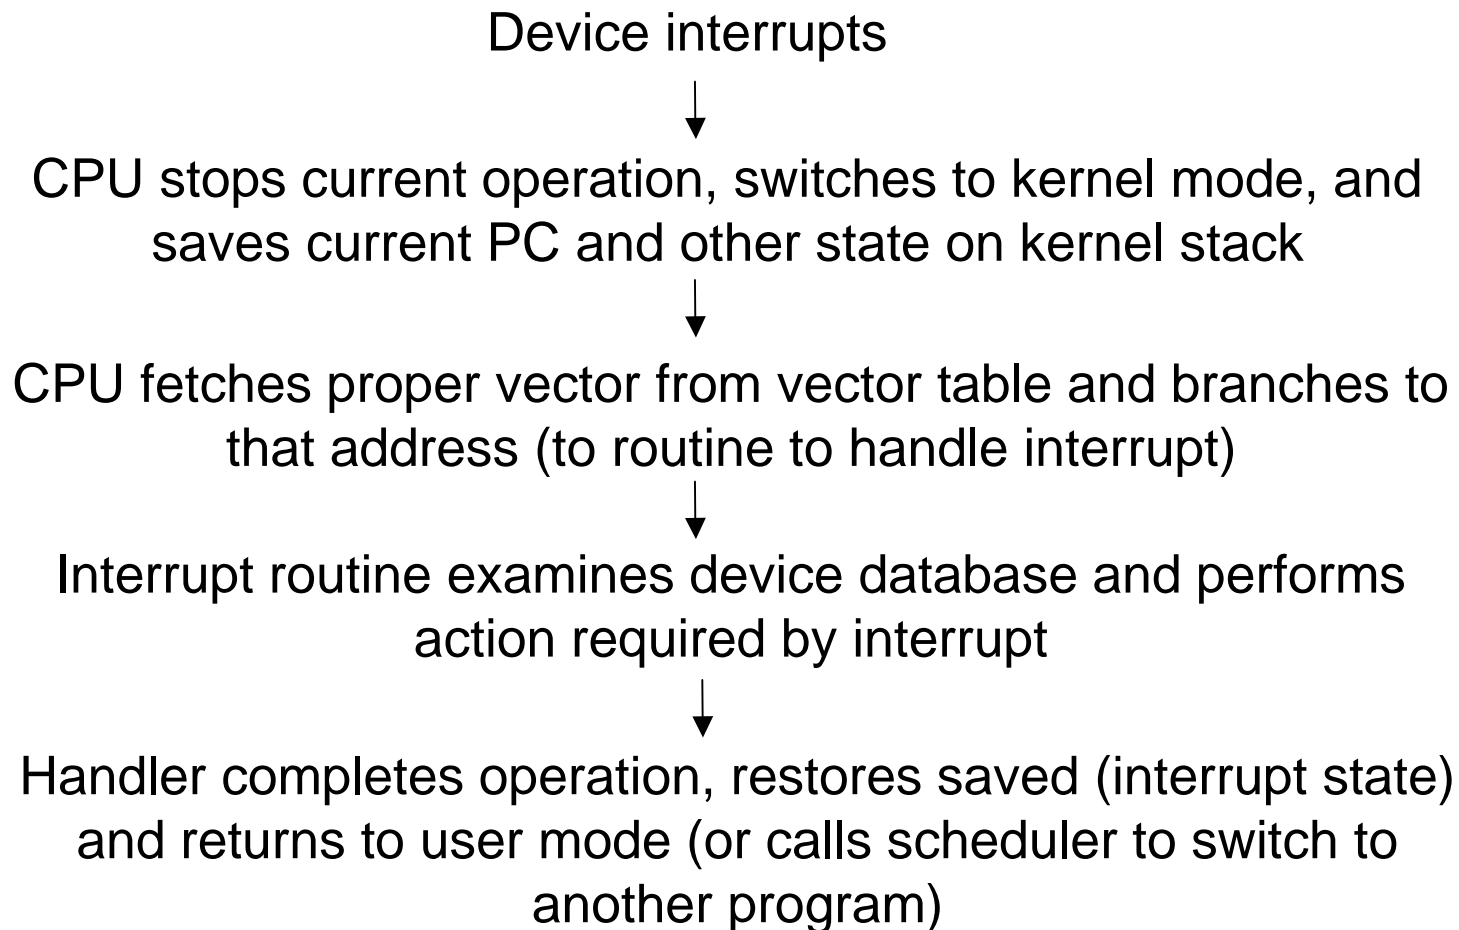
I/O Control

n How does the kernel notice an I/O has finished?

