Hard disks in audio systems

- Hard disk principles and construction:
  see W.Intro 9.1-9.12, WArt chapter 10
- Types of data transfer
  (I/O, DMA, bus master, shared memory)
- Interprocess communication (IPC)

  Polling and interrupts
- Disk caches and performance

From A-D converter to sound file

Smallest possible system

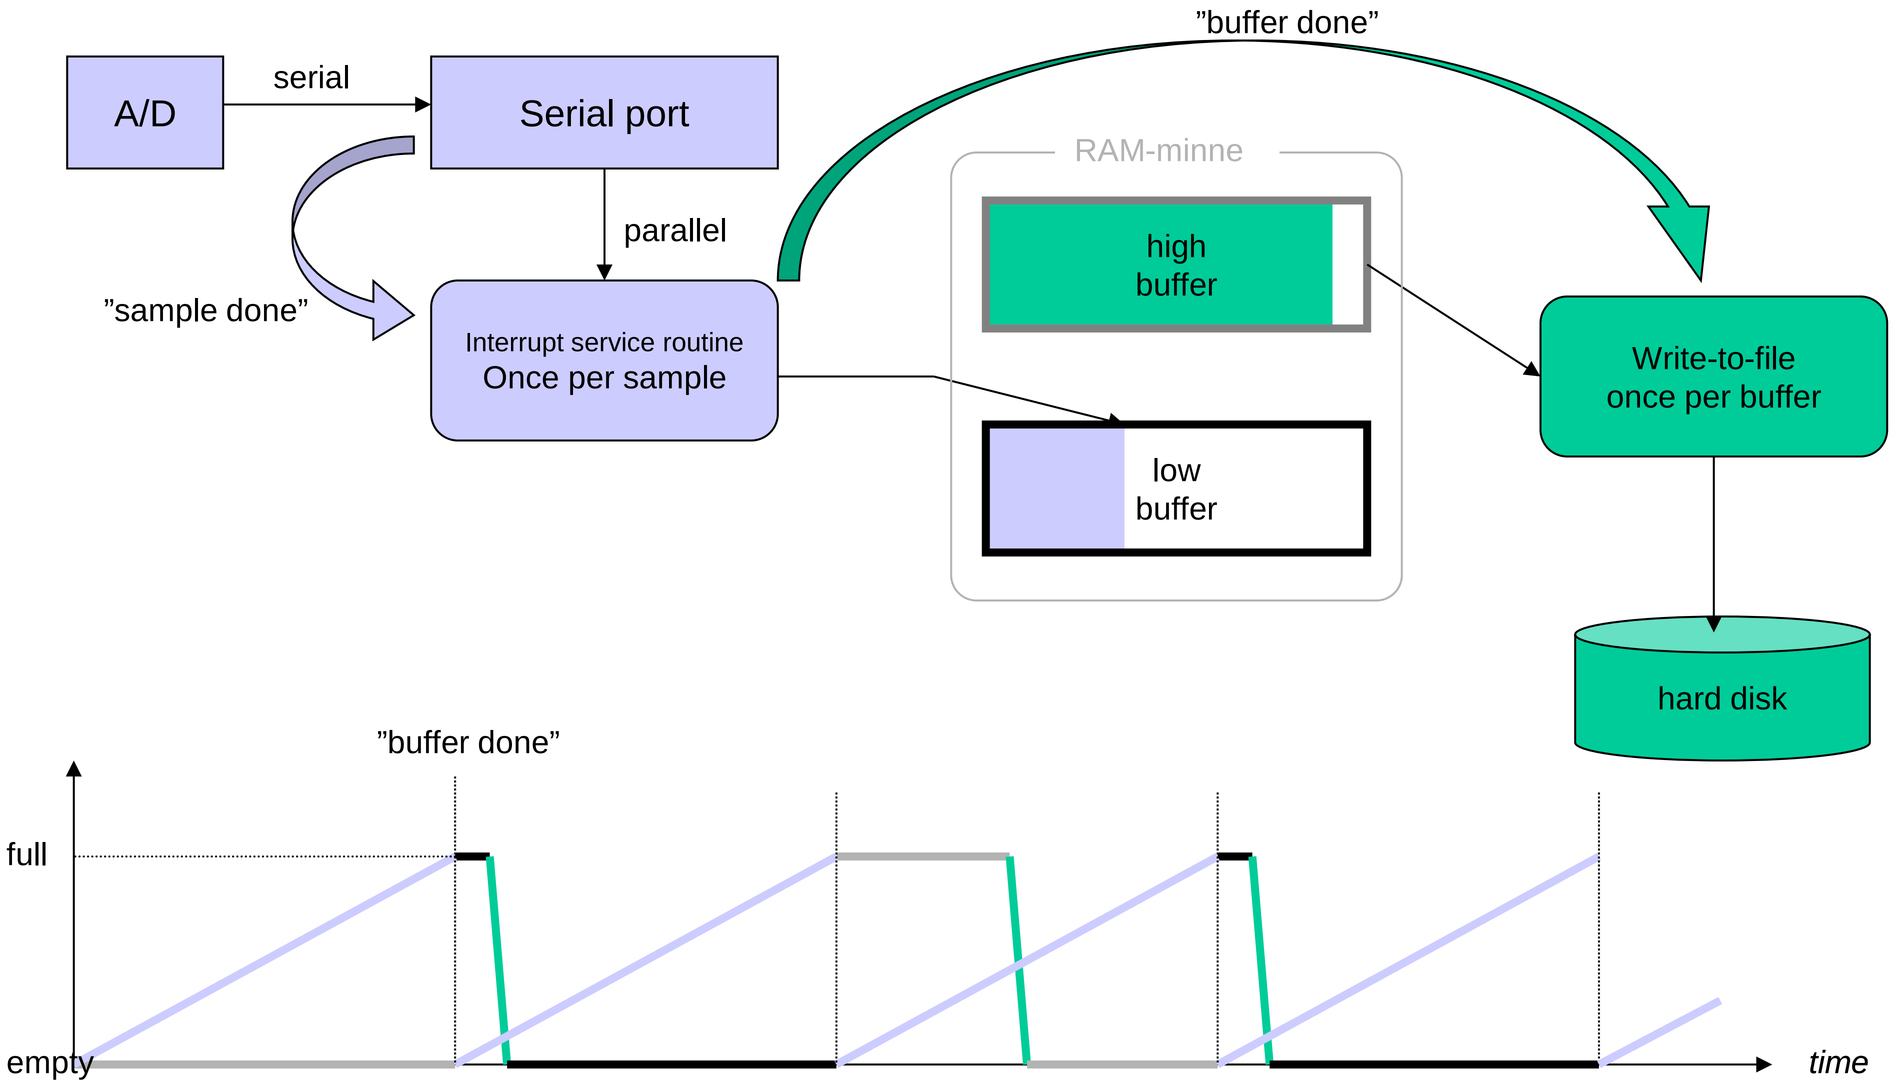

Components of a hard disk storage system

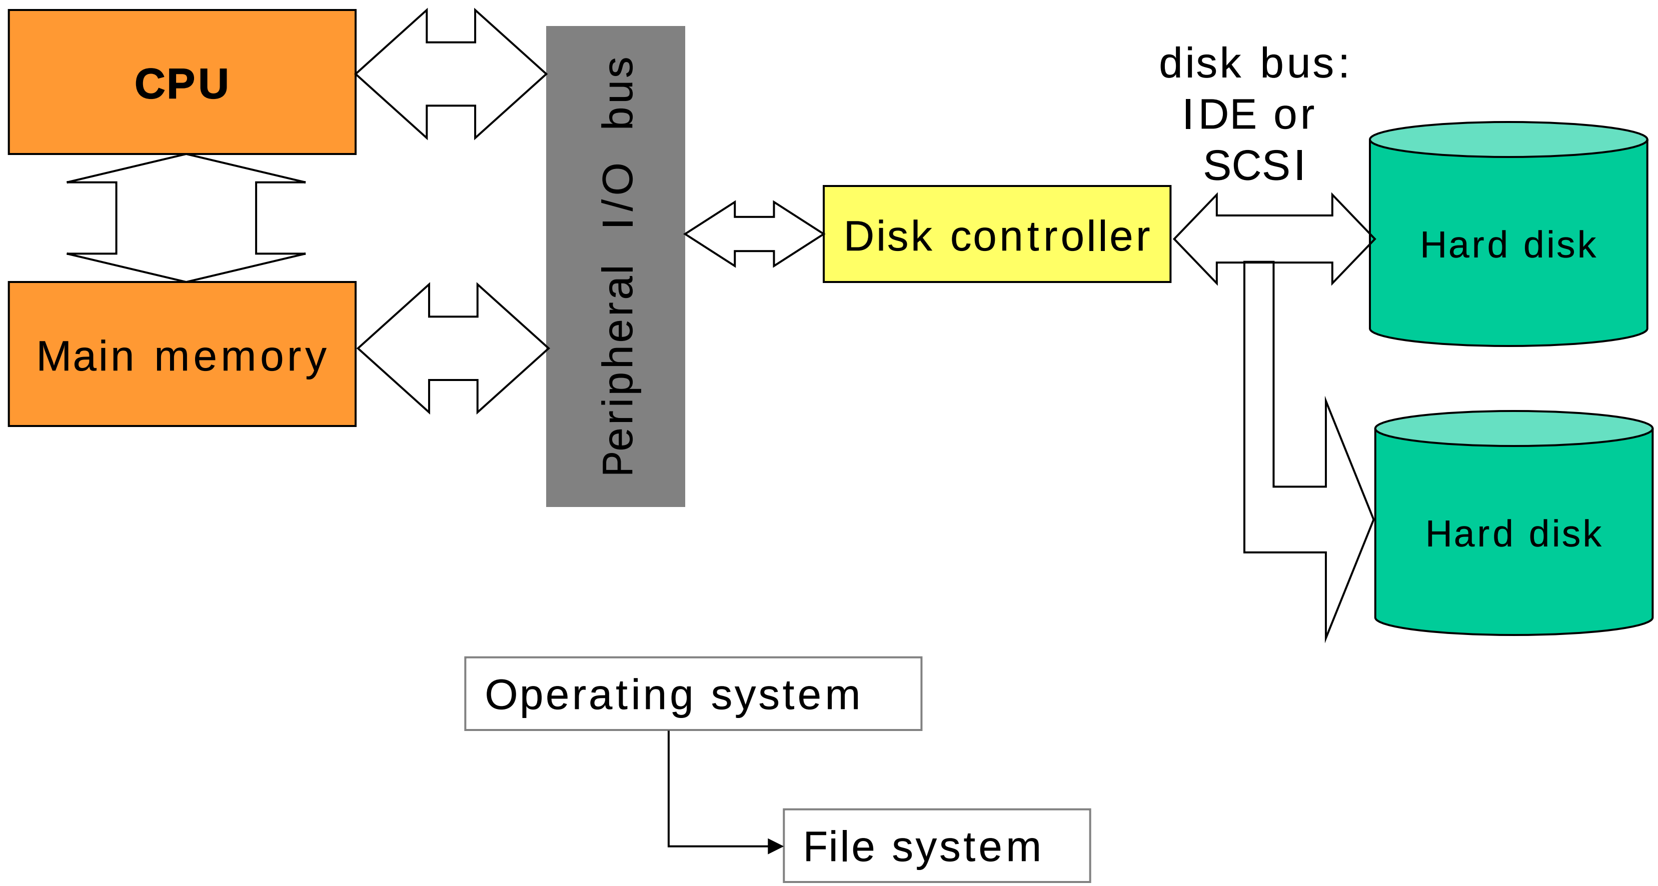

IBM 305

History

The IBM RAMAC 305 was the first commercial computer that used magnetic disk storage. IBM introduced it on September 4, **1956**. RAMAC stood for "Random Access Method of Accounting and Control." The 305 was one of the last vacuum tube computers that IBM built. The IBM 350 disk system stored 5 million characters. It had **fifty 24-inch diameter** disks. Two independent access arms moved up and down to select a disk and in and out to select a recording track, all under servo control. Several improved models were added in the 1950s. The IBM RAMAC 305 system with 350 disk storage leased for $3,200 per month. More than 1000 systems were built. Production ended in 1961 and the 305 was withdrawn in 1969.

From Wikipedia, the free encyclopedia.

