

Chapter 13: I/O Systems - Objectives

- ▶ Explore the structure of an operating system's I/O subsystem
- ▶ Discuss the principles of I/O hardware and its complexity
- ▶ Provide details of the performance aspects of I/O hardware and software




I/O Hardware

- ▶ Incredible variety of I/O devices
- ▶ Common concepts
 - **Port**
 - **Bus (daisy chain or shared direct access)**
 - **Controller (host adapter)**
- ▶ I/O instructions control devices
- ▶ Devices have addresses, used by
 - Direct I/O instructions
 - **Memory-mapped I/O**



A Typical PC Bus Structure



Polling

- ▶ Determines state of device
 - command-ready
 - busy
 - Error
- ▶ **Busy-wait** cycle to wait for I/O from device

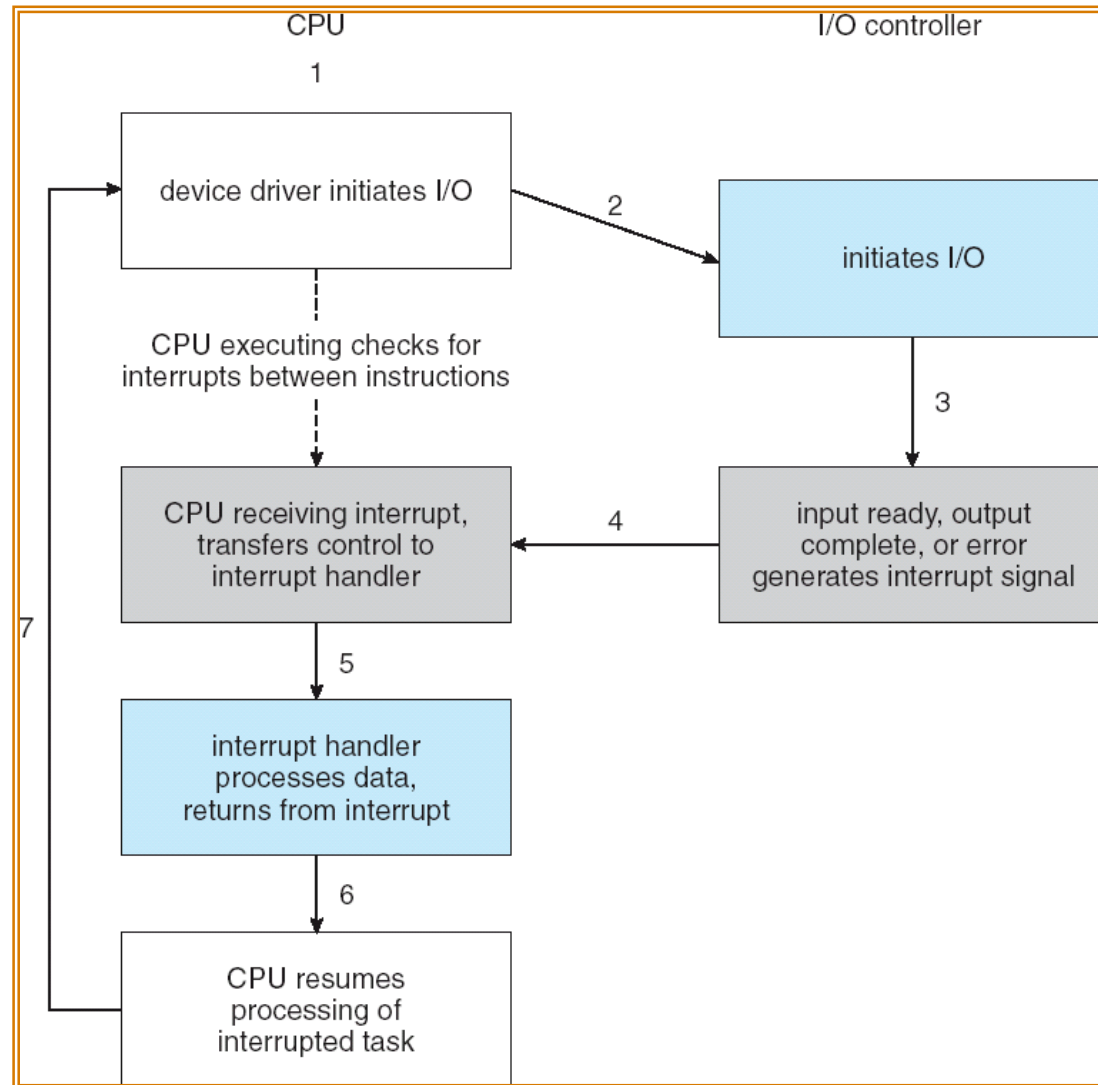


Interrupts

- ▶ **CPU Interrupt-request line** triggered by I/O device
- ▶ **Interrupt handler** receives interrupts
- ▶ **Maskable** to ignore or delay some interrupts
- ▶ Interrupt vector to dispatch interrupt to correct handler
 - Based on priority
 - Some **nonmaskable**
- ▶ Interrupt mechanism also used for exceptions