- ü Polling
- ü Hardware interrupt



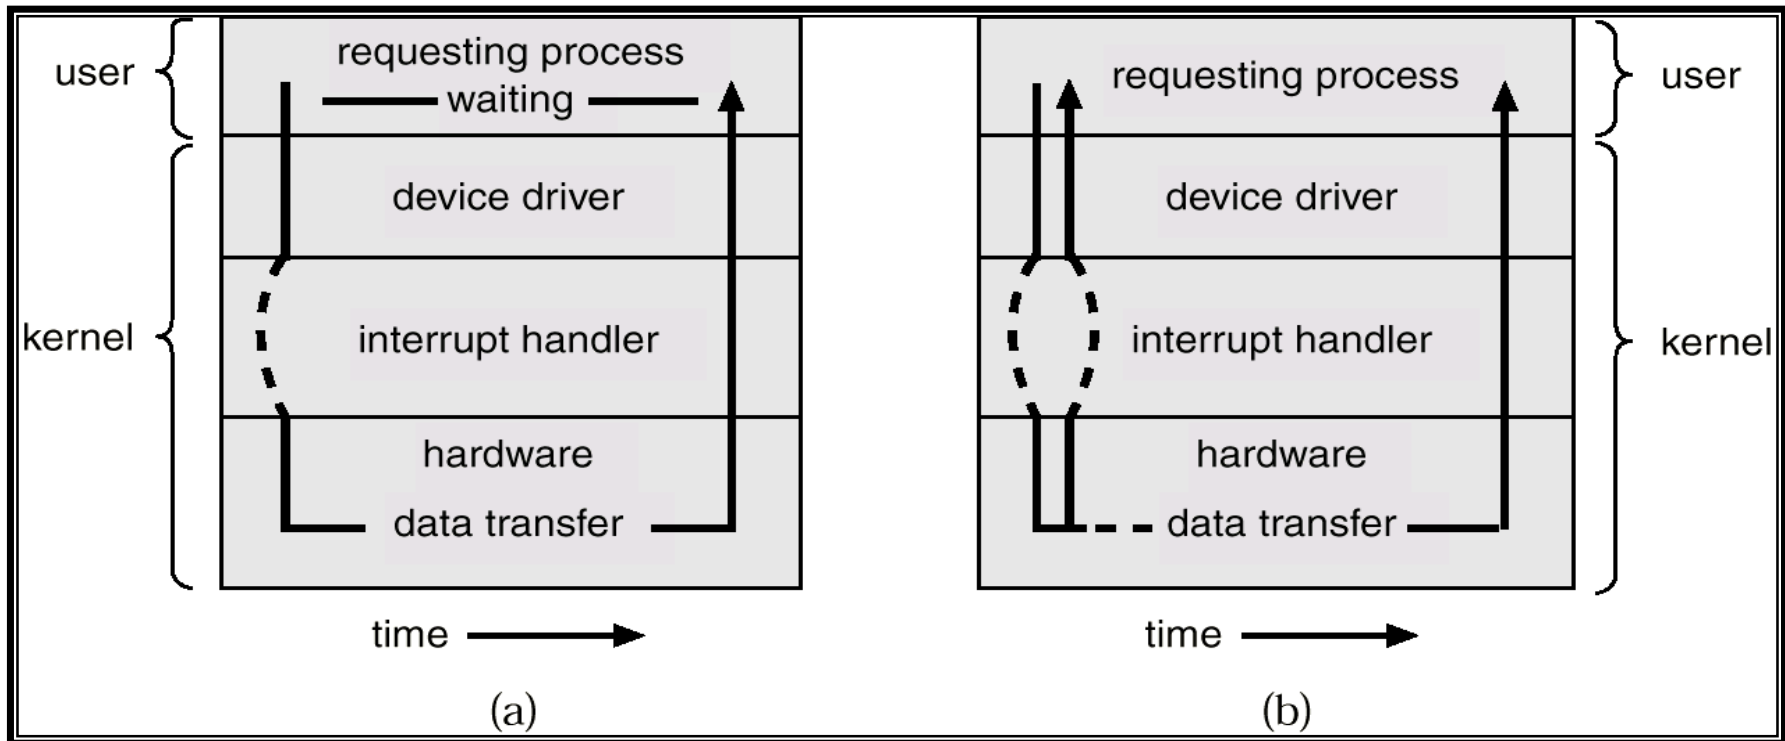
I/O Control (Cont'd)



