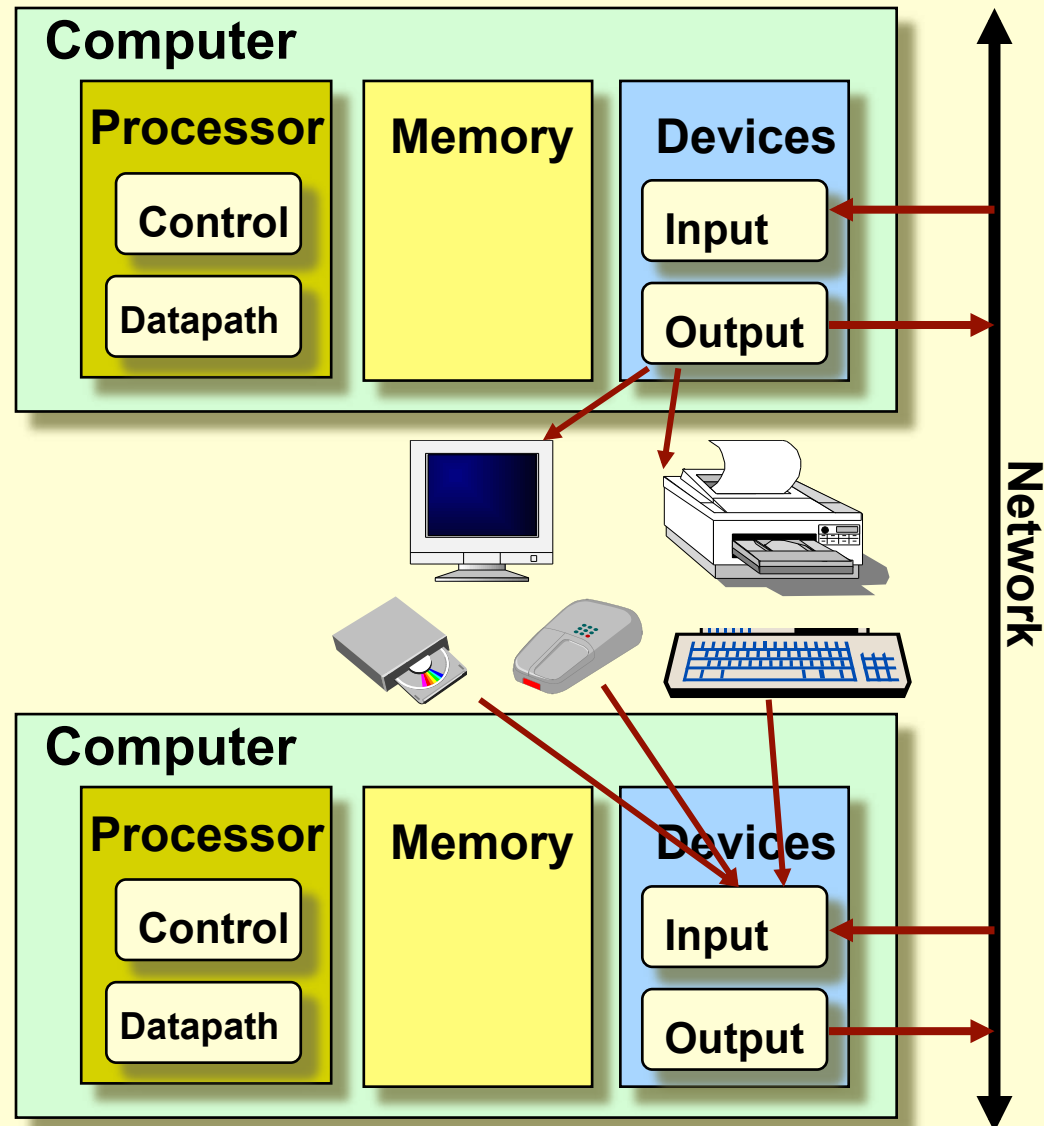


CMSC 611: Advanced Computer Architecture

I/O & Storage

Input/Output

- I/O Interface
 - Device drivers
 - Device controller
 - Service queues
 - Interrupt handling
- Design Issues
 - Performance
 - Expandability
 - Standardization
 - Resilience to failure
- Impact on Tasks
 - Blocking conditions
 - Priority inversion
 - Access ordering



Impact of I/O on System Performance

Suppose we have a benchmark that executes in 100 seconds of elapsed time, where 90 seconds is CPU time and the rest is I/O time. If the CPU time improves by 50% per year for the next five years but I/O time does not improve, how much faster will our program run at the end of the five years?

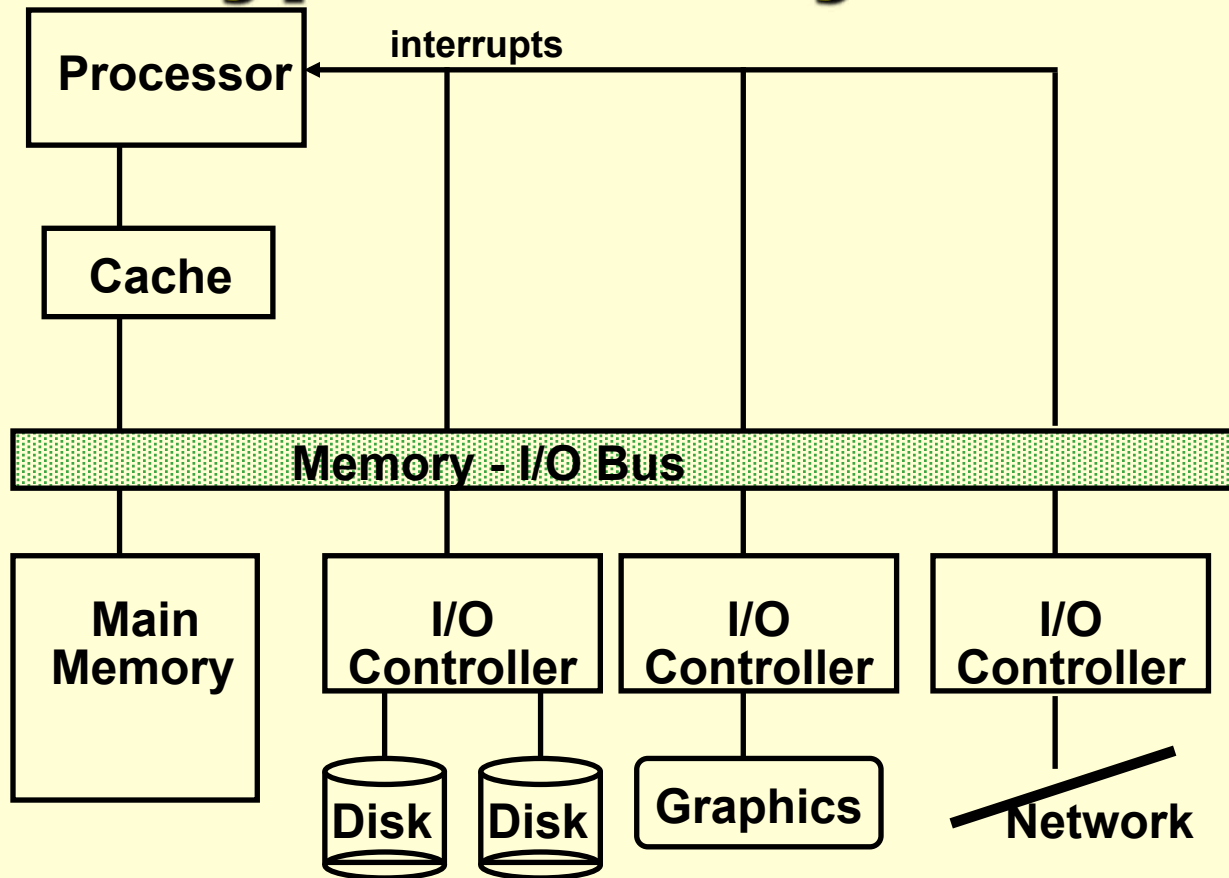
Answer: Elapsed Time = CPU time + I/O time

