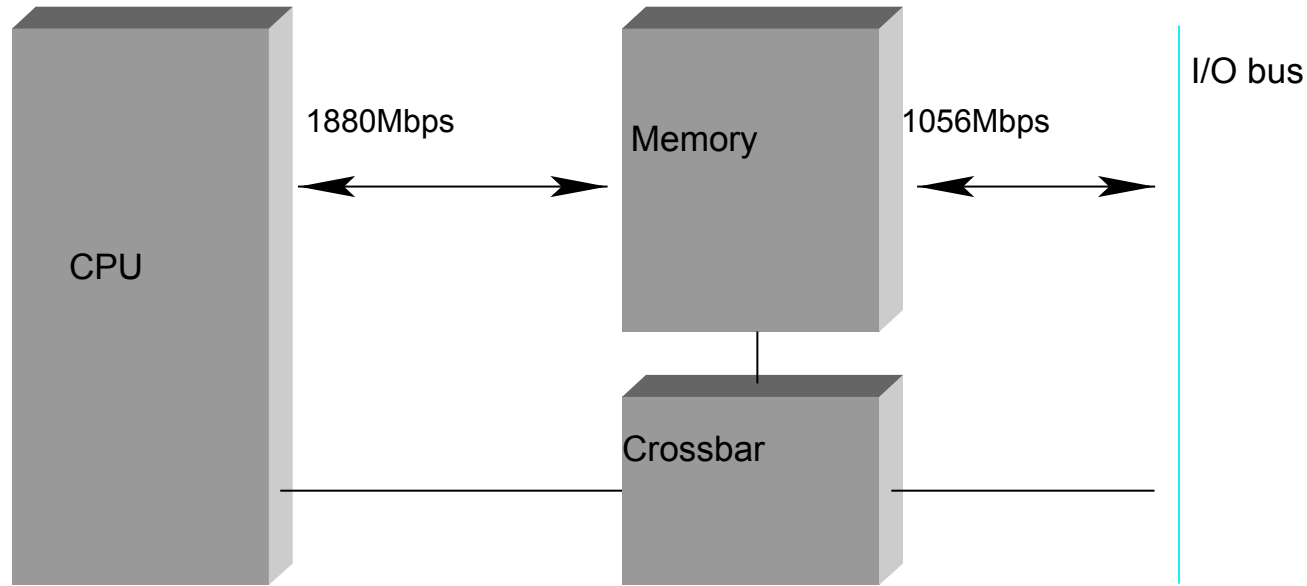


11 I/O Systems

chapter 12 of the book

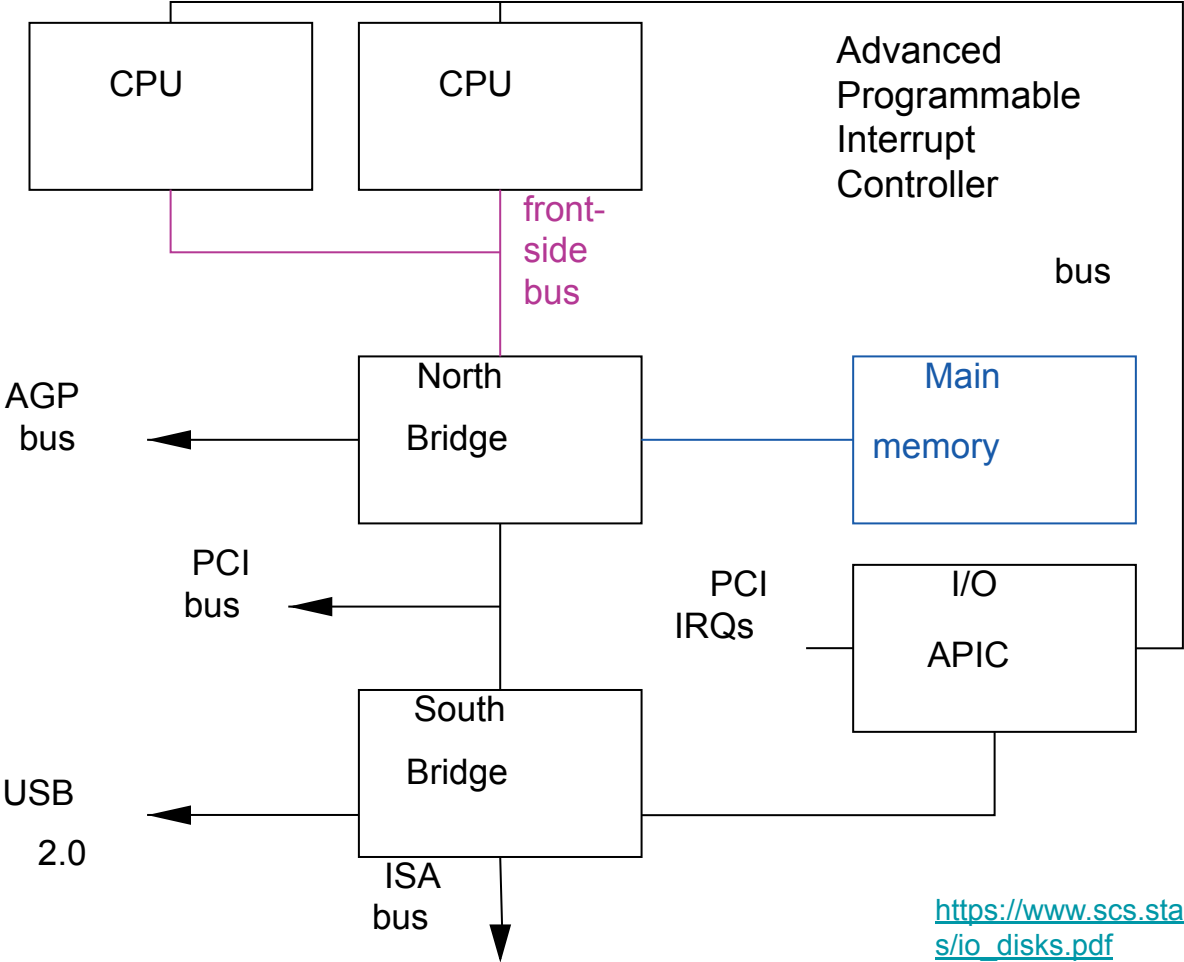
- Overview
- I/O Hardware
- Application I/O Interface
- Kernel I/O Subsystem
- Transforming I/O Requests to Hardware Operations
- STREAMS
- Performance

Memory and I/O buses



- CPU accesses physical memory over a bus
- Devices access memory over I/O bus with DMA
- Devices can appear to be a region of memory

Realistic ~2005 PC architecture



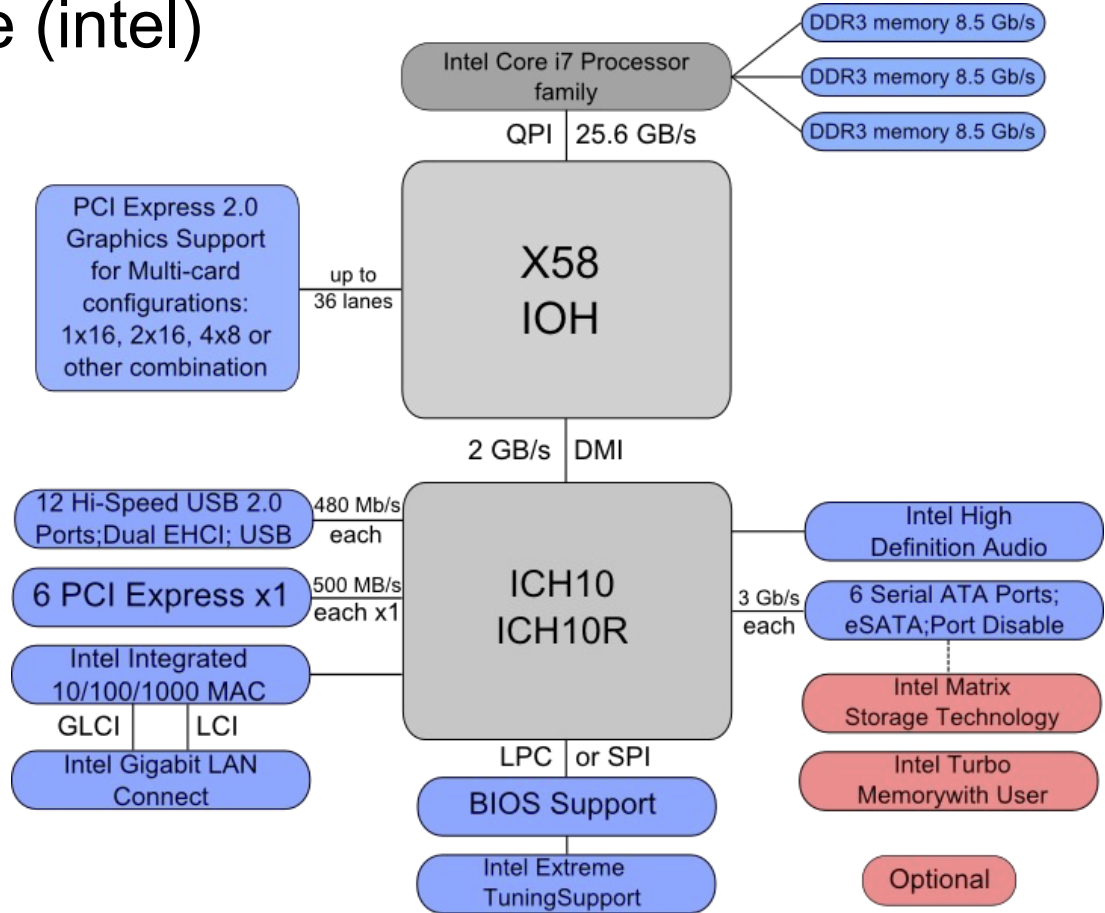
https://www.scs.stanford.edu/24wi-cs212/notes/io_disks.pdf