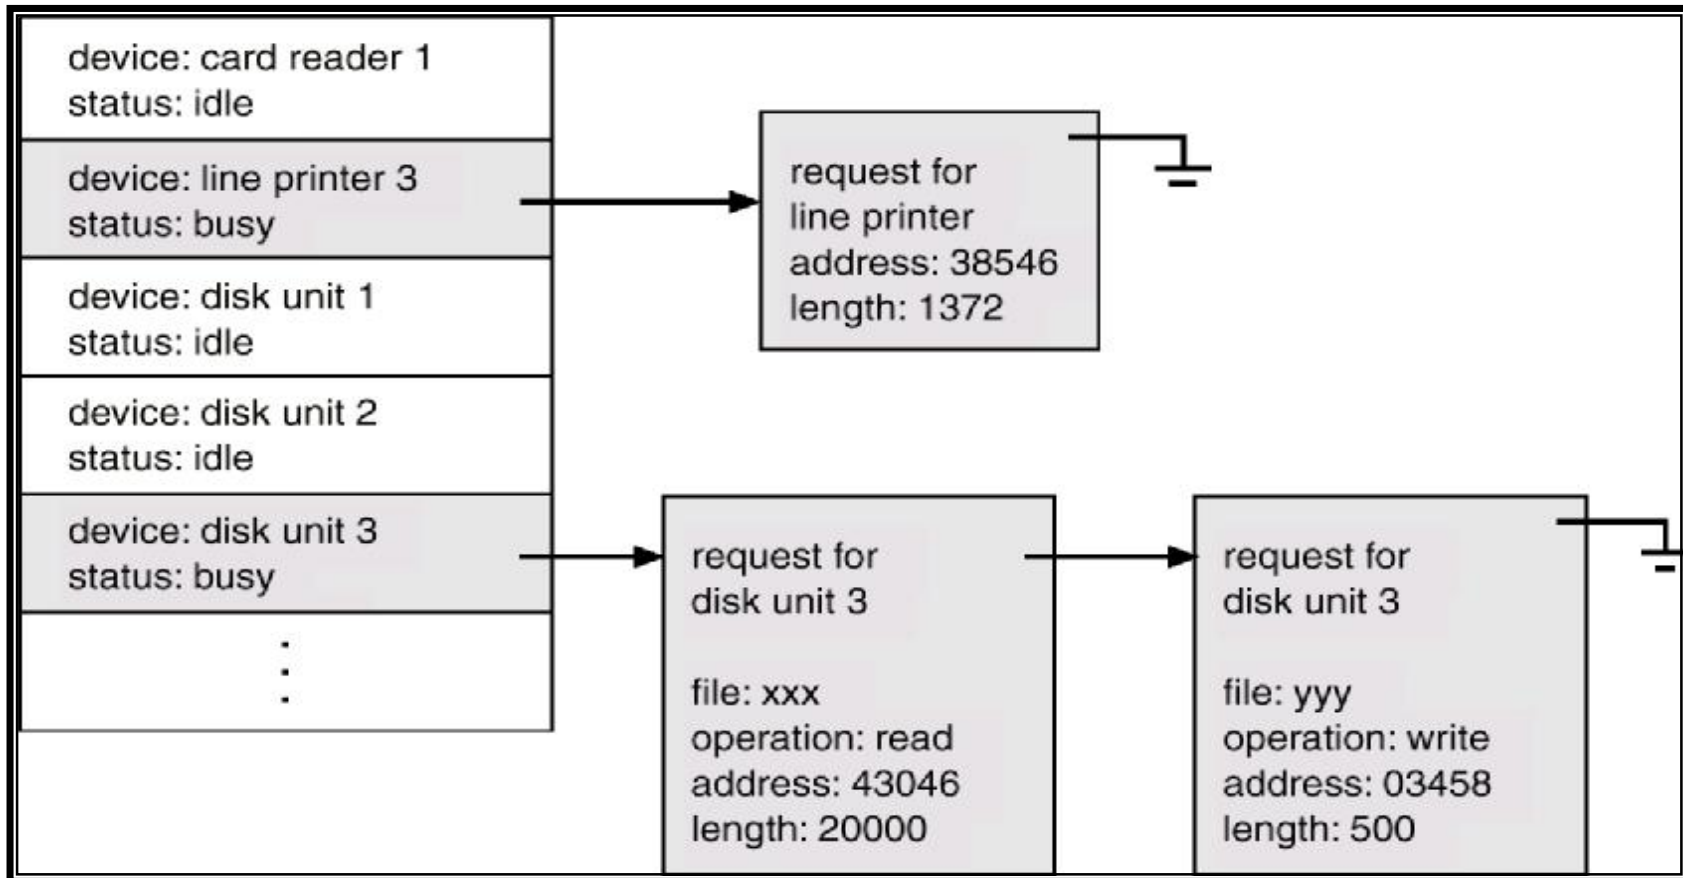
Two I/O Methods

Synchronous/ Blocking

Asynchronous/ Non-blocking



Device-Status Table



Direct Memory Access Structure

- n Used for high-speed I/O devices able to transmit information at close to memory speeds
- n Device controller transfers blocks of data from buffer storage directly to main memory without CPU intervention
- n Only one interrupt is generated per block, rather than the one interrupt per byte

DMA

n Data Transfer Modes in I/O

ü Programmed I/O (PIO)

- § By special I/O instructions
- § Memory-mapped I/O

ü DMA (Direct Memory Access)

- § Used for high-speed I/O devices able to transmit information at close to memory speeds
- § Device controller transfers blocks of data from buffer storage directly to main memory without CPU intervention.
- § Only an interrupt is generated per block


Storage Structure

- n Main memory – only large storage media that the CPU can access directly
 - ü von Neumann architecture vs. Harvard architecture
 - ü SRAM vs. DRAM
 - ü DDR, QDR, RDRAM

- n Secondary storage – extension of main memory that provides large nonvolatile storage capacity

- n Magnetic disks – rigid metal or glass platters covered with magnetic recording material
 - ü Disk surface is logically divided into *tracks*, which are subdivided into *sectors*
 - ü The *disk controller* determines the logical interaction between the device and the computer

Moving-Head Disk Mechanism



A Modern Disk Drive

n IBM Deskstar 120GXP (120GB)

- ü Disks/Head: 3/6
- ü Max. areal density: 29.7 Gbits/sq.inch
- ü Max. recording density: 524K BPI (bits/inch)
- ü Track density: 56.7K TPI (tracks/inch)
- ü Seek time: 8.5ms average
- ü Rotational latency: 7200rpm (8.3ms/rotation)
- ü Max. media transfer rate: 592 Mbits/sec
- ü Max. interface transfer rate: 100MB/sec
- ü Sustained data rate: 48 to 23 MB/sec

Storage Hierarchy


n Storage systems organized in hierarchy

- ü Speed
- ü Cost
- ü Volatility

n *Caching*

- ü copying information into faster storage system
- ü main memory can be viewed as a last *cache* for secondary storage

