



Advanced Parallel Architecture

Lesson 4 bis



Annalisa Massini - 2014/2015



Internal Memory

RAM

- ▶ Many memory types are random access
 - ▶ individual words of memory are directly accessed through wired-in addressing logic
- ▶ The most common is referred to as *random-access memory* (RAM)
 - ▶ misuse of the term, because also other types of memory are random access
 - ▶ One distinguishing characteristic of RAM is that it is possible both to read data from the memory and to write new data into the memory easily and rapidly
 - ▶ Both the reading and writing are accomplished through the use of electrical signals

Semiconductor Memory Types

Memory Type	Category	Erasure	Write Mechanism	Volatility
Random-access memory (RAM)	Read-write memory	Electrically, byte-level	Electrically	Volatile
Read-only memory (ROM)	Read-only memory	Not possible	Masks	Nonvolatile
Programmable ROM (PROM)			Electrically	
Erasable PROM (EPROM)	UV light, chip-level			
Electrically Erasable PROM (EEPROM)	Electrically, byte-level			
Flash memory	Electrically, block-level			

RAM

- ▶ RAM technology is divided into two technologies:
 - ▶ Dynamic
 - ▶ Static
- ▶ A **dynamic RAM** (DRAM) is made with cells that store data as charge on **capacitors**
- ▶ The presence or absence of charge in a capacitor is interpreted as a **binary 1 or 0**
- ▶ Because capacitors have a natural **tendency to discharge**, dynamic RAMs require periodic charge refreshing to maintain data storage → the term *dynamic*

RAM

- ▶ RAM technology is divided into two technologies:
 - ▶ Dynamic
 - ▶ Static
- ▶ A **static RAM** (SRAM) is a digital device that uses the same logic elements used in the processor
- ▶ Binary values are stored using traditional flip-flop logic-gate configurations
- ▶ A static RAM will hold its data as long as power is supplied to it

SRAM v DRAM

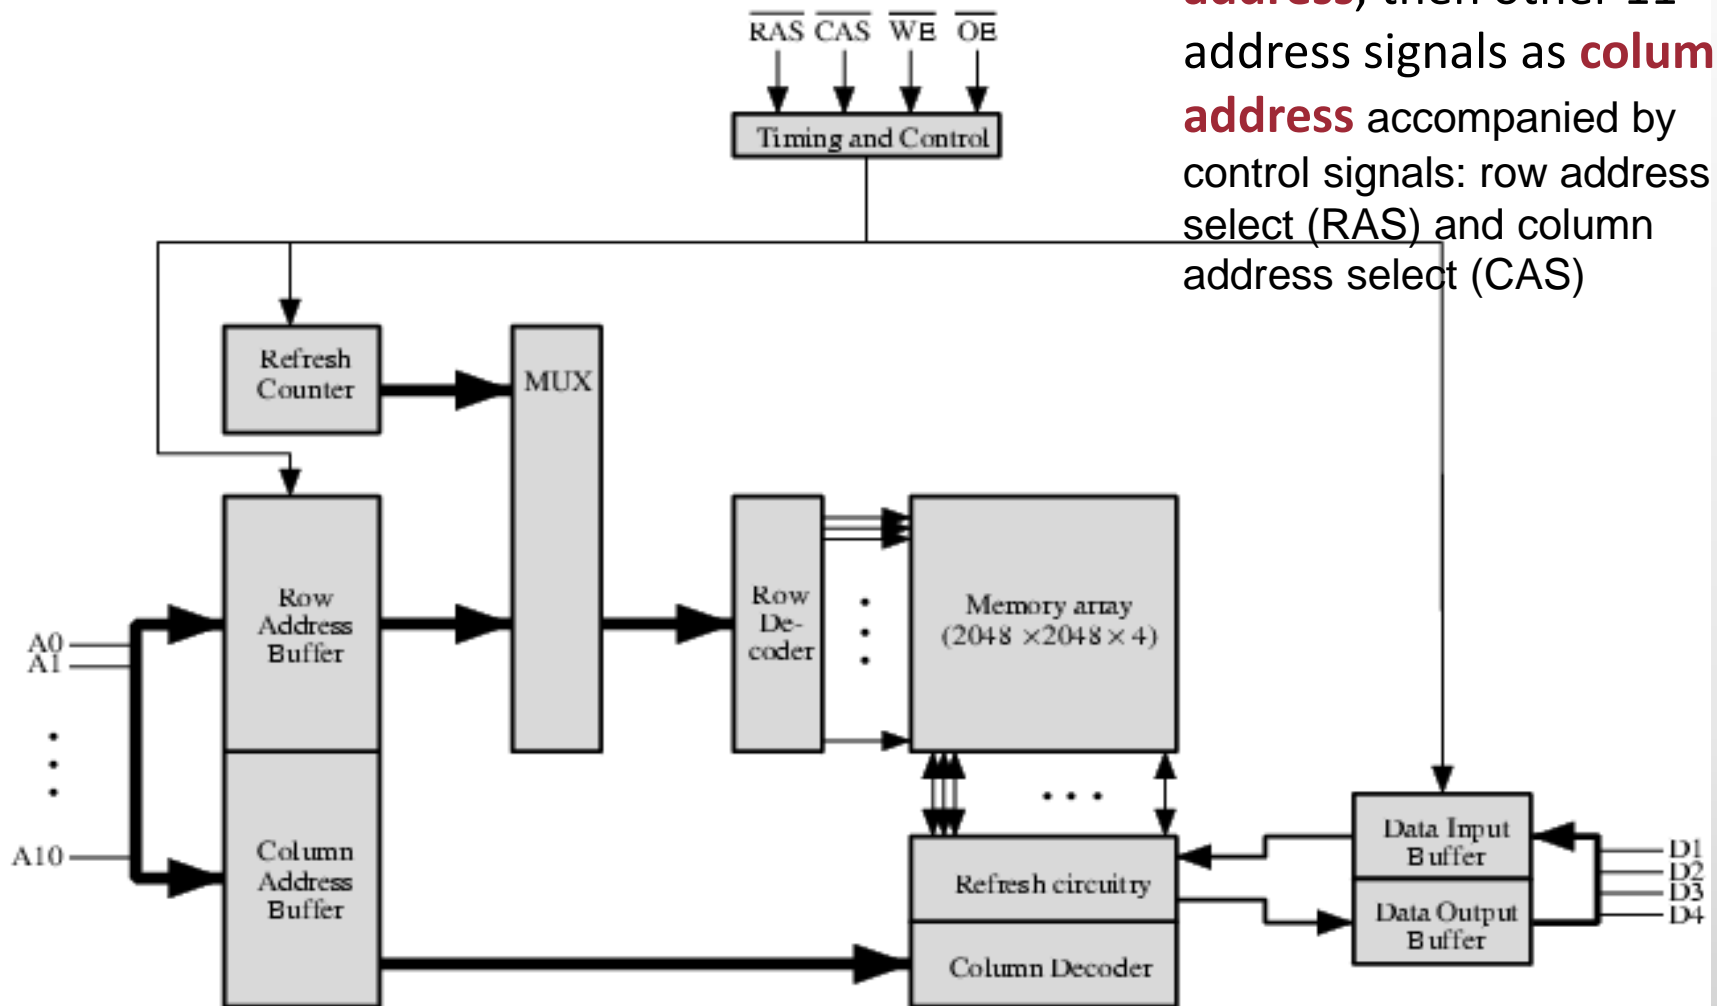
- ▶ Both volatile
 - ▶ Power needed to preserve data
- ▶ Dynamic cell
 - ▶ Simpler to build, smaller
 - ▶ More dense
 - ▶ Less expensive
 - ▶ Needs refresh circuitry
 - ▶ Larger memory units
- ▶ Static
 - ▶ Faster
 - ▶ Cache (both on and off chip)

Example of chip organization

- ▶ A 16Mbit chip can be organised as:
 - ▶ 1M of 16 bit words
 - ▶ 1-bit-per-chip that is 16 lots of 1Mbit chip with bit 1 of each word in chip 1 and so on
- ▶ For example a 16Mbit chip can be organised as a 2048 x 2048 x 4bit array
 - ▶ Reduces number of address pins
 - ▶ Multiplex row address and column address
 - ▶ 11 pins to address ($2^{11}=2048$)
 - ▶ Adding one more pin doubles range of values so x4 capacity

Typical 16 Mb DRAM (4M x 4)

11 address signals are passed to the chip as **row address**, then other 11 address signals as **column address** accompanied by control signals: row address select (RAS) and column address select (CAS)