After n years	CPU time	I/O time	Elapsed time	% I/O time
0	90 Seconds	10 Seconds	100 Seconds	10%
1	$\frac{90}{1.5} = 60$ Seconds	10 Seconds	70 Seconds	14%
2	$\frac{60}{1.5} = 40$ Seconds	10 Seconds	50 Seconds	20%
3	$\frac{40}{1.5} = 27$ Seconds	10 Seconds	37 Seconds	27%
4	$\frac{27}{1.5} = 18$ Seconds	10 Seconds	28 Seconds	36%
5	$\frac{18}{1.5} = 12$ Seconds	10 Seconds	22 Seconds	45%

Over five years:

CPU improvement = $90/12 = 7$. **BUT** System improvement = $100/22 = 4.5$

Typical I/O System



- The connection between the I/O devices, processor, and memory are usually called (local or internal) bus
- Communication among the devices and the processor use both protocols on the bus and interrupts

I/O Device Examples

<u><i>Device</i></u>	<u><i>Behavior</i></u>	<u><i>Partner</i></u>	<u><i>Data Rate (KB/sec)</i></u>
Keyboard	Input	Human	0.01
Mouse	Input	Human	0.02
Line Printer	Output	Human	1.00
Floppy disk	Storage	Machine	50.00
Laser Printer	Output	Human	100.00
Optical Disk	Storage	Machine	500.00
Magnetic Disk	Storage	Machine	5,000.00
Network-LAN	Input or Output	Machine	20 – 1,000.00
Graphics Display	Output	Human	30,000.00

Disk History

Data density in
Mbit/square inch

Model 3340 hard disk
1973

1.7

140



Model 3370
1979

7.7

2,300



Capacity of Unit
Shown in Megabytes

Model 3390
1989

62.5

60,000



Travelstar VP
1997

1,450

1,600




Travelstar 8GS
1997

3,090

8,100




Organization of a Hard Magnetic Disk



- Typical numbers (depending on the disk size):
 - 500 to 2,000 tracks per surface
 - 32 to 128 sectors per track
 - A sector is the smallest unit that can be read or written to
- Traditionally all tracks have the same number of sectors:
 - Constant bit density: record more sectors on the outer tracks
 - Recently relaxed: constant bit size, speed varies with track location

Magnetic Disk Operation



- Cylinder: all the tracks under the head at a given point on all surface
- Read/write is a three-stage process:
 - Seek time
 - position the arm over proper track
 - Rotational latency
 - wait for the sector to rotate under the read/write head
 - Transfer time
 - transfer a block of bits (sector) under the read-write head
- Average seek time
 - $(\sum \text{time for all possible seeks}) / (\# \text{ seeks})$
 - Typically in the range of 8 ms to 12 ms
 - Due to locality of disk reference, actual average seek time may only be 25% to 33% of the advertised number

Magnetic Disk Characteristic

- Rotational Latency:
 - Most disks rotate at 5,400 to 10,000 RPM
 - Approximately 11 ms to 6 ms per revolution, respectively
 - An average latency to the desired information is halfway around the disk:
 - 5.5 ms at 5400 RPM, 3 ms at 10000 RPM
- Transfer Time is a function of :
 - Transfer size (usually a sector): 1 KB / sector
 - Rotation speed: 5400 RPM to 10000 RPM
 - Recording density: bits per inch on a track
 - Diameter: typical diameter ranges from 2.5 to 5.25"
 - Typical values ~500MB per second

Example

Calculate the access time for a disk with 512 byte/sector and 12 ms advertised seek time. The disk rotates at 5400 RPM and transfers data at a rate of 4MB/sec. The controller overhead is 1 ms. Assume that the queue is idle (so no service time)

Answer:

$$\begin{aligned}\text{Disk Access Time} &= \text{Seek time} + \text{Rotational Latency} + \text{Transfer time} \\ &\quad + \text{Controller Time} + \text{Queuing Delay} \\ &= 12 \text{ ms} + 0.5 / 5400 \text{ RPM} + 0.5 \text{ KB} / 4 \text{ MB/s} + 1 \text{ ms} + 0 \\ &= 12 \text{ ms} + 0.5 / 90 \text{ RPS} + 0.125 / 1024 \text{ s} + 1 \text{ ms} + 0 \\ &= 12 \text{ ms} + 5.5 \text{ ms} \quad + 0.1 \text{ ms} \quad + 1 \text{ ms} + 0 \\ &= 18.6 \text{ ms}\end{aligned}$$

ms

If real seeks are 1/3 the advertised seeks, disk access time would be 10.6 ms, with rotation delay contributing 50% of the access time!


