

# Input/Output

ICT 2203 Computer Architecture

*Partially based on Computer Organization and Architecture 8th Edition by William Stallings*

# Input Devices

- Keyboard
- Mouse
- Scanner
- Game Controller



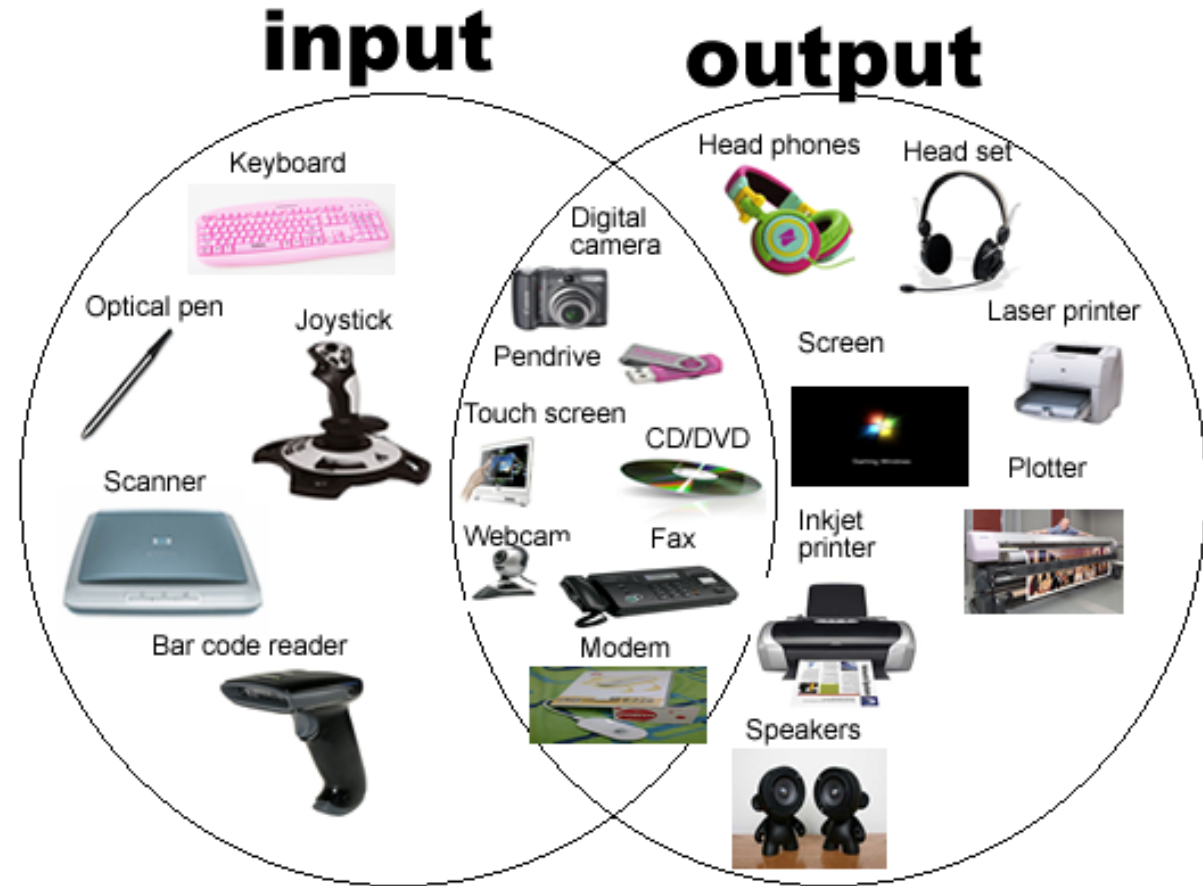
# Output Devices

- Monitor
- Printers
- Disk Drive
- Speakers



