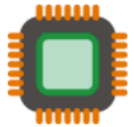

CASPER

CS305: Computer Architecture

Storage and I/O

https://www.cse.iitb.ac.in/~biswa/courses/CS305/main.html

https://www.cse.iitb.ac.in/~biswa/

10K Feet View

Storage devices are accessed differently than memory devices

- Access granularity • typically 4 – 8 KB
- Access is usually through a deeper, software stack
- Much higher access latencies
- Milliseconds vs ns
- Interfaces are slower than everything that we have seen so far
- SATA, SAS, PCIe
- Metrics for comparison • Latency, Bandwidth, IOPS

Magnetic disks/HDDs/

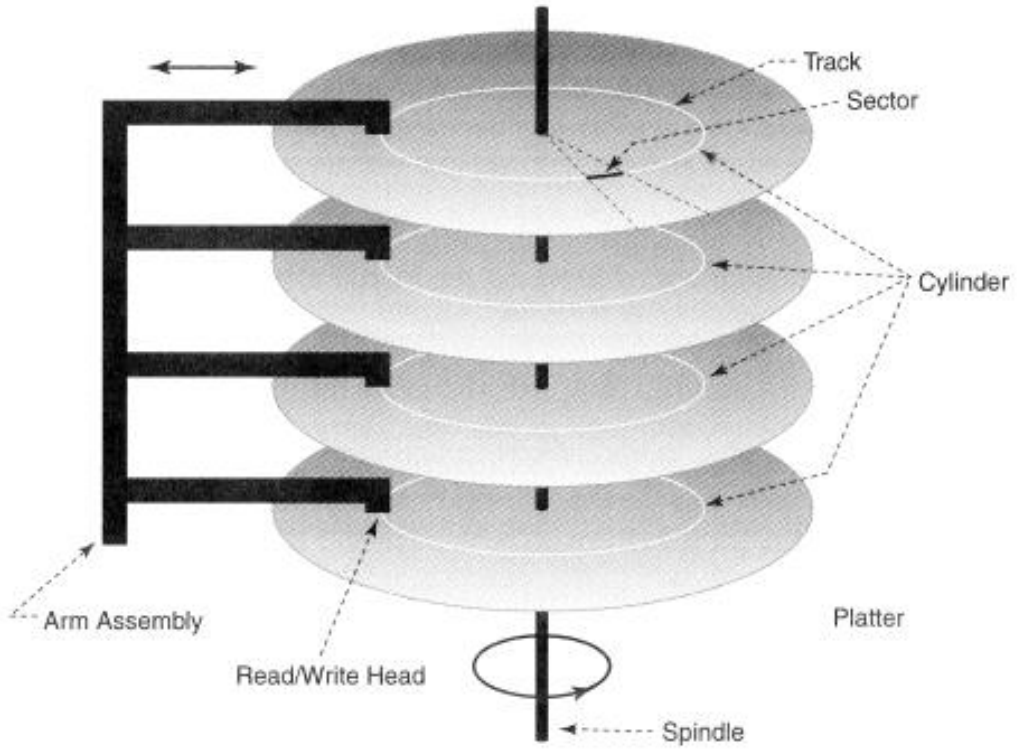

Magnetic Disk

A magnetic disk with platters (magnetic recording material on both sides

Each platter has concentric tracks (5 - 30K)

each track divided into sectors (512 B)

A movable arm holds the heads for each disk surface and moves them in tandem

Disk Read/Write

Slide courtesy
https://www.snia.org/sites/default/files/ESF/Performance_Benchmarking_3_Block_Components_FinalPPT.pptx

Disk Latency

Seek latency: Time to move the arm to the correct track; takes 5 to 12 ms on average; can take less if there is spatial locality

Rotational latency: time taken to rotate the correct sector under the head; typically, 2 ms (15,000 RPM)

Transfer time: time taken to transfer a block of bits out of the disk and is typically 100s MB/second

Other overheads, depending on disk design

Think about these latencies

Slide courtesy
https://www.snia.org/sites/default/files/ESF/Performance_Benchmarking_3_Block_Components_FinalPPT.pptx