Interrupt-Driven I/O Cycle

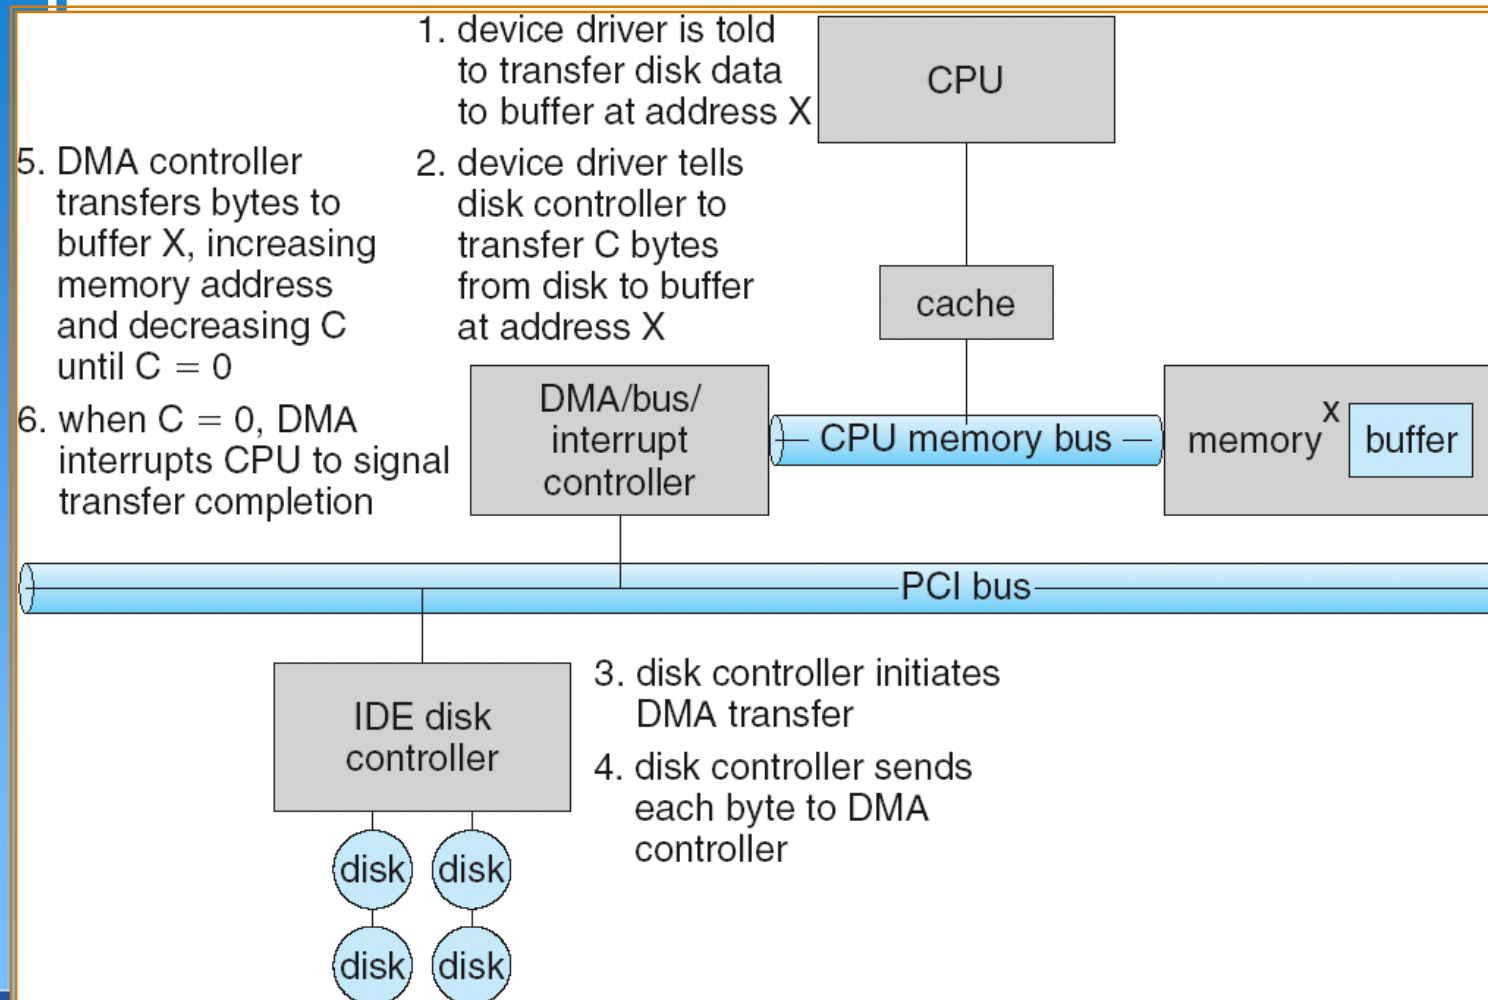


Direct Memory Access

- ▶ Used to avoid **programmed I/O** for large data movement
- ▶ Requires **DMA** controller
- ▶ Bypasses CPU to transfer data directly between I/O device and memory