# Input and Output devices

- Modem
- Network Interface Card
- Portable storage drives
- Touchscreen display

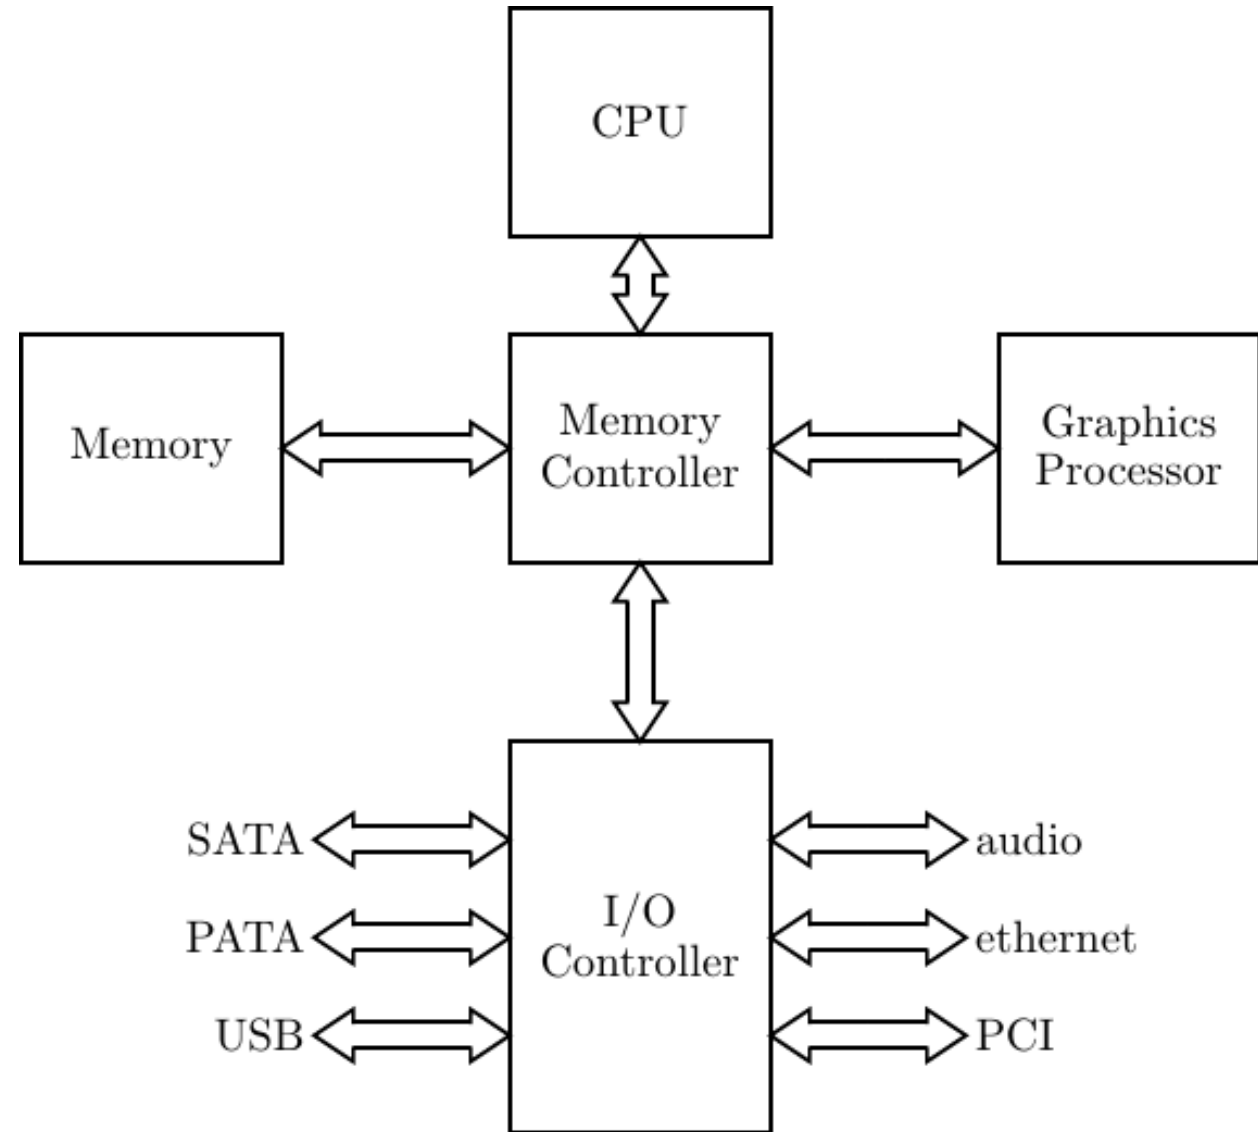


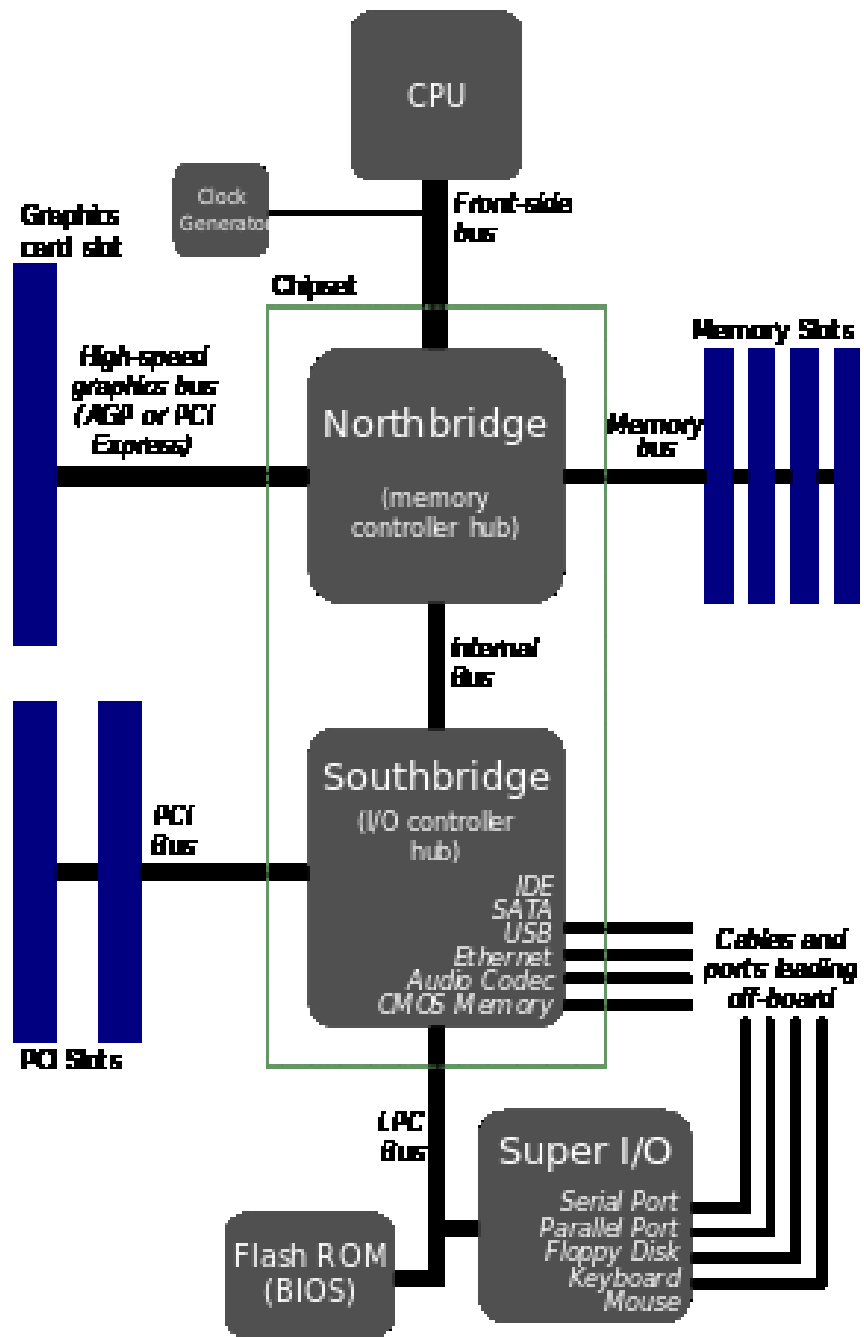
# Input/Output Problems

- Wide variety of peripherals
  - Delivering different amounts of data
  - At different speeds
  - In different formats