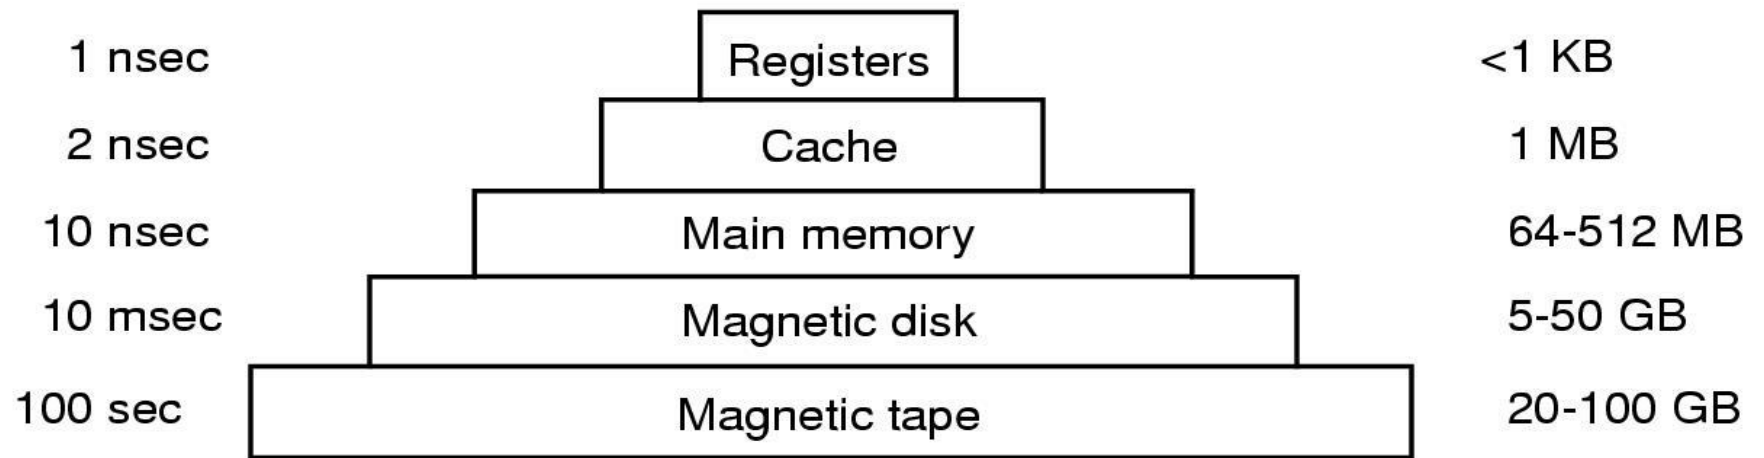
Storage-Device Hierarchy



Memory Hierarchy

Typical access time


Typical capacity



Caching

- n Use of high-speed memory to hold recently-accessed data
- n Requires a *cache management* policy
 - ü Write-through vs. Write-back
- n Caching introduces another level in storage hierarchy. This requires data that is simultaneously stored in more than one level to be *consistent*
 - ü Cache coherency

Migration of A From Disk to Register



Hardware Protection

- n Dual-Mode Operation
- n I/O Protection
- n Memory Protection
- n CPU Protection


Dual-Mode Operation

- n Sharing system resources requires operating system to ensure that an incorrect program cannot cause other programs to execute incorrectly

- n Provide hardware support to differentiate between at least two modes of operations
 1. *User mode*
 - § execution done on behalf of a user
 2. *Monitor mode (also kernel mode or system mode)*
 - § execution done on behalf of operating system

Dual-Mode Operation (Cont'd)

- n *Mode bit* added to computer hardware to indicate the current mode:
 - ü monitor (0) or user (1)
- n When an interrupt or fault occurs hardware switches to monitor mode




