

# CSC501

# Operating Systems Principles

---

## I/O Devices

# Previous Lectures

---

q Abstraction

Q Process

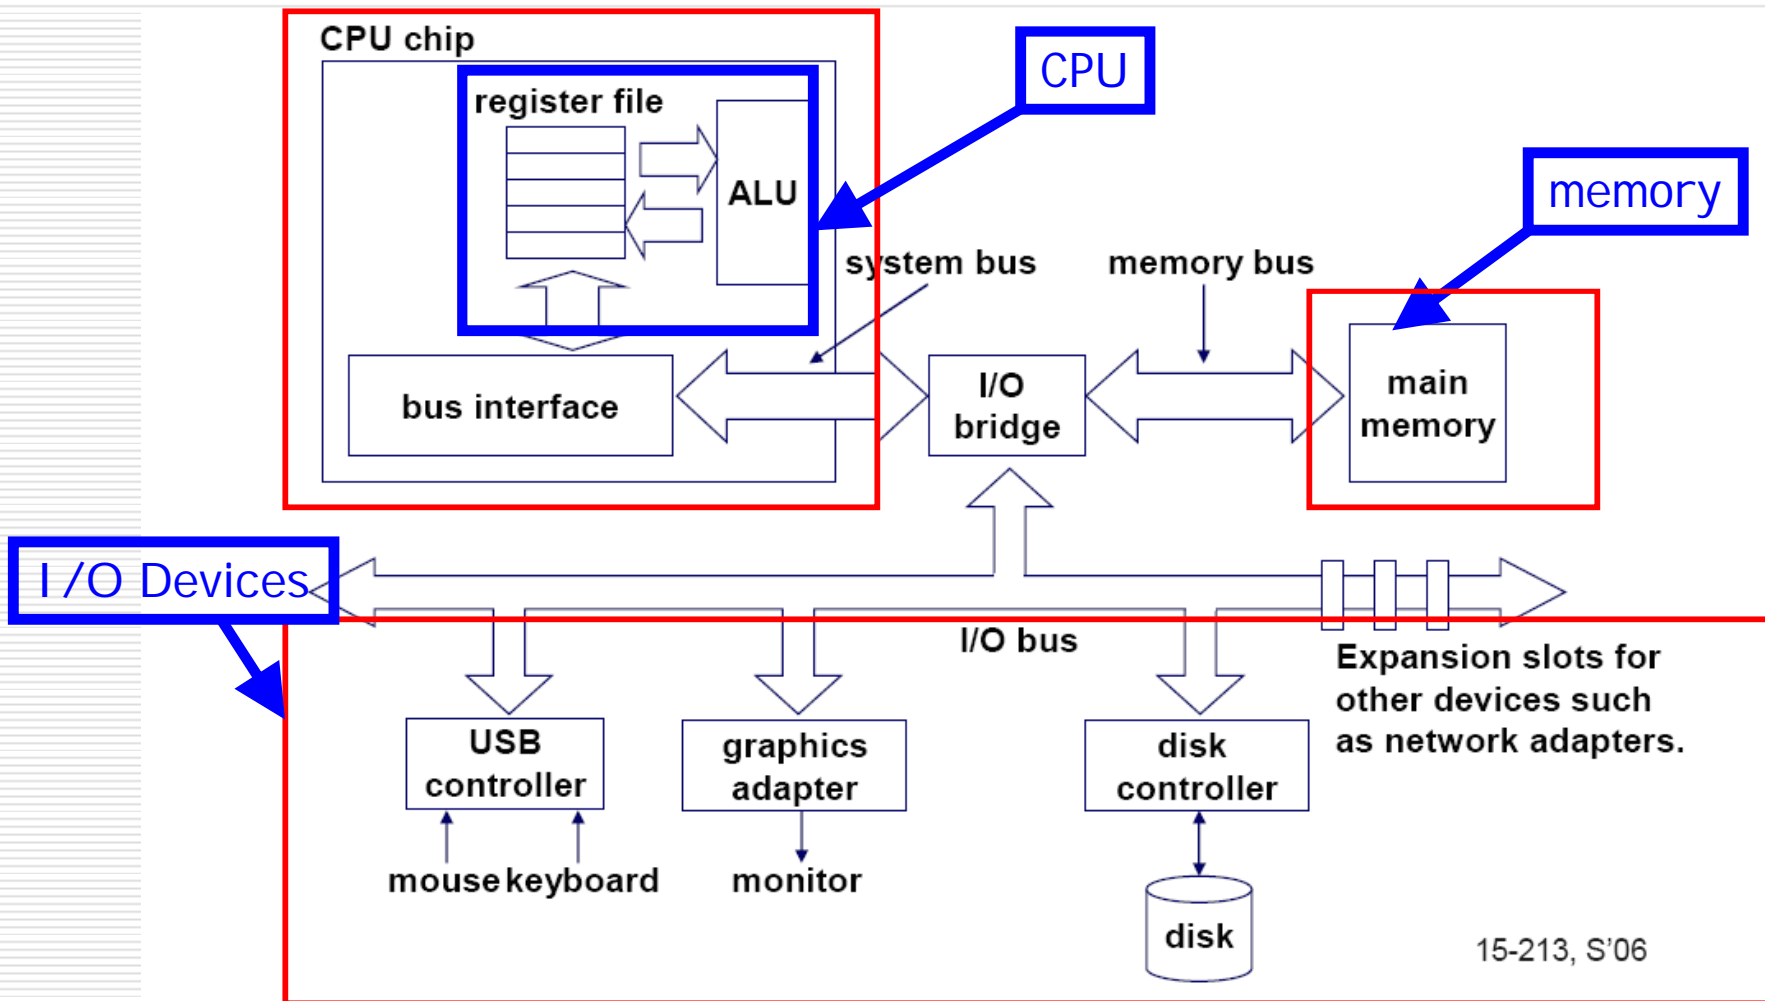
Q Memory

q Now

Q I/O Devices

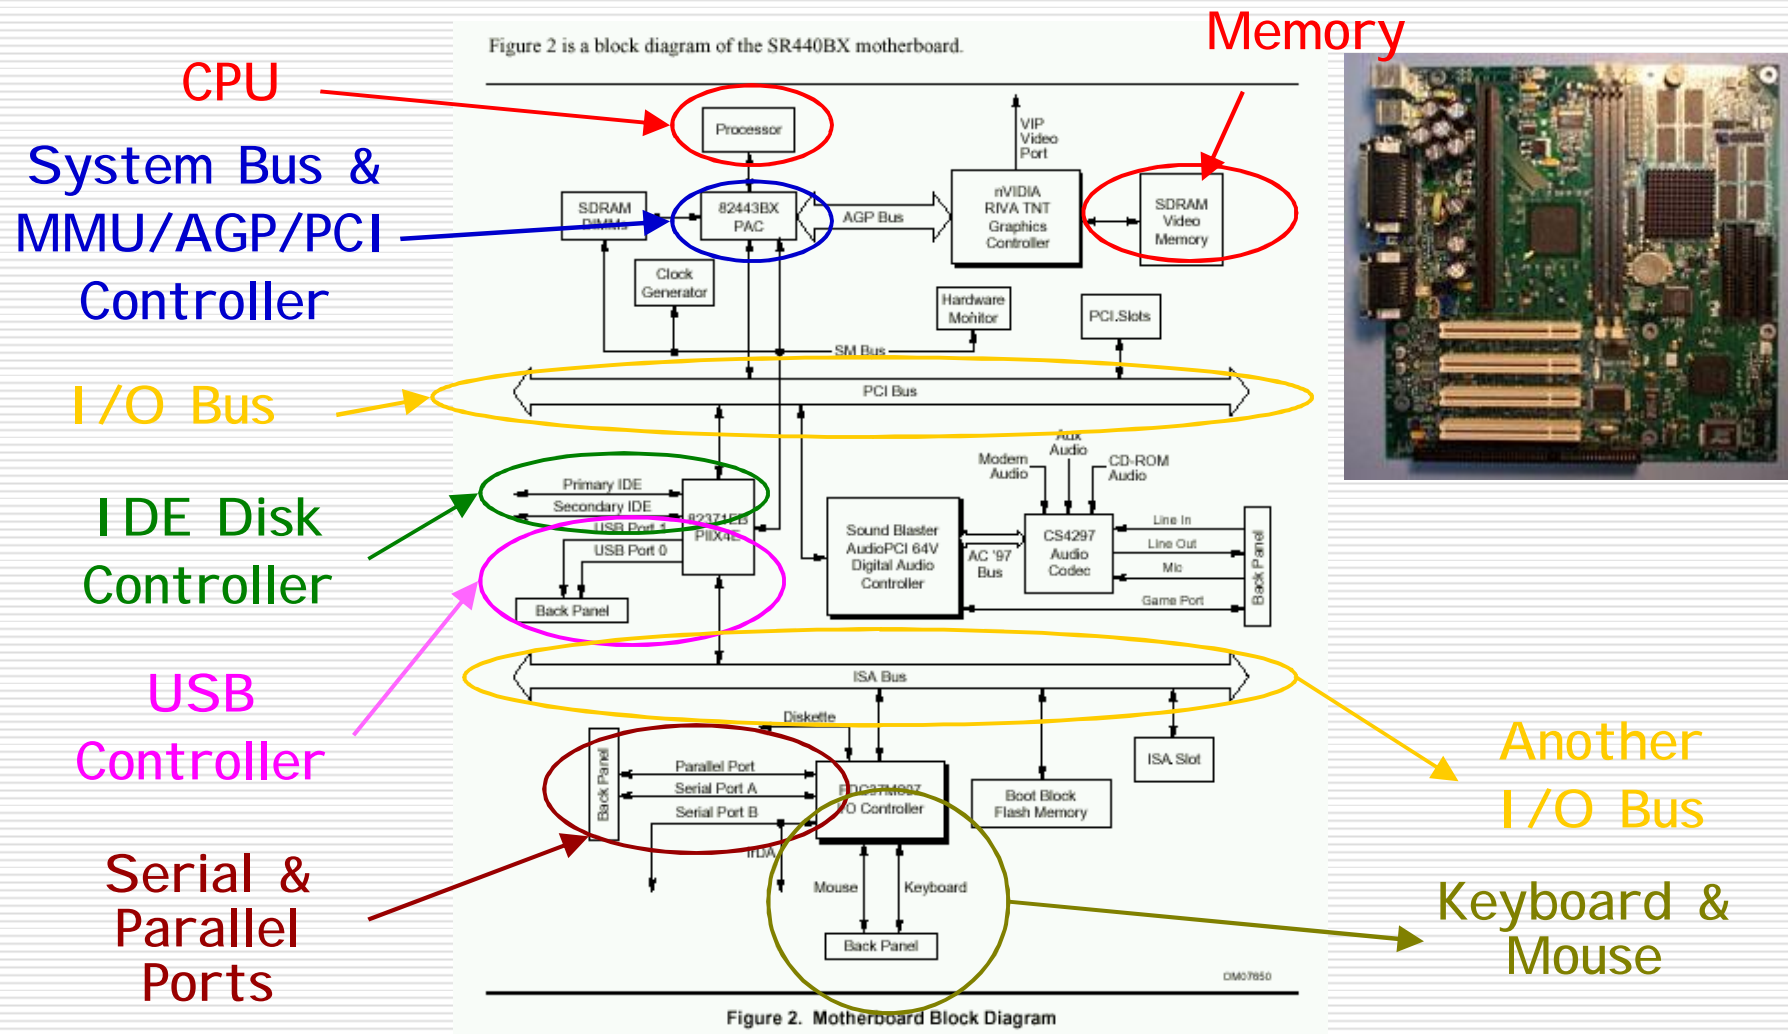
---

# Basic Computer Architecture



# Intel SR440BX Motherboard

Figure 2 is a block diagram of the SR440BX motherboard.



## OS Designer's Perspective

---

- q The largest, most complex subsystem in OS
  - q Most lines of code with the highest rate of code changes
    - Q Where OS engineers most likely to work
    - Q Difficult to test thoroughly
    - Q Big impact on performance and perception
    - Q Bigger impact on acceptability in market
-