Historical Trend

Characteristics	IBM 3090	IBM UltraStar	Integral 1820
Disk diameter (inches)	10.88	3.50	1.80
Formatted data capacity (MB)	22,700	4,300	21
MTTF (hours)	50,000	1,000,000	100,000
Number of arms/box	12	1	1
Rotation speed (RPM)	3,600	7,200	3,800
Transfer rate (MB/sec)	4.2	9-12	1.9
Power/box (watts)	2,900	13	2
MB/watt	8	102	10.5
Volume (cubic feet)	97	0.13	0.02
MB/cubic feet	234	33000	1050

Reliability and Availability

- Two terms that are often confused:
 - Reliability: Is anything broken?
 - Availability: Is the system still available to the user?
- Availability can be improved by adding hardware:
 - Example: adding ECC on memory
- Reliability can only be improved by:
 - Enhancing environmental conditions
 - Building more reliable components
 - Building with fewer components
 - Improve availability may come at the cost of lower reliability

Disk Arrays




- Increase potential throughput by having many disk drives:
 - Data is spread over multiple disk
 - Multiple accesses are made to several disks
- Reliability is lower than a single disk:
 - Reliability of N disks = Reliability of 1 Disk \div N
 - (50,000 Hours \div 70 disks = 700 hours)
 - Disk system MTTF: Drops from 6 years to 1 month
 - Arrays (without redundancy) too unreliable to be useful!
 - But availability can be improved by adding redundant disks (RAID):
 - Lost information can be reconstructed from redundant information


Manufacturing Advantages of Disk Arrays

Disk Product Families

Conventional:
4 disk
designs



Disk Array:
1 disk
design




Replace Small # of Large Disks with Large # of Small Disks!

Redundant Arrays of Disks

- Redundant Array of Inexpensive Disks (RIAD)
 - Widely available and used in today's market
 - Files are "striped" across multiple spindles
 - Redundancy yields high data availability despite low reliability
 - Contents of a failed disk is reconstructed from data redundantly stored in the disk array
 - Drawbacks include capacity penalty to store redundant data and bandwidth penalty to update a disk block
 - Different levels based on replication level and recovery techniques

RAID level	Failures survived	Data disks	Check disks
0 Non-redundant	0	8	0
1 Mirrored	1	8	8
2 Memory-style ECC	1	8	4
3 Bit-interleaved parity	1	8	1
4 Block-interleaved	1	8	1
5 Block-interleaved distributed parity	1	8	1


RAID 1: Disk Mirroring/ Shadowing



- Each disk is fully duplicated onto its "shadow"
- Very high availability can be achieved
- Bandwidth sacrifice on write: Logical write = two physical writes
- Reads may be optimized
- Most expensive solution: 100% capacity overhead

Targeted for high I/O rate , high availability environments

RAID 3: Parity Disk

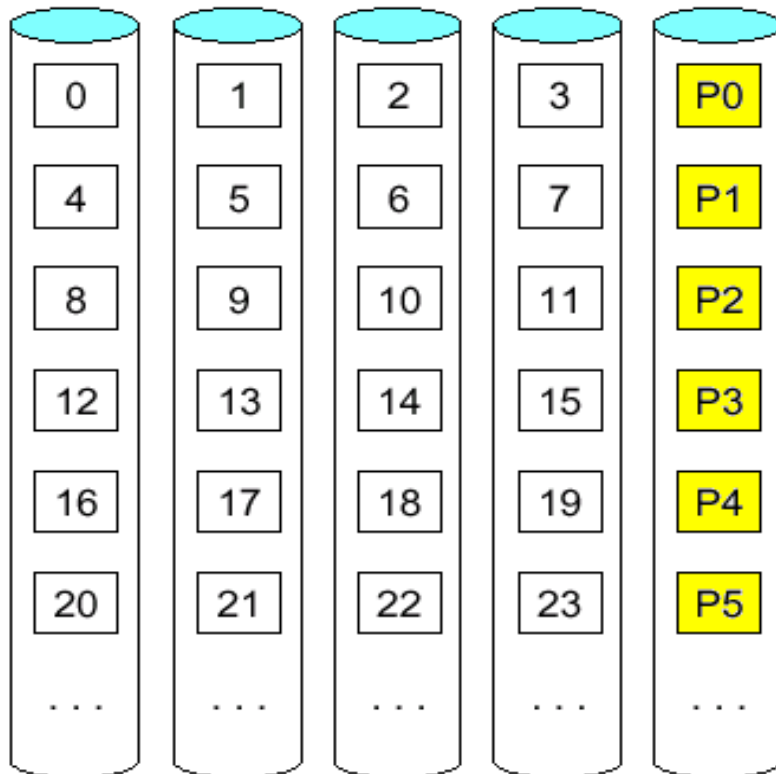


- ❑ Parity computed across recovery group to protect against hard disk failures
- ❑ 33% capacity cost for parity in this configuration: wider arrays reduce capacity costs, decrease expected availability, increase reconstruction time
- ❑ Arms logically synchronized, spindles rotationally synchronized (logically a single high capacity, high transfer rate disk)

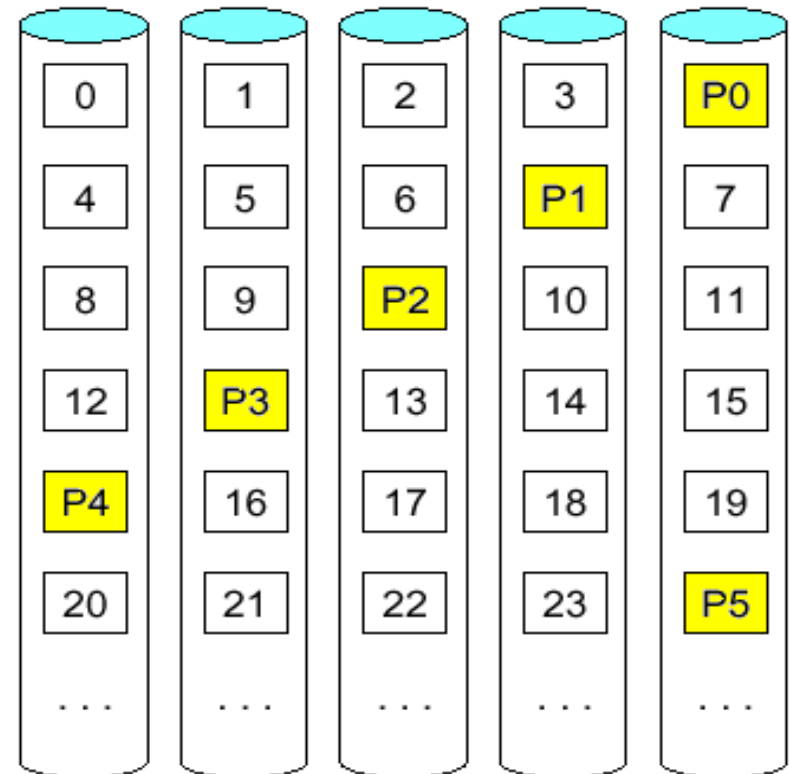
Targeted for high bandwidth applications: Scientific, Image Processing