Modern 3.5” drive

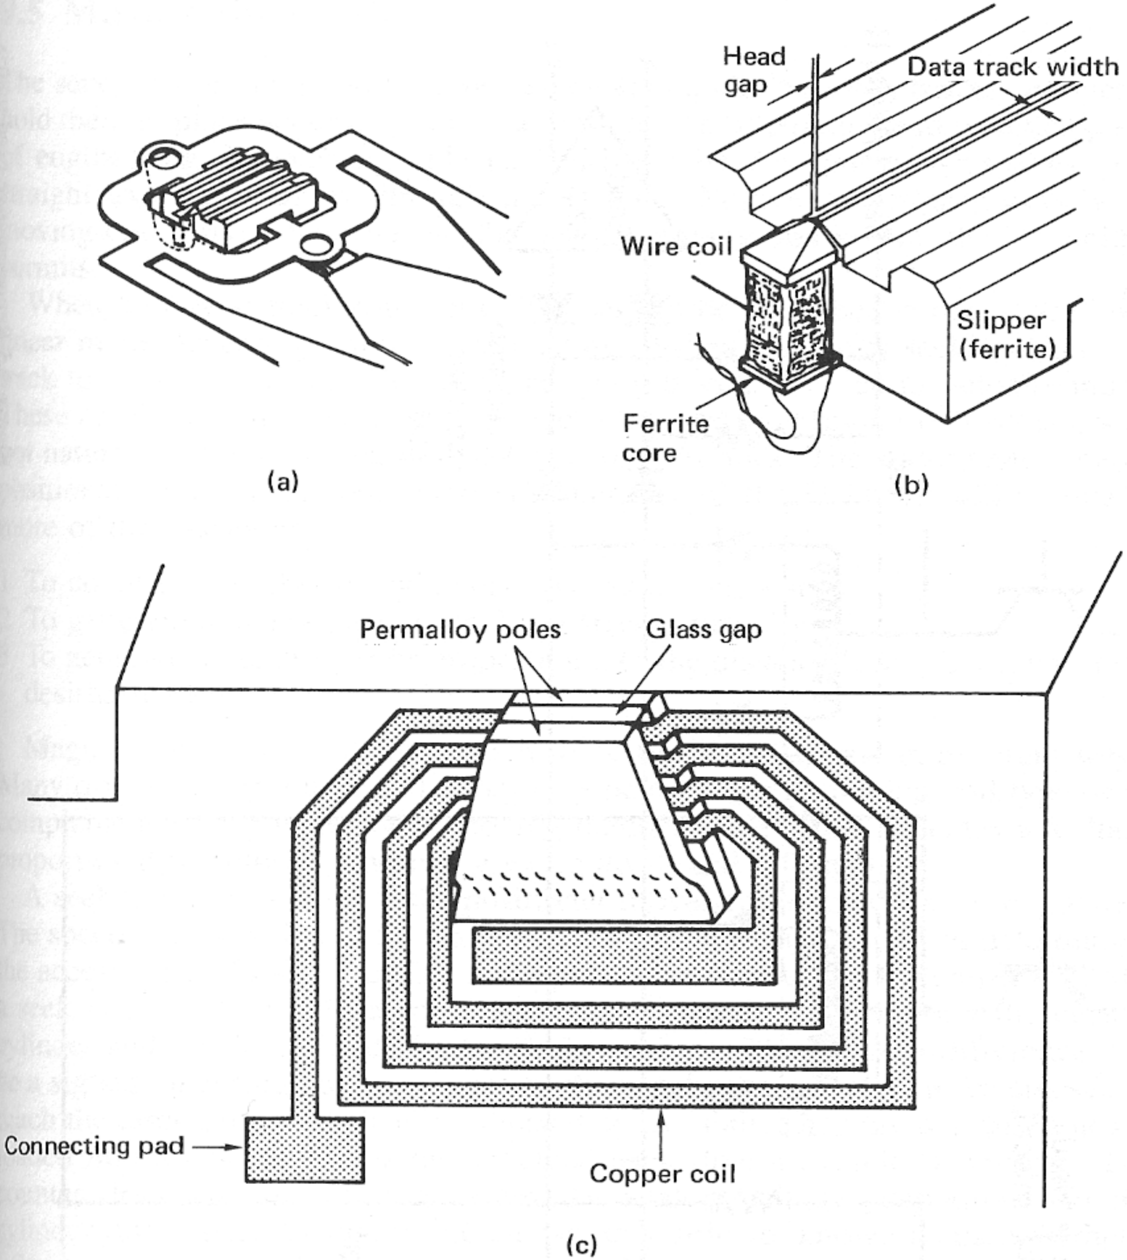

igure 9.4 (a) Winchester head construction showing large air bleed grooves. (b) Close-up of ipper showing magnetic circuit on trailing edge. (c) Thin-film head is fabricated on the end of he slipper using microcircuit technology.

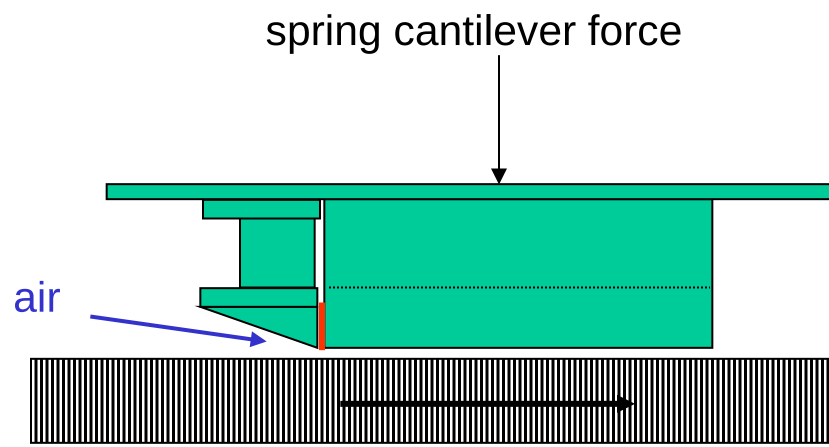

r = approx 3 cm
circumference = $2\pi \cdot 0{,}03 \approx 0{,}2$ m
7200 rpm = 120 rps
v = $0{,}2 \cdot 120 = 24$ m/s