Six Step Process to Perform DMA Transfer

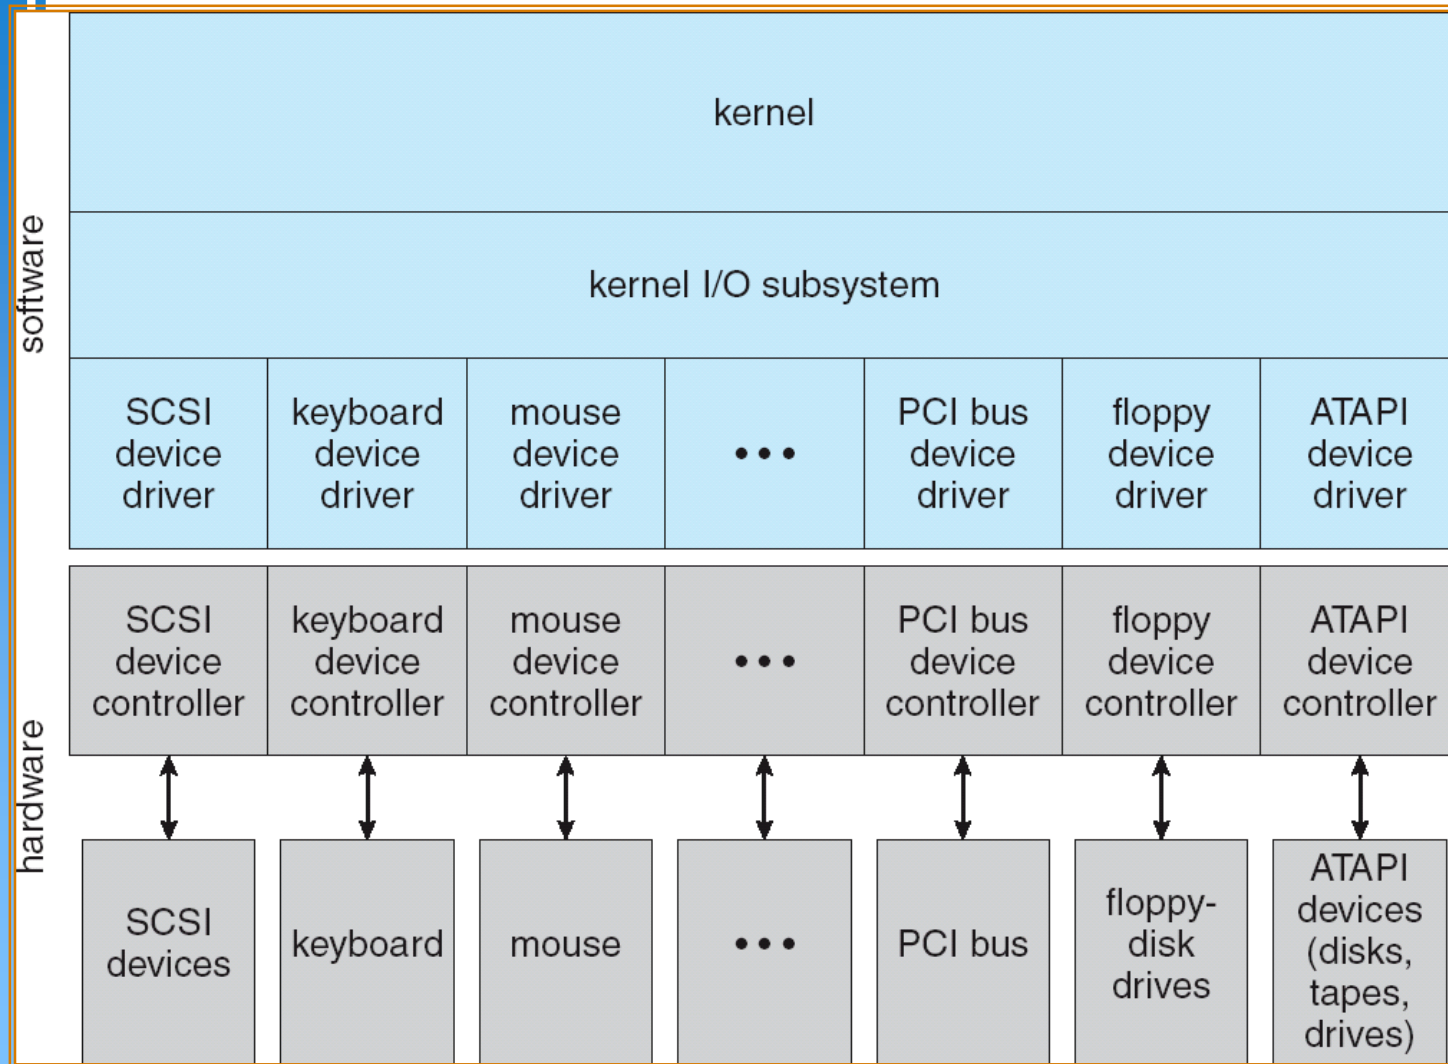


Application I/O Interface

- ▶ I/O system calls encapsulate device behaviors in generic classes
- ▶ Device-driver layer hides differences among I/O controllers from kernel
- ▶ Devices vary in many dimensions
 - **Character-stream or block**
 - **Sequential or random-access**
 - **Sharable or dedicated**
 - **Speed of operation**
 - **read-write, read only, or write only**



A Kernel I/O Structure



Characteristics of I/O Devices

aspect	variation	example
data-transfer mode	character block	terminal disk
access method	sequential random	modem CD-ROM
transfer schedule	synchronous asynchronous	tape keyboard
sharing	dedicated sharable	tape keyboard
device speed	latency seek time transfer rate delay between operations	
I/O direction	read only write only read–write	CD-ROM graphics controller disk



Block and Character Devices

- ▶ Block devices include disk drives
 - Commands include read, write, seek
 - Raw I/O or file-system access
 - Memory-mapped file access possible
- ▶ Character devices include keyboards, mice, serial ports
 - Commands include `get`, `put`
 - Libraries layered on top allow line editing