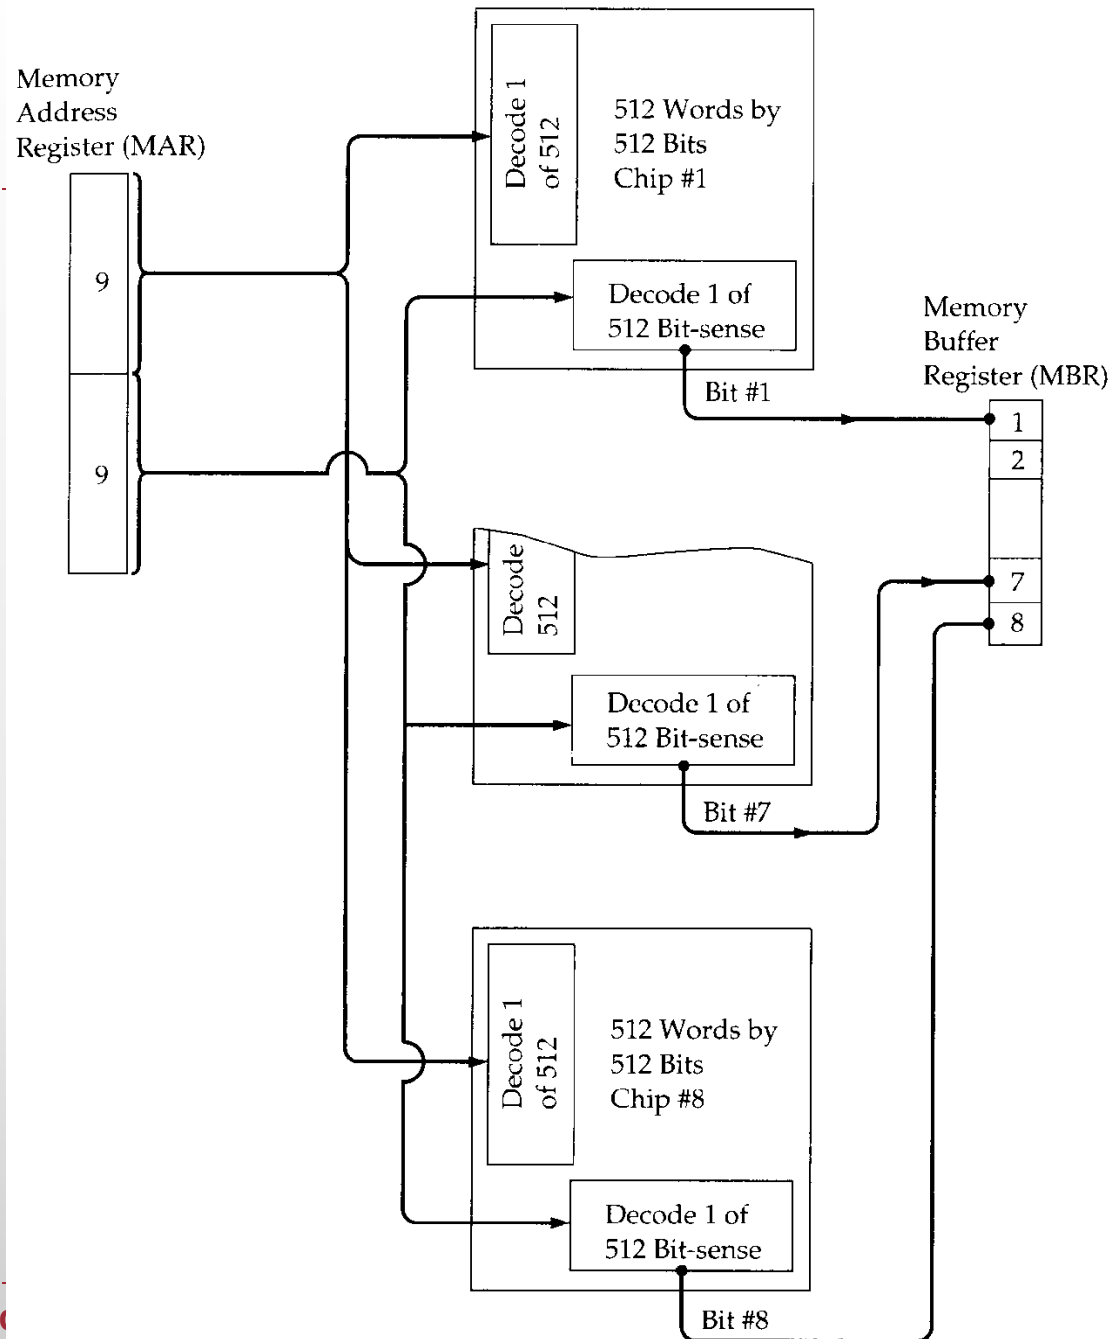


Example

How a memory module consisting of 256K 8-bit words could be organized

256K words → **18-bit address** supplied to the module from some external source (e.g., the address lines of a bus to which the module is attached)

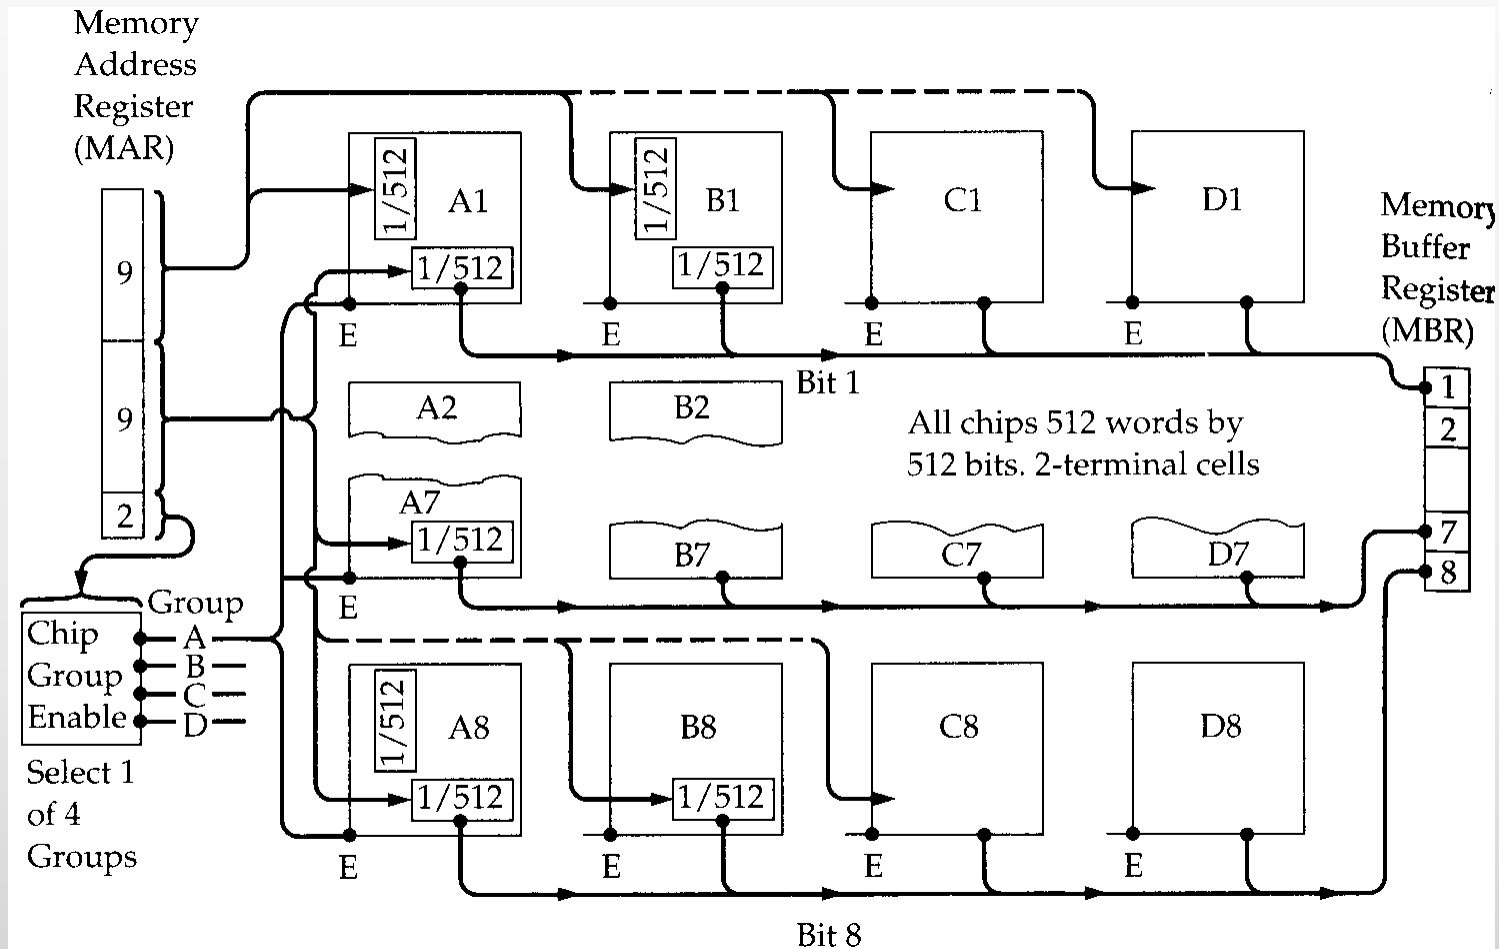
The address is presented to **8 256K 1-bit chips**, each of which provides the input/output of 1 bit



1MByte Module Organisation

Possible organization of a memory consisting of **1M word** by **8 bits per word**

Four columns of chips, each column containing **256K words** arranged as before



Interleaved Memory

- ▶ **Main memory** is composed of a collection of **DRAM memory chips** that can be grouped together to form a *memory bank*
- ▶ It is possible to organize the memory banks in a way known as **interleaved memory**
- ▶ Each bank is independently able
 - ▶ to service a memory read or write request
- ▶ a system with **K banks** can service **K requests simultaneously**, increasing memory read or write rates by a factor of K
- ▶ If **consecutive words** of memory are stored **in different banks**, then the transfer of a block is speeded up

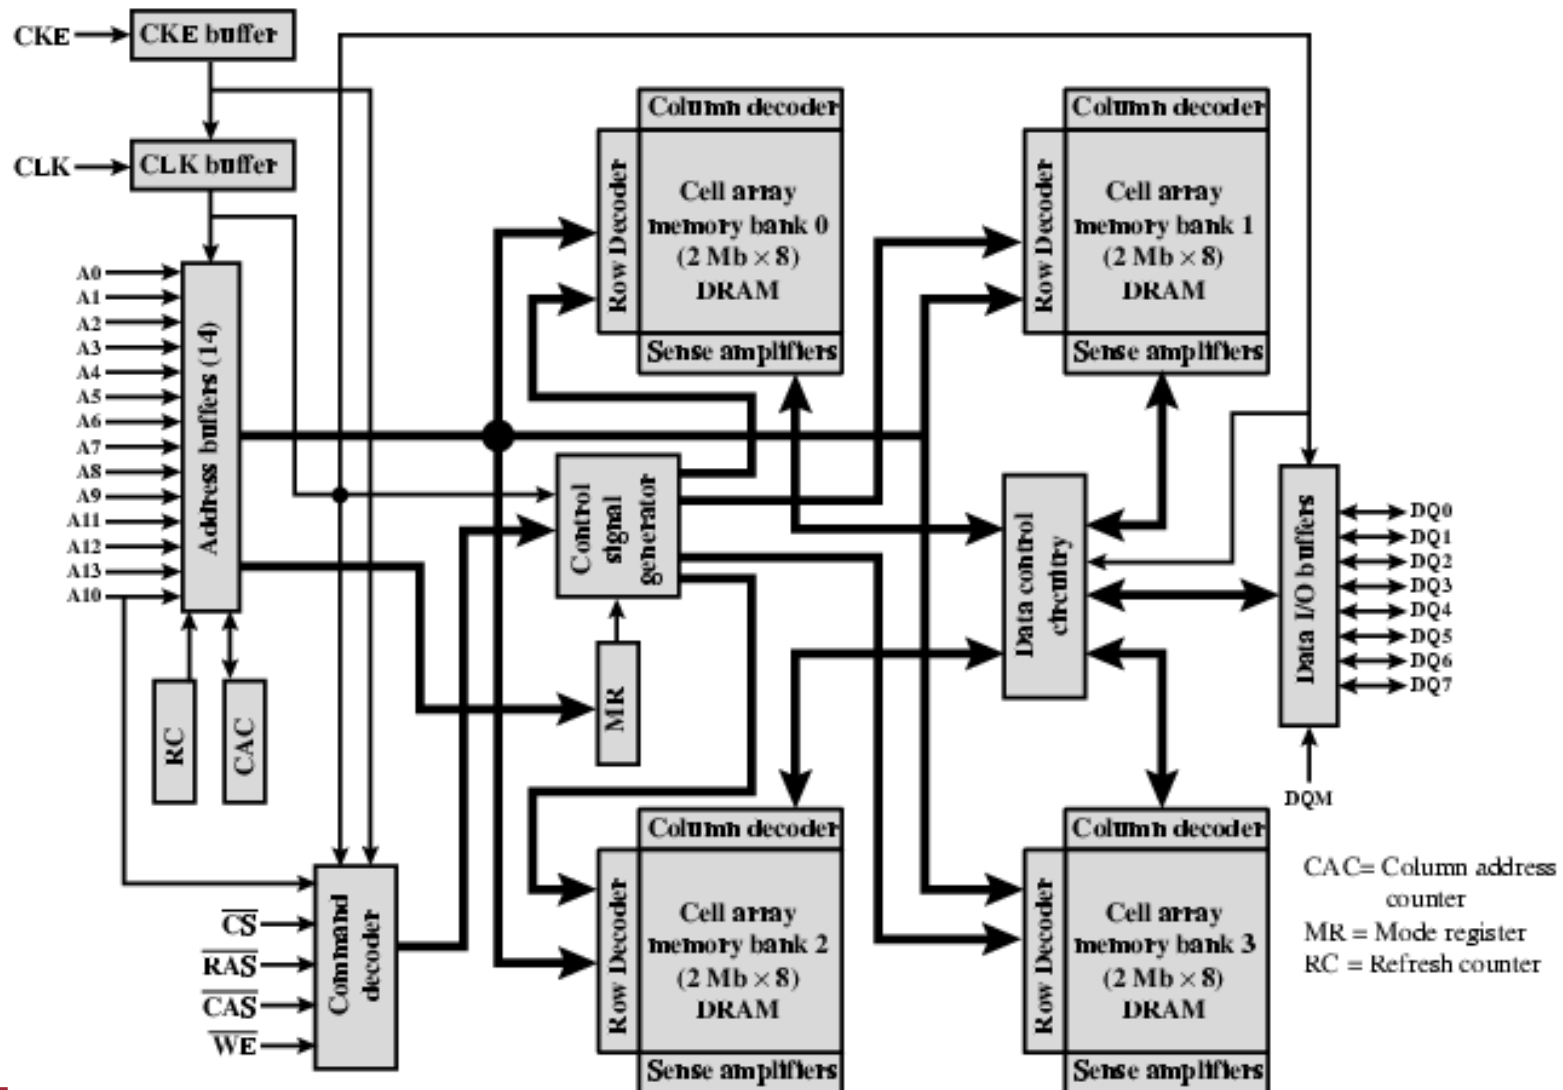
Advanced DRAM Organization

- ▶ Basic DRAM same since first RAM chips
- ▶ Enhanced DRAM
 - ▶ Contains small SRAM as well
 - ▶ SRAM holds last line read (c.f. Cache!)
- ▶ Cache DRAM
 - ▶ Larger SRAM component
 - ▶ Use as cache or serial buffer