Block-Based Parity

- ❑ Block-based parity leads to more efficient read access compared to RAID 3
- ❑ Designating a parity disk allows recovery but will keep it idle in the absence of a disk failure
- ❑ RAID 5 distribute the parity block to allow the use of all disk and enhance parallelism of disk access



RAID 4



RAID 5

RAID 5+: High I/O Rate Parity

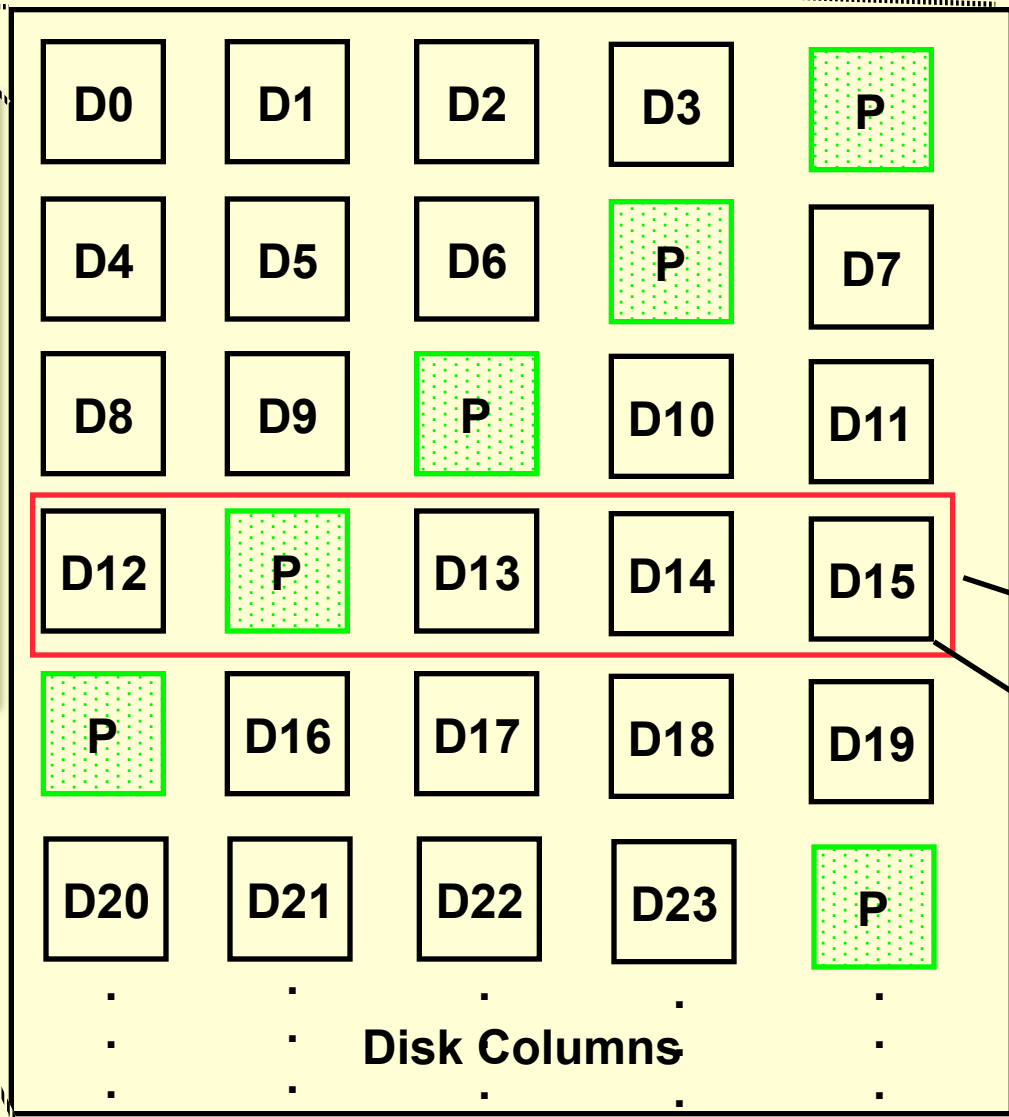