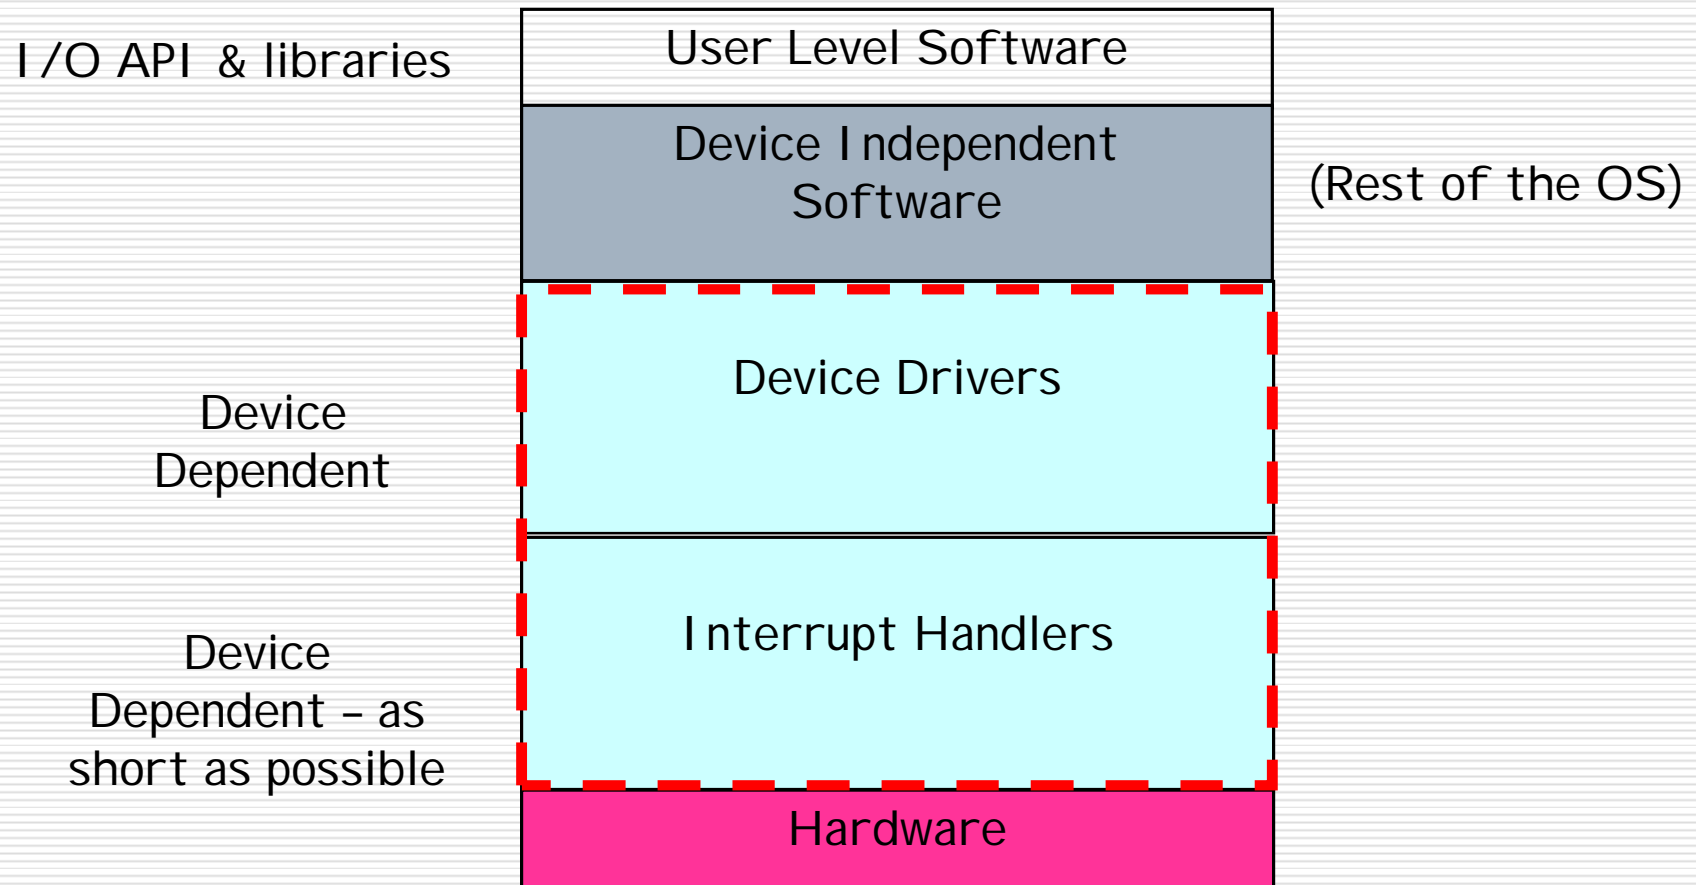
# Principles of I/O Software

---

- q **Efficiency** - Do not allow I/O operations to become system bottleneck, especially for slow devices
  
- q **Generality** - Provide a common API to access I/O devices
  - Q Uniform naming/interface
    - v open, close read, write...
  - Q Device independence
    - v isolate OS and application programs from device specific details and peculiarities
    - v isolate the impact of device errors, retry where possible
      - n Errors are abundant in I/O

# I/O Software "Stack"

---



# Communication with I/O Devices

---

- q Programmed I/O (a.k.a., polling)
- q Interrupt-Driven I/O
- q DMA-based I/O



## Programmed I/O (Polling)

---

**q** Used when device and controller are relatively **quick** to process an I/O operation

**Q** Device driver

- ✓ Gains access to device
- ✓ Initiates I/O operation
- ✓ Loops testing for completion of I/O operation
- ✓ If there are more I/O operations, repeat

**Any Example?**

**Q** Used in following kinds of cases

- ✓ Service interrupt time > Device response time
  - ✓ Device has no interrupt capability
  - ✓ Embedded systems where CPU has nothing else to do
-

# Interrupt-Driven I/O

---

q Interrupts occur on I/O events

- Q Operation completion
- Q Error or change of status

Any Example?

q Interrupt handling

- Q Step 1: Save the current machine state
- Q Step 2: Load the machine state for interrupt handling
- Q Step 3: Invoke the corresponding ISR
- Q Step 4: Resume the program execution

**Question:**

Where do we save the current machine states?