Synchronous DRAM (SDRAM)

- ▶ Access is synchronized with **an external clock**
- ▶ Address is presented to RAM
- ▶ RAM finds data (CPU waits in conventional DRAM)
- ▶ Since SDRAM moves data in time with system clock, CPU knows when data will be ready
- ▶ **CPU does not have to wait**, it can do something else
- ▶ Burst mode allows SDRAM to set up stream of data and fire it out in block
- ▶ DDR-SDRAM sends data twice per clock cycle (leading & trailing edge)

SDRAM



DDR SDRAM

- ▶ SDRAM can only send data once per clock
- ▶ Double-data-rate SDRAM can send data twice per clock cycle
 - ▶ Rising edge and falling edge

Solid-state drive (SSD)

- ▶ A **solid-state drive** is a **solid-state storage** device that uses **integrated circuit** assemblies as memory to store data **persistently**
- ▶ SSD technology primarily uses electronic interfaces compatible with traditional block input/output (I/O) hard disk drives (HDDs)
- ▶ New I/O interfaces (like SATA Express) have been designed to address specific requirements of the SSD technology
- ▶ SSDs have ***no moving mechanical components***
 - ▶ This distinguishes them from traditional electromechanical magnetic disks such as hard disk drives (HDDs) or floppy disks, which contain spinning disks and movable read/write heads

Solid-state drive (SSD)

- ▶ Compared with electromechanical disks, SSDs:
 - ▶ Are more resistant to physical shock
 - ▶ Are run silently
 - ▶ Have lower access time
 - ▶ Have lower latency
 - ▶ Are more expensive per unit of storage than HDDs (even if price has continued to decline)

Solid-state drive (SSD)

- ▶ Since 2009 NAND flash non-volatile memory
- ▶ Every SSD includes a controller:
 - ▶ that incorporates the electronics that bridge the NAND memory components to the host computer
 - ▶ embedded processor that executes firmware-level code → important factor of SSD performance

Read Only Memory (ROM)

- ▶ Permanent storage
 - ▶ Nonvolatile
- ▶ Microprogramming
- ▶ Library subroutines
- ▶ Systems programs (BIOS)
- ▶ Function tables

Types of ROM

- ▶ Written during manufacture
 - ▶ Very expensive for small runs
- ▶ Programmable (once)
 - ▶ PROM
 - ▶ Needs special equipment to program
- ▶ Read “mostly”
 - ▶ Erasable Programmable (EPROM)
 - ▶ Erased by UV
 - ▶ Electrically Erasable (EEPROM)
 - ▶ Takes much longer to write than read
 - ▶ Flash memory
 - ▶ Erase whole memory electrically



External Memory

Types of External Memory

- ▶ Magnetic Disk
 - ▶ RAID
 - ▶ Removable
- ▶ Optical
 - ▶ CD-ROM
 - ▶ CD-Recordable (CD-R)
 - ▶ CD-R/W
 - ▶ DVD
- ▶ Magnetic Tape

Magnetic Disk

- ▶ Disk substrate coated with magnetizable material
- ▶ Substrate used to be aluminium
- ▶ Now glass
 - ▶ Improved surface uniformity
 - ▶ Increases reliability
 - ▶ Reduction in surface defects
 - ▶ Reduced read/write errors
 - ▶ Lower flight heights (See later)
 - ▶ Better stiffness
 - ▶ Better shock/damage resistance

Read and Write Mechanisms

- ▶ Recording & retrieval via conductive coil called a head
- ▶ May be single read/write head or separate ones
- ▶ During read/write, head is stationary, platter rotates
- ▶ Write
 - ▶ Current through coil produces magnetic field
 - ▶ Pulses sent to head
 - ▶ Magnetic pattern recorded on surface below
 - ▶ Higher storage density and speed

Read and Write Mechanisms

- ▶ Read (traditional)

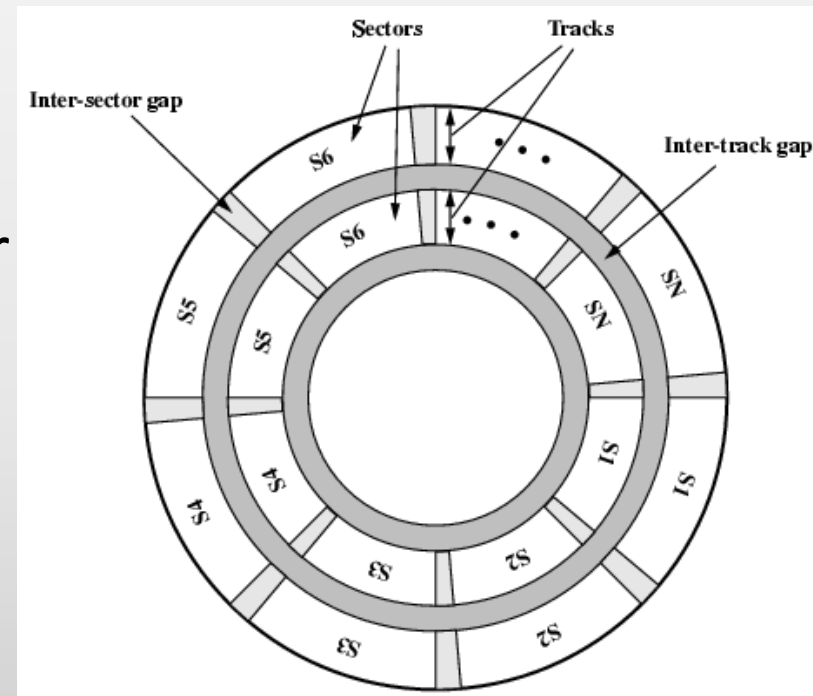
- ▶ Magnetic field moving relative to coil produces current
- ▶ Coil is the same for read and write

- ▶ Read (contemporary)

- ▶ Separate read head, close to write head
- ▶ Partially shielded magneto resistive (MR) sensor
- ▶ Electrical resistance depends on direction of magnetic field
- ▶ High frequency operation
 - ▶ Higher storage density and speed

Data Organization and Formatting

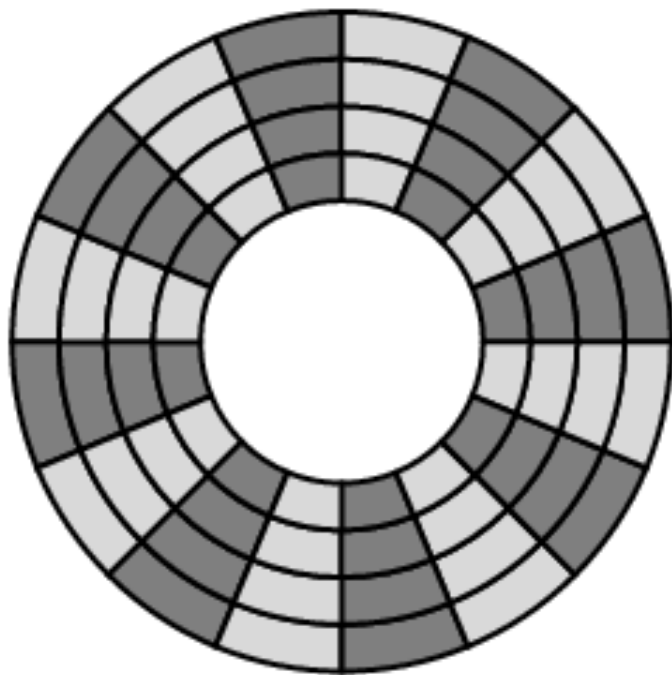
- ▶ Concentric rings or tracks
 - ▶ Gaps between tracks
 - ▶ Reduce gap to increase capacity
 - ▶ Same number of bits per track (variable packing density)
 - ▶ Constant angular velocity
- ▶ Tracks divided into sectors
- ▶ Minimum block size is one sector
- ▶ May have more than one sector per block



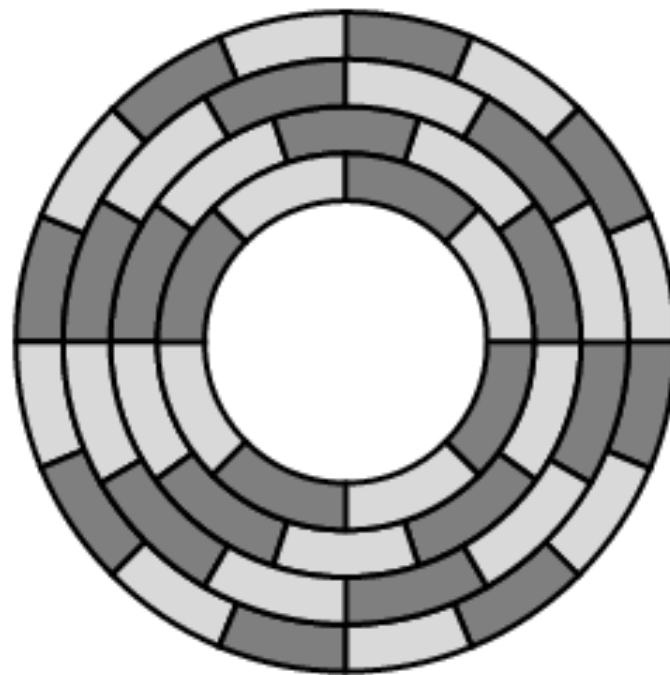
Disk Velocity

- ▶ Bit near centre of rotating disk passes fixed point slower than bit on outside of disk
- ▶ Increase spacing between bits in different tracks
- ▶ Rotate disk at constant angular velocity (CAV)
 - ▶ Gives pie shaped sectors and concentric tracks
 - ▶ Individual tracks and sectors addressable
 - ▶ Move head to given track and wait for given sector
 - ▶ Waste of space on outer tracks
 - ▶ Lower data density
- ▶ Can use zones to increase capacity
 - ▶ Each zone has fixed bits per track
 - ▶ More complex circuitry