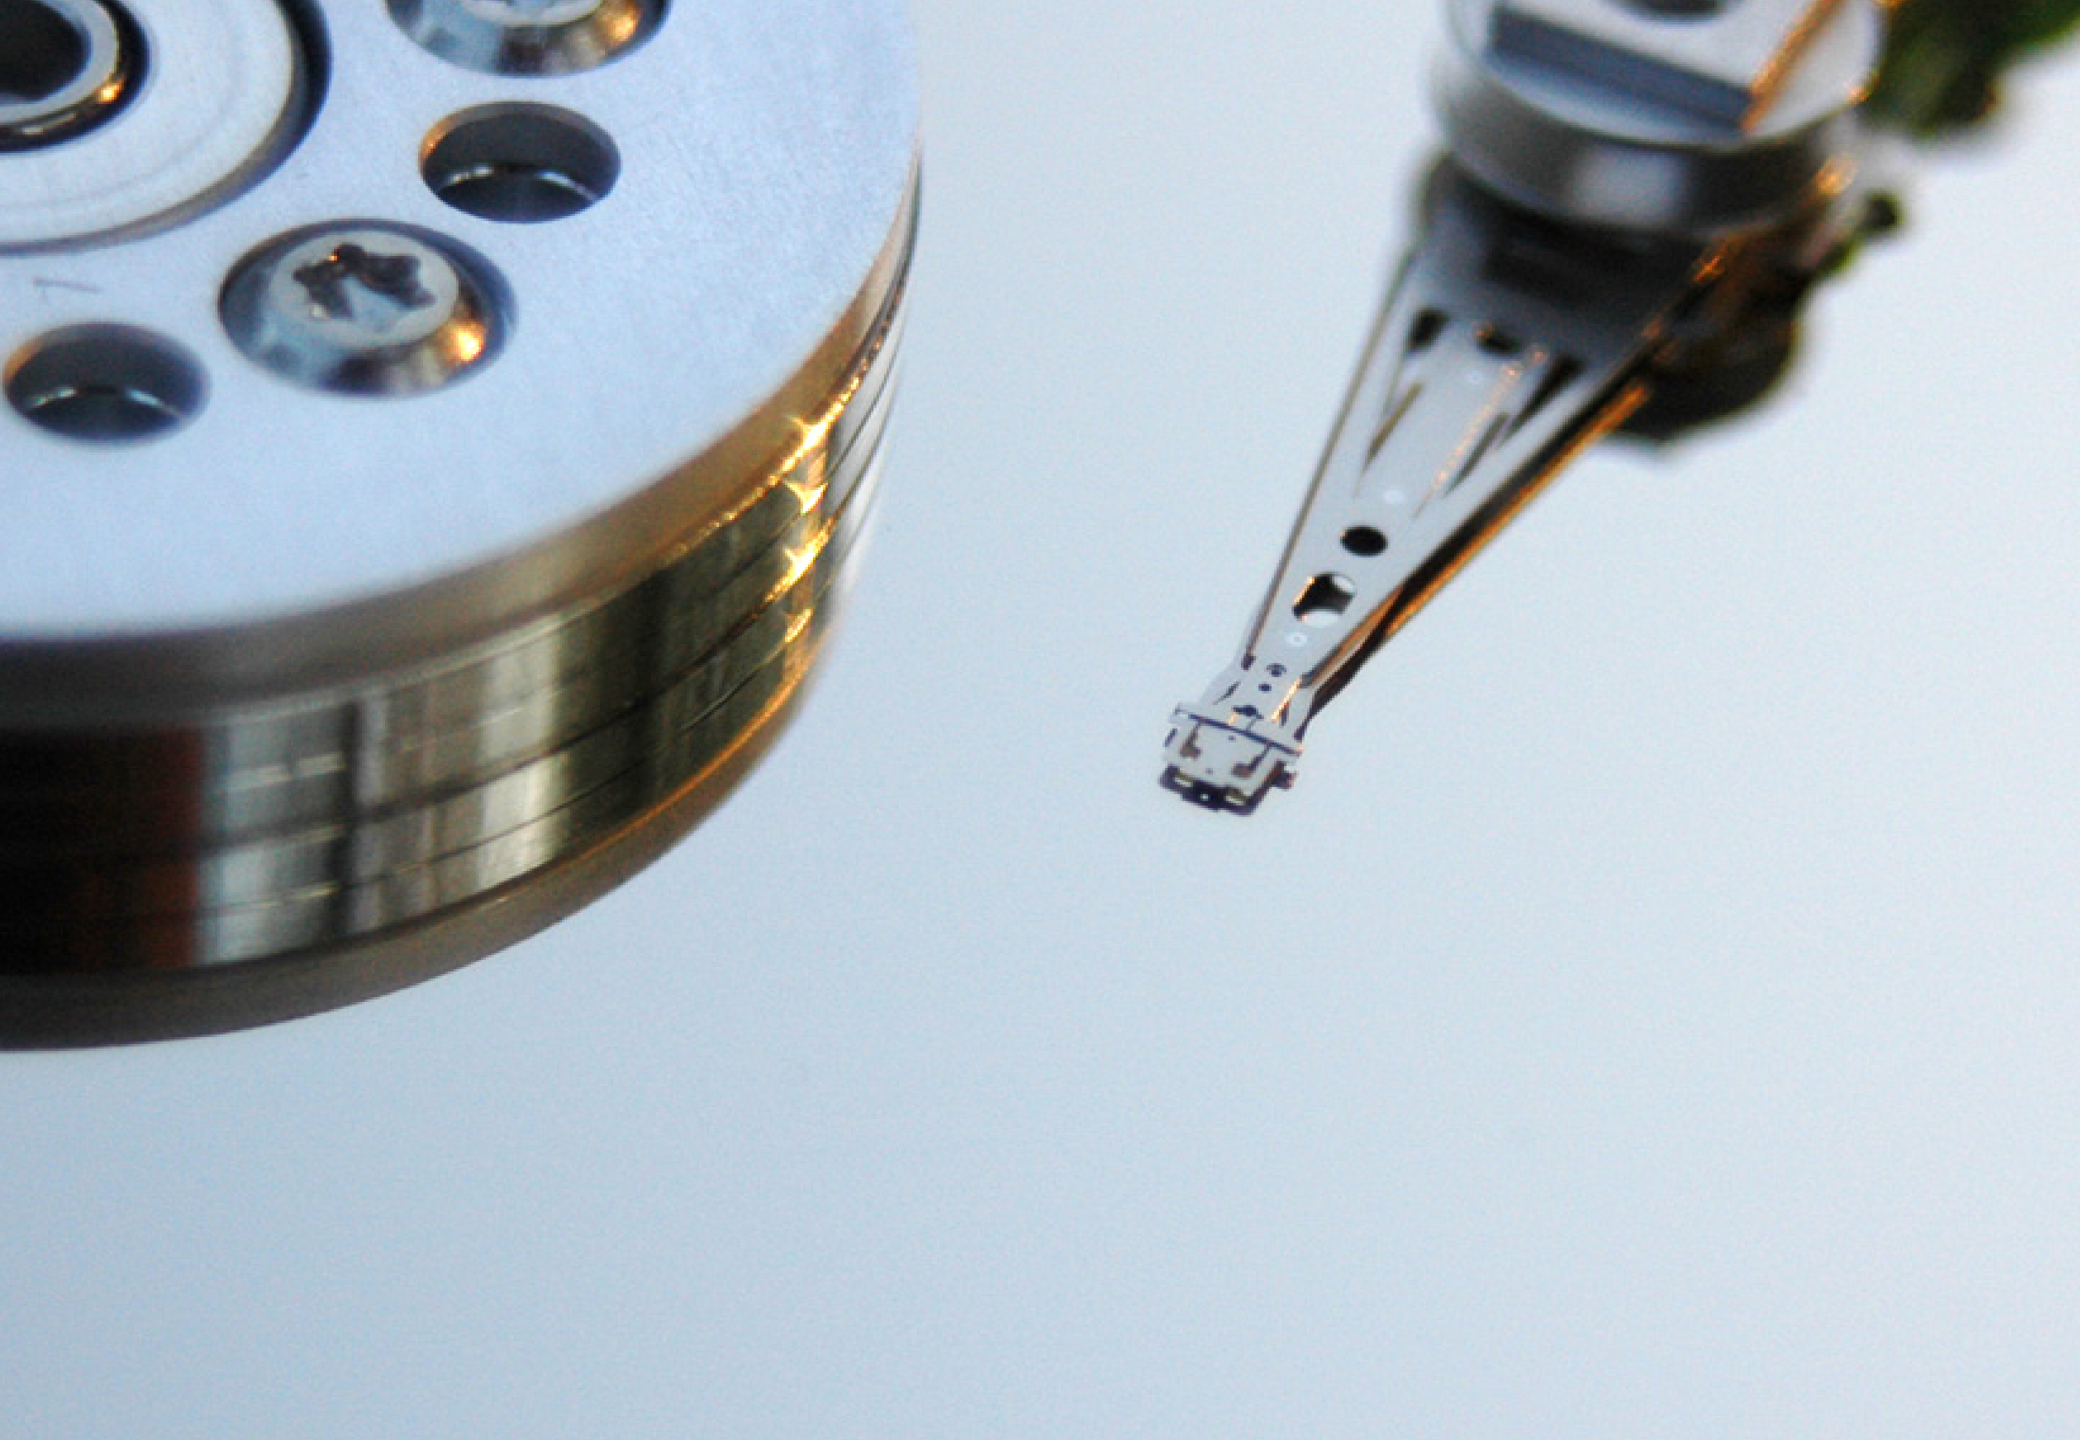

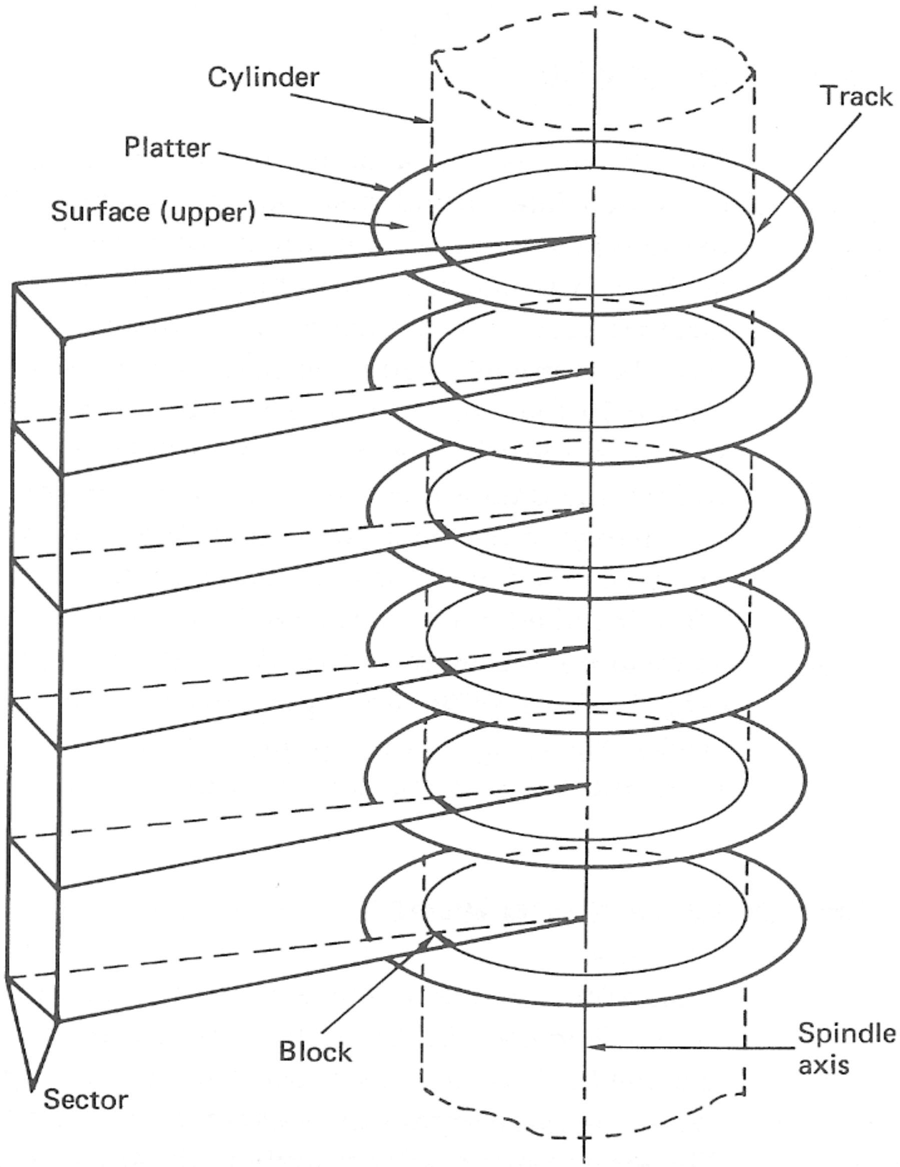

Figure 10.2 Disk terminology. Surface: one side of a platter. Track: path described on a surface by a fixed head. Cylinder: imaginary shape intersecting all surfaces at tracks of the same radius. Sector: angular subdivision of pack. Block: that part of a track within one sector. Each block has a unique cylinder, head and sector address.

platter = skiva
cylinder = cylinder
surface = yta
head = huvud
track = spår
sector = sektor
block = block