- All slower than CPU and RAM
- Need I/O modules


# I/O Module

- Interface to CPU and Memory
- Interface to one or more peripherals






South Bridge North Bridge



# I/O Module Function

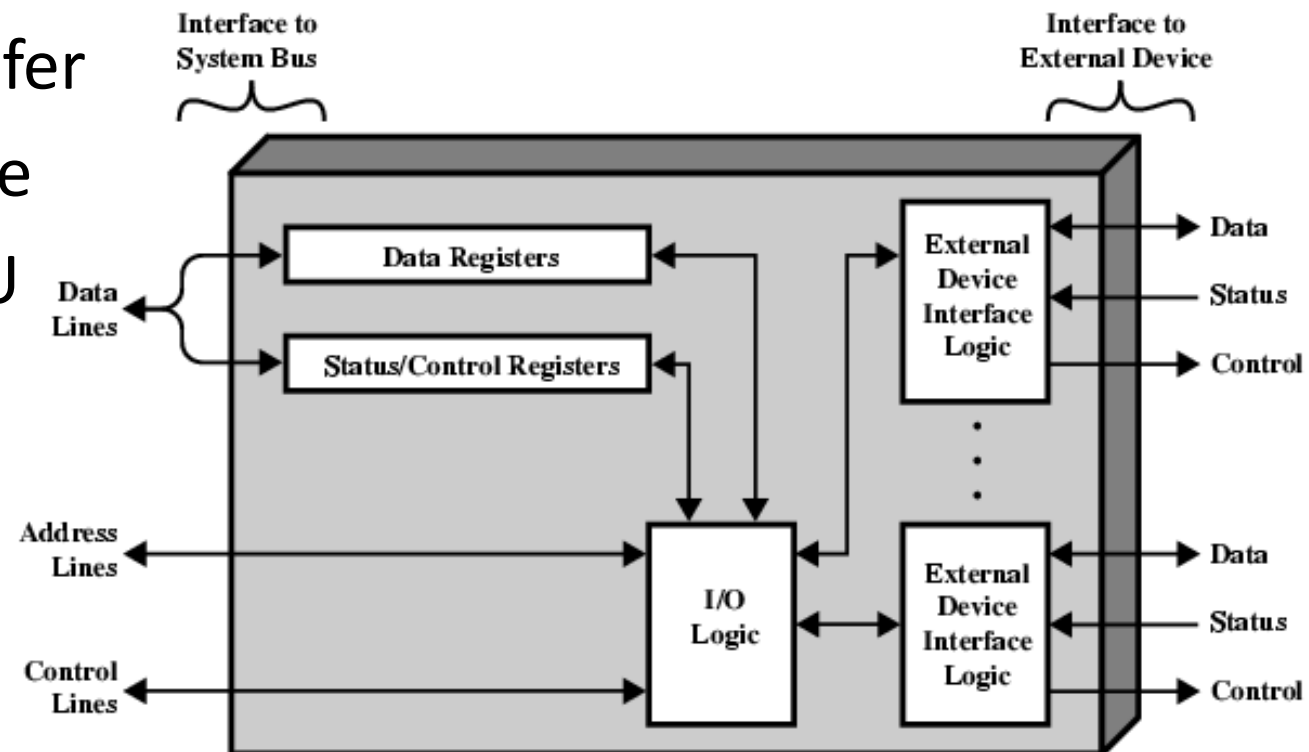
1. Control and Timing
2. CPU Communicating
3. Device Communication
4. Data Buffering
5. Error Detection