Privileged instructions can be issued only in monitor mode

I/O Protection

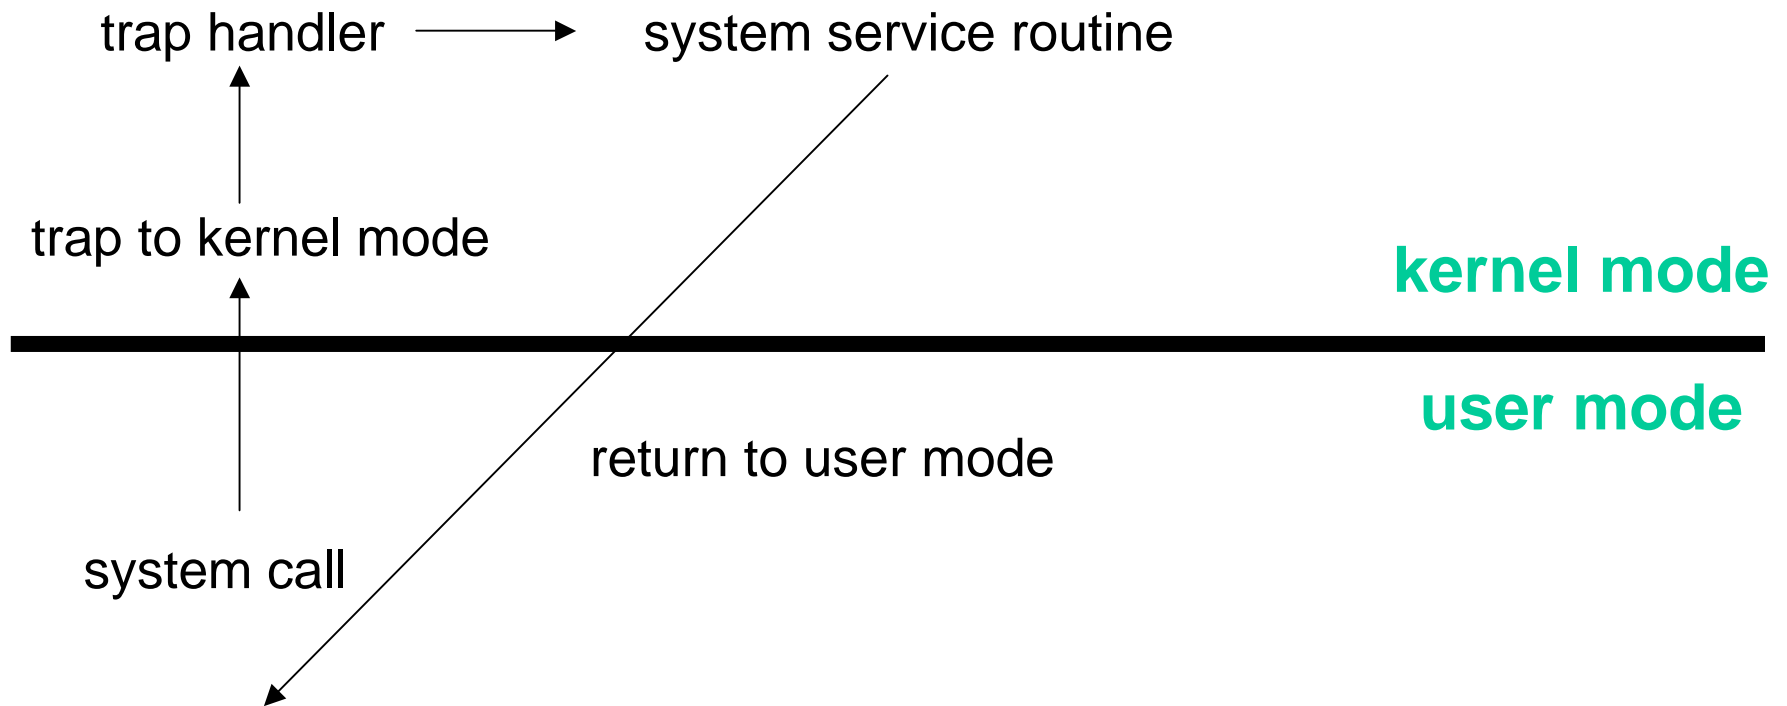
- n All I/O instructions are privileged instructions
- n Must ensure that a user program could never gain control of the computer in monitor mode (i.e., a user program that, as part of its execution, stores a new address in the interrupt vector)