A logical write becomes four physical I/Os

Independent writes possible because of interleaved parity

Reed-Solomon Codes ("Q") for protection during reconstruction

Targeted for mixed applications



Increasing Logical Disk Addresses

Stripe

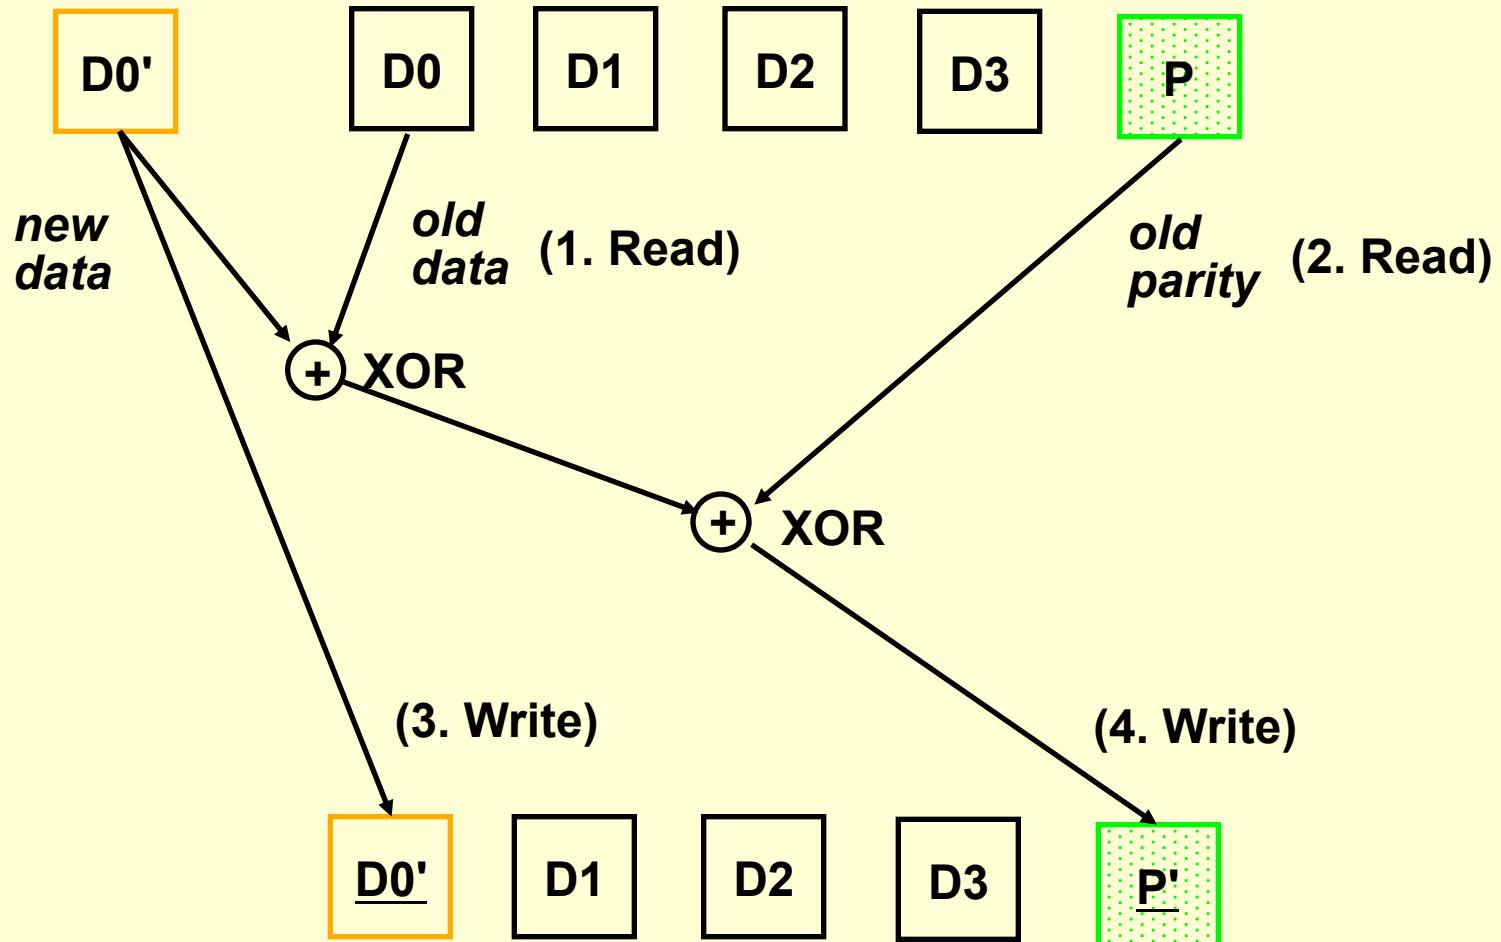
Stripe Unit

Disk Columns

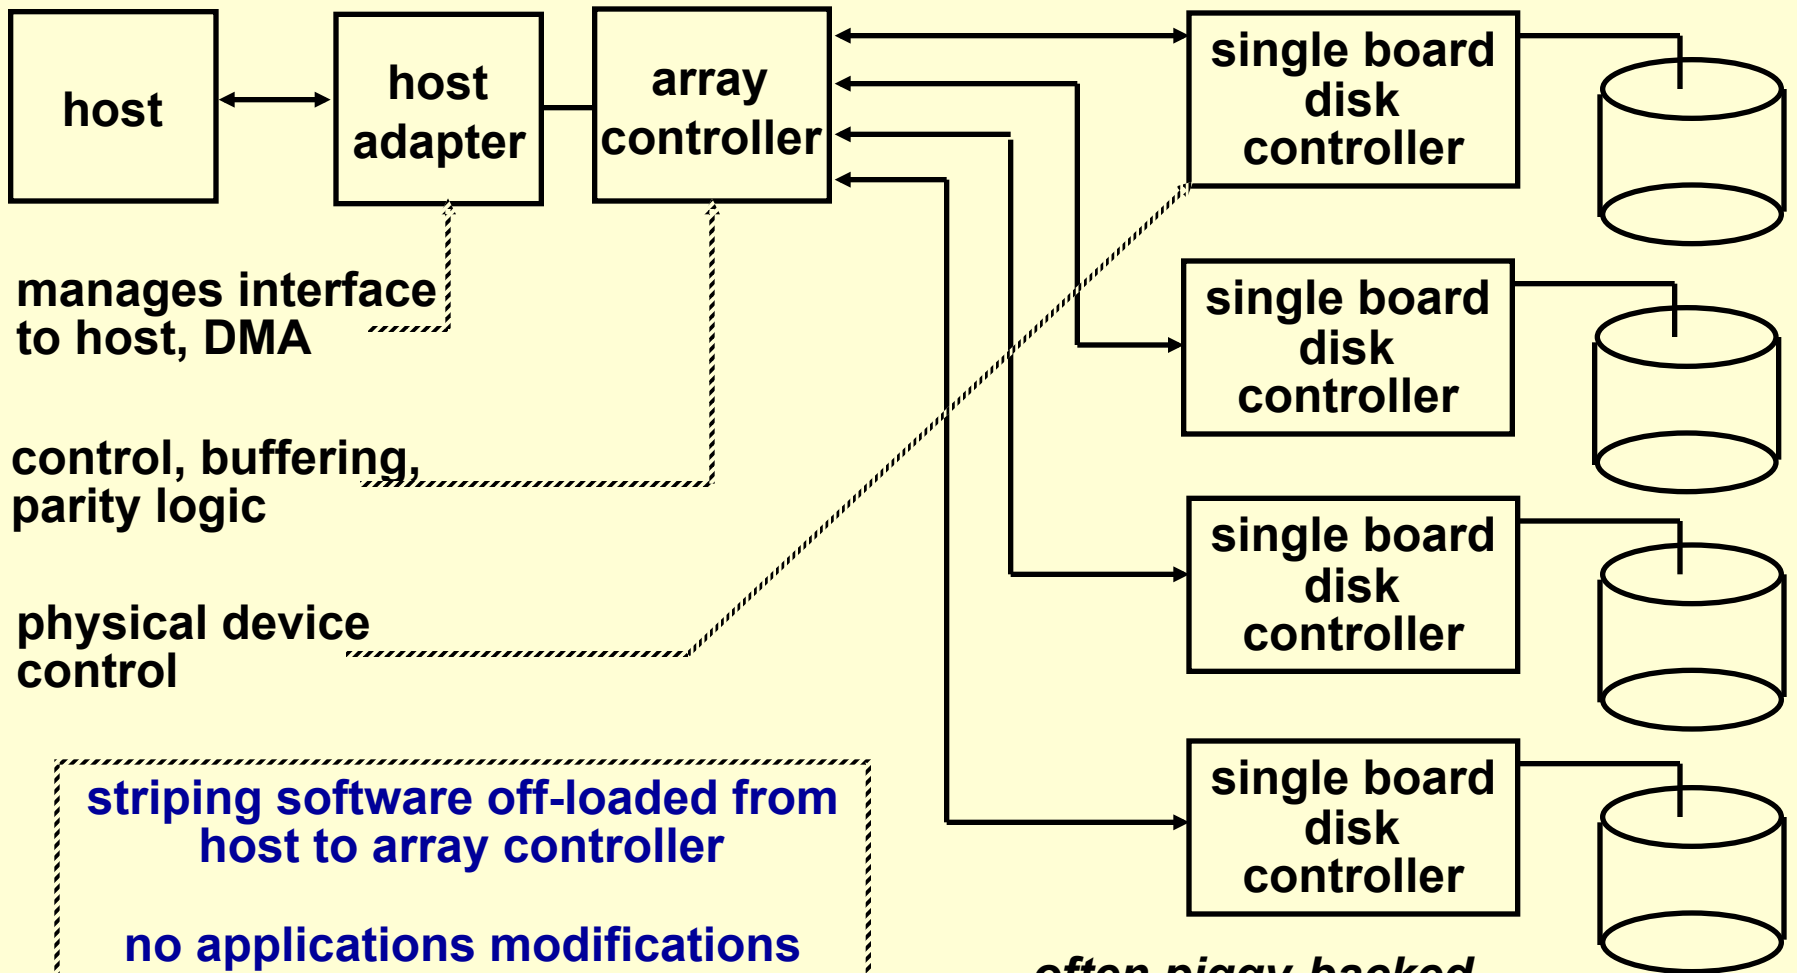