# I/O Steps

- CPU checks I/O module device status
- I/O module returns status
- If ready, CPU requests data transfer
- I/O module gets data from device
- I/O module transfers data to CPU
- Variations for output, DMA, etc.



# I/O Module Decisions

- Hide or reveal device properties to CPU
- Support multiple or single device
- Control device functions or leave for CPU
- Also O/S decisions
  - e.g. Unix treats everything it can as a file

# Input Output Techniques

- Programmed I/O
  - The CPU issues a command then waits for I/O operations to be complete.
  - The CPU is faster than the I/O module then method is wasteful.
- Interrupt Driven I/O
  - The CPU issues commands then proceeds with its normal work until interrupted by I/O device on completion of its work.
- DMA (Direct Memory Access)
  - Memory and I/O Module exchange data without involvement of CPU.
  - CPU grants I/O module authority to read from or write to memory without involvement.
  - DMA module controls exchange of data between main memory and the I/O device.

# Addressing I/O Devices

- Under programmed I/O data transfer is very like memory access (CPU viewpoint)
- Each device given unique identifier
- CPU commands contain identifier (address)

# I/O Mapping

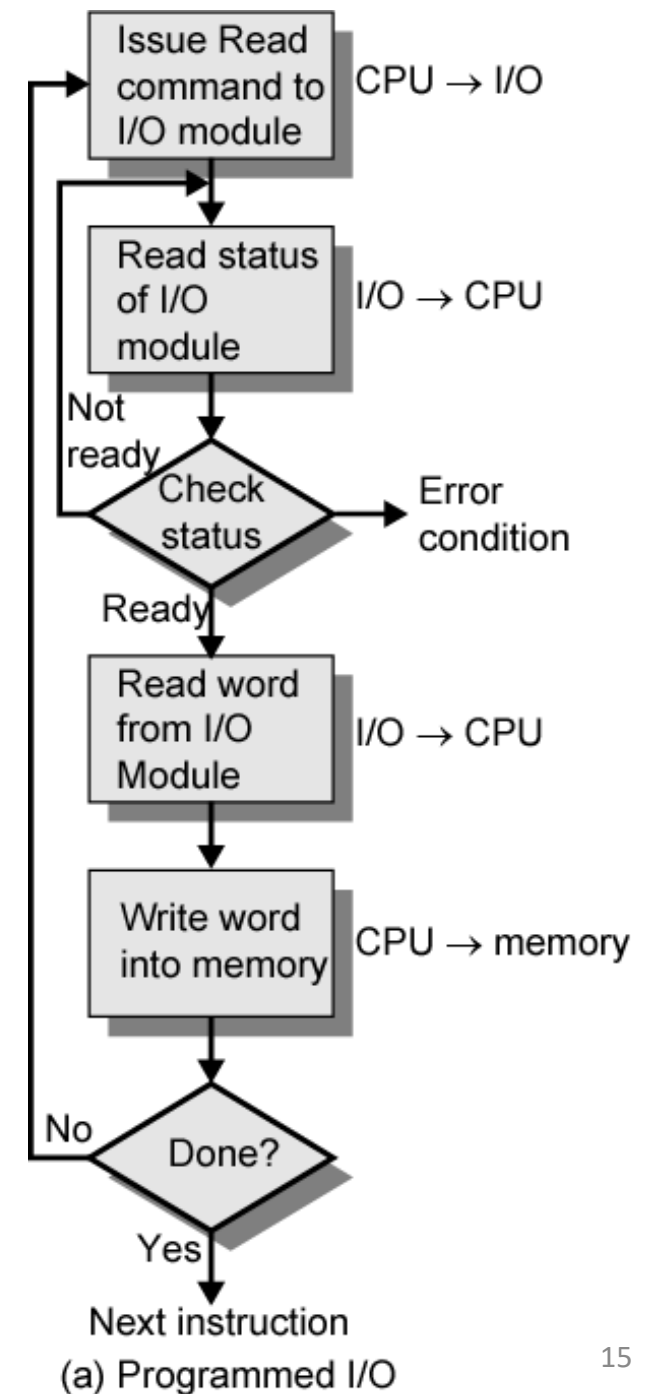
- Memory mapped I/O
  - Devices and memory share an address space
  - I/O looks just like memory read/write
  - No special commands for I/O
    - Large selection of memory access commands available
- Isolated I/O
  - Separate address spaces
  - Need I/O or memory select lines
  - Special commands for I/O
    - Limited set