Use of A System Call to Perform I/O



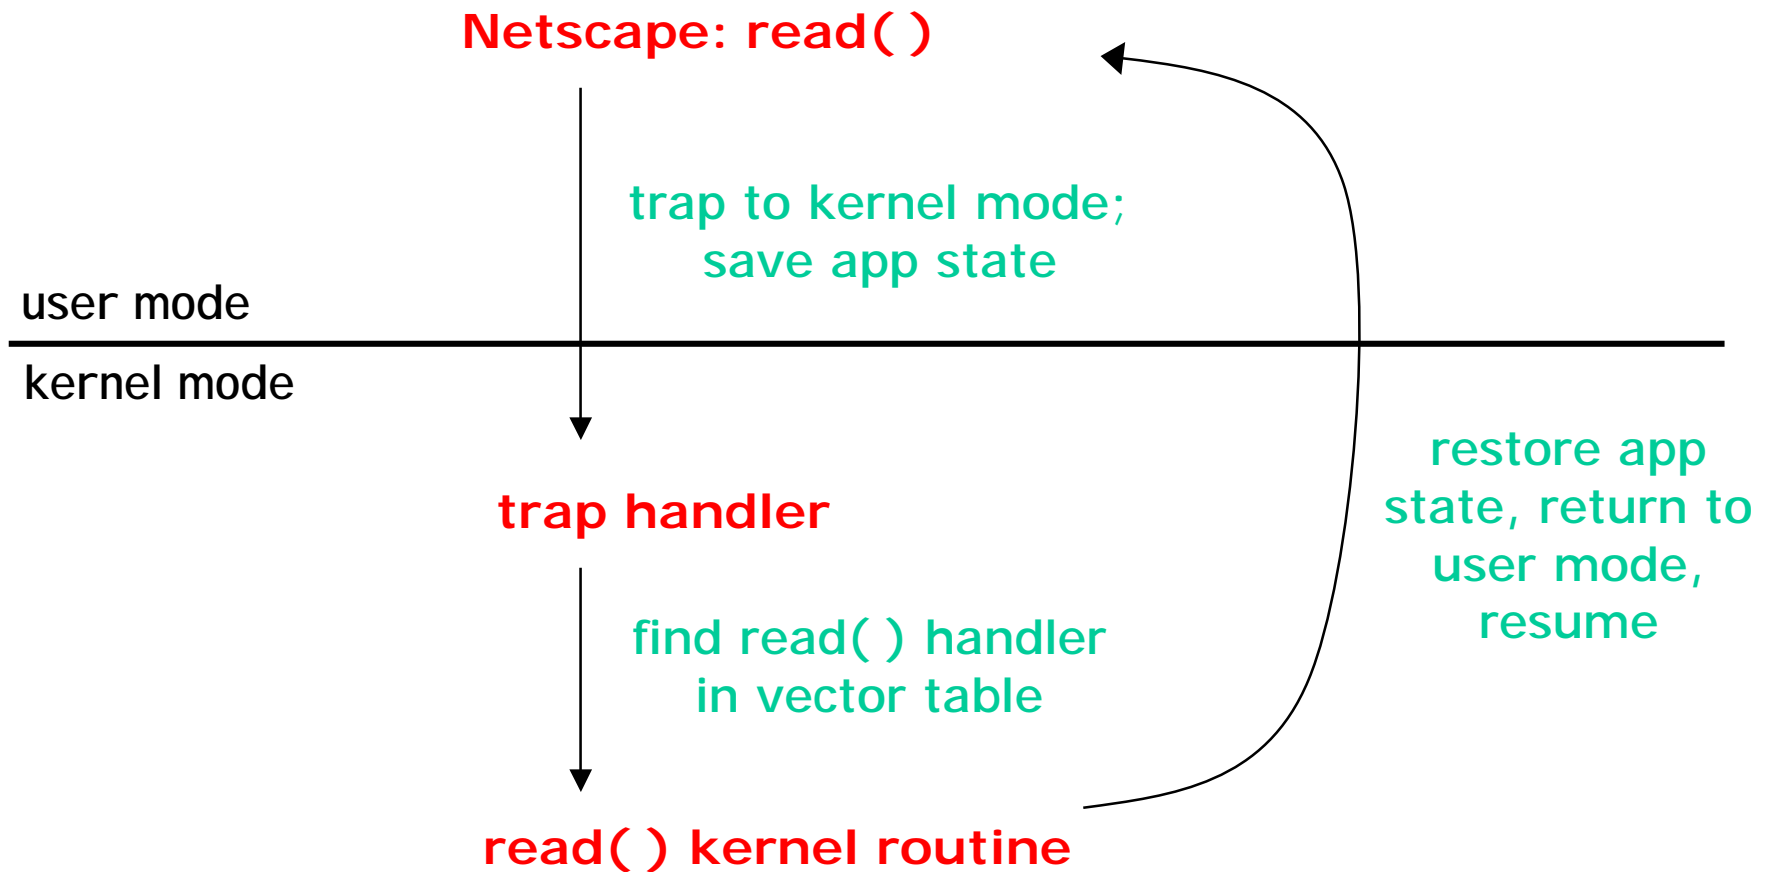
System Call

OS Kernel



User Programs

System Call (Cont'd)



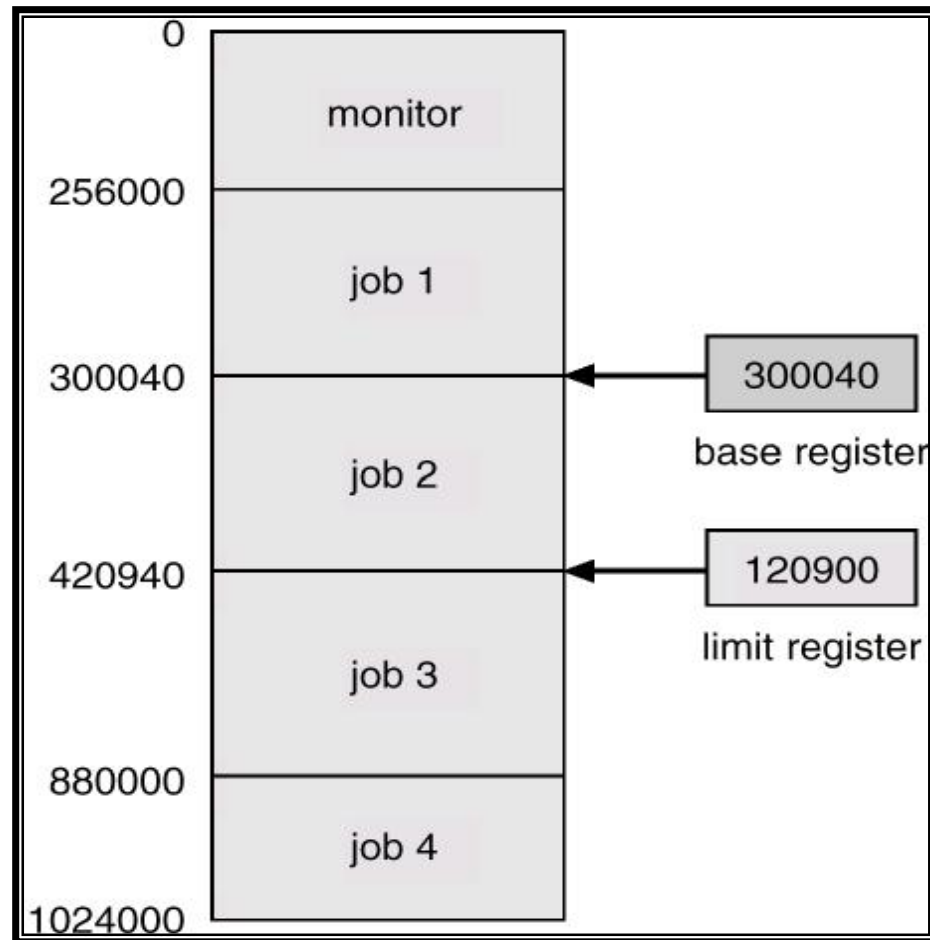
Memory Protection

- n Must provide memory protection at least for the interrupt vector and the interrupt service routines