---

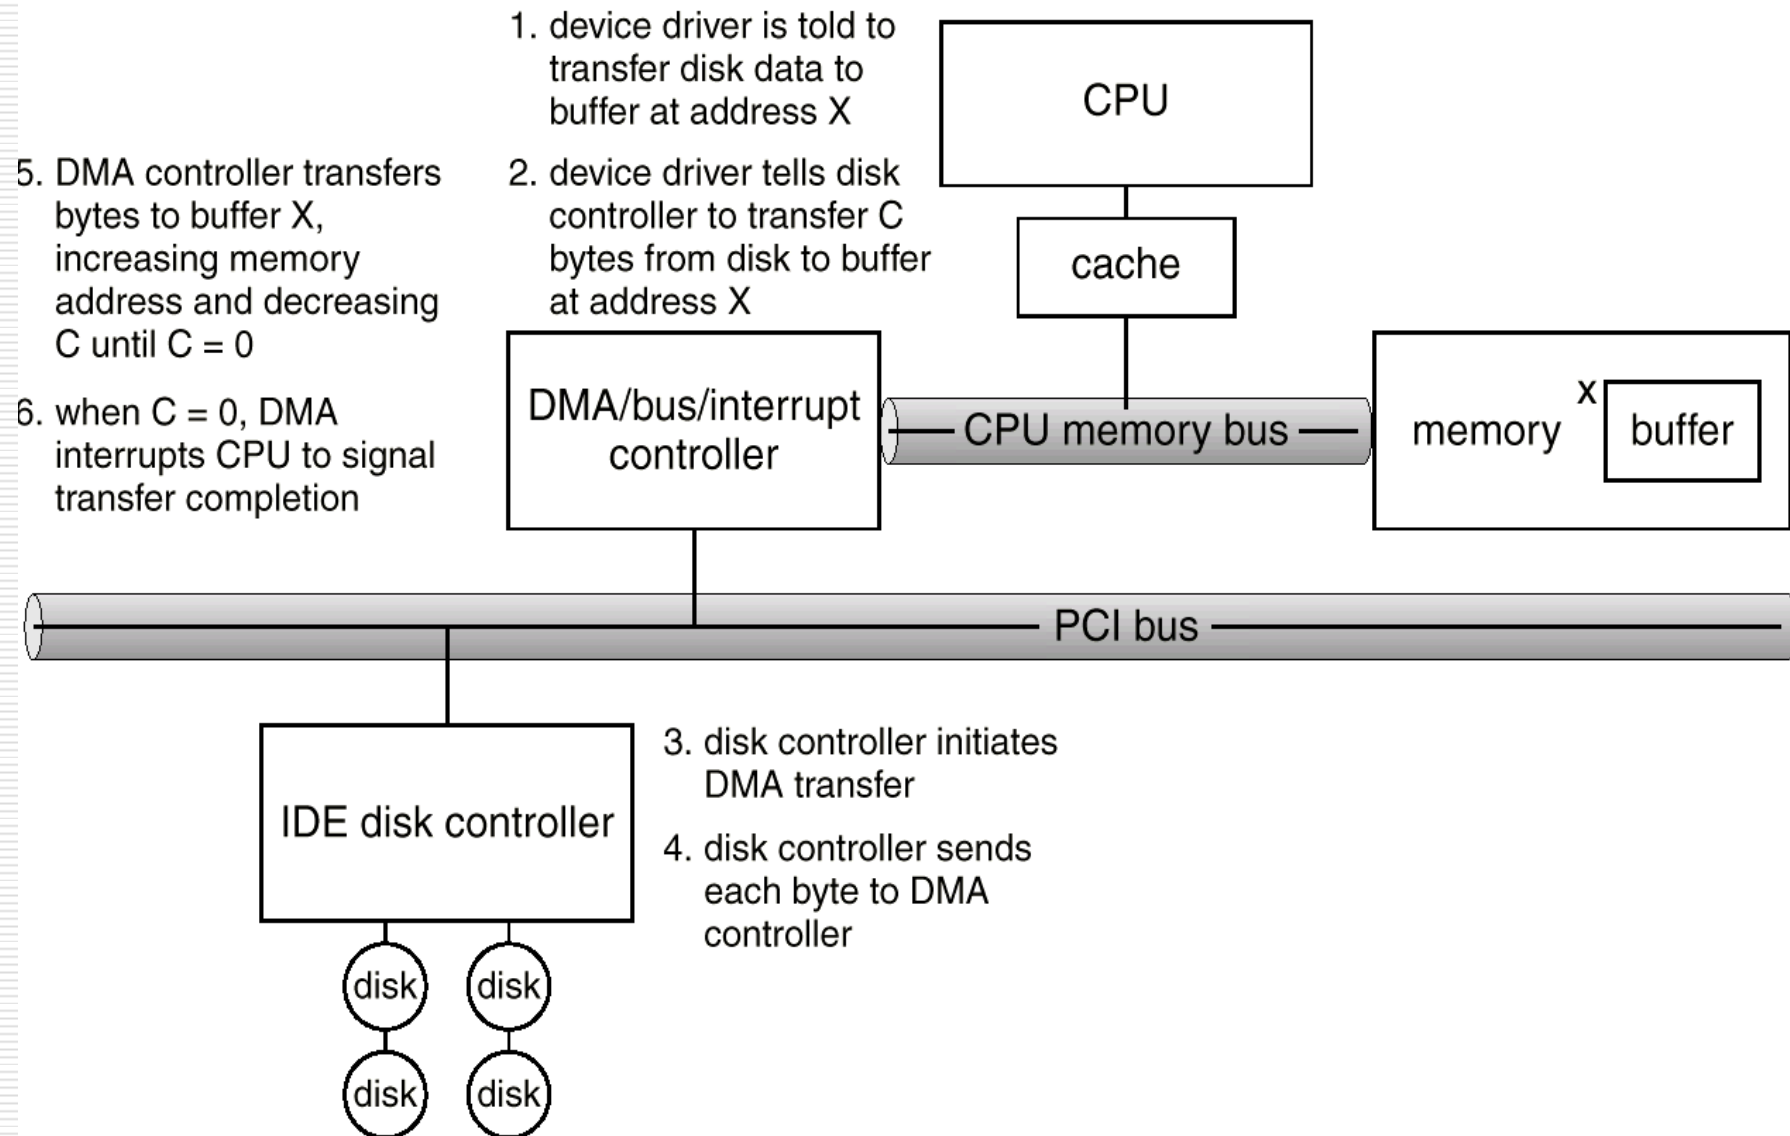
# DMA-Based I/O

---

- q Direct Memory Access
  - Q Device read/write directly from/to memory
  - Q Transfer from memory to device typically initiated from CPU
  - Q Transfer from device to memory can be initiated by the device or the CPU

Any Example?

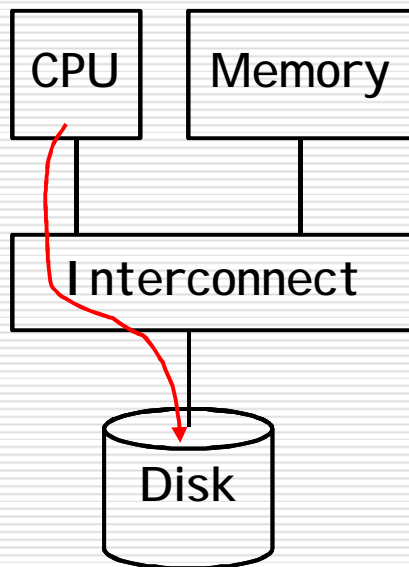
# DMA-Based I/O



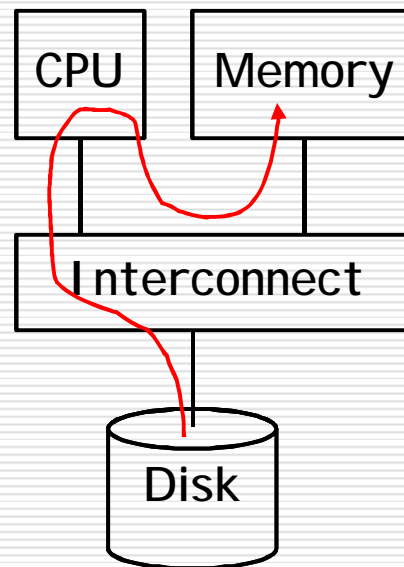
# Communication with I/O Devices

---


- q Comparison: Polling, Interrupt, and DMA
  - Q Polling and Interrupt = word at a time
  - Q DMA = block at a time



Polling



Interrupt



DMA

---

**Question:**

How should the device driver be designed?

# Device Driver

---

- q Hides the device specifics from the above layers
    - Q Supporting a common API
  - q Translates logical I/O into device I/O
    - Q E.g., logical disk blocks into {head, track, sector}
    - Q Performs data buffering and scheduling of I/O operations
  - q Structure
    - Q Several **synchronous** entry points: device initialization, queue I/O requests, state control, read/write
    - Q An **asynchronous** entry point to handle interrupts
-

# Device Driver – A Typical Design

---

- q OS and driver communication: **OS-dependent**
  - Q Commands and data between OS and drivers
    - v Meet the interface specs of the device independent layer
    - v Utilize the facilities supplied by the OS – buffers, error codes, etc.
  