# Programmed I/O

- CPU has direct control over I/O
  - Sensing status
  - Read/write commands
  - Transferring data
- CPU waits for I/O module to complete operation
- Wastes CPU time

# Programmed I/O - Basic Operation

- CPU requests I/O operation
- I/O module performs operation
- I/O module sets status bits
- CPU checks status bits periodically
- I/O module does not inform CPU directly
- I/O module does not interrupt CPU
- CPU may wait or come back later



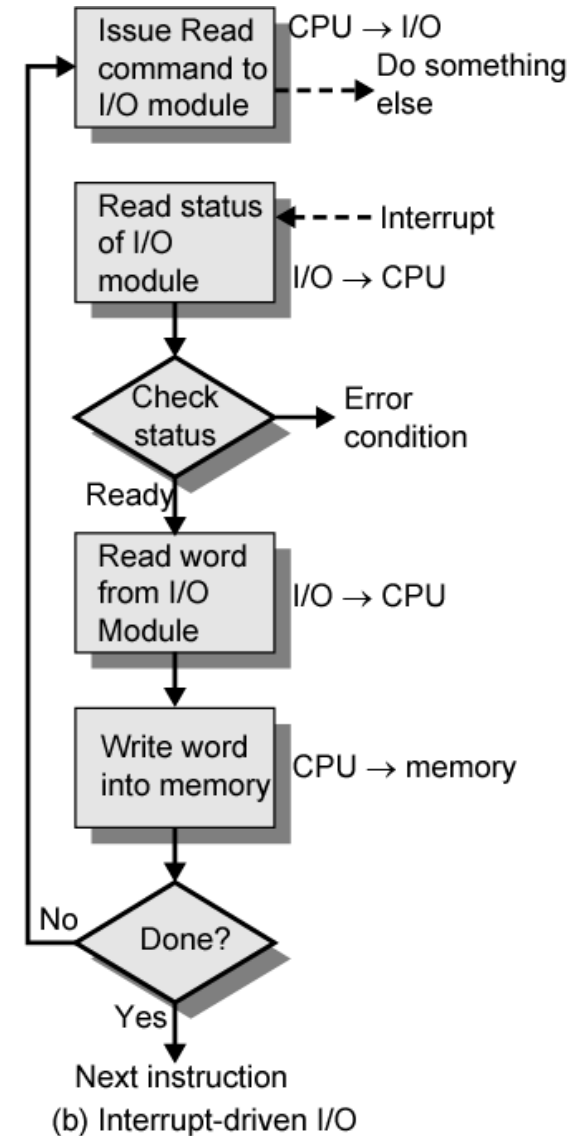
# Interrupt Driven I/O

- Overcomes CPU waiting
- No repeated CPU checking of device
- I/O module interrupts when ready