Disk Layout Methods Diagram



(a) Constant angular velocity

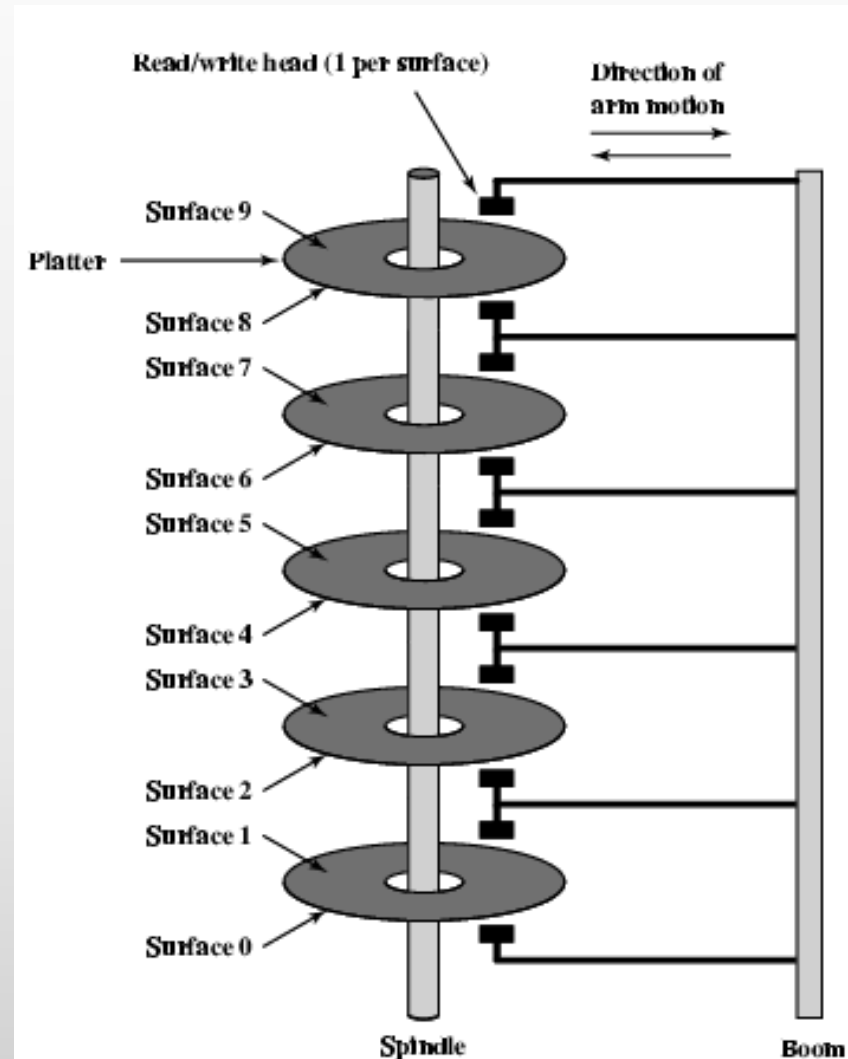


(b) Multiple zoned recording

Must be able to identify start of track and sector
Format disk (Additional information not available to user)

Multiple Platter

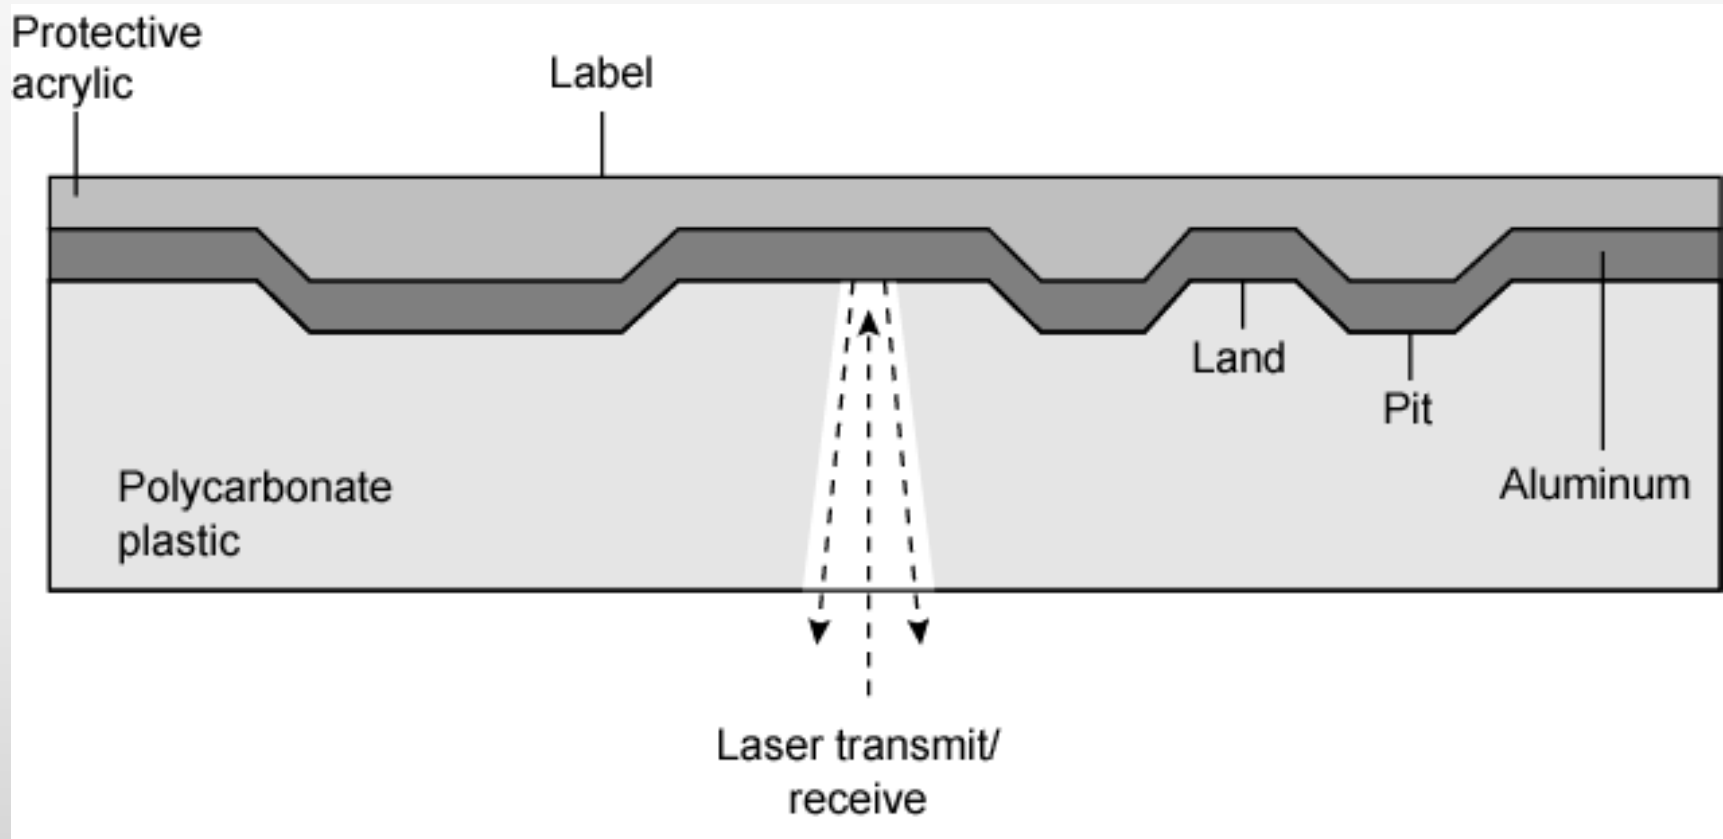
- ▶ One head per side
- ▶ Heads are joined and aligned
- ▶ Aligned tracks on each platter form cylinders
- ▶ Data is striped by cylinder
 - ▶ reduces head movement
 - ▶ Increases speed (transfer rate)



Optical Storage CD-ROM

- ▶ Originally for audio
- ▶ 650Mbytes giving over 70 minutes audio
- ▶ Polycarbonate coated with highly reflective coat, usually aluminium
- ▶ Data stored as pits
- ▶ Read by reflecting laser
- ▶ Constant packing density
- ▶ Constant linear velocity

CD Operation



Other Optical Storage

- ▶ CD-Recordable (CD-R)