Problems of Small Writes

RAID-5: Small Write Algorithm

1 Logical Write = 2 Physical Reads + 2 Physical Writes



Subsystem Organization



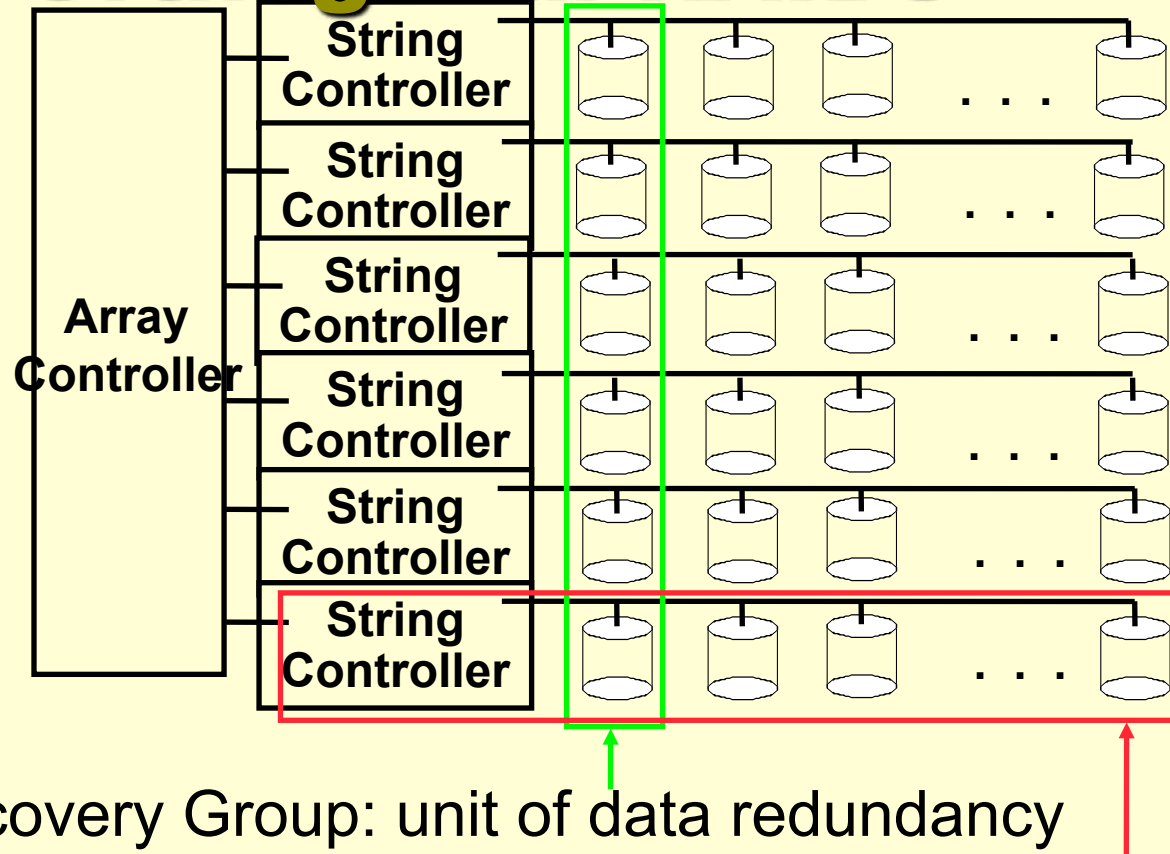
striping software off-loaded from host to array controller

no applications modifications

no reduction of host performance

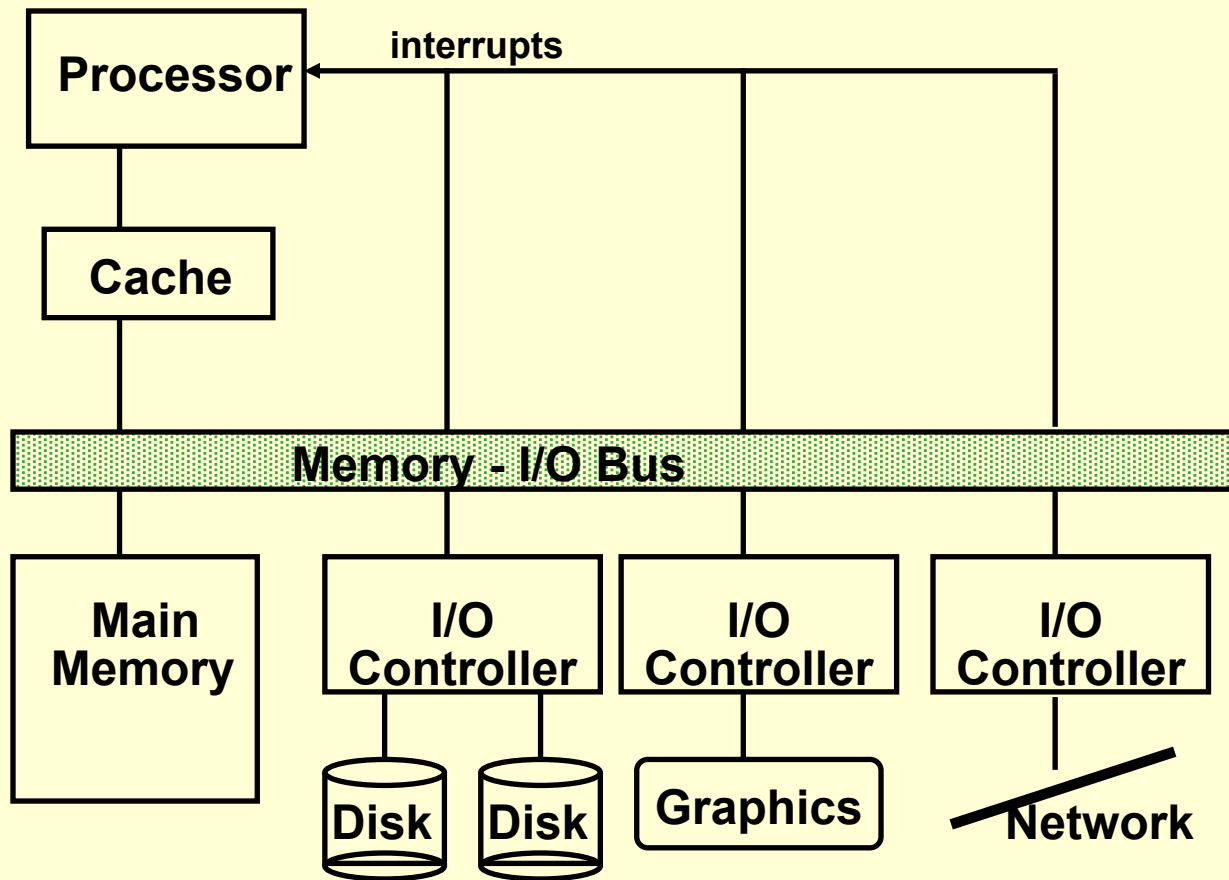
often piggy-backed in small format devices