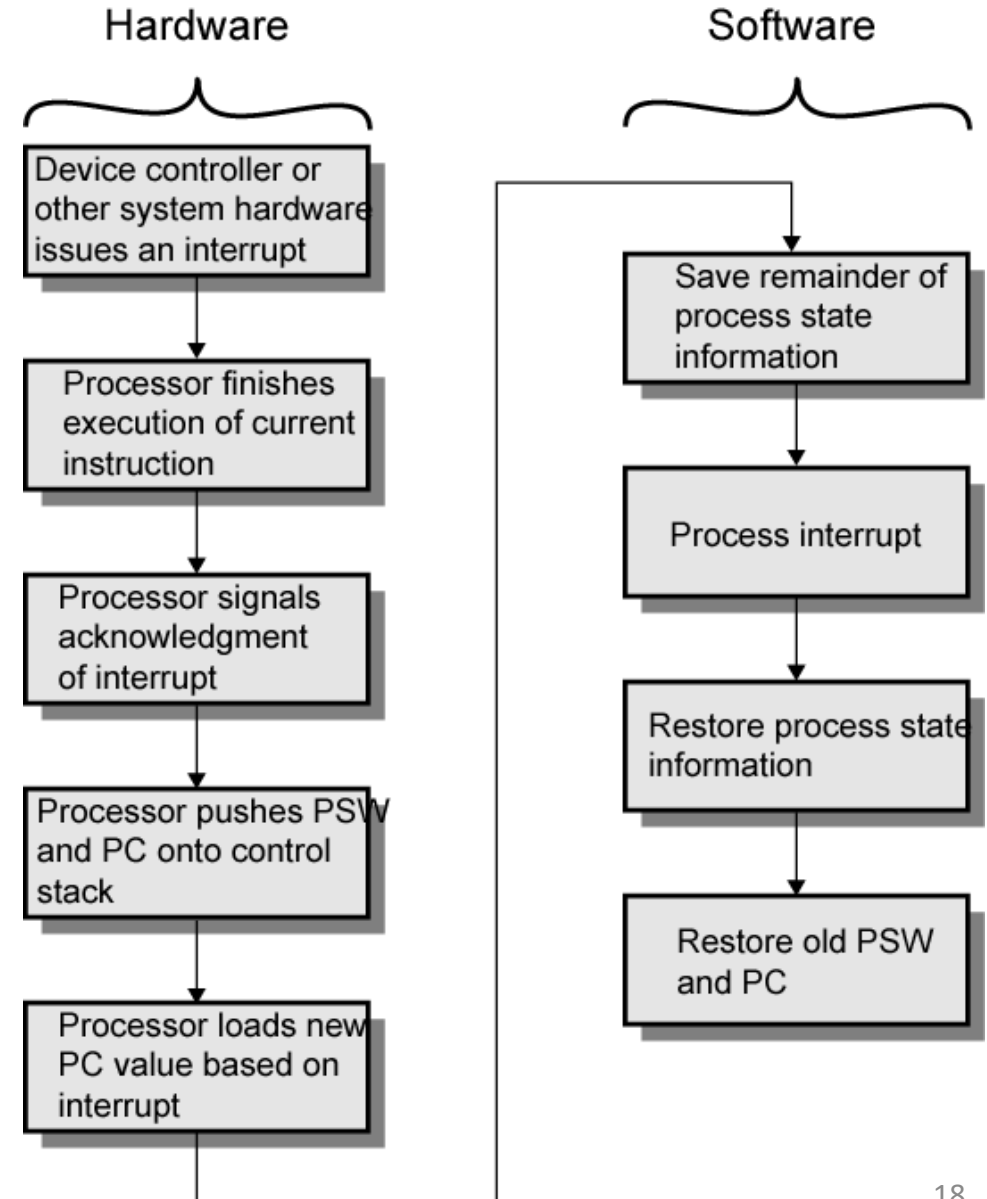
# Interrupt Driven I/O - Basic Operation

- CPU issues read command
- I/O module gets data from peripheral whilst CPU does other work
- I/O module interrupts CPU
- CPU requests data
- I/O module transfers data



# Simple Interrupt Processing

- Issue read command
- Do other work
- Check for interrupt at end of each instruction cycle
- If interrupted:-
  - Save context (registers)
  - Process interrupt
    - Fetch data & store



# Design Issues

- How do you identify the module issuing the interrupt?
- How do you deal with multiple interrupts?
  - i.e. an interrupt handler being interrupted

# Identifying Interrupting Module

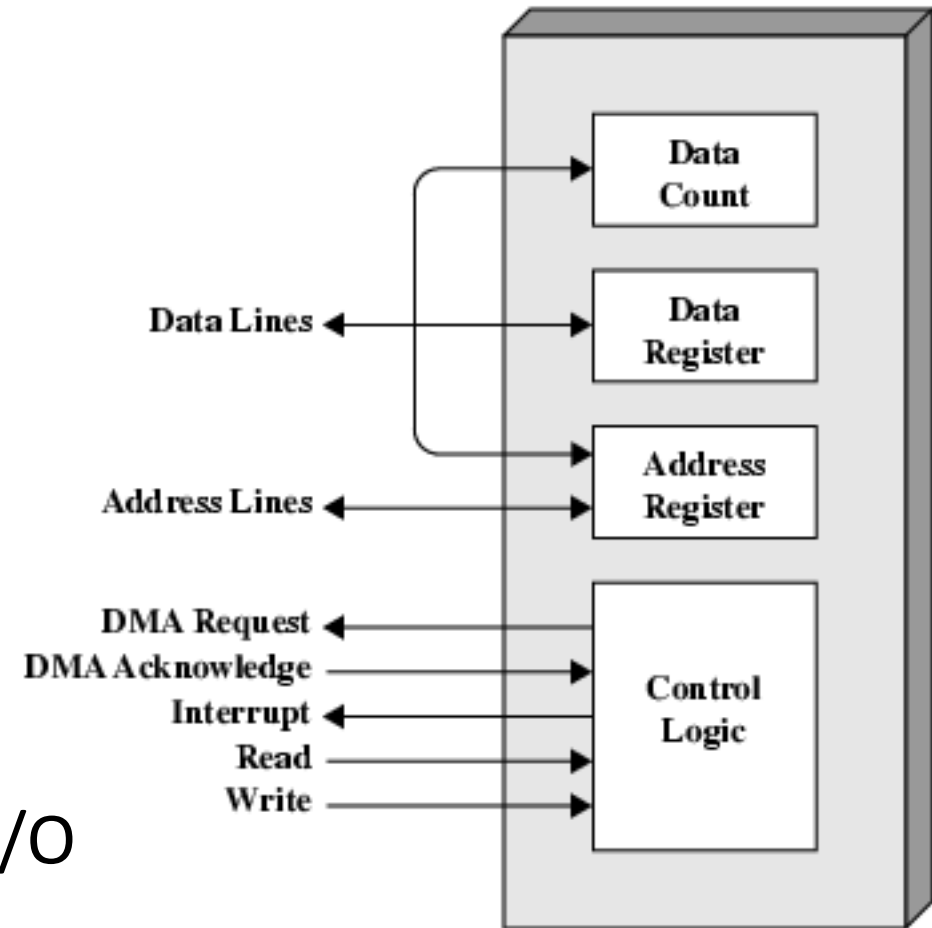
- Different line for each module
  - PC
  - Limits number of devices
- Software poll
  - CPU asks each module in turn
  - Slow
- Daisy Chain or Hardware poll
  - Interrupt Acknowledge sent down a chain
  - Module responsible places vector on bus
  - CPU uses vector to identify handler routine
- Bus Master
  - Module must claim the bus before it can raise interrupt
  - e.g. PCI & SCSI