- ▶ WORM
- ▶ Now affordable
- ▶ Compatible with CD-ROM drives

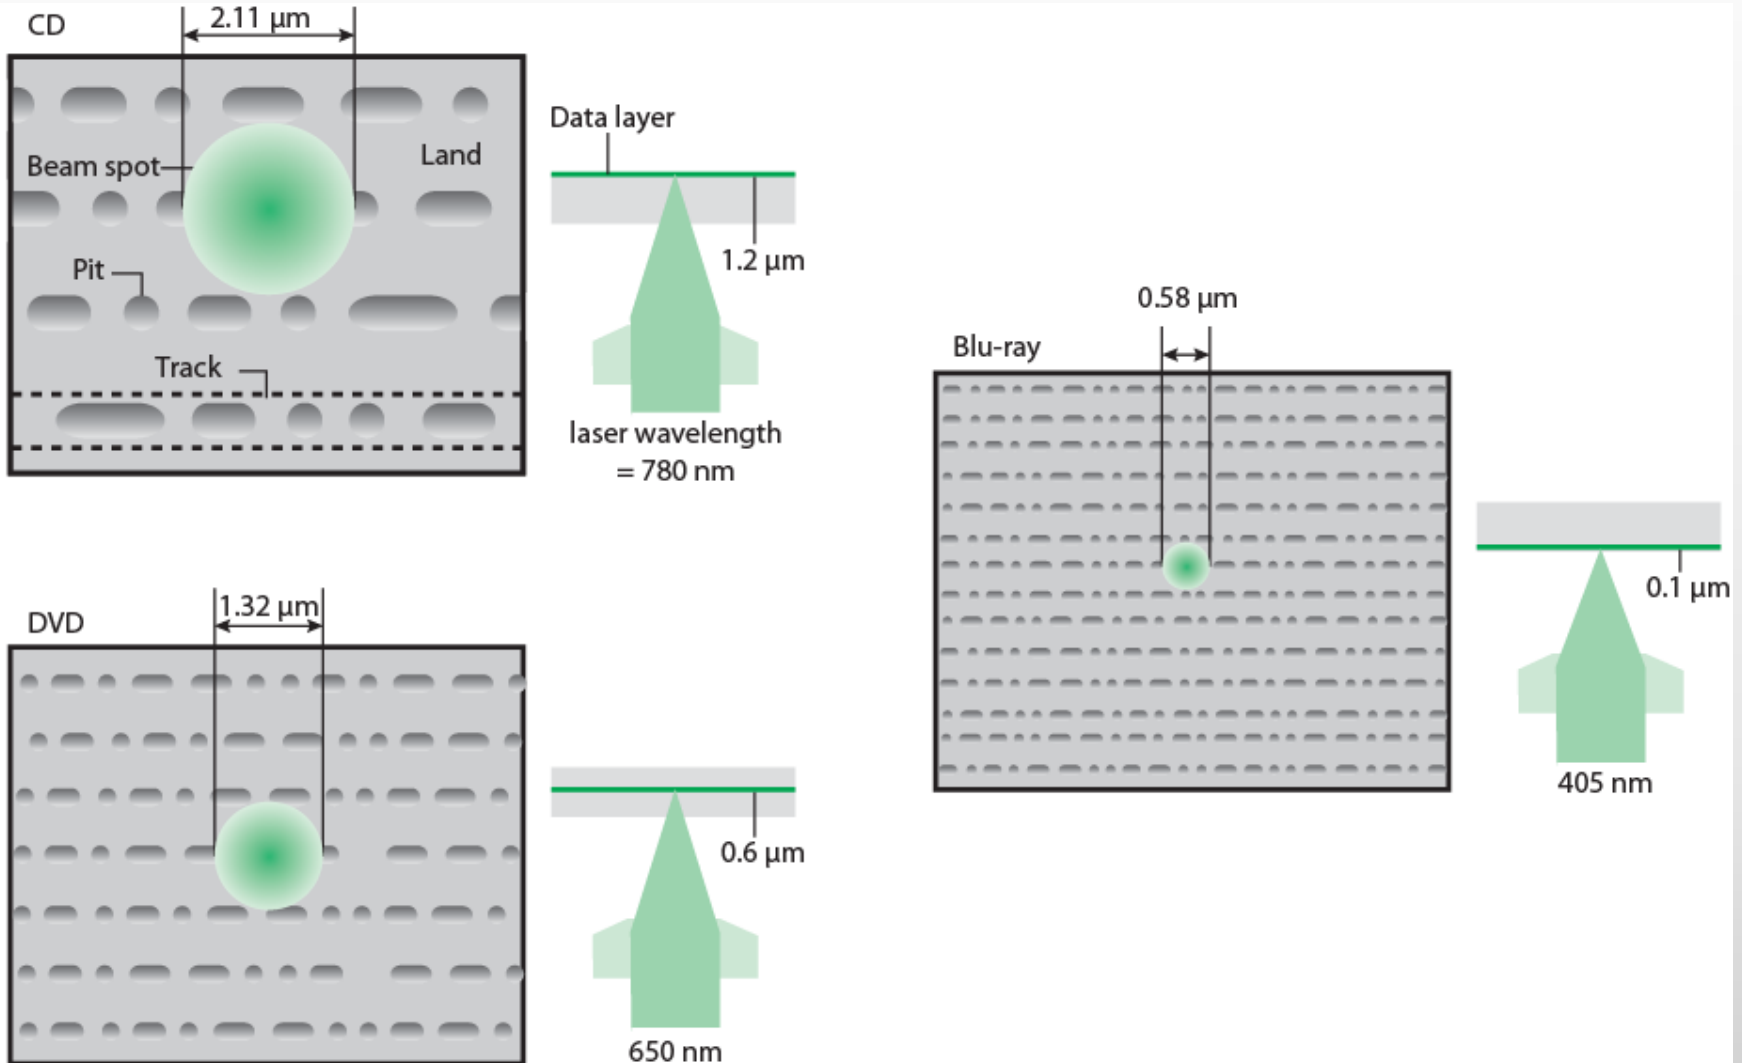
- ▶ CD-RW

- ▶ Erasable
- ▶ Getting cheaper
- ▶ Mostly CD-ROM drive compatible
- ▶ Phase change
 - ▶ Material has two different reflectivities in different phase states

DVD and High Definition Optical Disks

- ▶ DVD - Digital Versatile Disk
 - ▶ Multi-layer with very high capacity (4.7G per layer)
- ▶ High Definition Optical Disk
 - ▶ Much higher capacity than DVD
 - ▶ Shorter wavelength laser (Blue-violet range)
 - ▶ Smaller pits
- ▶ HD-DVD
 - ▶ 15GB single side single layer
- ▶ Blue-ray
 - ▶ Data layer closer to laser (Tighter focus, less distortion, smaller pits)
 - ▶ 25GB on single layer

Optical Memory Characteristics



Magnetic Tape

- ▶ Serial access
- ▶ Slow
- ▶ Very cheap
- ▶ Backup and archive
- ▶ Linear Tape-Open (LTO) Tape Drives
 - ▶ Developed late 1990s
 - ▶ Open source alternative to proprietary tape systems



Input Output

Input/Output Problems

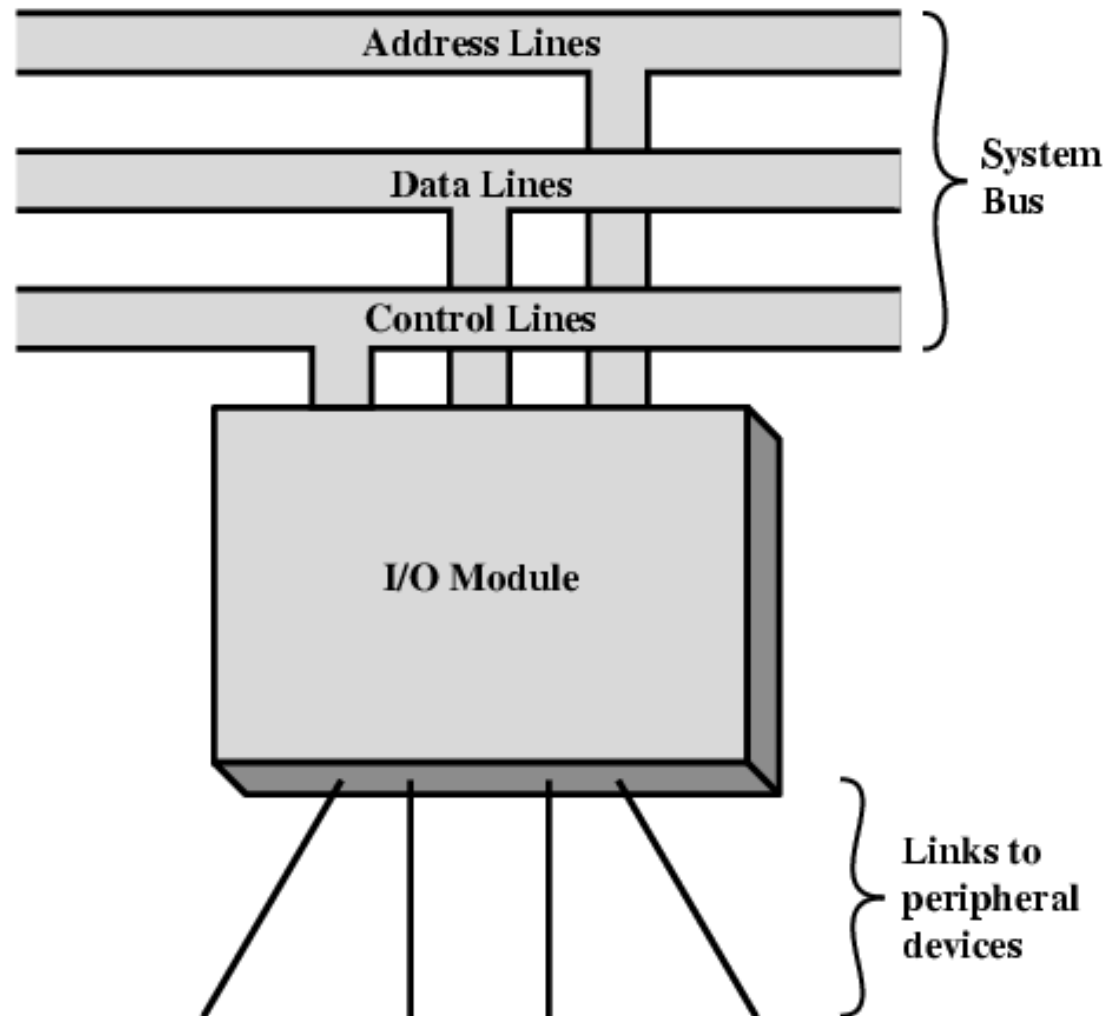
Input/Output Problems

- ▶ 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

Input/Output module

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

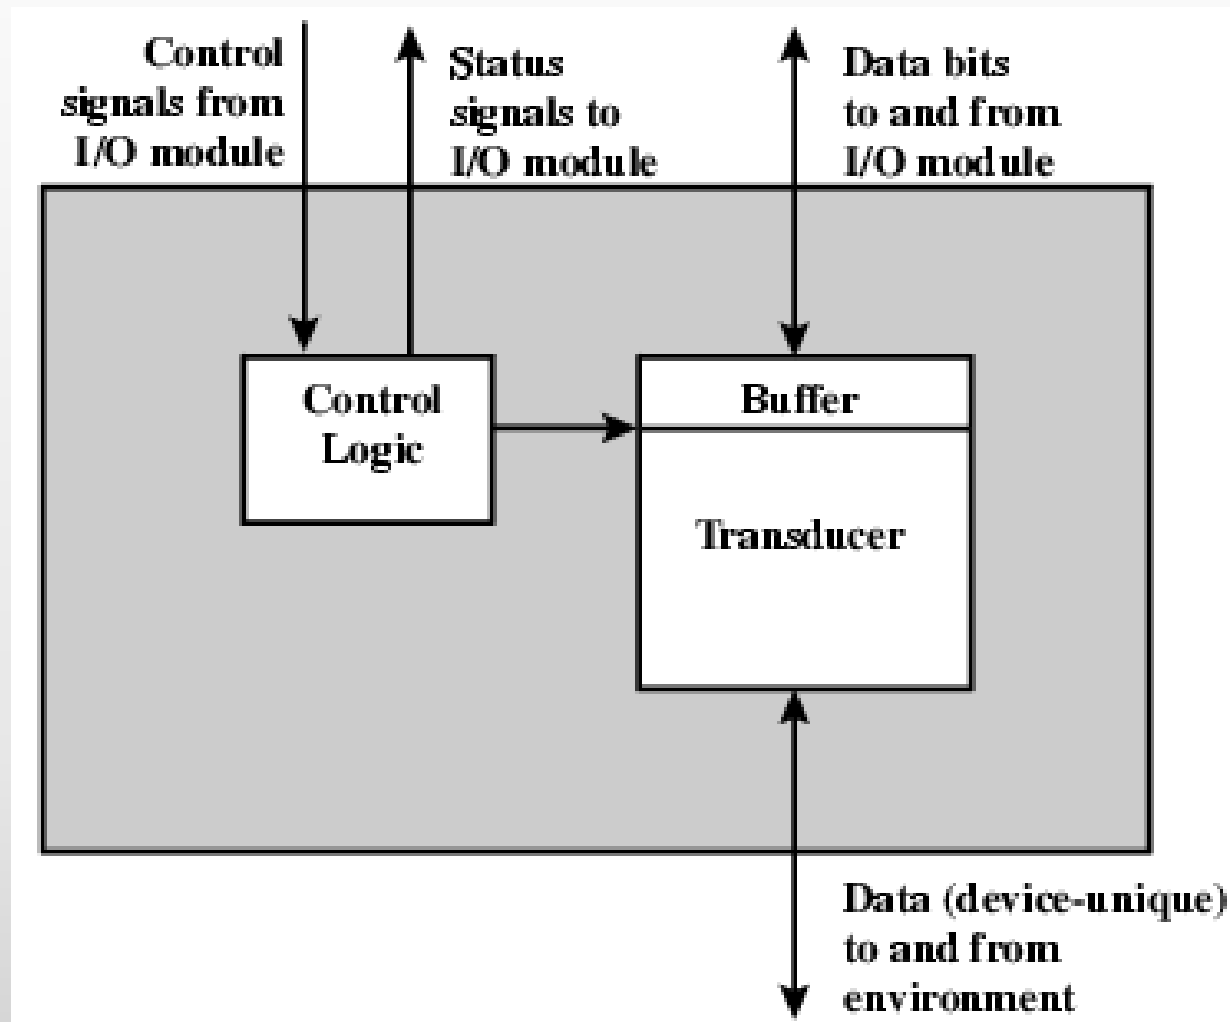
Generic Model of I/O Module



External Devices

- ▶ Human readable
 - ▶ Screen, printer, keyboard
- ▶ Machine readable
 - ▶ Monitoring and control
- ▶ Communication
 - ▶ Modem
 - ▶ Network Interface Card (NIC)