Modern PC architecture (intel)

IOH-I/O hub

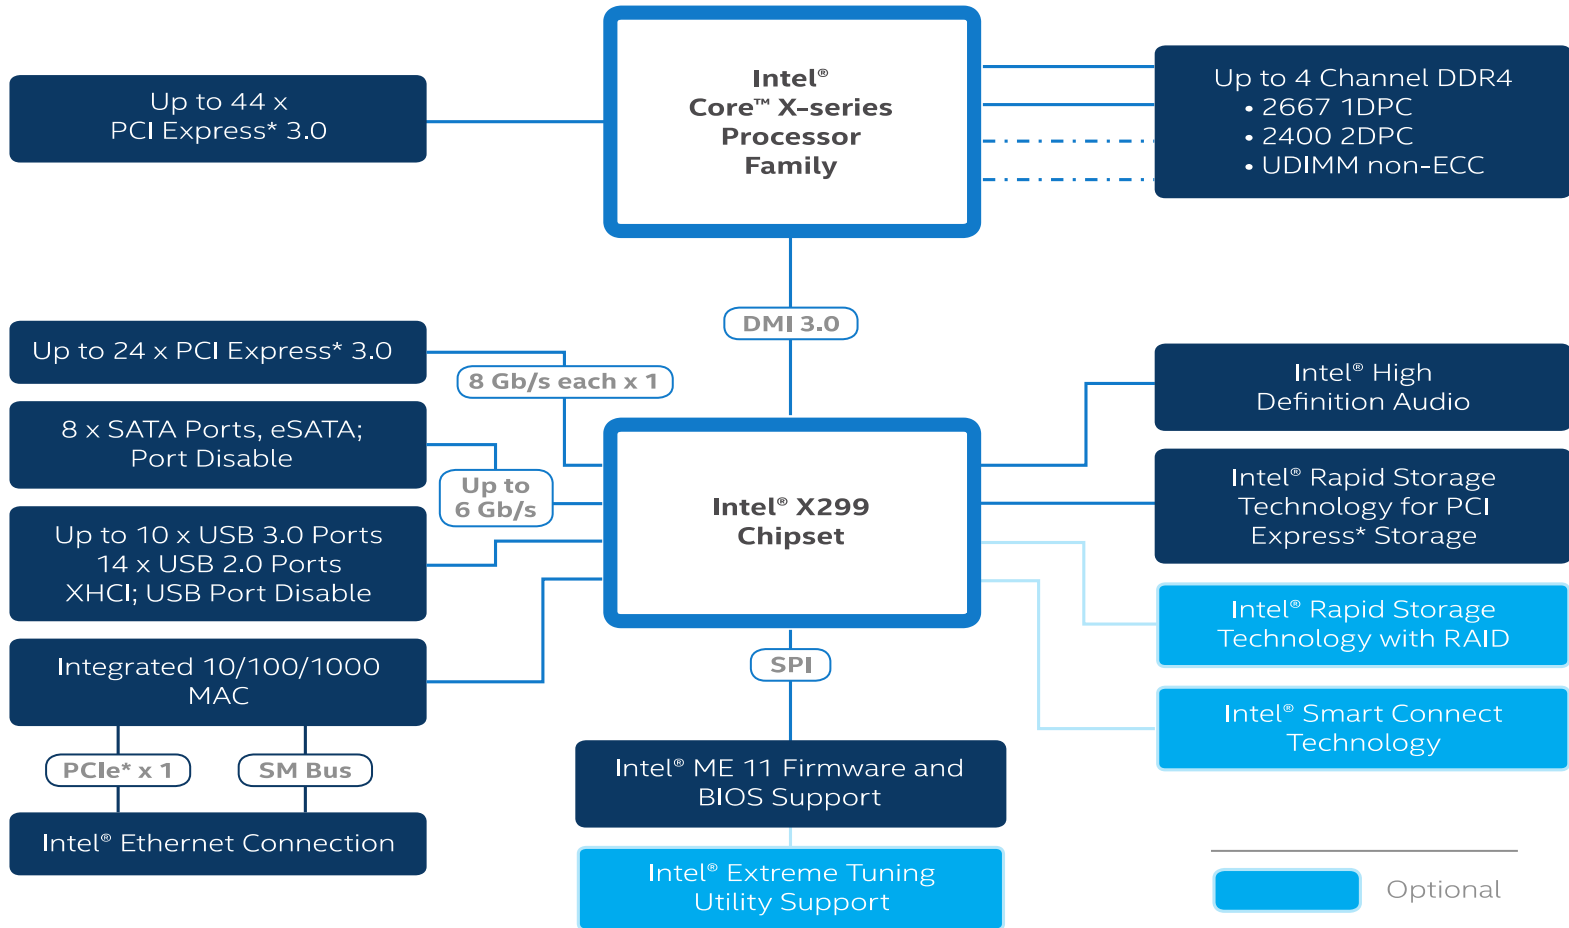
ICH = I/O controller Hub

QPI = [Intel QuickPath Interconnect](https://www.intel.com/technology/qpi/) - Wikipedia




https://www.scs.stanford.edu/24wi-cs212/notes/io_disks.pdf
https://en.wikipedia.org/wiki/Intel_X58

CPU now entirely subsumes IOH [intel]



AMD EPYC is essentially an SoC



4094 pins: both memory controller and 128 lanes PCIe
directly on chip!

What is memory

SRAM – Static RAM

- Like two NOT gates circularly wired input-to-output
- 4–6 transistors per bit, actively holds its value
- Very fast, used to cache slower memory

DRAM – Dynamic RAM

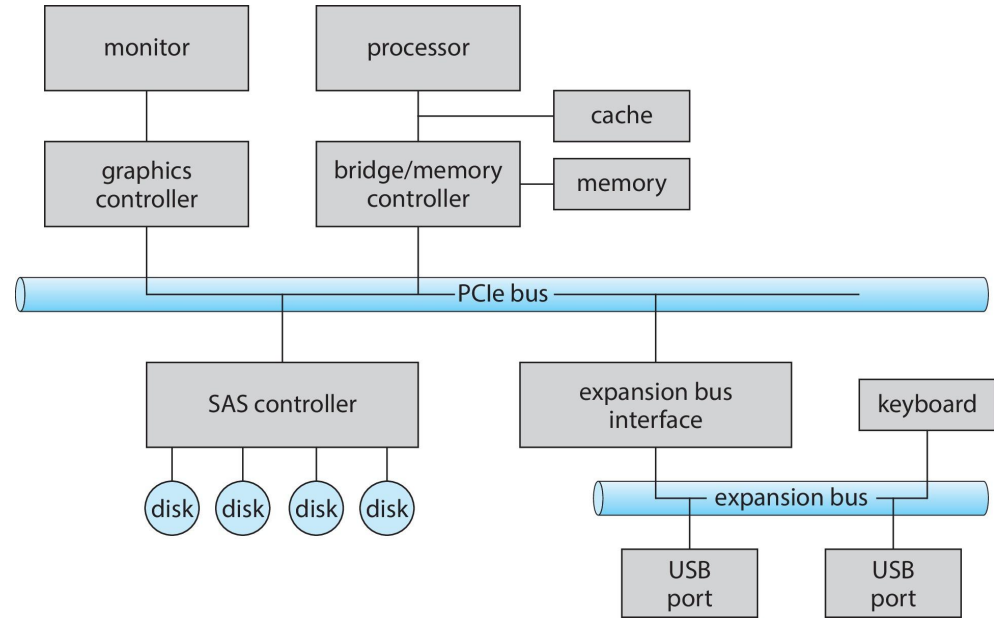
- A capacitor + gate, holds charge to indicate bit value
- 1 transistor per bit – extremely dense storage
- Charge leaks – need slow comparator to decide if bit 1 or 0
- Must rewrite charge after reading, and periodically refresh

VRAM – “Video RAM”

- - Dual ported DRAM, can write while someone else reads

What is I/O bus? E.g., PCI:A Typical PC Bus Structure

- Incredible variety of I/O devices
 - Storage
 - Transmission
 - Human-interface
- Common concepts – signals from I/O devices interface with computer
 - **Port** – connection point for device
 - **Bus - daisy chain** or shared direct access
 - **PCI** bus common in PCs and servers, PCI Express (**PCIe**)
 - **expansion bus** connects relatively slow devices
 - **Serial-attached SCSI (SAS)** common disk interface

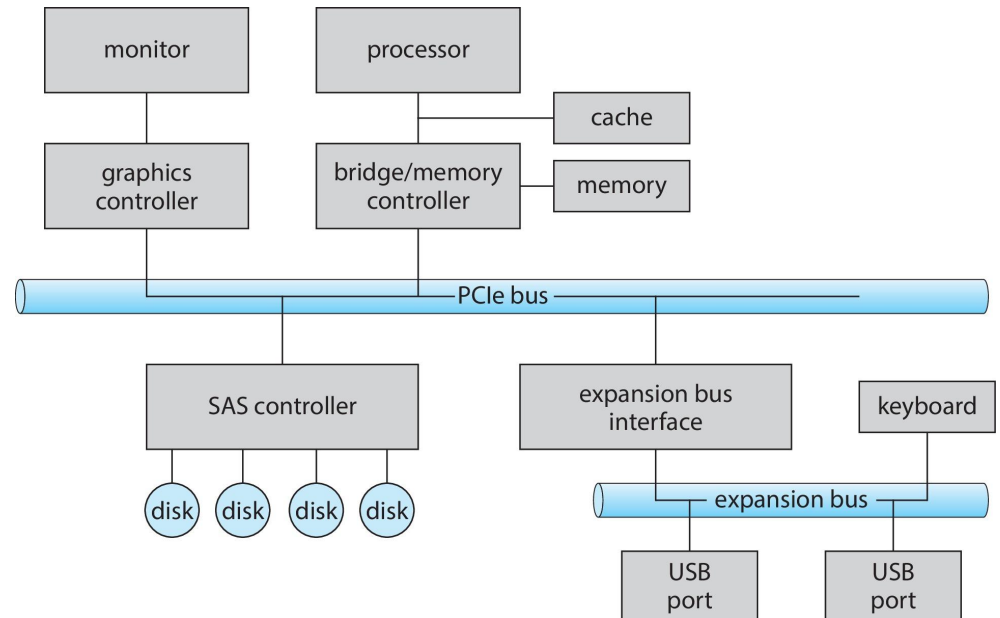


I/O Hardware (Cont.)