System Availability: Orthogonal RAIDs

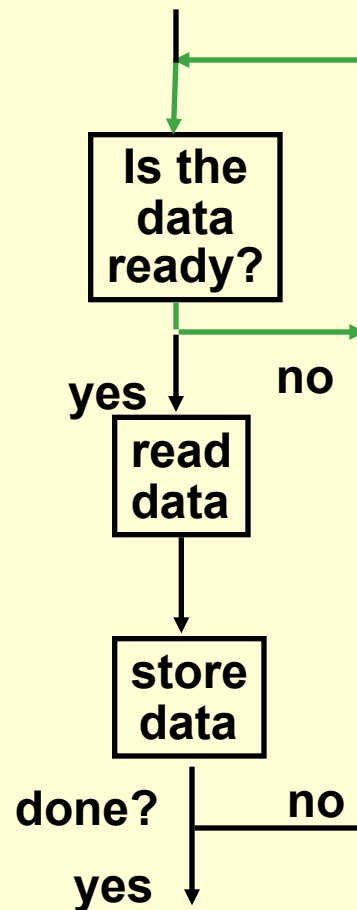



- Data Recovery Group: unit of data redundancy
- Redundant Support Components: fans, power supplies, controller, cables
- End to End Data Integrity: internal parity protected data paths

I/O Control



Polling: Programmed I/O

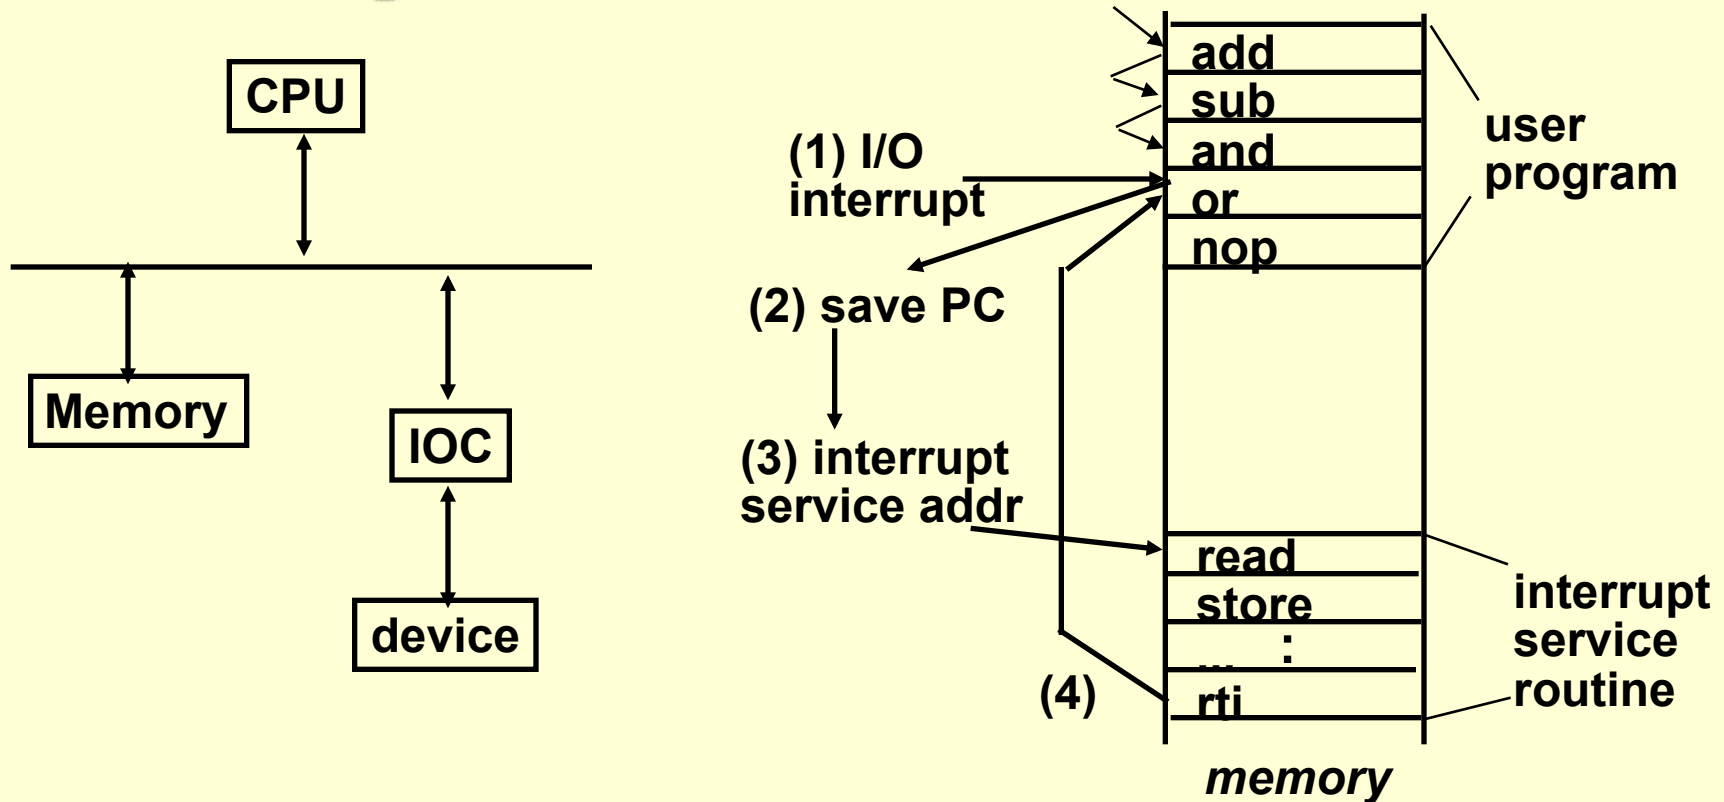


**busy wait loop
not an efficient
way to use the CPU
unless the device
is very fast!**

**but checks for I/O
completion can be
dispersed among
computation
intensive code**

- **Advantage:**
 - Simple: the processor is totally in control and does all the work
- **Disadvantage:**
 - Polling overhead can consume a lot of CPU time