- q Driver and hardware comm.: **Device-dependent**
  - Q Commands and data between driver and hardware
    - v Actions during interrupt handling

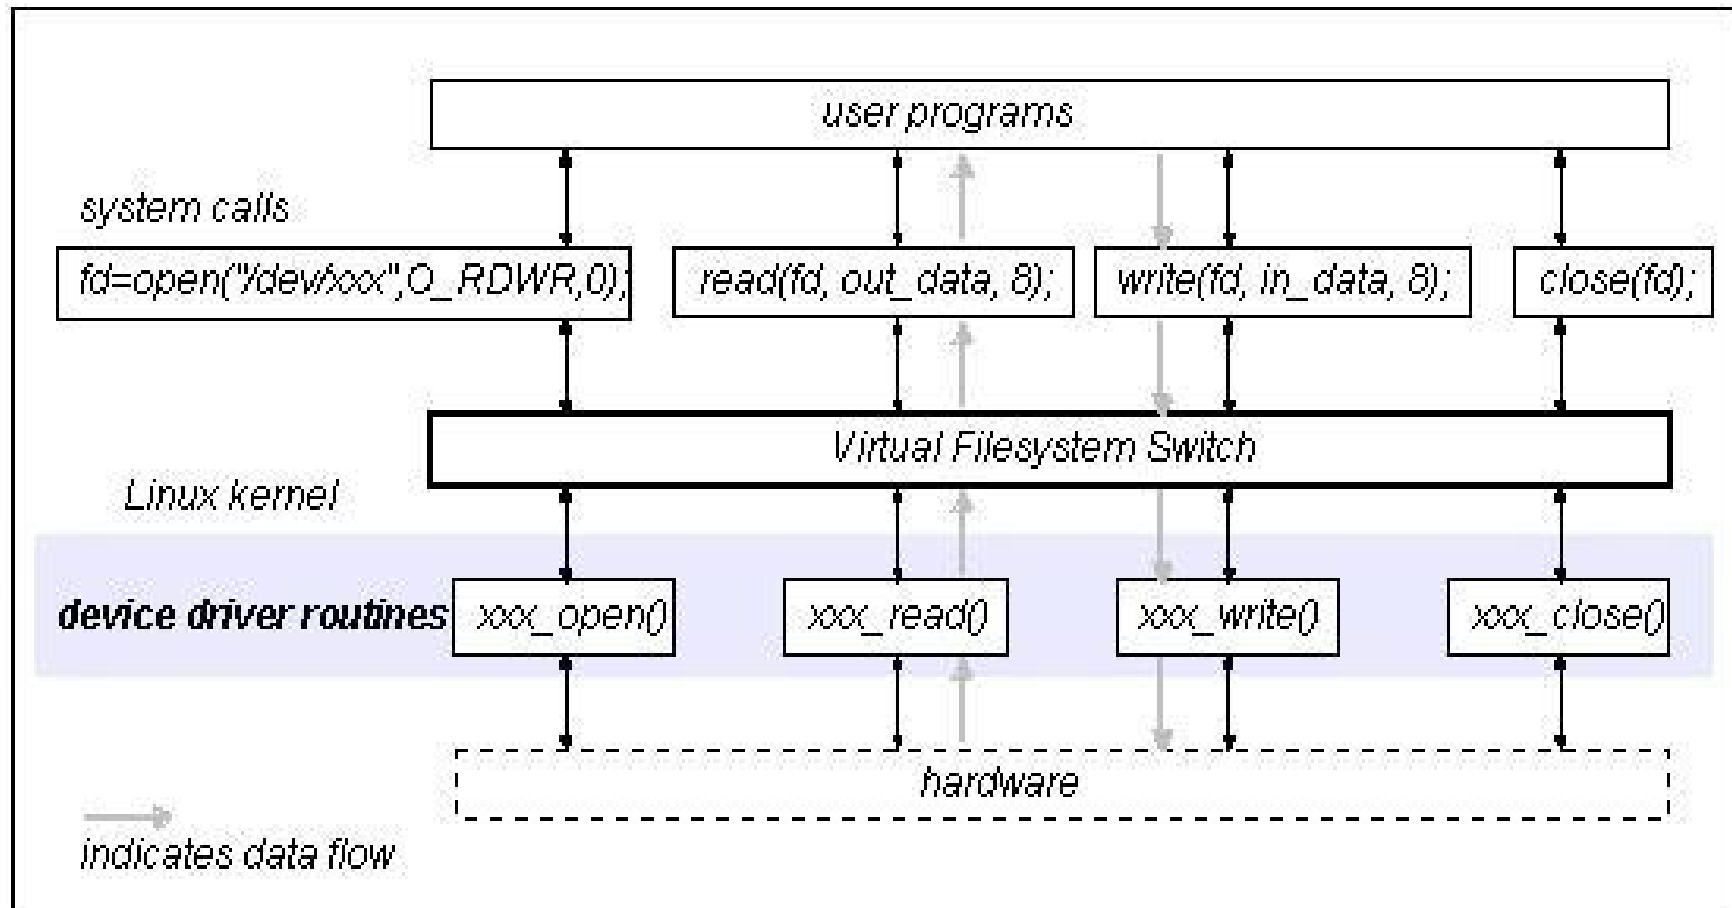


# Device Driver – A Typical Design

---

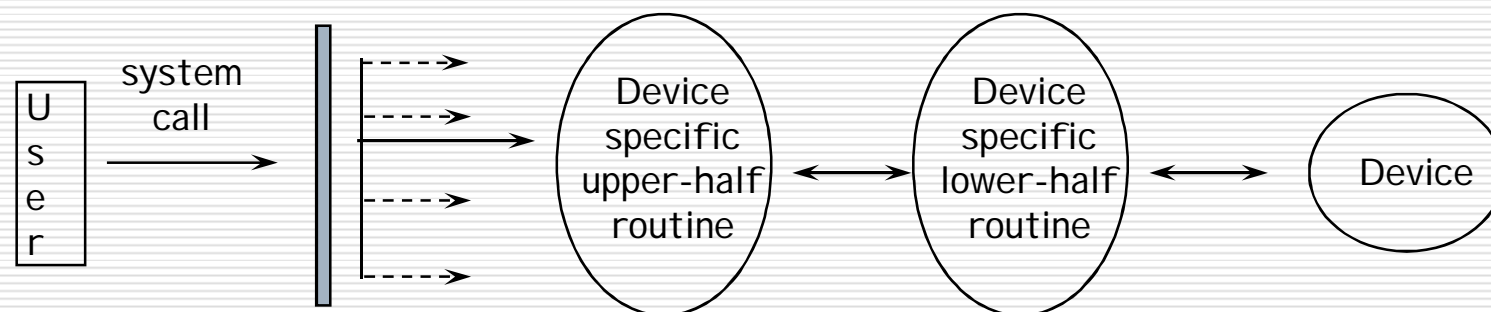
- q Driver operations
  - Q Initialize devices
  - Q Interpret commands from OS
  - Q Schedule multiple outstanding requests
  - Q Manage data transfer
  - Q Accept and process interrupts
  - Q ...
- q Typically use **open-close-read-write** paradigm
  - Q Support *open, close, read, write, seek* functions
  - Q Use *ioctl* function as escape mechanism

# Uniform API (e.g., in Linux)



## Case Study w/ Xinu

- q Code divided into two parts
  - Q Upper Half
    - ✓ Defines abstract operations
    - ✓ Executed by application to request I/O
  - Q Lower Half
    - ✓ Device-specific interrupt routine
    - ✓ Executed by arbitrary (current) process
    - ✓ Invoked when operation completes



Device switch table devtab[ ]

# Case Study w/ Xinu


- q Device switch table
  - Q One row per device
  - Q One column per operation
  - Q Entry points to a procedure

*device* ↓

*operation* →

	open	read	write	
CONSOLE	&ttyopen	&ttyread	&ttywrite	
SERIAL0	&ionull	&comread	&comwrite	
SERIAL1	&ionull	&comread	&comwrite	...
ETHER	&ethopen	&ethread	&ethwrite	

⋮



Device Switch Table

```

int      dvnum;
long int (*dvinit)();
long int (*dvshutdown)();
long int (*dvopen)();
long int (*dvclose)();
long int (*dvread)();
long int (*dvwwrite)();
long int (*dvseek)();
long int (*dvgetc)();
long int (*dvputc)();
long int (*dvcntl)();
long int (*dviint)();
long int (*dvoint)();
long int dvcsr;
long int dvivec;
long int dvovec;