External Device Block Diagram



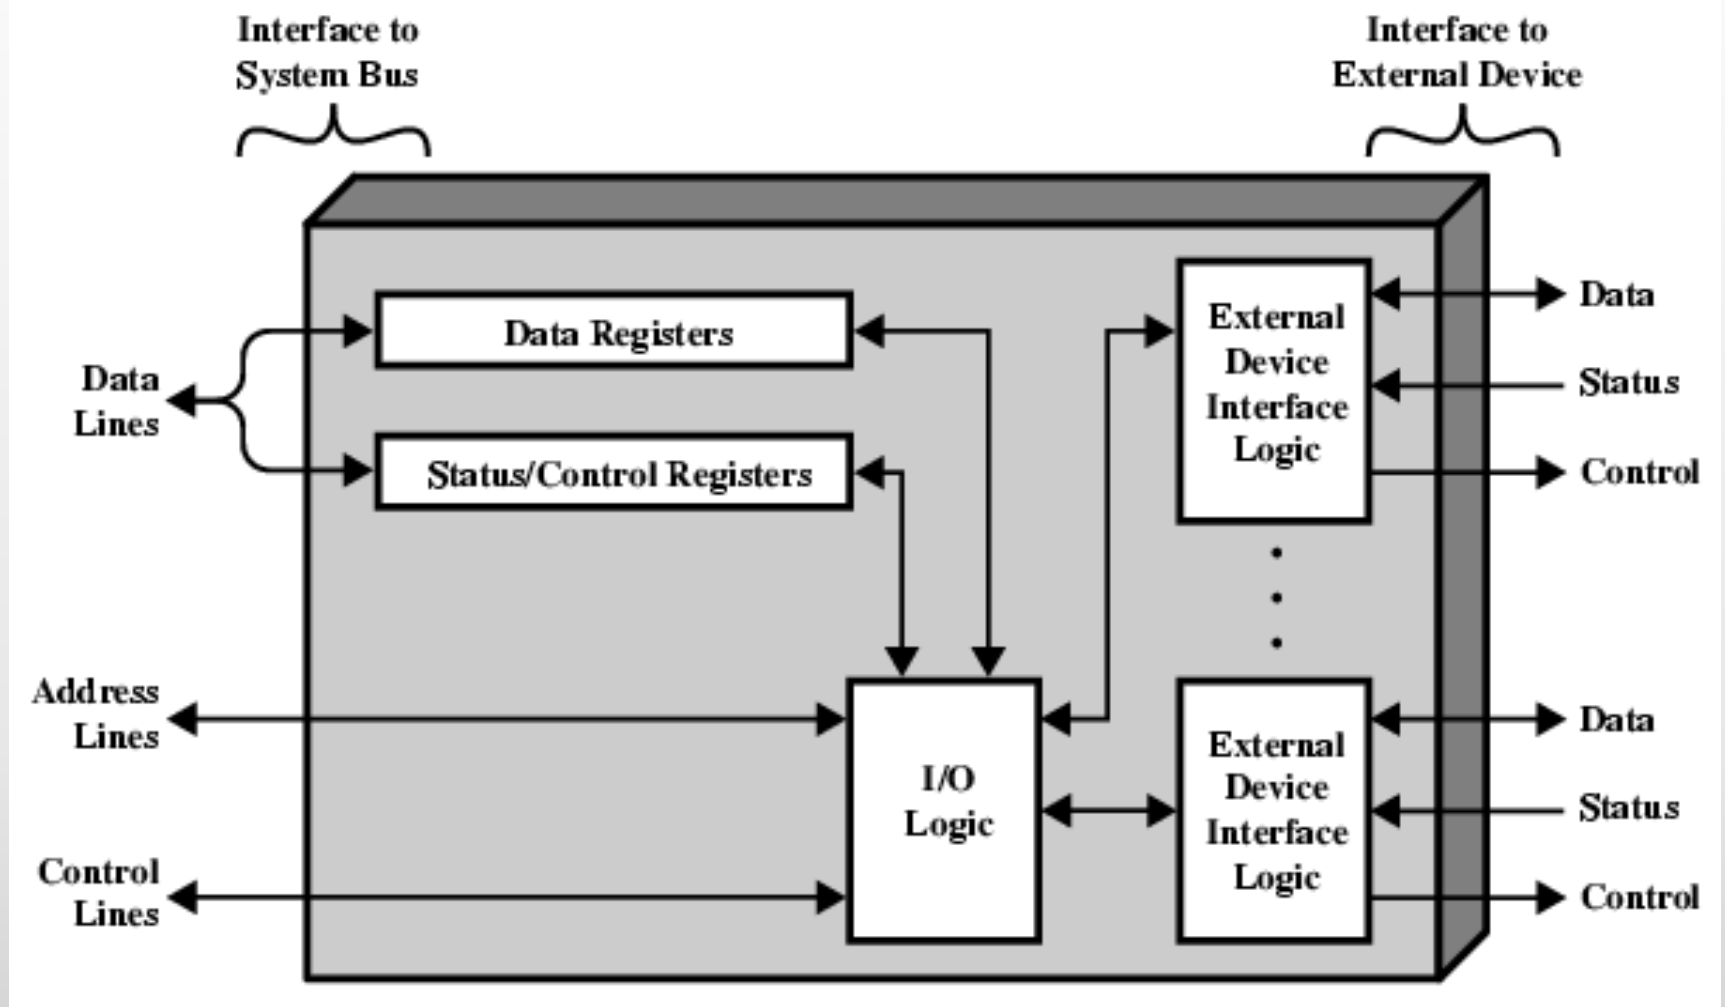
I/O Module Function

- ▶ Control & Timing
- ▶ CPU Communication
- ▶ Device Communication
- ▶ Data Buffering
- ▶ 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 Diagram



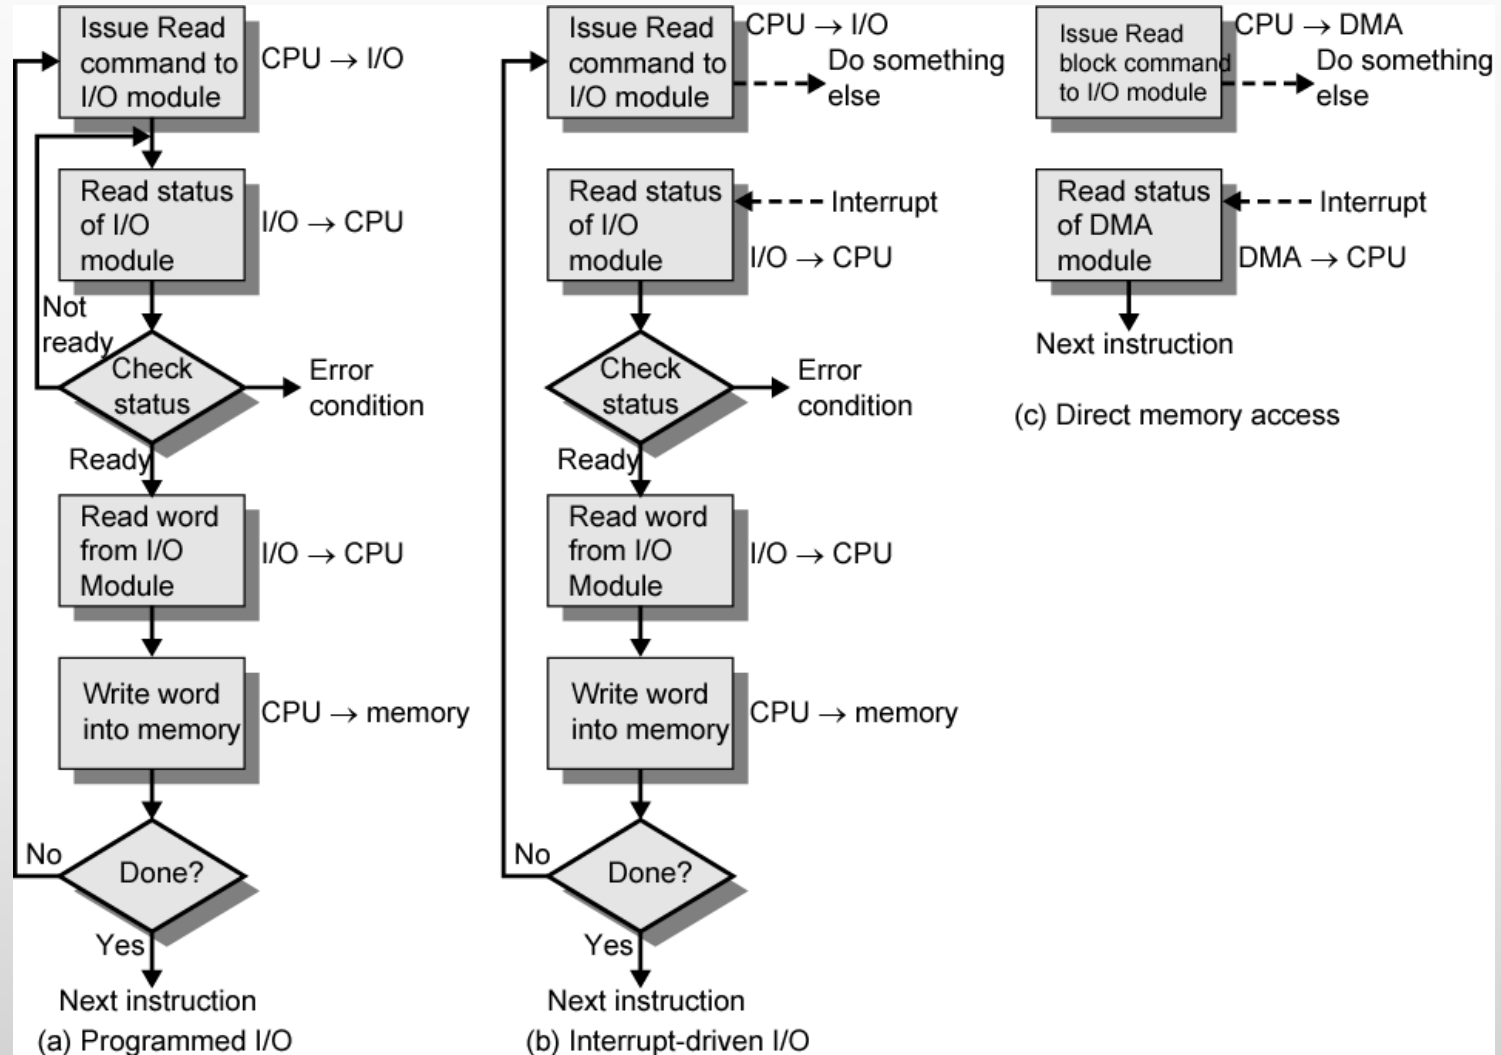
I/O Module Decisions

- ▶ 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
- ▶ Interrupt driven
- ▶ Direct Memory Access (DMA)

Three Techniques for Input of a Block of Data



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 - detail

- ▶ 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

I/O Commands

- ▶ CPU issues address
 - ▶ Identifies module (& device if >1 per module)
- ▶ CPU issues command
 - ▶ Control - telling module what to do
 - ▶ e.g. spin up disk
 - ▶ Test - check status
 - ▶ e.g. power? Error?
 - ▶ Read/Write
 - ▶ Module transfers data via buffer from/to 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