# Multiple Interrupts

- Each interrupt line has a priority
- Higher priority lines can interrupt lower priority lines
- If bus mastering only current master can interrupt

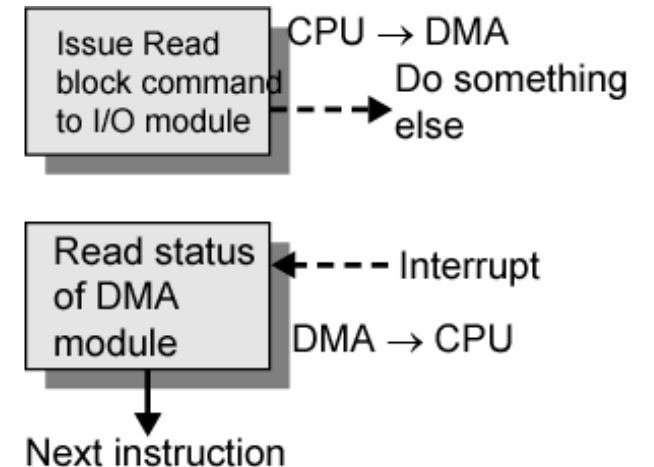
# Direct Memory Access

- Interrupt driven and programmed I/O require active CPU intervention
  - Transfer rate is limited
  - CPU is tied up
- DMA is the answer
  - Additional Module (hardware) on bus
  - DMA controller takes over from CPU for I/O



# DMA Operation - Basic Operation

- CPU tells DMA controller:
  - Whether to Read or Write
  - Address of the device to be addressed
  - Starting address of memory block for data
  - Amount of data to be transferred
- CPU carries on with other work
- DMA controller deals with transfer
- DMA controller sends interrupt when finished