char     *dvioblk;
int      dvminor;

```

Ptrs to device specific high-level routines

Ptr to device control block

device switch table entry ( struct devsw)

The structure of the control block varies from device to device. The control block contains information associated with the device.

## Case Study w/ Xinu

---

- q Device switch table
  - Q Device-independent operations “generic”
  - Q Given operation may not apply to given device
    - v Example: *seek on keyboard or network*
  - Q All entries in device switch table must be valid

**Any solution?**

## Case Study w/ Xinu

---

- q Device switch table
  - Q Device-independent operations “generic”
  - Q Given operation may not apply to given device
    - v Example: *seek on keyboard or network*
  - Q All entries in device switch table must be valid
- q Special Entries
  - Q *ionull*
    - v Used for innocuous operation, always returns *OK*
  - Q *ioerr*
    - v Used for incorrect operation, always returns *SYSERR*

## Example Device Driver *tty* in Xinu

---

- q Used for I /O to console
- q Separate input and output buffers
- q Use semaphores to synchronize
  
- q Upper-half routines
  - Q `ttyinit`, `ttyopen`, `ttyclose`, `ttyread`, `ttywrite`, `ttyputc`, `ttygetc`, `ttypcntl`
- q Lower-half routines
  - Q `ttyiin` (input interrupt handler)
  - Q `ttyoin` (output interrupt handler)



## Example Device Driver *tty* in Xinu

---

- q Actions during a character output
  - Q Upper-half
    - v Waits for buffer space
    - v Deposits character
    - v Triggers device
  - Q Lower-half
    - v Extracts character and pass to device
    - v Signals space semaphore

# Next Lecture

---

- q Hard Disk
  - Q Scheduling and RAID