block address =
cylinder, head, sector

Multi-platter versus single platter

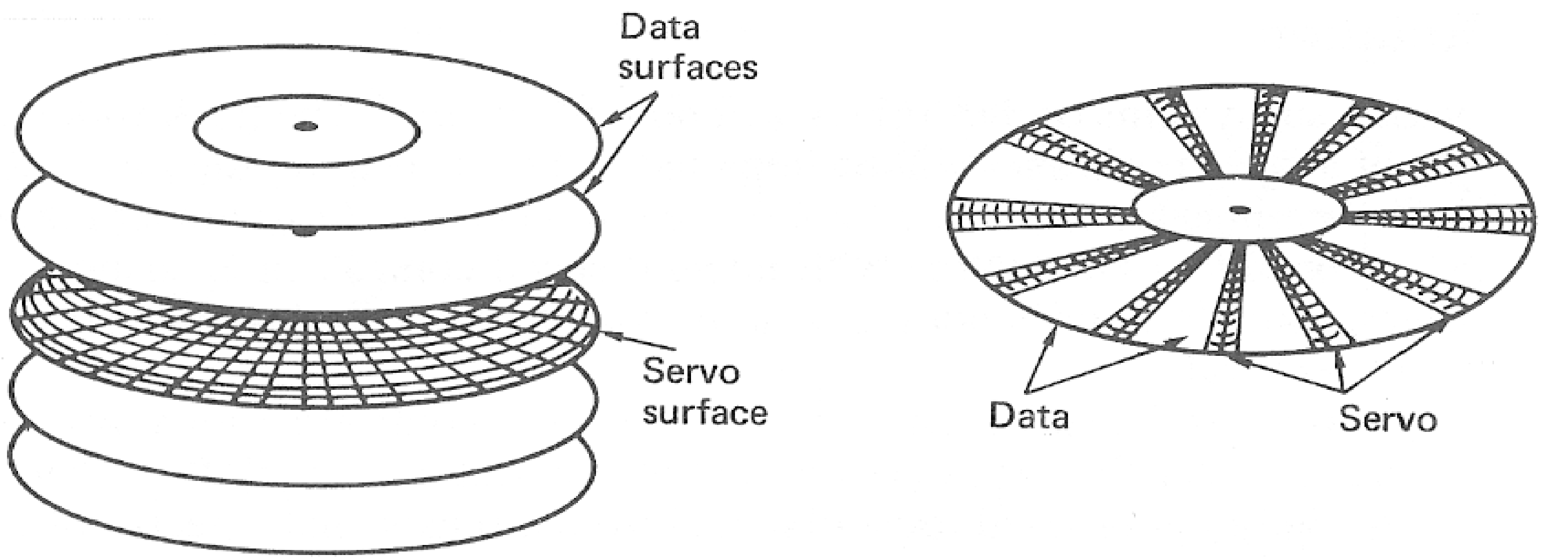

Figure 9.8 In a multiplatter disk pack, one surface is dedicated to servo information. In a single platter, the servo information is embedded in the data on the same surfaces.


1. The head is here
2. Direction of actuator positioning: Seek Time
3. Direction of disc rotation: Latency
4. The data needed is here

Primary performance factors

- Seek time

  = time for radial movement of the head (typ. 9 ms)
- Disk latency

  = average time for the desired sector to arrive under the head (always ½ revolution -> 4.17 ms at 7200 rpm)
- Bit transfer rate

Subsidiary performance factors

- Access time = seek time + disk latency
- External data rate

Contiguous files, or not?

- Cluster allocation is not always sequential
- With time, files become fragmented across non-contiguous sectors – more head seeks, longer time.
- ”Defragment” the hard drive periodically.
  OS might do this automatically when idling (good or bad)?
- Some real-time operating systems have function calls that allow the programmer to request files to be contiguous on disk. Files are pre-created with the desired length.

Architecture

There are four main mechanisms for data transfer

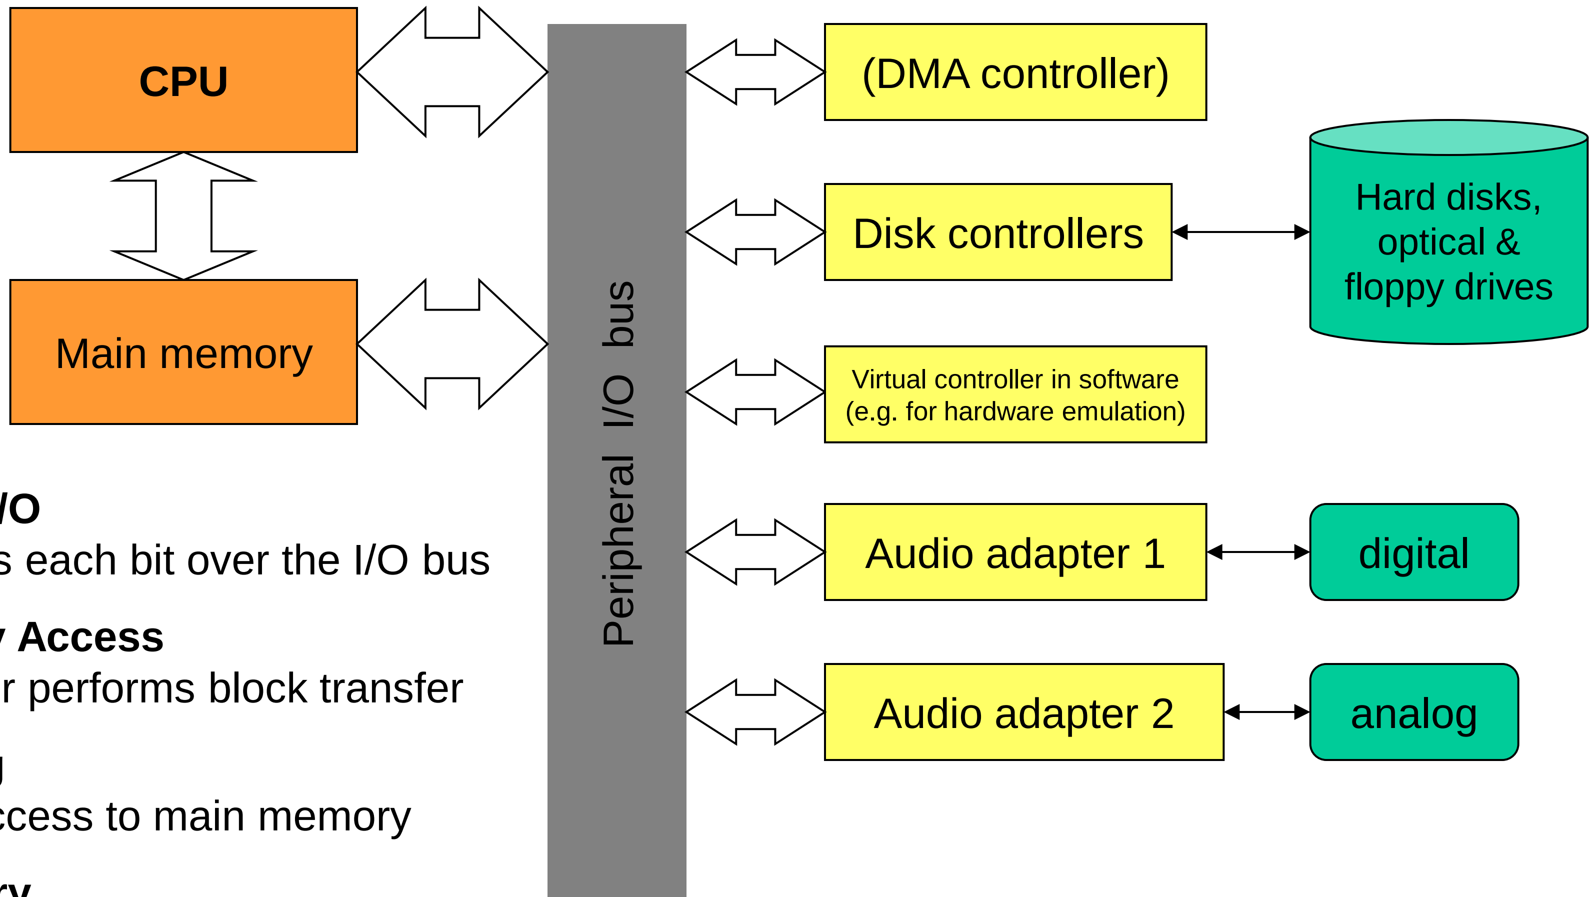

1. **Programmed I/O**
   – CPU transfers each bit over the I/O bus
2. **Direct Memory Access**
   - DMA controller performs block transfer
3. **Bus mastering**
   – device has access to main memory
4. **Shared memory**
   - CPU has access to memory on device