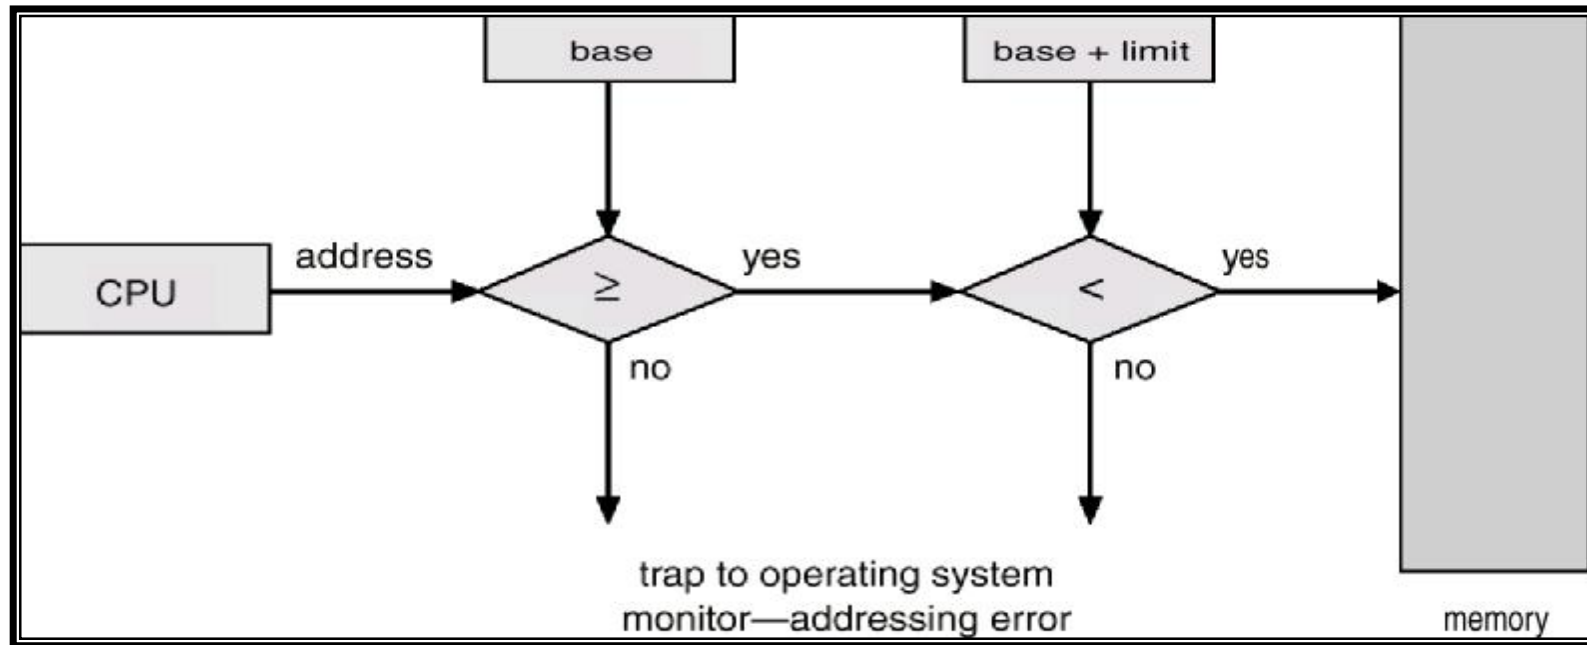
- n In order to have memory protection, add two registers that determine the range of legal addresses a program may access:
 - ü **Base register** – holds the smallest legal physical memory address
 - ü **Limit register** – contains the size of the range

- n Memory outside the defined range is protected

Use of A Base and Limit Register



Hardware Address Protection



Hardware Protection

- n When executing in monitor mode, the operating system has unrestricted access to both monitor and user's memory
- n The load instructions for the *base* and *limit* registers are privileged instructions

Memory Protection

n MMU (Memory Management Unit)

- ü Memory management hardware provides more sophisticated memory protection mechanisms
 - § base and limit registers
 - § page table pointers, page protection, TLBs
 - § virtual memory
 - § segmentation
- ü Manipulation of memory management hardware are protected (privileged) operations

CPU Protection

- n *Timer* – interrupts computer after specified period to ensure operating system maintains control
 - ü Timer is decremented every clock tick
 - ü When timer reaches the value 0, an interrupt occurs
- n Timer commonly used to implement time sharing
- n Time also used to compute the current time
- n Load-timer is a privileged instruction

Timers

- n How does the OS take control of CPU from the running programs?
 - ü Use a hardware timer that generates a periodic interrupt
 - ü The timer interrupt transfers control back to OS
 - ü The OS preloads the timer with a time to interrupt. (“quantum”)
 - § 10ms for Linux
 - § Cf) time slice or time quantum
 - ü The timer is privileged.
 - § Only the OS can load it

Synchronization

n Problems

- ü Interrupt can occur at any time and may interfere with the interrupted code
- ü OS must be able to synchronize concurrent processes