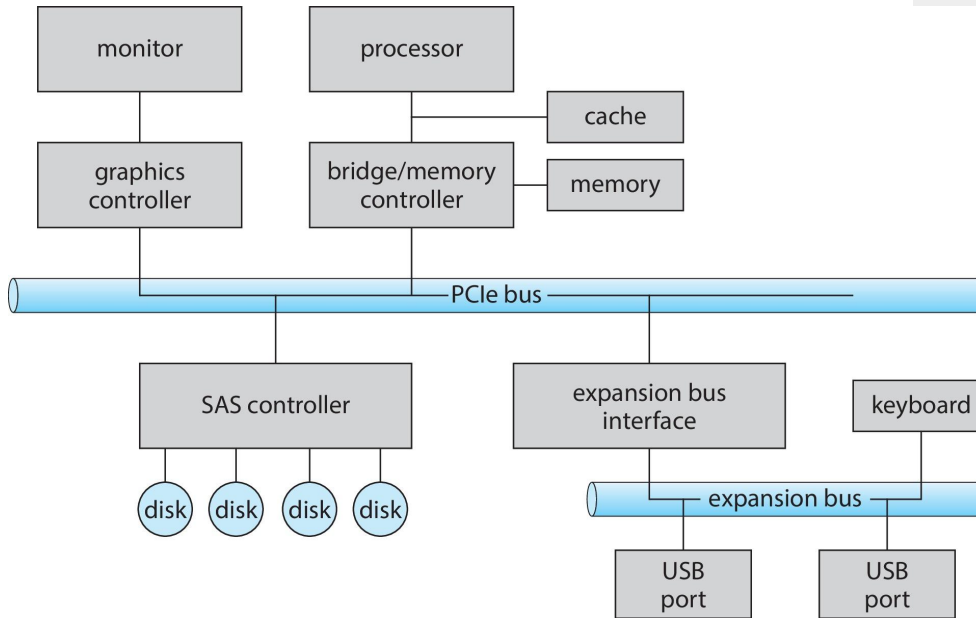
- **Controller (host adapter)** – electronics that operate port, bus, device
 - Sometimes integrated
 - Sometimes separate circuit board (host adapter)
 - Contains processor, microcode, private memory, bus controller, etc.
 - Some talk to per-device controller with bus controller, microcode, memory, etc.

Device drivers

- I/O management is a major component of operating system design and operation
 - Important aspect of computer operation
 - I/O devices vary greatly
 - Various methods to control them
 - Performance management
 - New types of devices frequent
- Ports, busses, device controllers connect to various devices
- **Device drivers** encapsulate device details
 - Present uniform device-access interface to I/O subsystem



Communicating with a device



- communicating with devices **through registers**
 - Data-in register, data-out register, status register, control register
 - Typically 1-4 bytes, or FIFO buffer
- communicating with devices **through memory-mapped I/O**
 - a certain portion of the processor's address space is mapped to the device
 -

communicating with devices through registers

data-in register is read by the host to get input from the device.

data-out register is written by the host to send output.

status register has bits read by the host to ascertain the status of the device, such as idle, ready for input, busy, error, transaction complete, etc.

control register has bits written by the host to issue commands or to change settings of the device such as parity checking, word length, or full- versus half-duplex operation.

I/O address range (hexadecimal)	device
000-00F	DMA controller
020-021	interrupt controller
040-043	timer
200-20F	game controller
2F8-2FF	serial port (secondary)
320-32F	hard-disk controller
378-37F	parallel port
3D0-3DF	graphics controller
3F0-3F7	diskette-drive controller
3F8-3FF	serial port (primary)

Device I/O Port Locations on PCs (partial)

x86 I/O instructions

```
static inline uint8_t
inb(uint16_t port){
    uint8_t data;
    asm volatile("inb %w1, %b0" : "=a"(data) : "Nd"(port));
    return data;
}

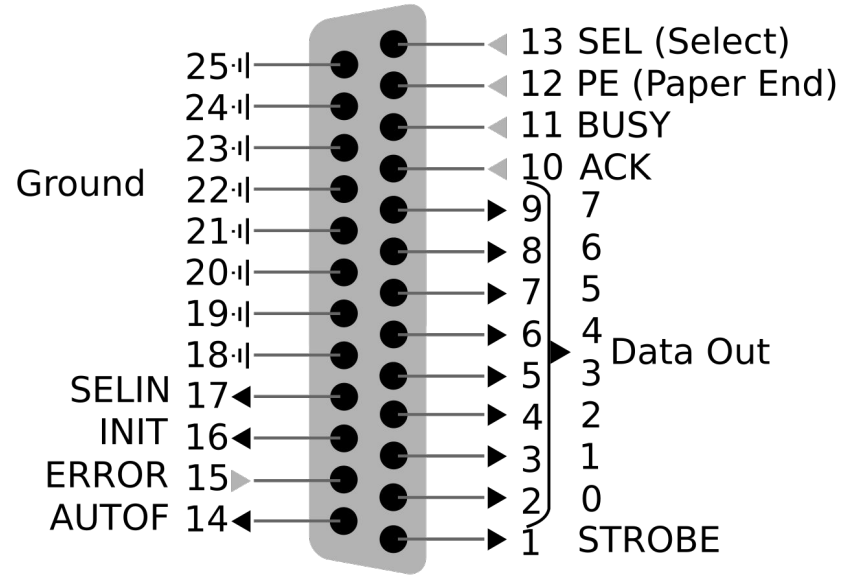
static inline void
outb(uint16_t port, uint8_t data){
    asm volatile("outb %b0, %w1" : : "a"(data), "Nd"(port));
}

static inline void
insw(uint16_t port, void *addr, size_t cnt){
    asm volatile("rep insw" : "+D"(addr), "+c"(cnt) : "d"(port) : "memory");
}
```

https://www.scs.stanford.edu/24wi-cs212/notes/io_disks.pdf

<https://man7.org/linux/man-pages/man2/outb.2.html>

Example LPT1






Simple hardware has three control registers:

Bit-to-pin mapping for the Standard Parallel Port (SPP):

Address	MSB							LSB
	Bit: 7	6	5	4	3	2	1	0
Base (Data port)	Pin: 9	8	7	6	5	4	3	2
Base+1 (Status port)	Pin: ~11	10	12	13	15			
Base+2 (Control port)	Pin:				~17	16	~14	~1

~ indicates a hardware inversion of the bit.



Writing a byte to parallel port

```
/* https : // wiki.osdev.org/Parallel_port */
void sendbyte(uint8_t byte) {
    /* Wait until BSY bit is 1. */
    while ((inb(0x379) & 0x80) == 0)
        delay();
    /* Put the byte we wish to send on pins D7-0. */
    outb(0x378, byte);
    /* Pulse STR (strobe) line to inform the printer
     * that a byte is available */
    uint8_t ctrlval = inb(0x37a);
    outb(0x37a, ctrlval | 0x01);
    delay();
    outb(0x37a, ctrlval);
}
```

[Parallel port - OSDev Wiki](https://www.osdev.org/wiki/Parallel_port)

https://www.scs.stanford.edu/24wi-cs212/notes/io_disks.pdf

IDE disk driver

```
void IDE_ReadSector(int disk, int off, void *buf){
    outb(0x1F6, disk == 0 ? 0xE0 : 0xF0); // Select Drive
    IDEWait();
    outb(0x1F2, 1);           // Read length (1 sector = 512 B)
    outb(0x1F3, off);        // LBA(linear block address) low
    outb(0x1F4, off >> 8);   // LBA mid
    outb(0x1F5, off >> 16);  // LBA high
    outb(0x1F7, 0x20);       // Read command
    insw(0x1F0, buf, 256);   // Read 256 words
}
```

```
void IDEWait() {
    // Discard status 4 times
    inb(0x1F7); inb(0x1F7); inb(0x1F7); inb(0x1F7);
    // Wait for status BUSY flag to clear
    while ((inb(0x1F7) & 0x80) != 0)
        ;
}
```

https://www.scs.stanford.edu/24wi-cs212/notes/ido_disks.pdf

Summary of I/O instructions

- in/out instructions slow and clunky
 - Instruction format restricts what registers you can use
 - Only allows 2^{16} different port numbers
 - Per-port access control turns out not to be useful (any port access allows you to disable all interrupts)

Alternative use **memory mapped I/O**

- Devices can achieve same effect with physical addresses, e.g.:

```
volatile int32_t *device_control
= (int32_t *) (0xc0100 + PHYS_BASE);
*device_control = 0x80;
int32_t status = *device_control;
```

- OS must map physical to virtual addresses, ensure non-cacheable
- Assign physical addresses at boot to avoid conflicts. PCI:
 - Slow/clunky way to access configuration registers on device
 - Use that to assign ranges of physical addresses to device