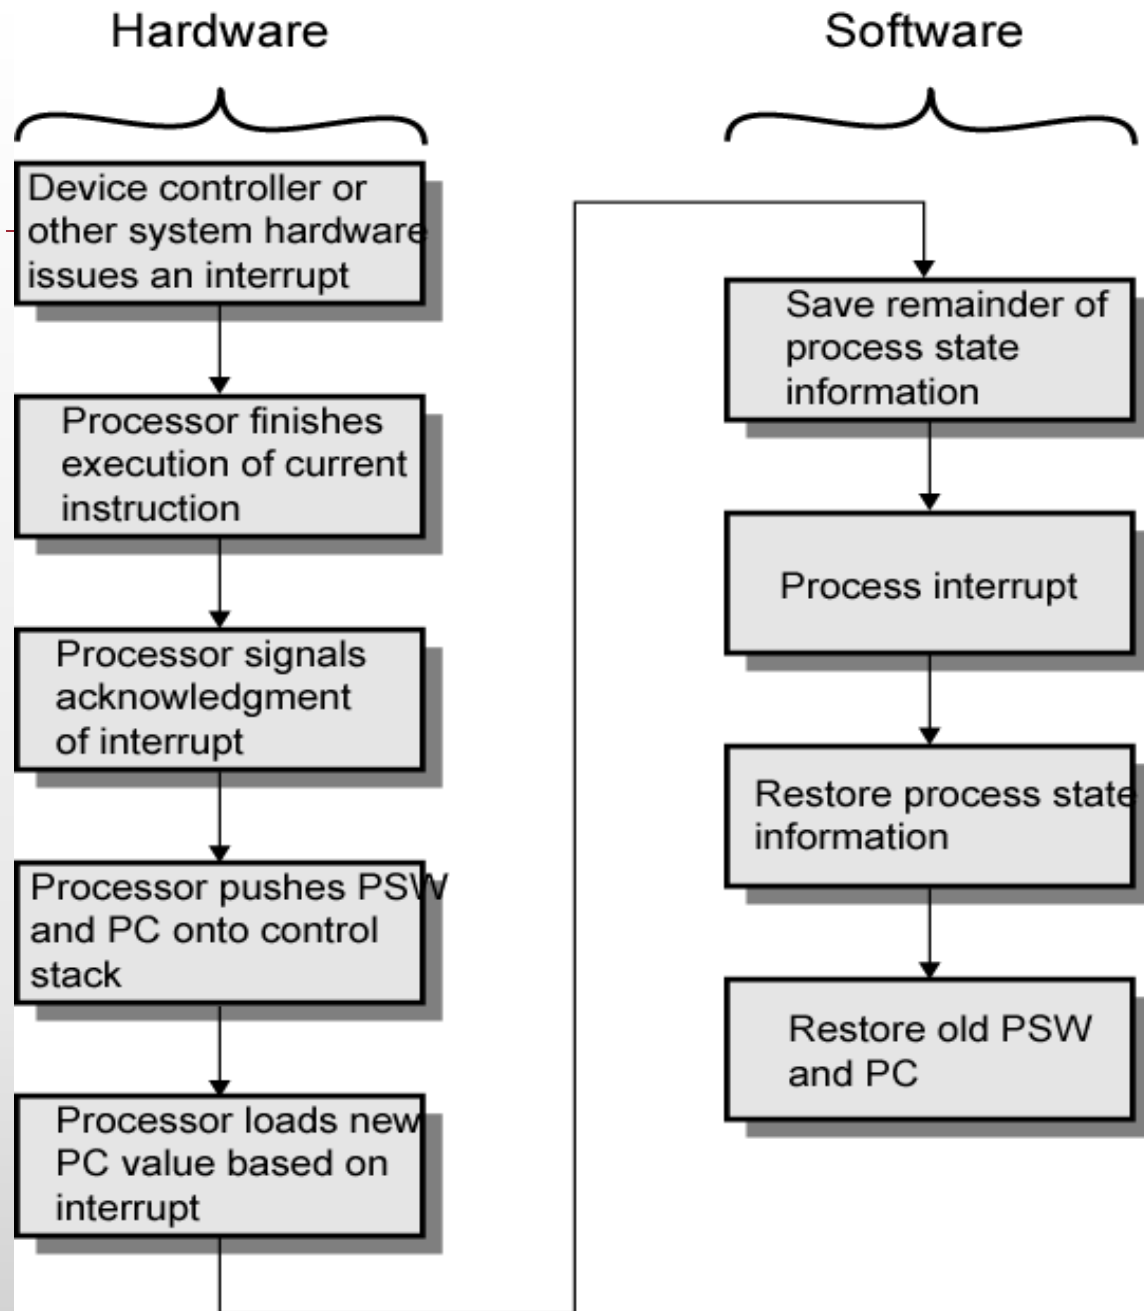
Interrupt Driven I/O

- ▶ 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

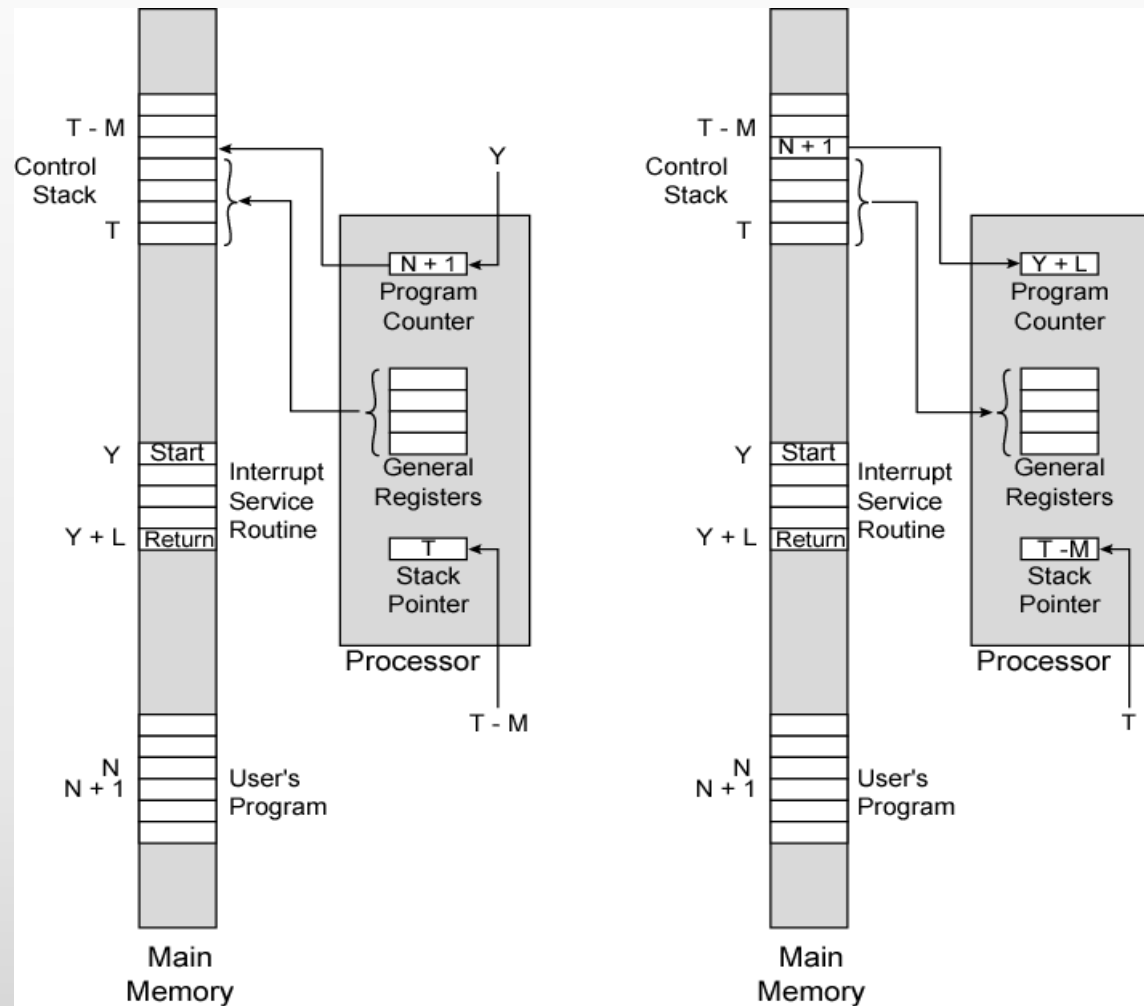
Simple Interrupt Processing



CPU Viewpoint

- ▶ 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
- ▶ See Operating Systems notes