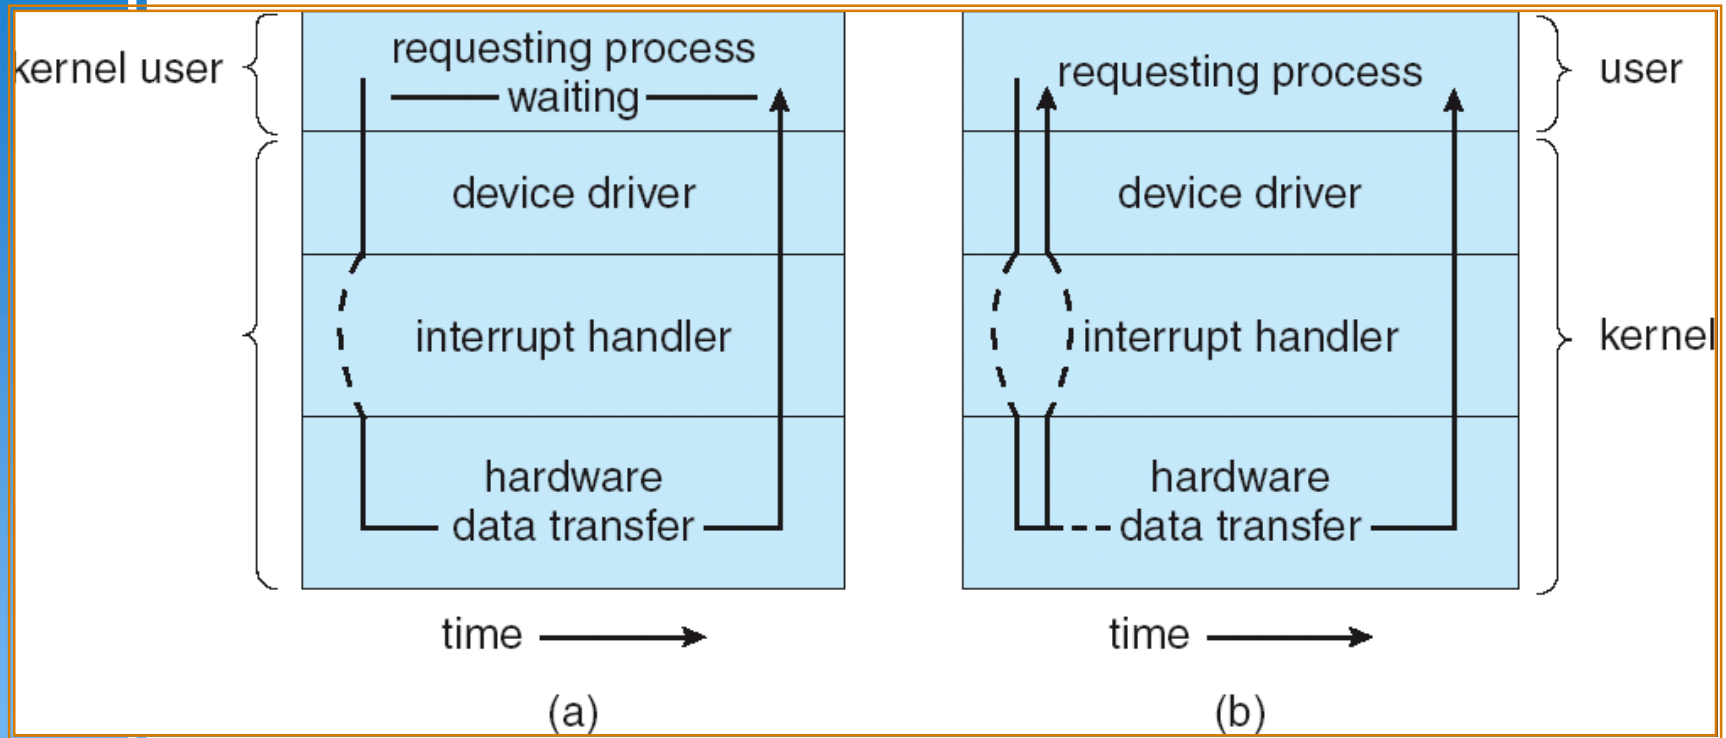


Blocking and Nonblocking I/O

- ▶ **Blocking** - process suspended until I/O completed
 - Easy to use and understand
 - Insufficient for some needs
- ▶ **Nonblocking** - I/O call returns as much as available
 - User interface, data copy (buffered I/O)
 - Implemented via multi-threading
 - Returns quickly with count of bytes read or written
- ▶ **Asynchronous** - process runs while I/O executes
 - Difficult to use
 - I/O subsystem signals process when I/O completed



Two I/O Methods



Synchronous

Asynchronous



Kernel I/O Subsystem

- ▶ Scheduling
 - Some I/O request ordering via per-device queue
 - Some OSs try fairness
- ▶ Buffering - store data in memory while transferring between devices
 - To cope with device speed mismatch
 - To cope with device transfer size mismatch
 - To maintain “copy semantics”



Kernel I/O Subsystem

- ▶ **Caching** - fast memory holding copy of data
 - Always just a copy
 - Key to performance
- ▶ **Spooling** - hold output for a device
 - If device can serve only one request at a time
 - i.e., Printing
- ▶ **Device reservation** - provides exclusive access to a device
 - System calls for allocation and deallocation
 - Watch out for deadlock



Error Handling

- ▶ OS can recover from disk read, device unavailable, transient write failures
- ▶ Most return an error number or code when I/O request fails
- ▶ System error logs hold problem reports



I/O Protection

- ▶ User process may accidentally or purposefully attempt to disrupt normal operation via illegal I/O instructions
 - All I/O instructions defined to be privileged
 - I/O must be performed via system calls
 - Memory-mapped and I/O port memory locations must be protected too



Kernel Data Structures

- ▶ Kernel keeps state info for I/O components, including open file tables, network connections, character device state
- ▶ Many, many complex data structures to track buffers, memory allocation, “dirty” blocks
- ▶ Some use object-oriented methods and message passing to implement I/O



I/O Requests to Hardware Operations

- ▶ Consider reading a file from disk for a process:
 - Determine device holding file
 - Translate name to device representation
 - Physically read data from disk into buffer
 - Make data available to requesting process
 - Return control to process



Improving Performance

- ▶ Reduce number of context switches
- ▶ Reduce data copying
- ▶ Reduce interrupts by using large transfers, smart controllers, polling
- ▶ Use DMA
- ▶ Balance CPU, memory, bus, and I/O performance for highest throughput