Memory Mapped I/O

Memory-mapped I/O

- Device data and command registers mapped to processor address space
 - (or you can say) a certain portion of the processor's address space is mapped to the device
- communications occur by reading and writing directly to/from those memory areas.
- suitable for devices which must move large quantities of data quickly
 - such as graphics cards.
- This can be used in combination with traditional registers:
 - For example, graphics cards still use registers for control information such as setting the video mode.
- A potential problem exists,
 - what if a process is allowed to write directly to the address space used by a memory-mapped I/O device.
- **Note this is NOT the same as DMA**

Direct memory access (DMA buffers)

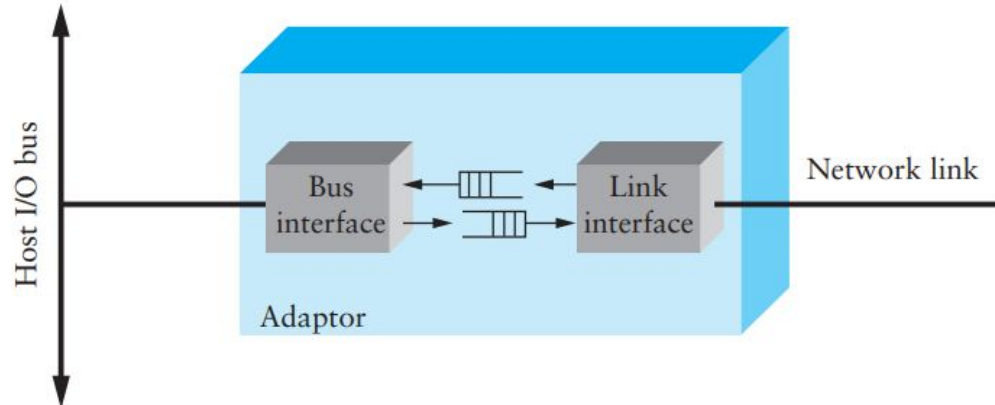
- It is wasteful to tie up the CPU transferring data in and out of registers one byte at a time.
 - Instead this work can be off-loaded to a special processor, known as the Direct Memory Access, DMA, Controller.
- **Idea:** only use CPU to transfer control requests, not data
- Include list of buffer locations in main memory
 - Device reads list and accesses buffers through DMA
 - Descriptions sometimes allow for scatter/gather I/O



Direct Memory Access

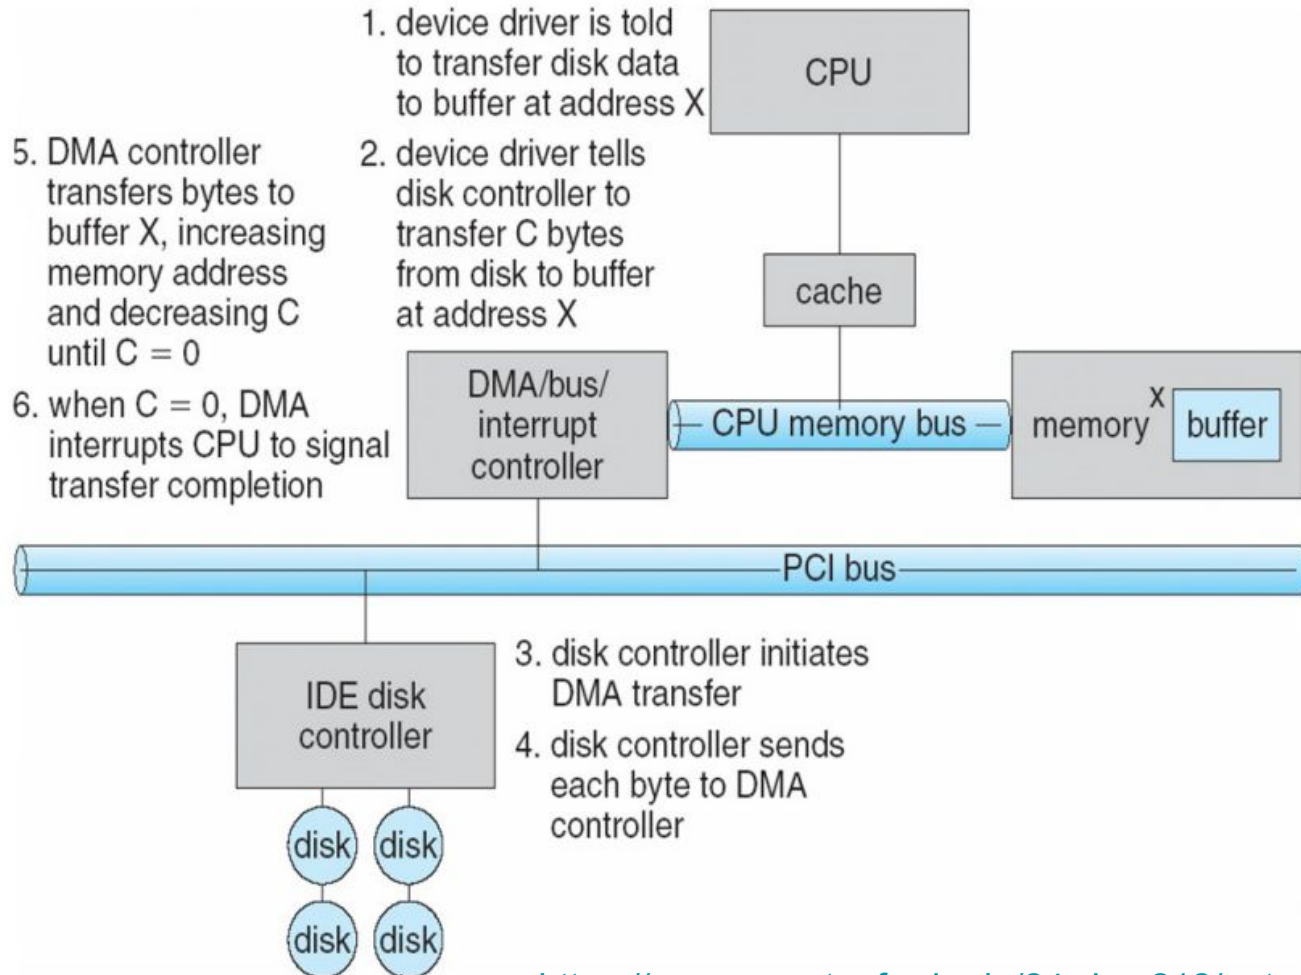
- Used to avoid **programmed I/O** (one byte at a time) for large data movement
- Requires **DMA** controller
- Bypasses CPU to transfer data directly between I/O device and memory
- OS writes DMA command block into memory
 - Source and destination addresses
 - Read or write mode
 - Count of bytes
 - Writes location of command block to DMA controller
 - Bus mastering of DMA controller – grabs bus from CPU
 - **Cycle stealing** from CPU but still much more efficient
 - When done, interrupts to signal completion
- Version that is aware of virtual addresses can be even more efficient - **DVMA**
 - possible to transfer from one mem to another without main memory chips

Example: Network Interface Card



- Link interface talks to wire/fiber/antenna
 - Typically does framing, link-layer CRC
- FIFOs on card provide small amount of buffering
- Bus interface logic uses DMA to move packets to and from buffers in main memory

Example: IDE disk read w. DMA



How should driver synchronize with card?

- **Device driver provides several entry points to kernel**
 - **Reset, ioctl, output, interrupt, read, write, strategy . . .**
- **How should driver synchronize with card?**
 - **E.g., Need to know when transmit buffers free or packets arrive**
 - **Need to know when disk request complete**
- **Two approach**
 - **pooling**
 - **interrupt driven devices**

Polling

Summary:

- Sent a packet? Loop asking card when buffer is free
- Waiting to receive? Keep asking card if it has packet
- Disk I/O? Keep looping until disk ready bit set

Details for each byte of I/O

1. **repeatedly** read **busy bit** from status register **until 0**
2. Host sets read or write bit and if write copies data into data-out register
3. Host sets command-ready bit
4. Controller sets busy bit, executes transfer
5. Controller clears busy bit, error bit, command-ready bit when transfer done

- Step 1 is **busy-wait** cycle to wait for I/O from device
 - Reasonable if device is fast
 - But inefficient if device slow
 - CPU switches to other tasks?
 - But if miss a cycle data overwritten / lost

Polling

Disadvantages

- Can't use CPU for anything else while polling
- Schedule poll in future?
 - High latency to receive packet or process disk block bad for response time

Polling can be **very fast and efficient** when

- both the **device** and the **controller** are **fast**
- and there is **significant data** to transfer.

It becomes **inefficient**, when

- the host must wait a long time in the busy loop for the device
- or frequent checks need to be made for data that is infrequently there.

How to be more efficient if status is non-zero **infrequently**?

Interrupt driven devices

- Asks card to interrupt CPU on events
 - Interrupts allow devices to notify the CPU
 - when they have data to transfer
 - or when an operation is complete

- This allows the CPU to perform other tasks when no I/O transfers need an immediate attention.

Interrupt handler runs at high priority

Asks card what happened:

- e.g. for NIC xmit buffer free, new packet

This is what most general-purpose OSes do

Interrupt driven devices

- Asks card to interrupt CPU on events
 - Interrupts allow devices to notify the CPU
 - when they have data to transfer
 - or when an operation is complete
- This allows the CPU to perform other tasks when no I/O transfers need an immediate attention.

CPU Interrupt-request line triggered by I/O device

- Checked(sensed) by processor after each instruction

A device's controller **raises** an interrupt by asserting a signal on the interrupt request line.

The CPU catches the interrupt and dispatches the interrupt handler.

- The CPU performs
 - a state save,
 - and transfers control to the interrupt handler routine at a fixed address in memory.

The interrupt handler clears the interrupt by servicing the device.