Interrupt Driven Data Transfer



- **Advantage:**
 - User program progress is only halted during actual transfer
- **Disadvantage:** special hardware is needed to:
 - Cause an interrupt (I/O device)
 - Detect an interrupt (processor)
 - Save the proper states to resume after the interrupt (processor)

I/O Interrupt vs. Exception

- An I/O interrupt is just like the exceptions except:
 - An I/O interrupt is asynchronous
 - Further information needs to be conveyed
 - Typically exceptions are more urgent than interrupts
- An I/O interrupt is asynchronous with respect to instruction execution:
 - I/O interrupt is not associated with any instruction
 - I/O interrupt does not prevent any instruction from completion
 - You can pick your own convenient point to take an interrupt
- I/O interrupt is more complicated than exception:
 - Needs to convey the identity of the device generating the interrupt
 - Interrupt requests can have different urgencies:
 - Interrupt request needs to be prioritized
 - Priority indicates urgency of dealing with the interrupt
 - high speed devices usually receive highest priority

Direct Memory Access

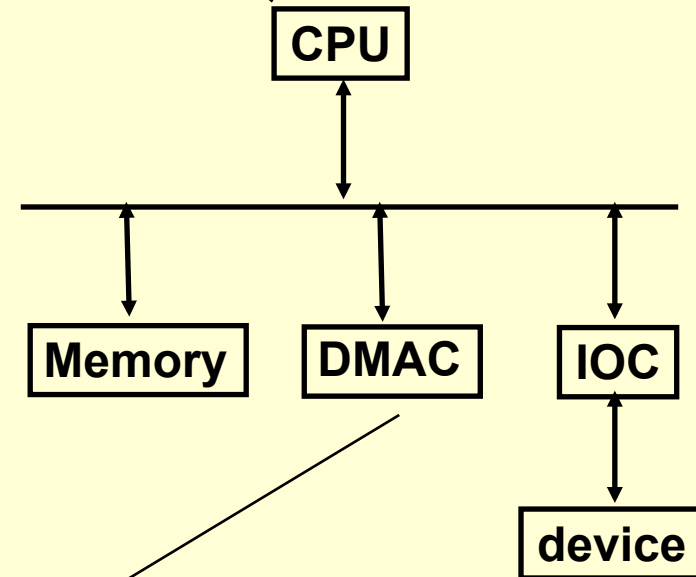
- Direct Memory Access (DMA):

- External to the CPU
- Use idle bus cycles (*cycle stealing*)
- Act as a master on the bus
- Transfer blocks of data to or from memory without CPU intervention
- Efficient for large data transfer, e.g. from disk
- ☞ Cache usage allows the processor to leave enough memory bandwidth for DMA

- How does DMA work?:

- CPU sets up and supply device id, memory address, number of bytes
- DMA controller (DMAC) starts the access and becomes bus master
- For multiple byte transfer, the DMAC increment the address
- DMAC interrupts the CPU upon completion

CPU sends a starting address, direction, and length count to DMAC. Then issues "start".



DMAC provides handshake signals for Peripheral Controller, and Memory Addresses and handshake signals for Memory.

For multiple bus system, each bus controller often contains DMA control logic

DMA Problems

① **With virtual memory systems:** (pages would have physical and virtual addresses)

- ➔ Physical pages re-mapping to different virtual pages during DMA operations
- ➔ Multi-page DMA cannot assume consecutive addresses

Solutions:

- ➔ Allow virtual addressing based DMA
 - ⇒ Add translation logic to DMA controller
 - ⇒ OS allocated virtual pages to DMA prevent re-mapping until DMA completes
- ➔ Partitioned DMA
 - ⇒ Break DMA transfer into multi-DMA operations, each is single page
 - ⇒ OS chains the pages for the requester

② **In cache-based systems:** (there can be two copies of data items)

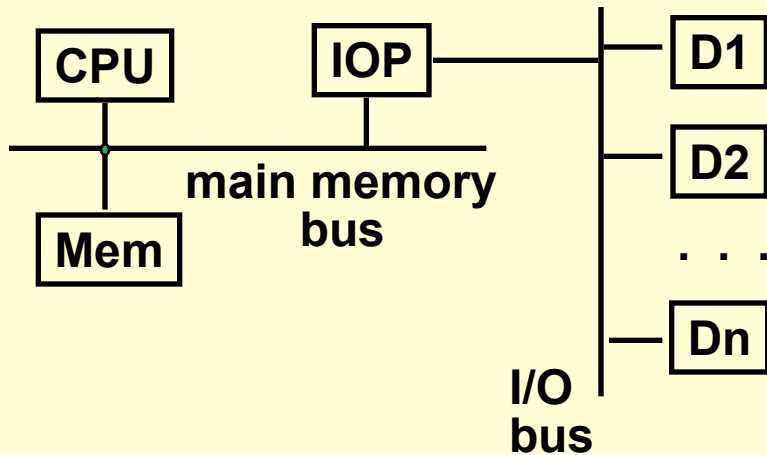
- ➔ Processor might not know that the cache and memory pages are different
- ➔ Write-back caches can overwrite I/O data or makes DMA to read wrong data

Solutions:


- ➔ Route I/O activities through the cache
 - ⇒ Not efficient since I/O data usually is not demonstrating temporal locality
- ➔ OS selectively invalidates cache blocks before I/O read or force write-back prior to I/O write
 - ⇒ Usually called cache *flushing* and requires hardware support

DMA allows another path to main memory with no cache and address translation

I/O Processor



- ➔ An I/O processor (IOP) offload the CPU
- ➔ Some processors, e.g. Motorola 860, include special purpose IOP for serial communication



Device to/from memory transfers are controlled by the IOP directly.

IOP steals memory cycles.