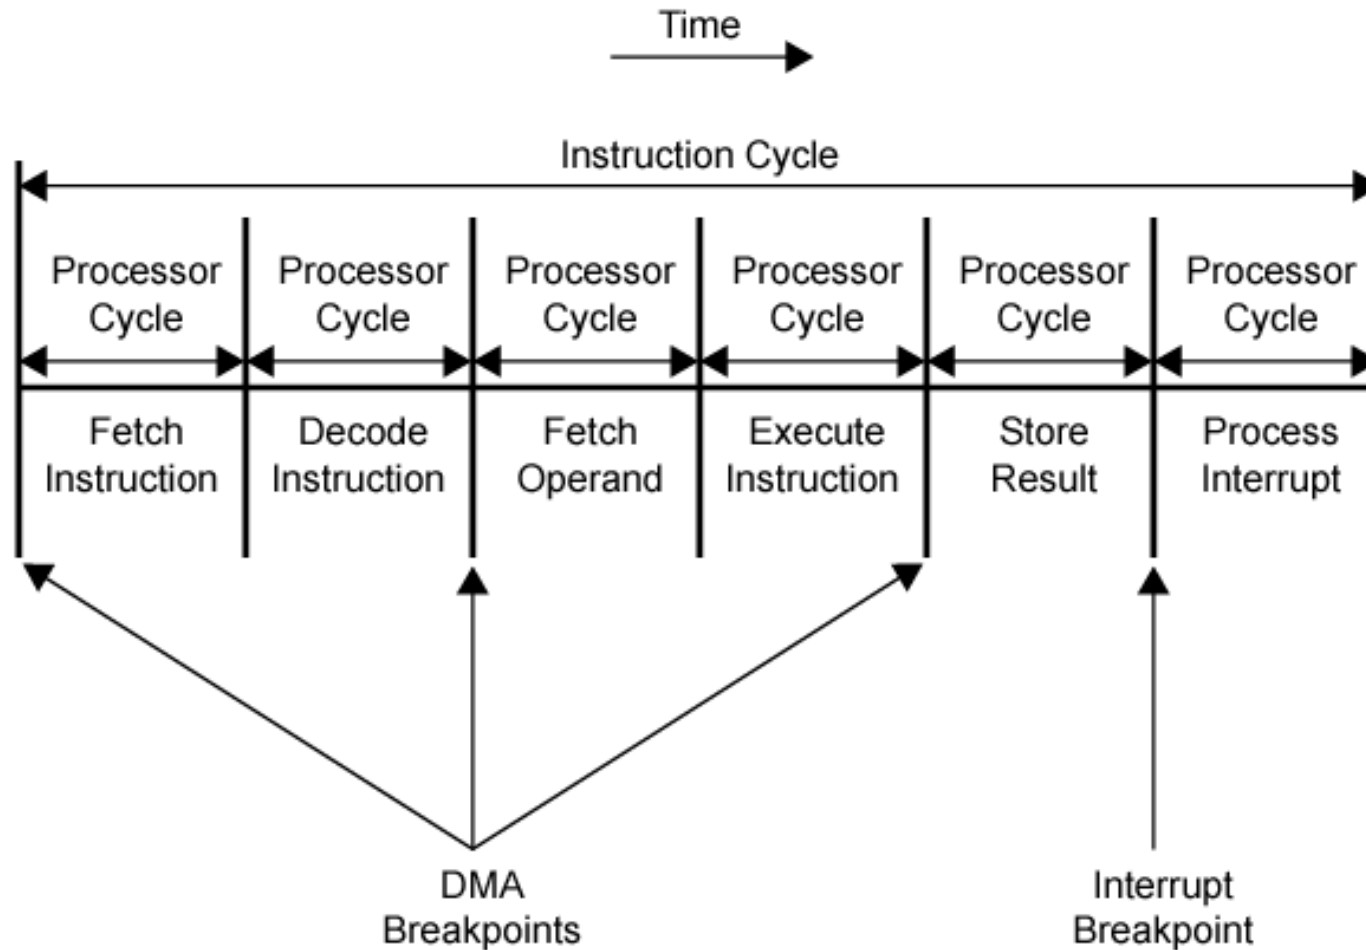
(c) Direct memory access

# DMA Transfer Cycle Stealing


- DMA controller takes over bus for a cycle
- Transfer of one word of data
- Not an interrupt
  - CPU does not switch context
- CPU suspended just before it accesses bus
  - i.e. before an operand or data fetch or a data write
- Slows down CPU but not as much as CPU doing transfer



# DMA and Interrupt Breakpoints During an Instruction Cycle




# DMA Configurations (1)



- Single Bus, Detached DMA controller
- Each transfer uses bus twice
  - I/O to DMA then DMA to memory
- CPU is suspended twice


## DMA Configurations (2)



(b) **Single-bus, Integrated DMA-I/O**

- Single Bus, Integrated DMA controller
- Controller may support >1 device
- Each transfer uses bus once
  - DMA to memory
- CPU is suspended once


# DMA Configurations (3)



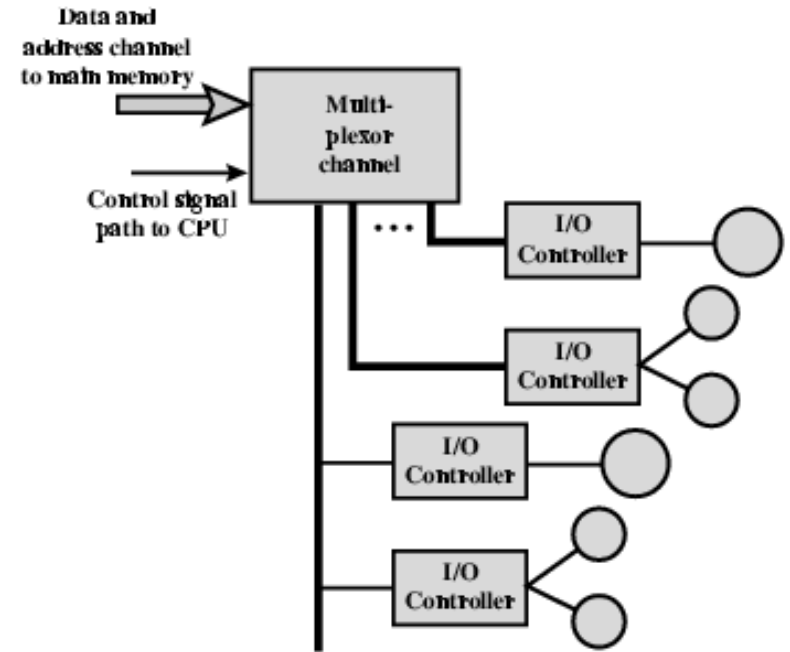
- Separate I/O Bus
- Bus supports all DMA enabled devices
- Each transfer uses bus once
  - DMA to memory
- CPU is suspended once

# I/O Channels

- I/O devices getting more sophisticated
- e.g. 3D graphics cards
- CPU instructs I/O controller to do transfer
- I/O controller does entire transfer
- Improves speed
  - Takes load off CPU
  - Dedicated processor is faster



(a) Selector



(b) Multiplexor

# Next:

Interfacing