- The interrupt handler
 - determines the cause of the interrupt,
 - performs the necessary processing,
 - performs a state restore,
 - and executes a **return from interrupt** instruction to return control to the CPU.

Illustration of Interrupt-driven I/O cycle

CPU Interrupt-request line triggered by I/O device

- Checked(sensed) by processor after each instruction

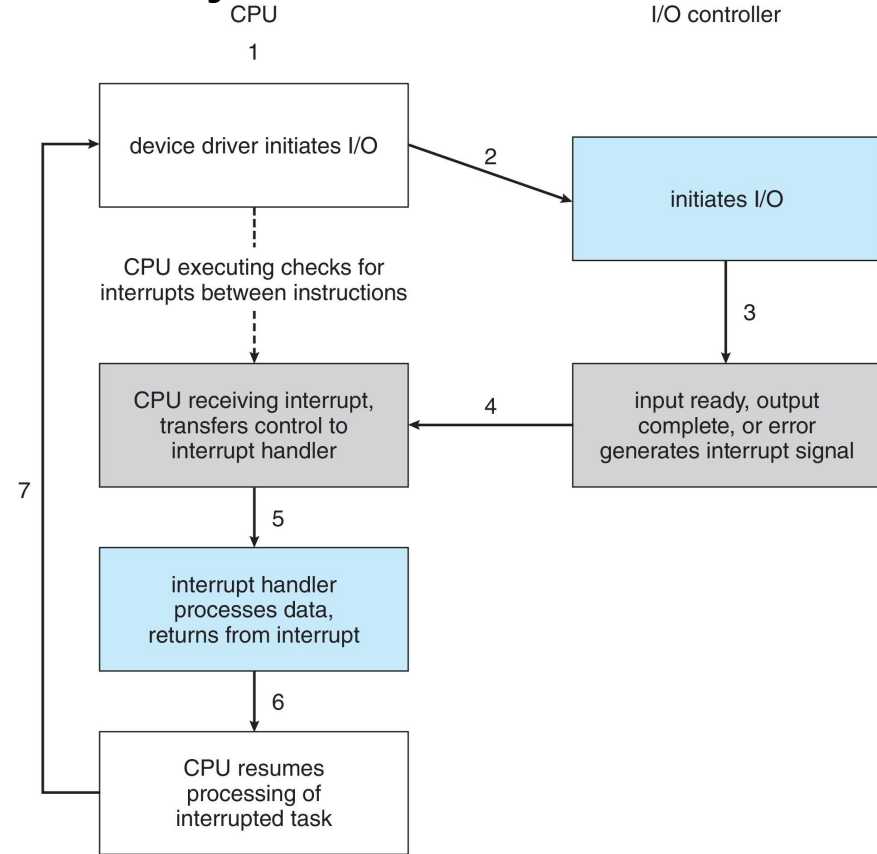
A device's controller **raises** an interrupt by asserting a signal on the interrupt request line.

The CPU catches the interrupt and dispatches the interrupt handler.

- The CPU performs
 - a state save,
 - and transfers control to the interrupt handler routine at a fixed address in memory.

The interrupt handler clears the interrupt by servicing the device.

- The interrupt handler
 - determines the cause of the interrupt,
 - performs the necessary processing,
 - performs a state restore,
 - and executes a **return from interrupt** instruction to return control to the CPU.



The previous example does not deal with the following Issues in modern computing

1. The need to defer interrupt handling during critical processing,
2. The need to determine which interrupt handler to invoke, without having to poll all devices to see which one needs attention, and
3. The need for multi-level interrupts, so the system can differentiate between high- and low-priority interrupts for proper response.

Interrupt controller on modern architectures

- two **Interrupt-request lines**
 - **Maskable** to ignore or delay some interrupts
 - **Non-maskable** for critical error conditions
- **Interrupt vector**
 - holds the addresses of routines prepared to process specific interrupts.
 - to dispatch interrupt to correct handler
 - Context switch at start and end
 - Based on priority
 - Some **nonmaskable**
 - **Interrupt chaining** if more than one device at same interrupt number

vector number	description
0	divide error
1	debug exception
2	null interrupt
3	breakpoint
4	INTO-detected overflow
5	bound range exception
6	invalid opcode
7	device not available
8	double fault
9	coprocessor segment overrun (reserved)
10	invalid task state segment
11	segment not present
12	stack fault
13	general protection
14	page fault
15	(Intel reserved, do not use)
16	floating-point error
17	alignment check
18	machine check
19–31	(Intel reserved, do not use)
32–255	maskable interrupts

Intel Pentium processor event-vector table.

Reminder: Exceptions

- Interrupt mechanism also used for **exceptions**
 - Terminate process, crash system due to hardware error
- Page fault executes when memory access error
- System call executes via **trap** to trigger kernel to execute request
- Multi-CPU systems can process interrupts concurrently
 - If operating system designed to handle it
- Used for time-sensitive processing, frequent, must be fast

Latency

- Stressing interrupt management because even single-user systems manage hundreds or interrupts per second and servers hundreds of thousands
- For example, a quiet macOS desktop generated 23,000 interrupts over 10 seconds

```
Fri Nov 25 13:55:59                                0:00:10
                SCHEDULER      INTERRUPTS
-----
total_samples          13          22998

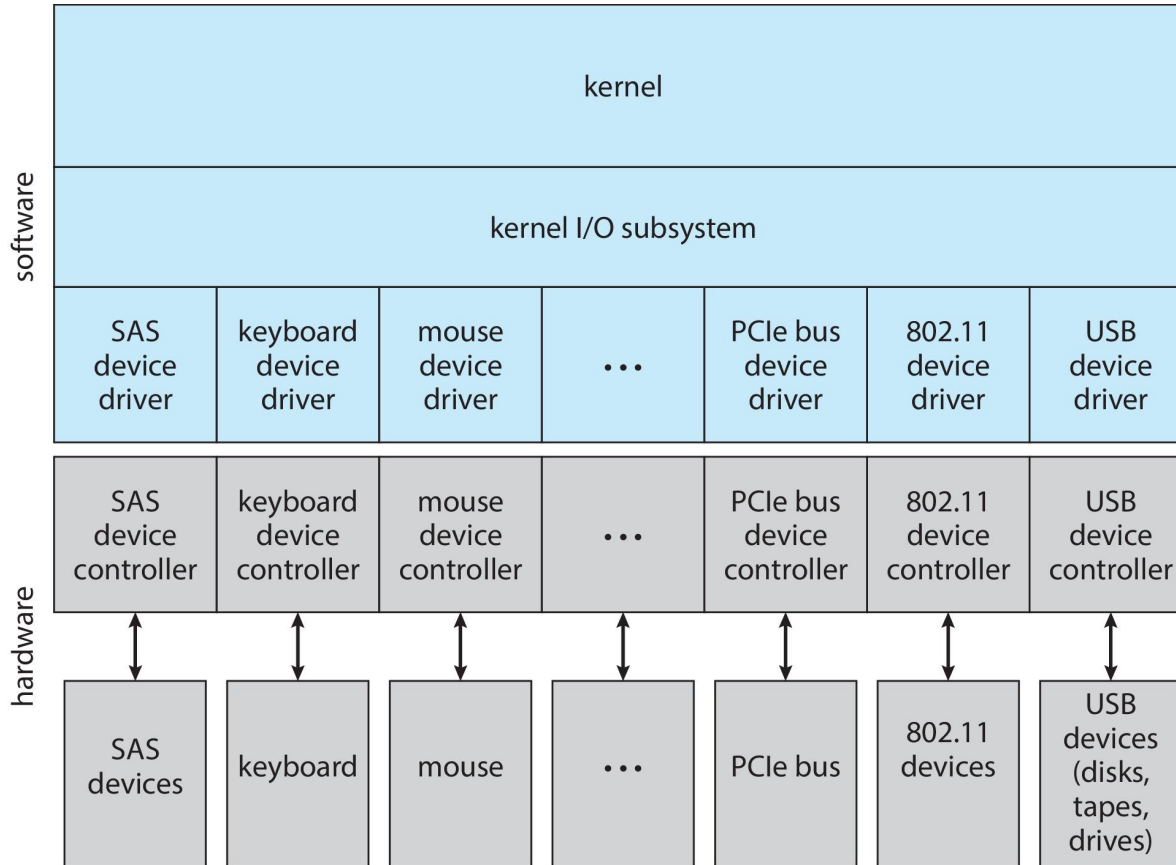
delays < 10 usecs      12          16243
delays < 20 usecs       1           5312
delays < 30 usecs       0            473
delays < 40 usecs       0            590
delays < 50 usecs       0             61
delays < 60 usecs       0            317
delays < 70 usecs       0              2
delays < 80 usecs       0              0
delays < 90 usecs       0              0
delays < 100 usecs      0              0
total < 100 usecs      13          22998
```

on linux
\$ cat /proc/interrupts

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
- New devices talking already-implemented protocols need no extra work
- Each OS has its own I/O subsystem structures and device driver frameworks
- Devices vary in many dimensions
 - **Character-stream** or **block**
 - **Sequential** or **random-access**
 - **Synchronous** or **asynchronous** (or both)
 - **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

Characteristics of I/O Devices (Cont.)