Changes in Memory and Registers for an Interrupt



(a) Interrupt occurs after instruction at location N

(b) Return from interrupt

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

Identifying Interrupting Module

- ▶ **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

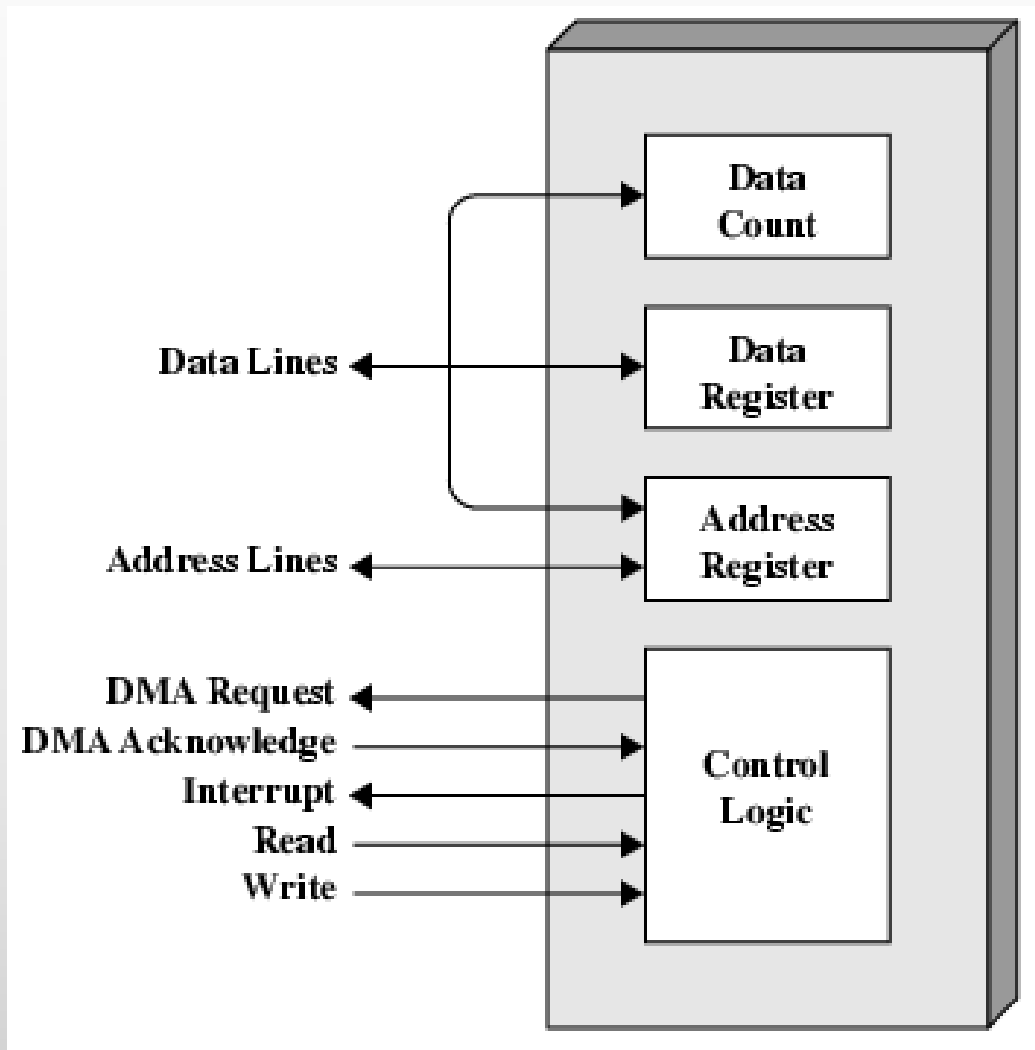
Multiple Interrupts

- ▶ 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

Typical DMA Module Diagram



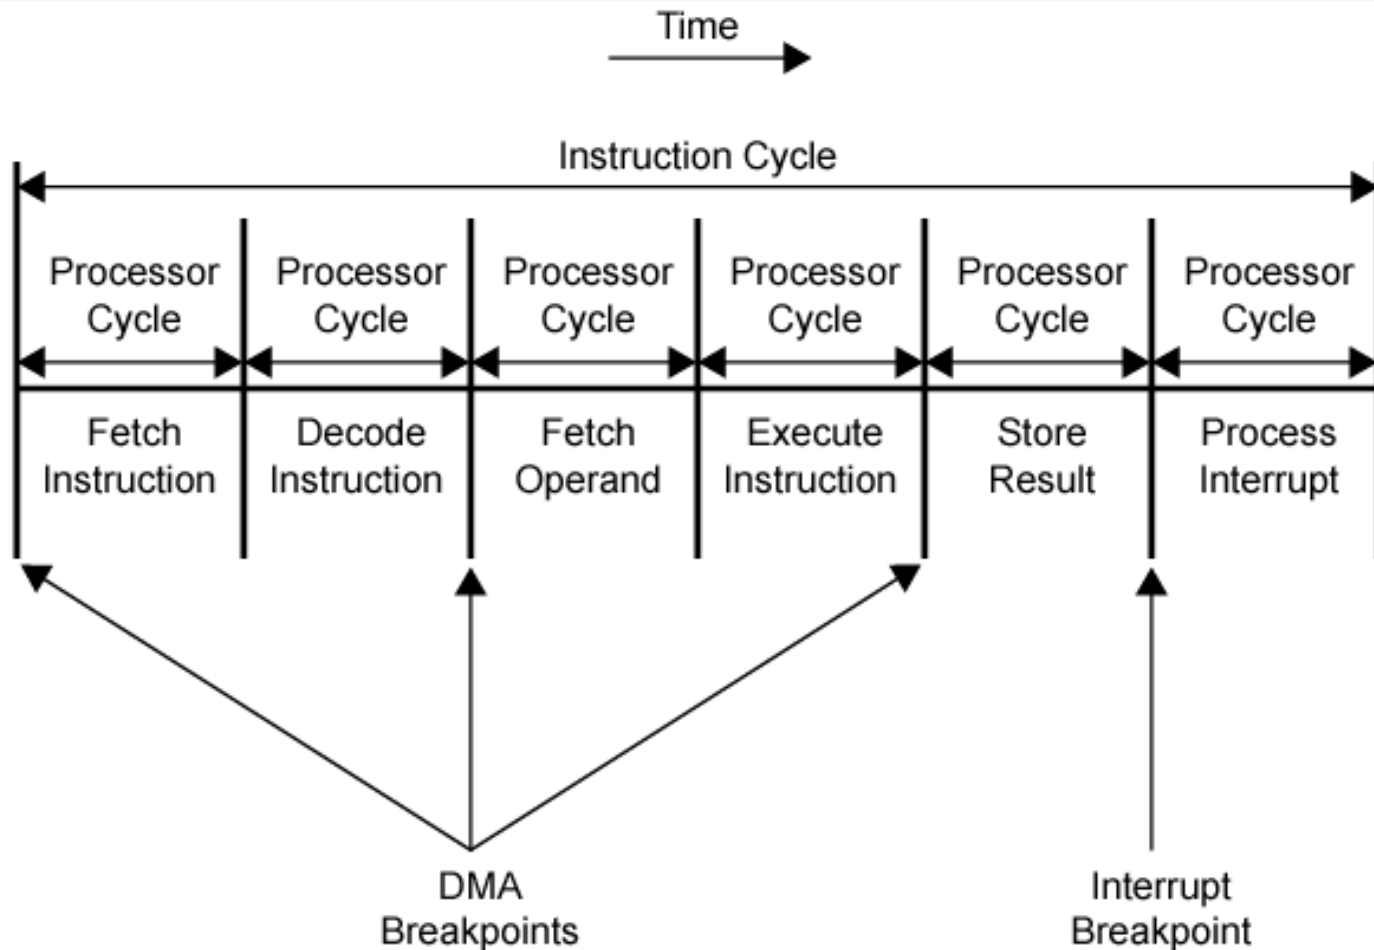
DMA Operation

- ▶ CPU tells DMA controller:
 - ▶ Read/Write
 - ▶ Device address
 - ▶ 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

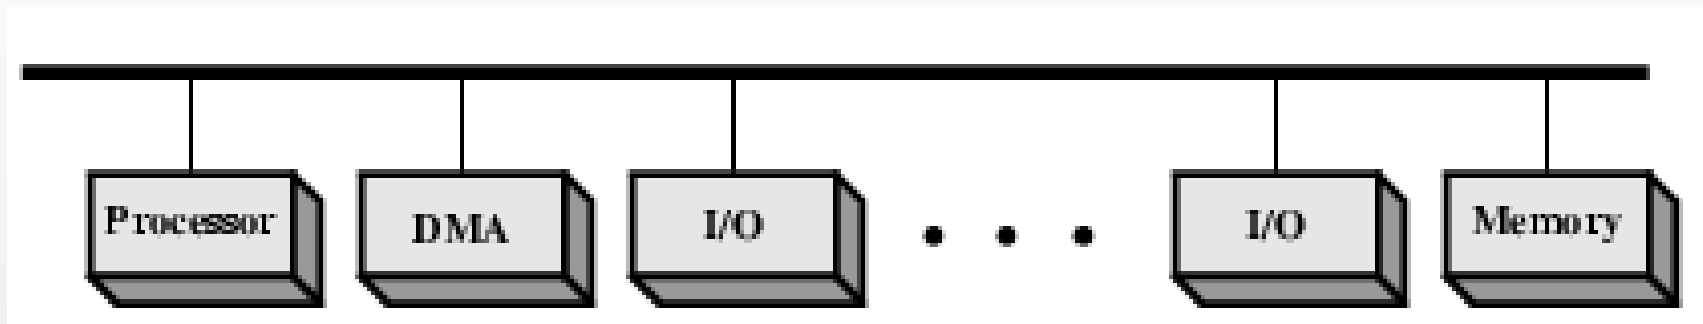
DMA Transfer - Cycle Stealing

- ▶ 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

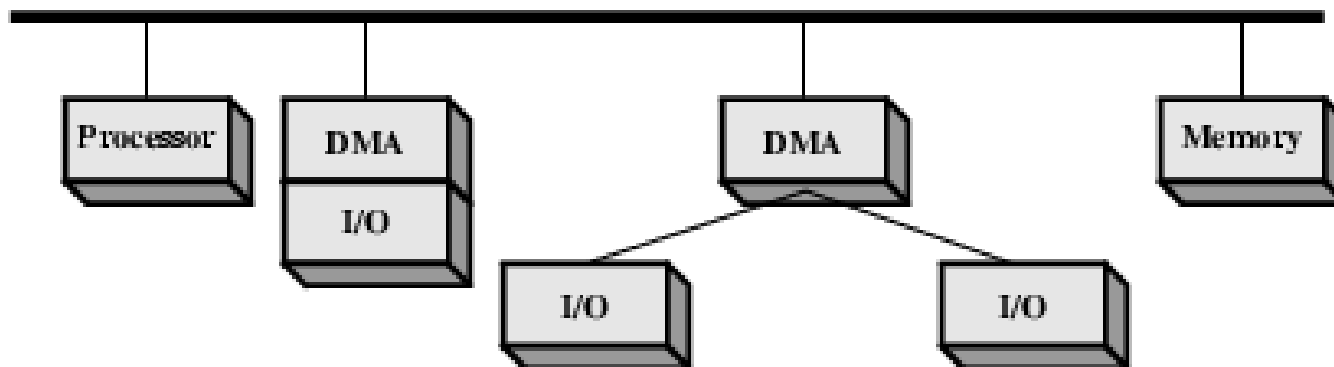


DMA Configurations (1)



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

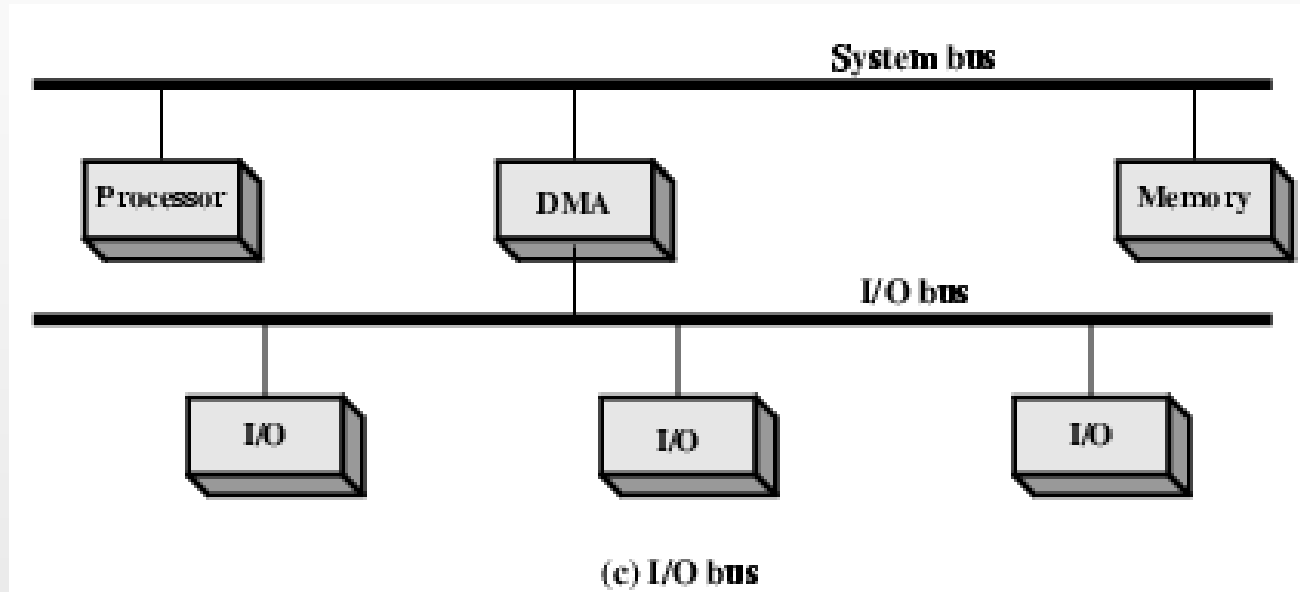
DMA Configurations (2)



(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

I/O Channel Architecture