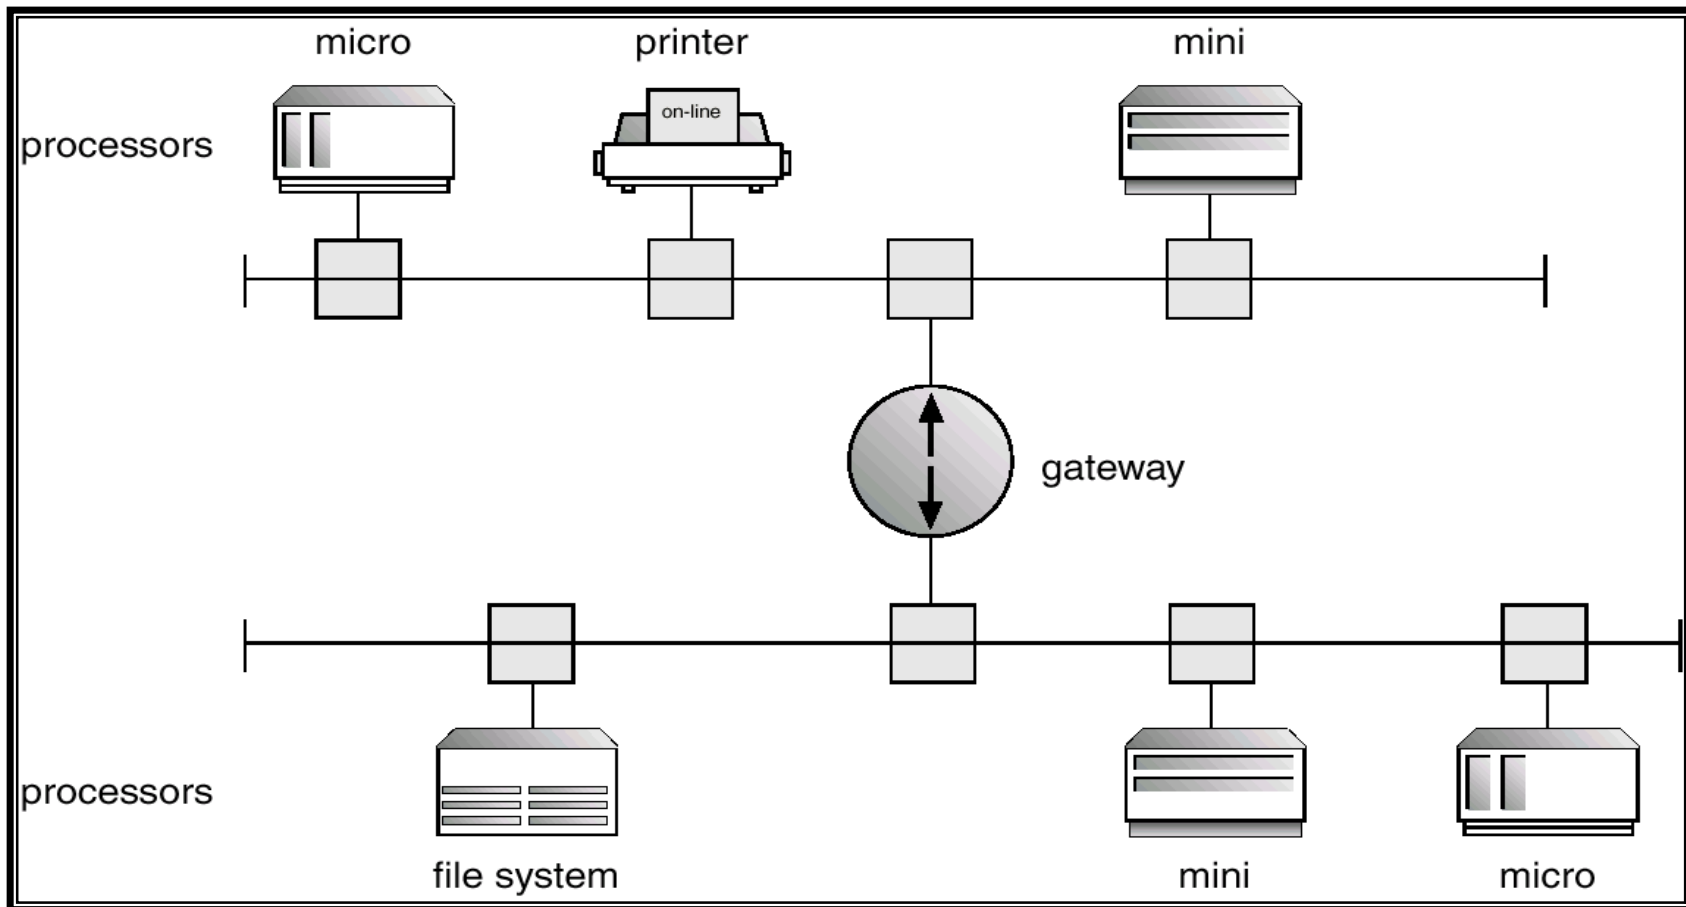
n Synchronization

- ü Turn off/on interrupts
- ü Use a special atomic instructions
 - § read-modify-write (e.g. INC, DEC)
 - § test-and-set

Network Structure

- n Local Area Networks (LAN)
 - ü Relay, bridge, gateway, router, switch, hub
- n Wide Area Networks (WAN)

Local Area Network Structure



Wide Area Network Structure