- Subtleties of devices handled by device drivers
- Broadly I/O devices can be grouped by the OS into
 - **Block I/O**
 - **Character I/O (Stream)**
 - **Memory-mapped file access**
 - **Network sockets**
- For direct manipulation of I/O device specific characteristics, usually an escape / back door
 - Unix `ioctl()` call to send arbitrary bits to a device control register and data to device data register
- UNIX and Linux use tuple of “major” and “minor” device numbers to identify type and

```
brw-rw---- 1 root disk 8, 0 Mar 16 09:18 /dev/sda s 0-4)
brw-rw---- 1 root disk 8, 1 Mar 16 09:18 /dev/sda1
brw-rw---- 1 root disk 8, 2 Mar 16 09:18 /dev/sda2
brw-rw---- 1 root disk 8, 3 Mar 16 09:18 /dev/sda3
```

Block and Character Devices

- Block devices include disk drives
 - Commands include read, write, seek
 - **Raw I/O, direct I/O**, or file-system access
 - Memory-mapped file access possible
 - File mapped to virtual memory and clusters brought via demand paging
 - DMA
- Character devices include keyboards, mice, serial ports
 - Commands include **get () , put ()**
 - Libraries layered on top allow line editing

Network Devices

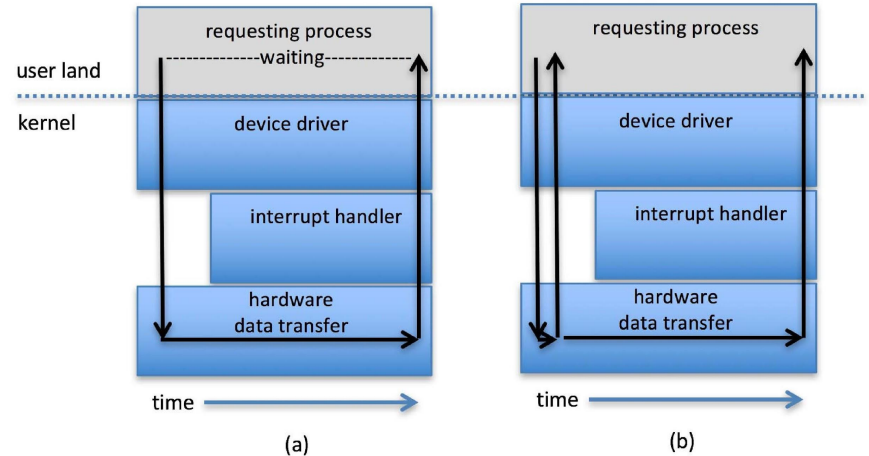
- Varying enough from block and character to have own interface
- Linux, Unix, Windows and many others include **socket** interface
 - Separates network protocol from network operation
 - Includes `select()` functionality
- Approaches vary widely (pipes, FIFOs, streams, queues, mailboxes)

Clocks and Timers

- Provide current time, elapsed time, timer
- Normal resolution about 1/60 second
- Some systems provide higher-resolution timers
- **Programmable interval timer** used for timings, periodic interrupts
- `ioctl()` (on UNIX) covers odd aspects of I/O such as clocks and timers

Nonblocking and Asynchronous 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
 - **select()** to find if data ready then **read()** or **write()** to transfer
- **Asynchronous** - process runs while I/O executes
 - Difficult to use
 - I/O subsystem signals process when I/O completed



Two I/O methods: (a) synchronous and (b) asynchronous.

Vectored I/O

- **Vectored I/O** allows one system call to perform multiple I/O operations
- For example, Unix **readve ()** accepts a vector of multiple buffers to read into or write from
- This scatter-gather method better than multiple individual I/O calls
 - Decreases context switching and system call overhead
 - Some versions provide atomicity
 - Avoid for example worry about multiple threads changing data as reads / writes occurring

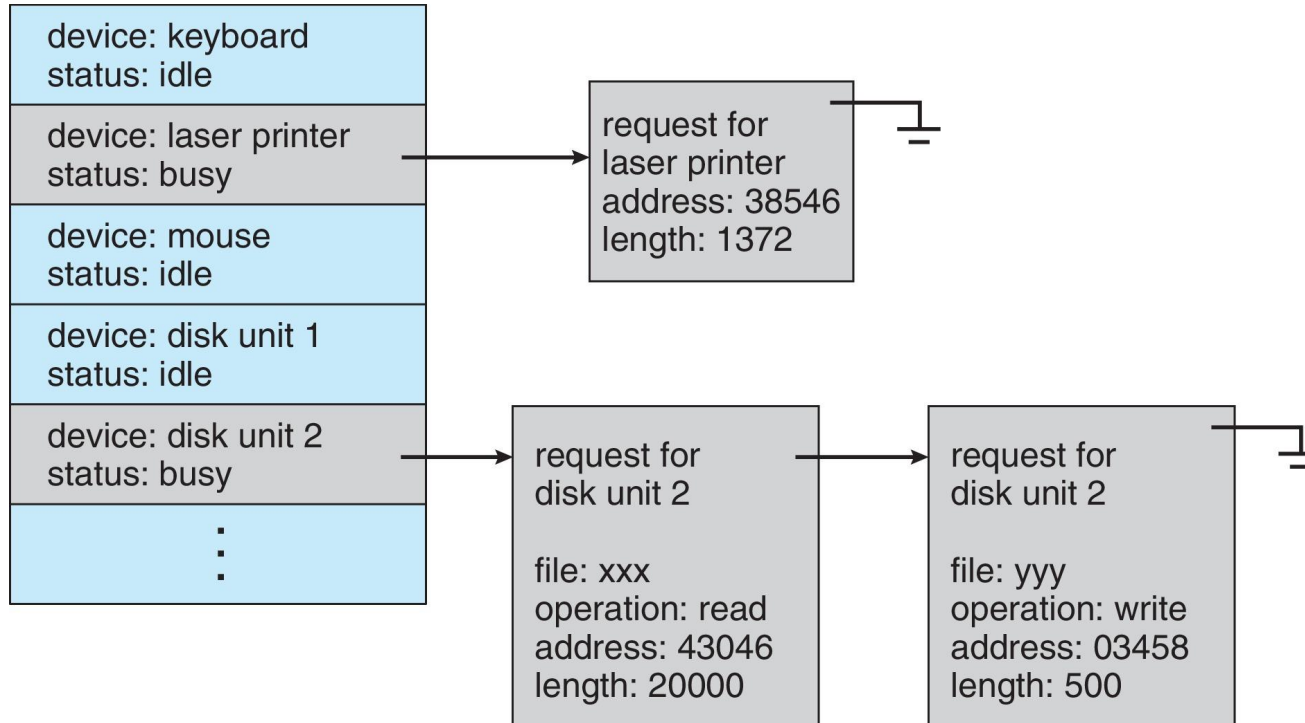
Kernel I/O Subsystem

- **I/O Scheduling**
 - Some I/O request ordering via per-device queue
 - Some OSs try fairness
 - Some implement Quality Of Service (i.e. IPQOS)
- **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”
 - **Double buffering** – two copies of the data
 - Kernel and user
 - Varying sizes
 - Full / being processed and not-full / being used
 - Copy-on-write can be used for efficiency in some cases

Kernel I/O Subsystem

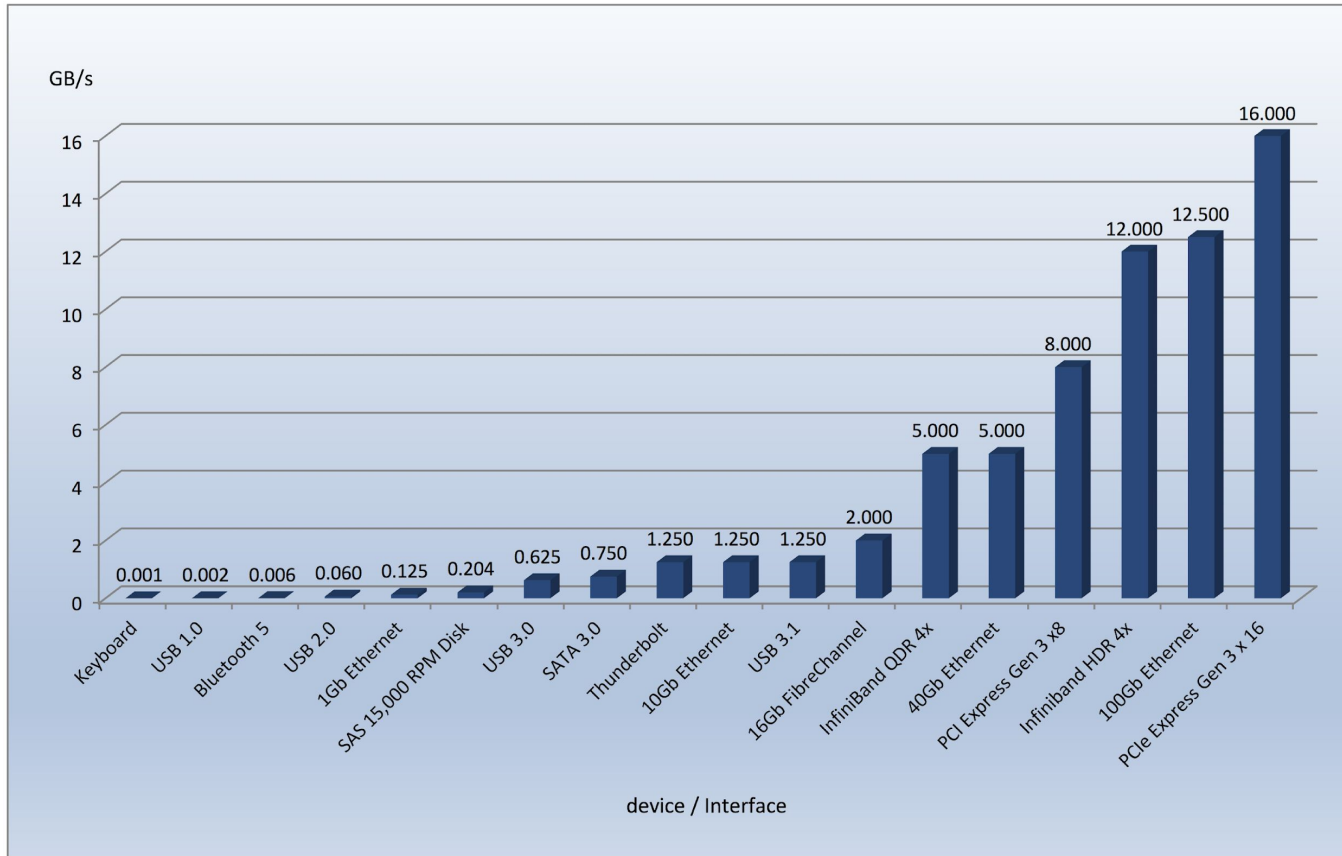
- **Caching** - faster device holding copy of data
 - Always just a copy
 - Key to performance
 - Sometimes combined with buffering
- **Spooling** - hold output for a device
 - If device can serve only one request at a time
 - i.e., Printing
- **Device reservation** - provides exclusive access to a device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

Device-status Table



When a kernel supports asynchronous I/O, it must be able to keep track of many I/O requests at the same time.