The world of buses/ports/IO devices

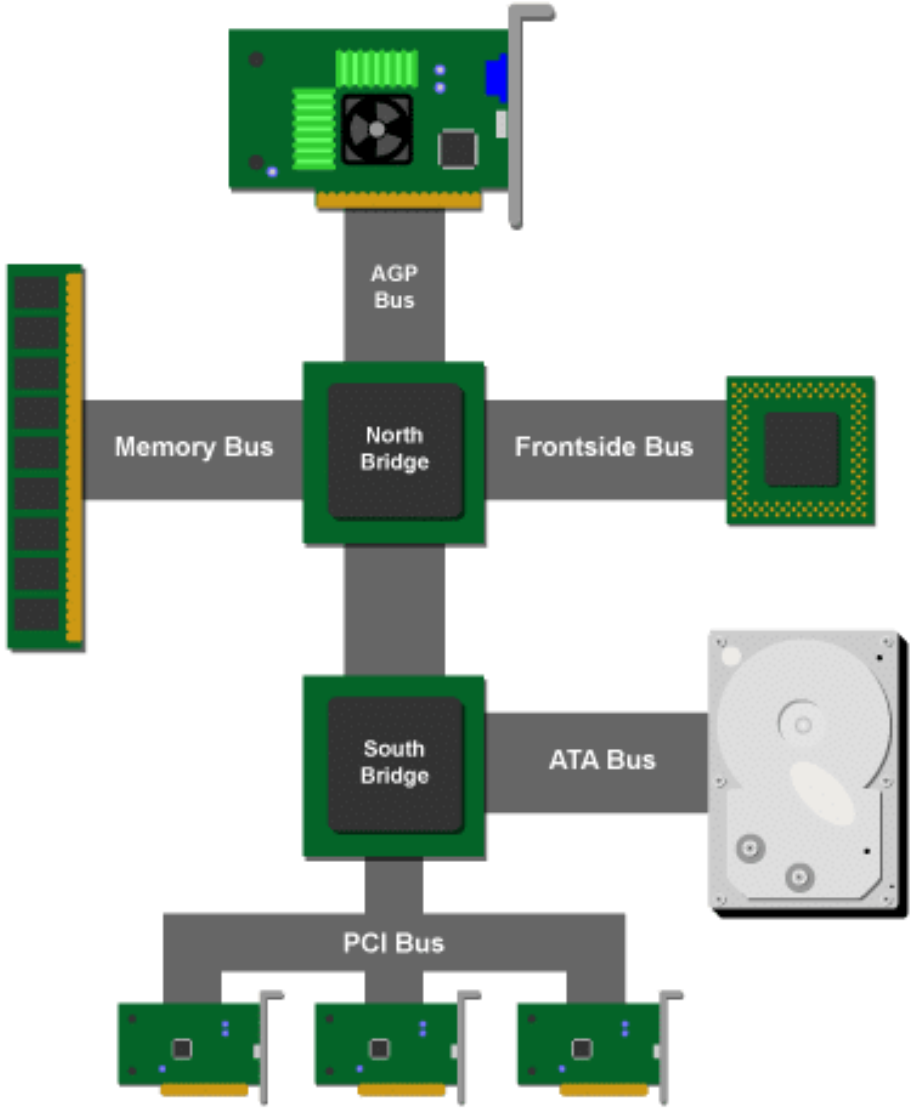

The motherboard

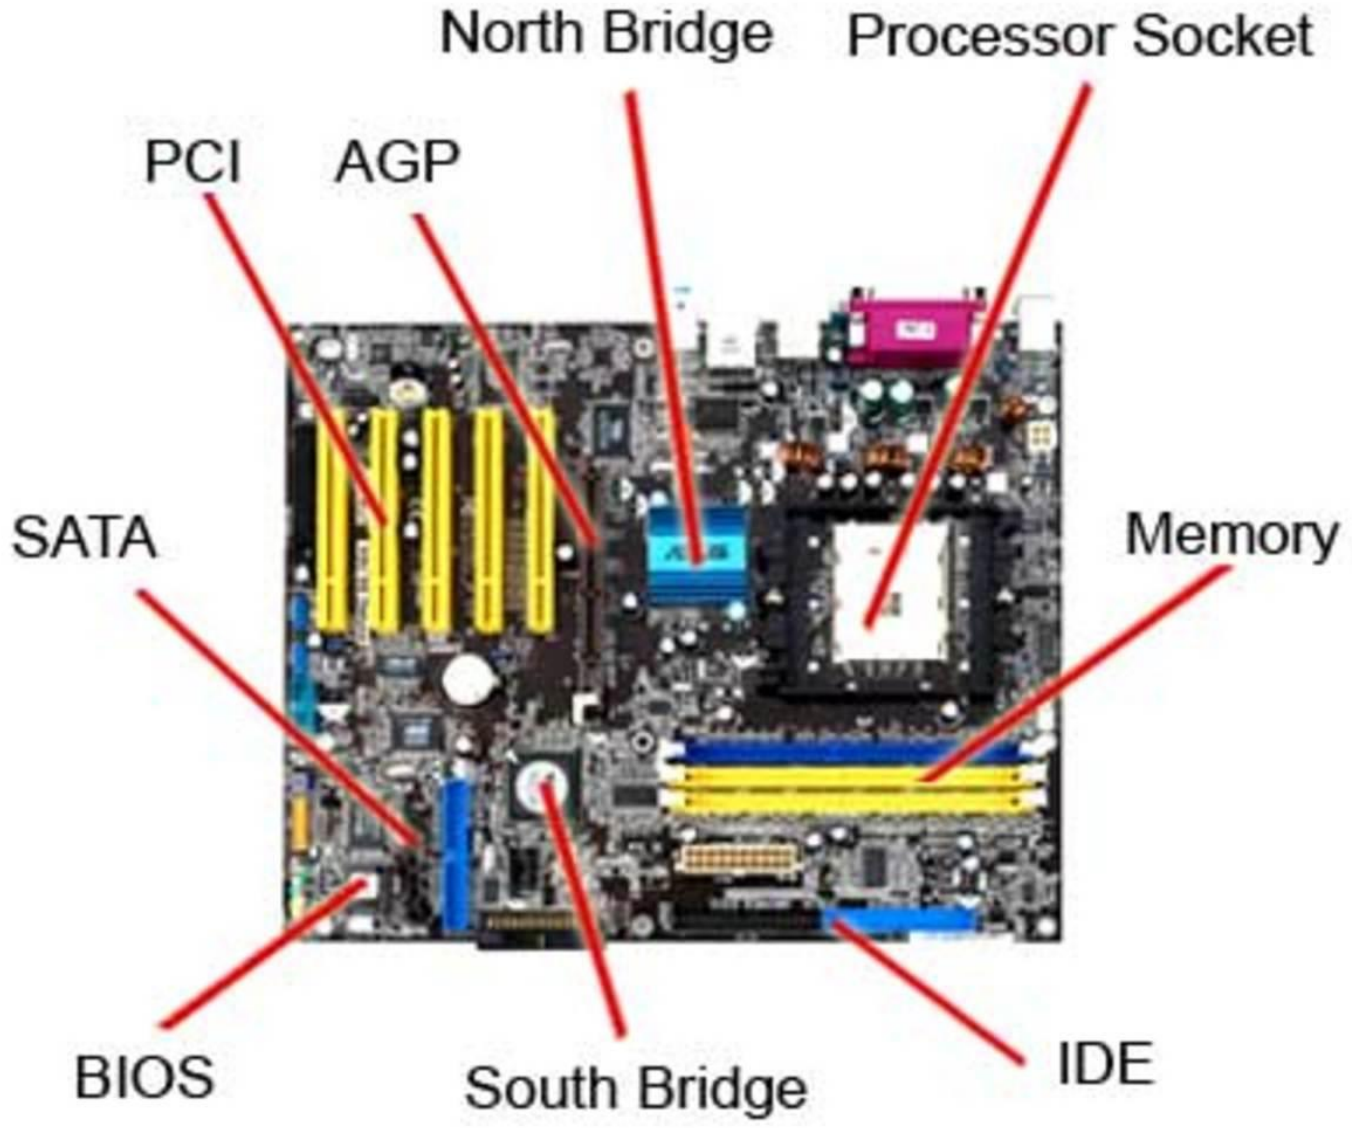

10K Feet View on I/O protocols

Slide courtesy
https://www.snia.org/sites/default/files/ESF/Performance_Benchmarking_3_Block_Components_FinalPPT.pptx

I/O Addressing and interaction with the processor

Memory Mapped I/O: I/O address is assigned to app's virtual address space. To access an I/O port, regular LOAD/STORE instructions can be used.

Interaction protocols:

Polling: Processor keeps on polling an I/O port address.

PIC: Programmable interrupt controller; buffers I/O requests for the processor

Still processor is in busy ☹

The DMA Approach

Direct Memory Access:

In polling and interrupt-based approaches, the processor has to execute thousands of instructions to move KBs of data to/from I/O to/from memory ☹