Common PC and Data-center I/O devices and Interface Speeds




Error Handling

- OS can recover from disk read, device unavailable, transient write failures
 - Retry a read or write, for example
 - Some systems more advanced – Solaris FMA, AIX
 - Track error frequencies, stop using device with increasing frequency of retry-able errors
- 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

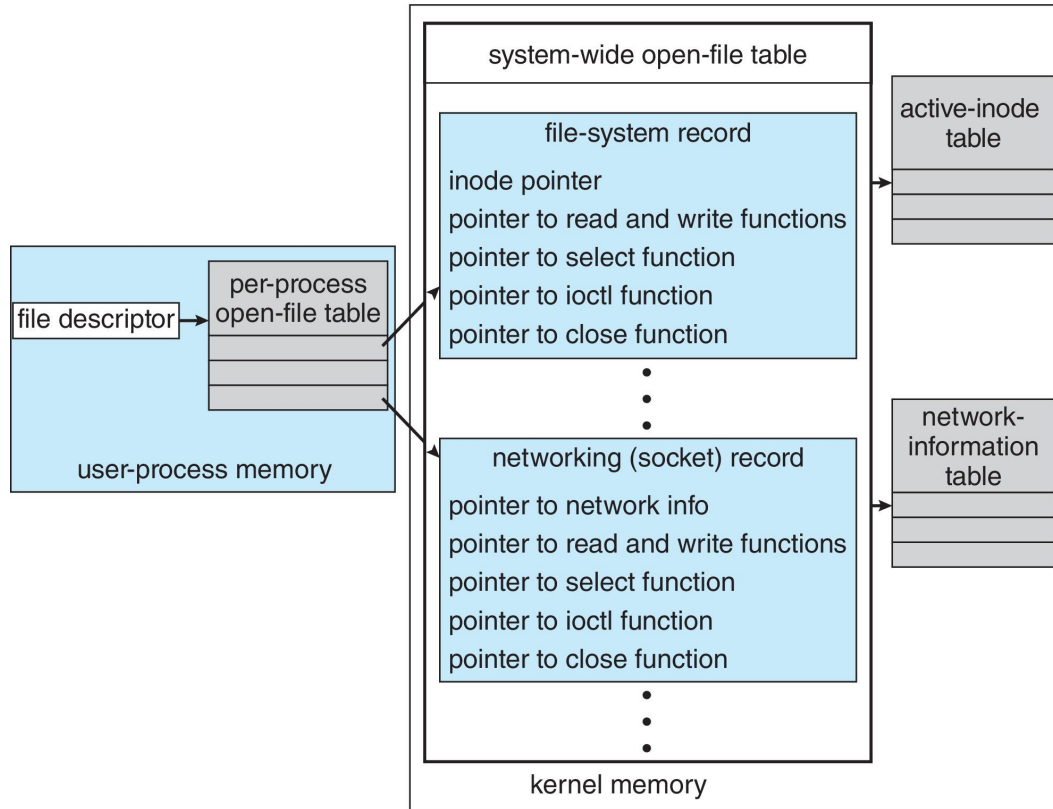
Use of a System Call to Perform I/O



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
 - Windows uses message passing
 - Message with I/O information passed from user mode into kernel
 - Message modified as it flows through to device driver and back to process
 - Pros / cons?

UNIX I/O Kernel Structure



Power Management

- Not strictly domain of I/O, but much is I/O related
- Computers and devices use electricity, generate heat, frequently require cooling
- OSes can help manage and improve use
 - Cloud computing environments move virtual machines between servers
 - Can end up evacuating whole systems and shutting them down
- Mobile computing has power management as first class OS aspect

Power Management (Cont.)

- For example, Android implements
 - Component-level power management
 - Understands relationship between components
 - Build device tree representing physical device topology
 - System bus -> I/O subsystem -> {flash, USB storage}
 - Device driver tracks state of device, whether in use
 - Unused component – turn it off
 - All devices in tree branch unused – turn off branch

Power Management (Cont.)

- For example, Android implements (Cont.)
 - Wake locks – like other locks but prevent sleep of device when lock is held
 - Power collapse – put a device into very deep sleep
 - Marginal power use
 - Only awake enough to respond to external stimuli (button press, incoming call)
- Modern systems use **advanced configuration and power interface (ACPI)** firmware providing code that runs as routines called by kernel for device discovery, management, error and power management

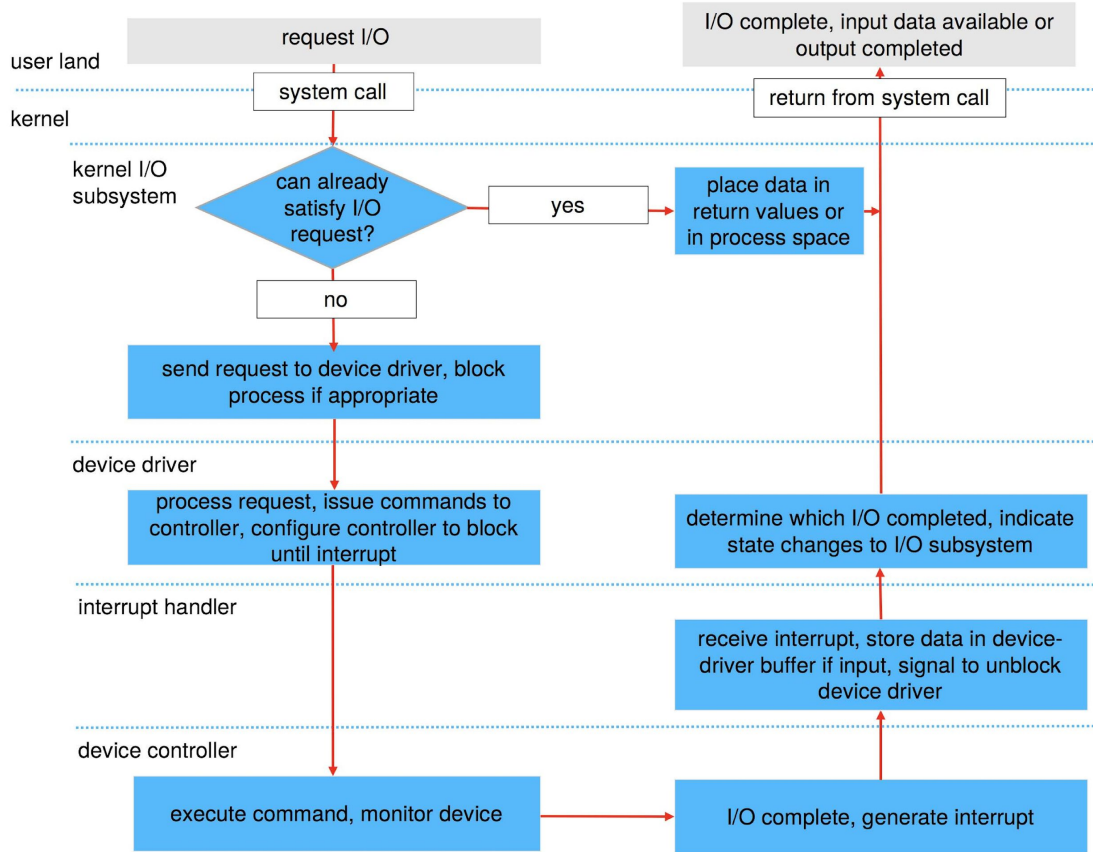
Kernel I/O Subsystem Summary

- In summary, the I/O subsystem coordinates an extensive collection of services that are available to applications and to other parts of the kernel
 - Management of the name space for files and devices
 - Access control to files and devices
 - Operation control (for example, a modem cannot seek())
 - File-system space allocation
 - Device allocation
 - Buffering, caching, and spooling
 - I/O scheduling
 - Device-status monitoring, error handling, and failure recovery
 - Device-driver configuration and initialization
 - Power management of I/O devices
- The upper levels of the I/O subsystem access devices via the uniform interface provided by the device drivers

Transforming 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


Life Cycle of An I/O Request



STREAMS

- **STREAM** – a full-duplex communication channel between a user-level process and a device in Unix System V and beyond
 - defines standard interfaces for char IO within kernel
- A STREAM consists of:
 - STREAM head interfaces with the user process
 - driver end interfaces with the device
 - zero or more STREAM modules between them
- Each module contains a **read queue** and a **write queue**
- Message passing is used to communicate between queues
 - **Flow control** option to indicate available or busy
- Asynchronous internally, synchronous where user process communicates with stream head

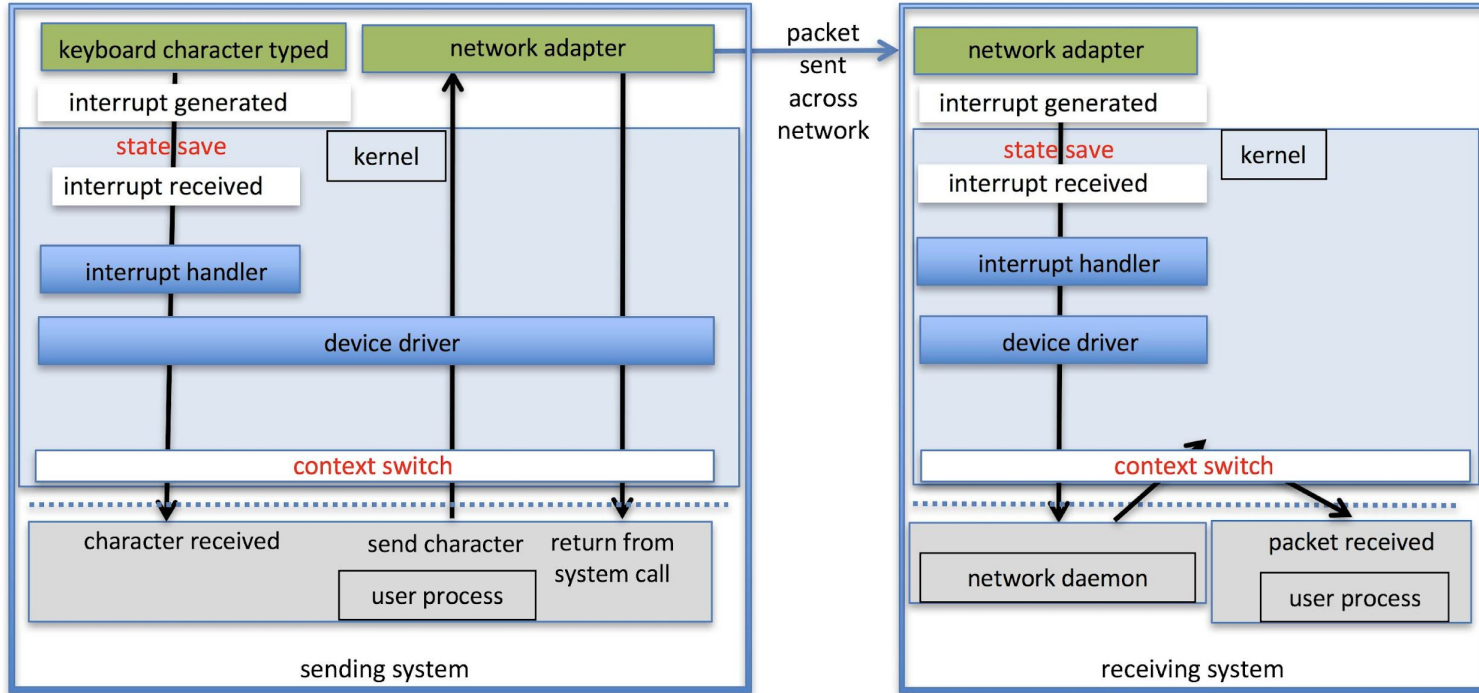
The STREAMS Structure in Unix



Performance

- I/O a major factor in system performance:
 - Demands CPU to execute device driver, kernel I/O code
 - Context switches due to interrupts
 - Data copying
 - Network traffic especially stressful

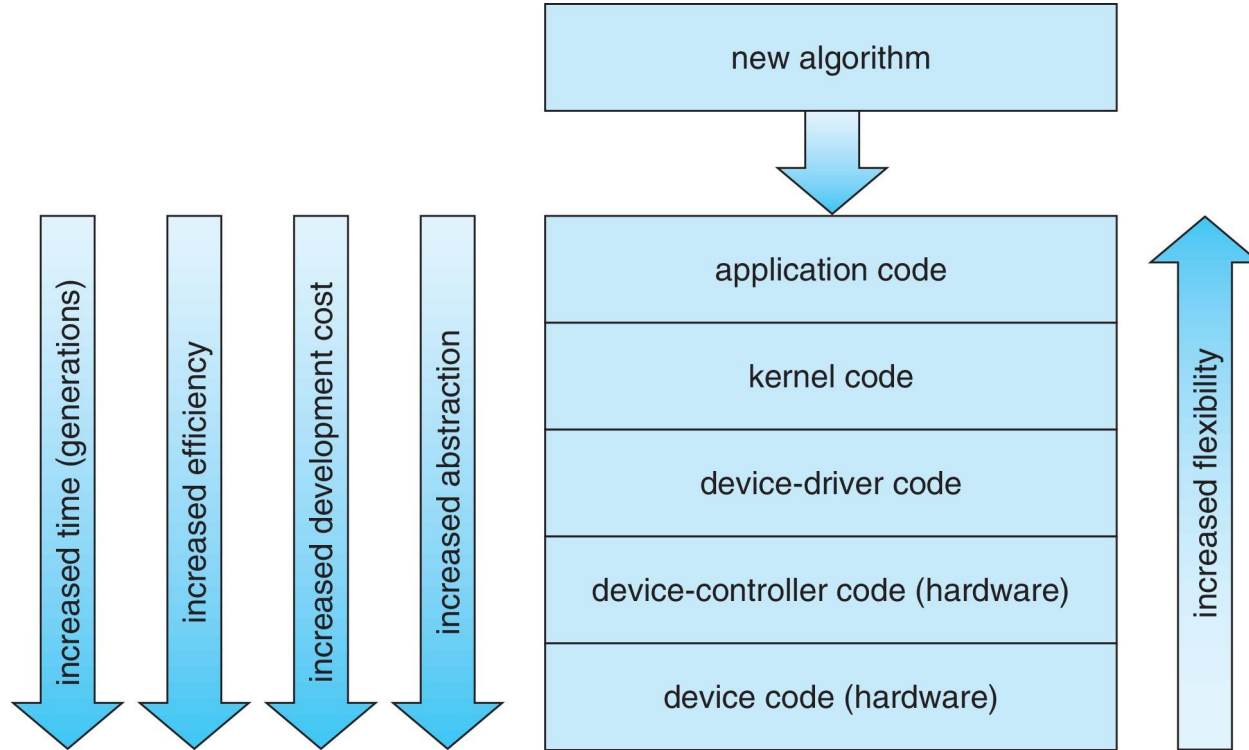
Intercomputer Communications



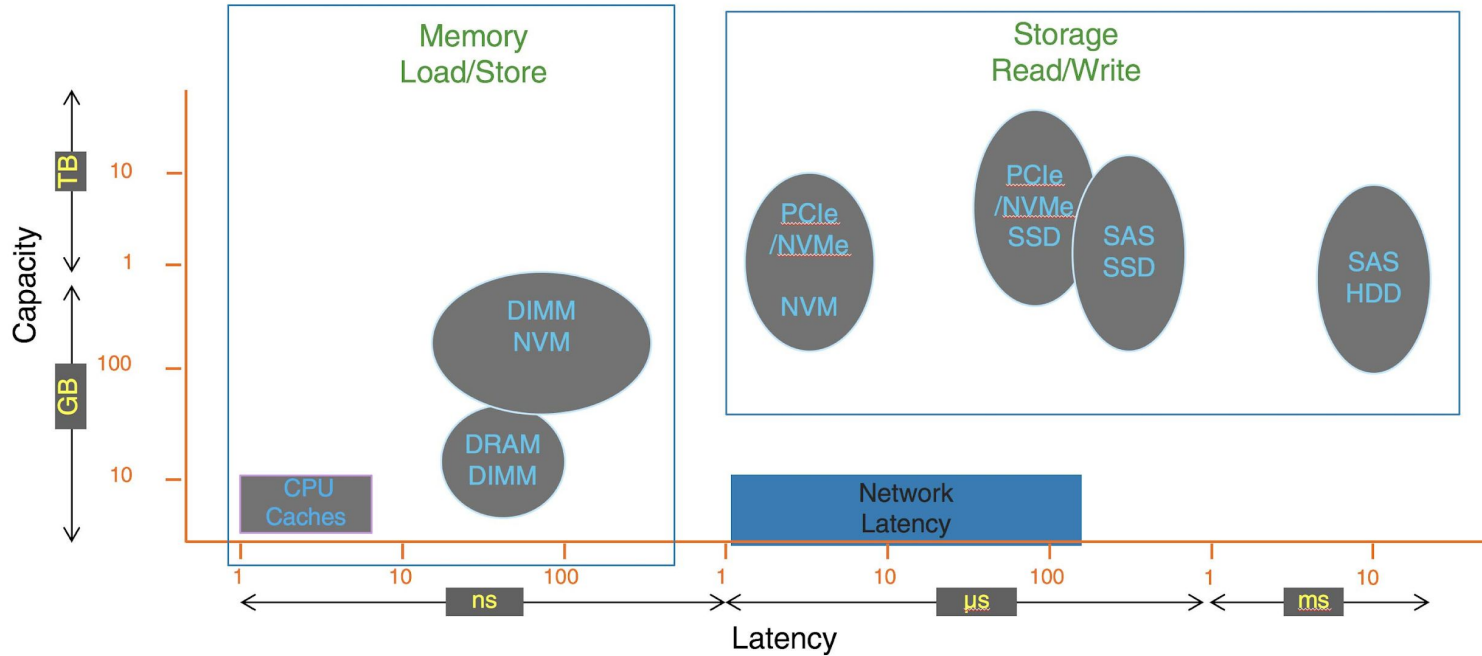
Improving Performance

- Reduce number of context switches
- Reduce data copying
- Reduce interrupts by using large transfers, smart controllers, polling
- Use DMA
- Use smarter hardware devices
- Balance CPU, memory, bus, and I/O performance for highest throughput
- Move user-mode processes / daemons to kernel threads

Device-Functionality Progression



I/O Performance of Storage (and Network Latency)



End of Chapter 12