DMA controller: Another peripheral

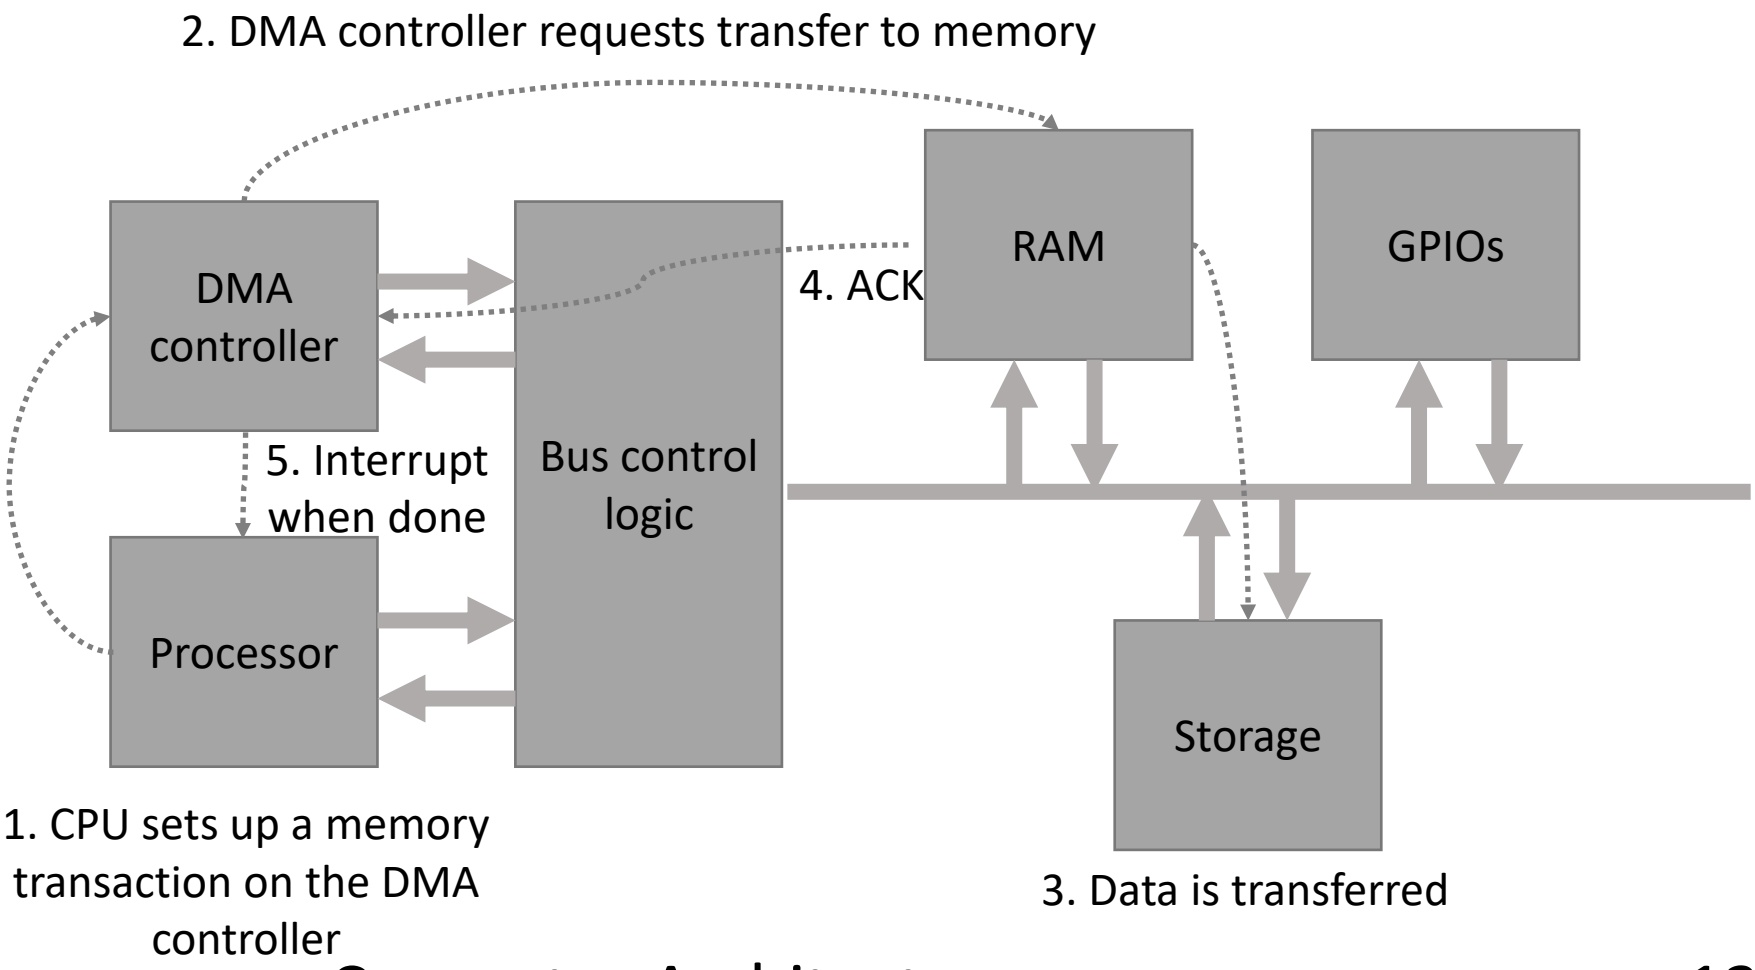

Two modes: Burst and cycle-stealing

Burst mode: Transfers all the data in one go. Becomes the master of the bus (south bridge and the bus to DRAM) for the entire duration.

Cycle stealing: Transfers data only when the bus is free.

Npezié