Architecture

Special subsystems such as main memory may have a private bus

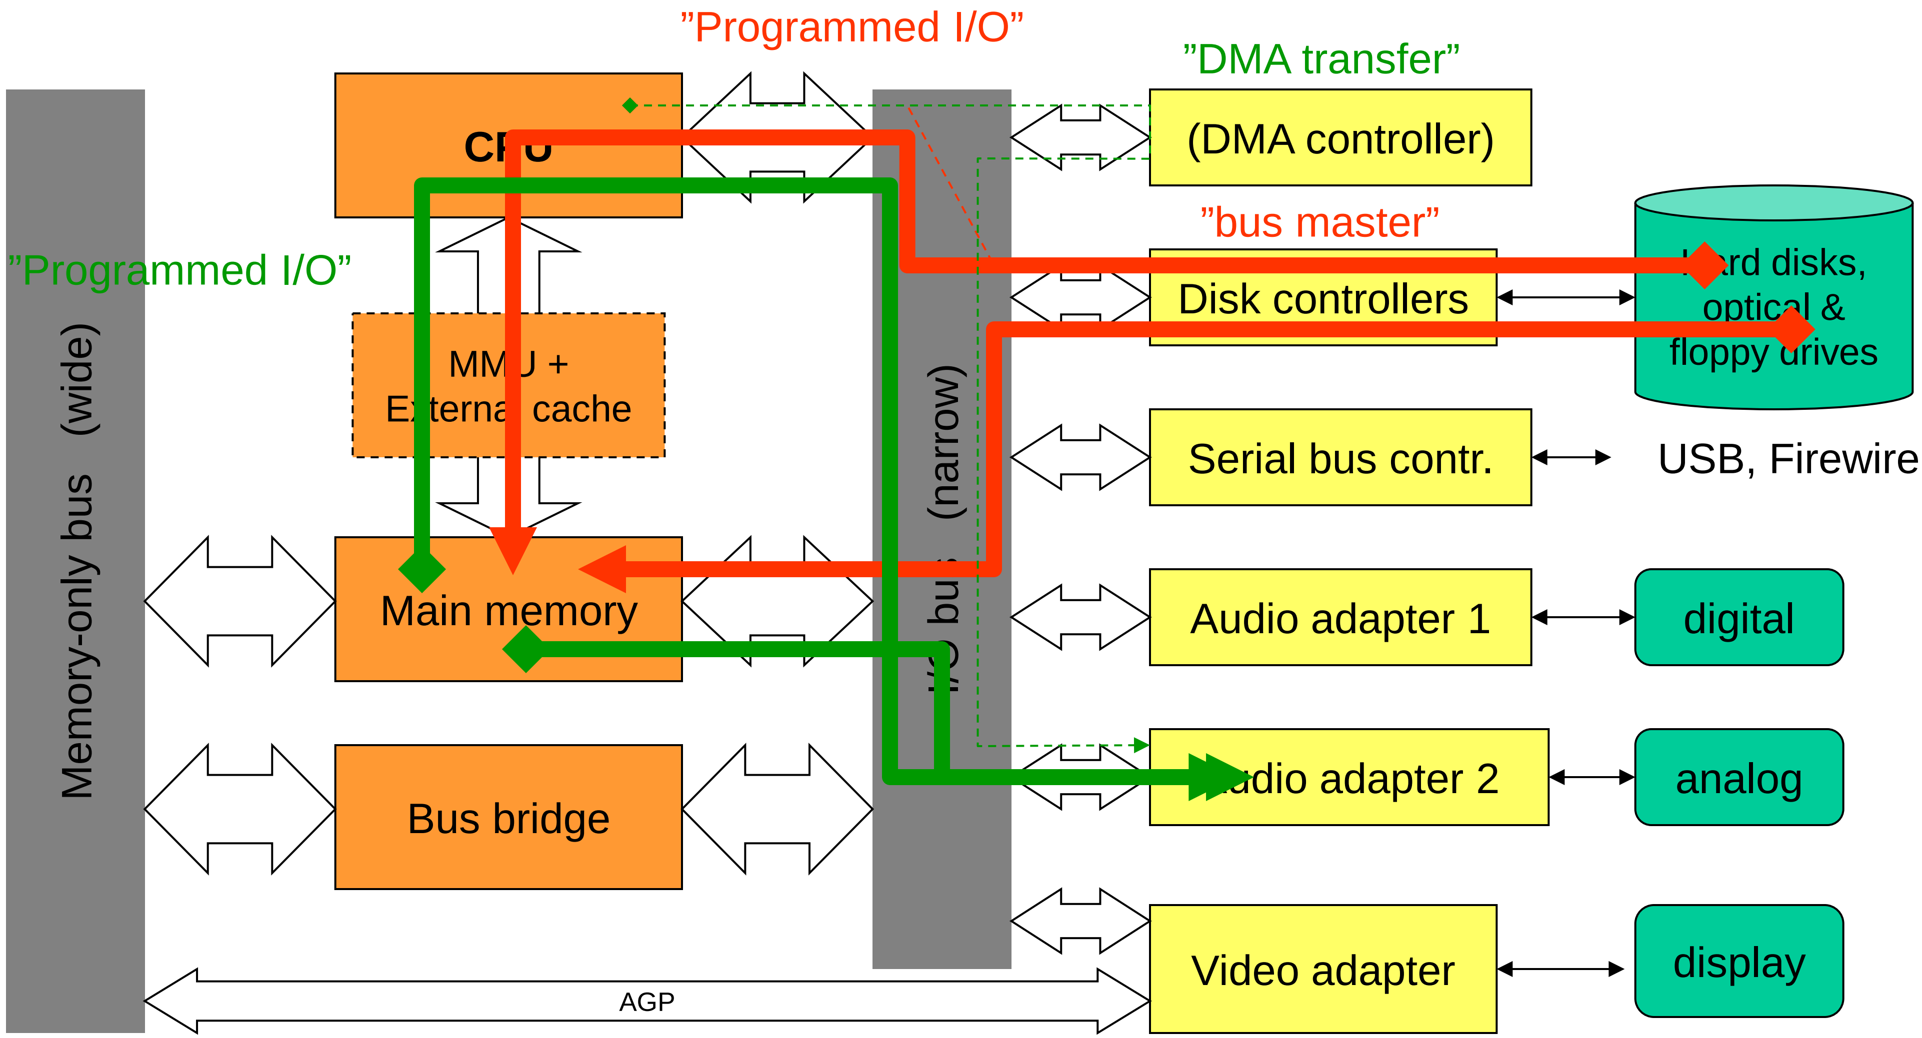

Timing in data streams

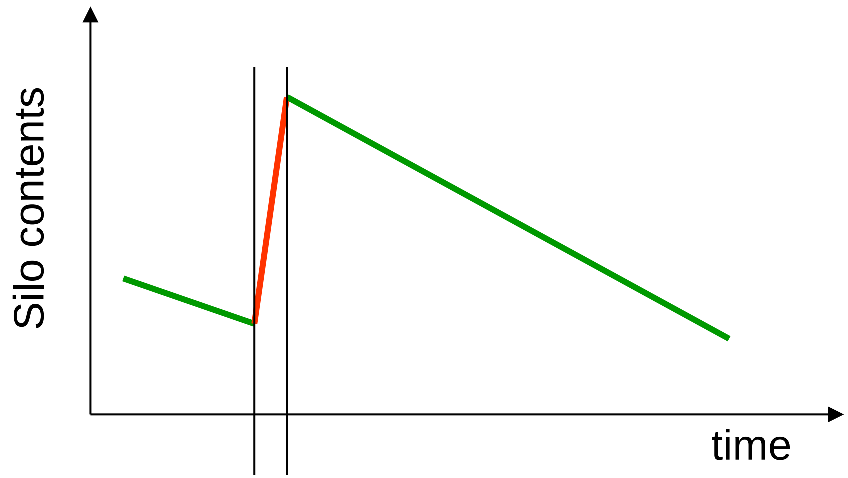

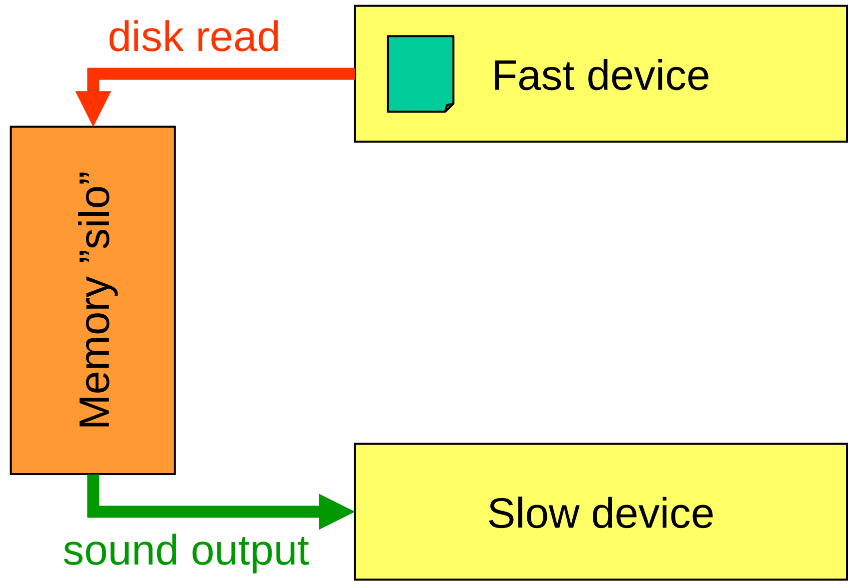

Dual buffering is used to maintain continuous throughput, in spite of system delays

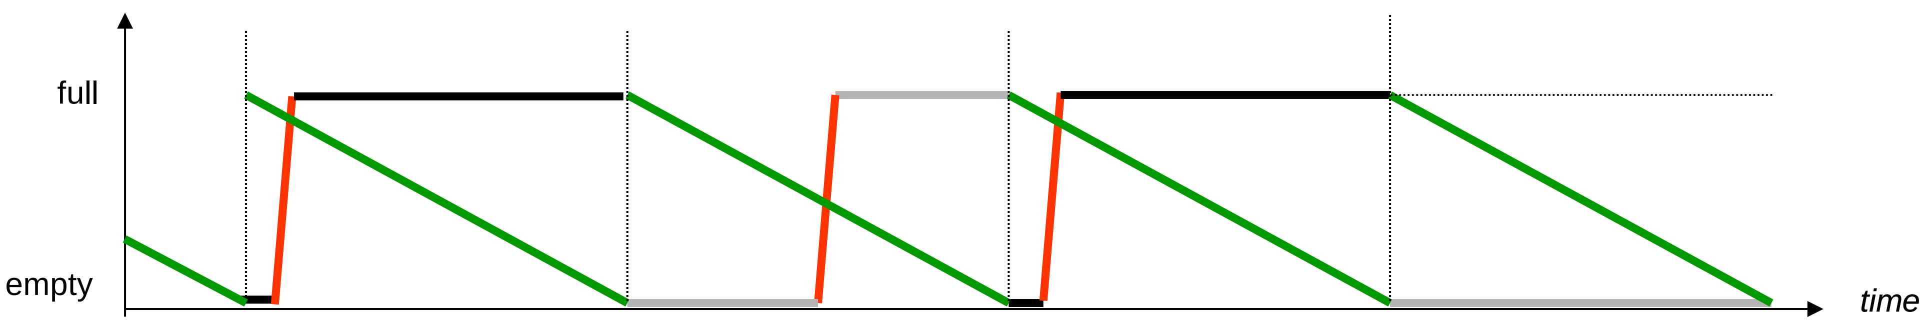

…but this requires two processes, or two processors

Multitasking and disk access

In a multitasking OS, the CPU is constantly switching between several threads or processes.
Hence several processes can be competing for access to the **hard disk**.

The OS **serialises** hardware access, making each process wait for its turn.

While the disk is seeking and turning,
the disk controller driver software can **yield** to other processes.

With several outstanding requests,
it can be quicker to process them out of order
(compare to an elevator).

Hardware subsystems involved in a "direct-to-disk" recording

A/D converter (synchronous)

- Controller or DSP on the sound card
  - Peripheral bus(-es) of the host computer
    - Host CPU
      - Host memory (RAM)
    - Host CPU
  - Peripheral bus
- Disk controller

Hard disk (burst operation, serialised access)

Virtual memory

A disk OS will typically implement virtual memory to make the RAM seem larger. Unused memory pages may be "swapped" or copied out to disk until a programme tries to access them.

Device programmers must ensure that real-time buffers are not "paged out" of RAM when the hardware interrupt occurs, or the delay to service will be too large.

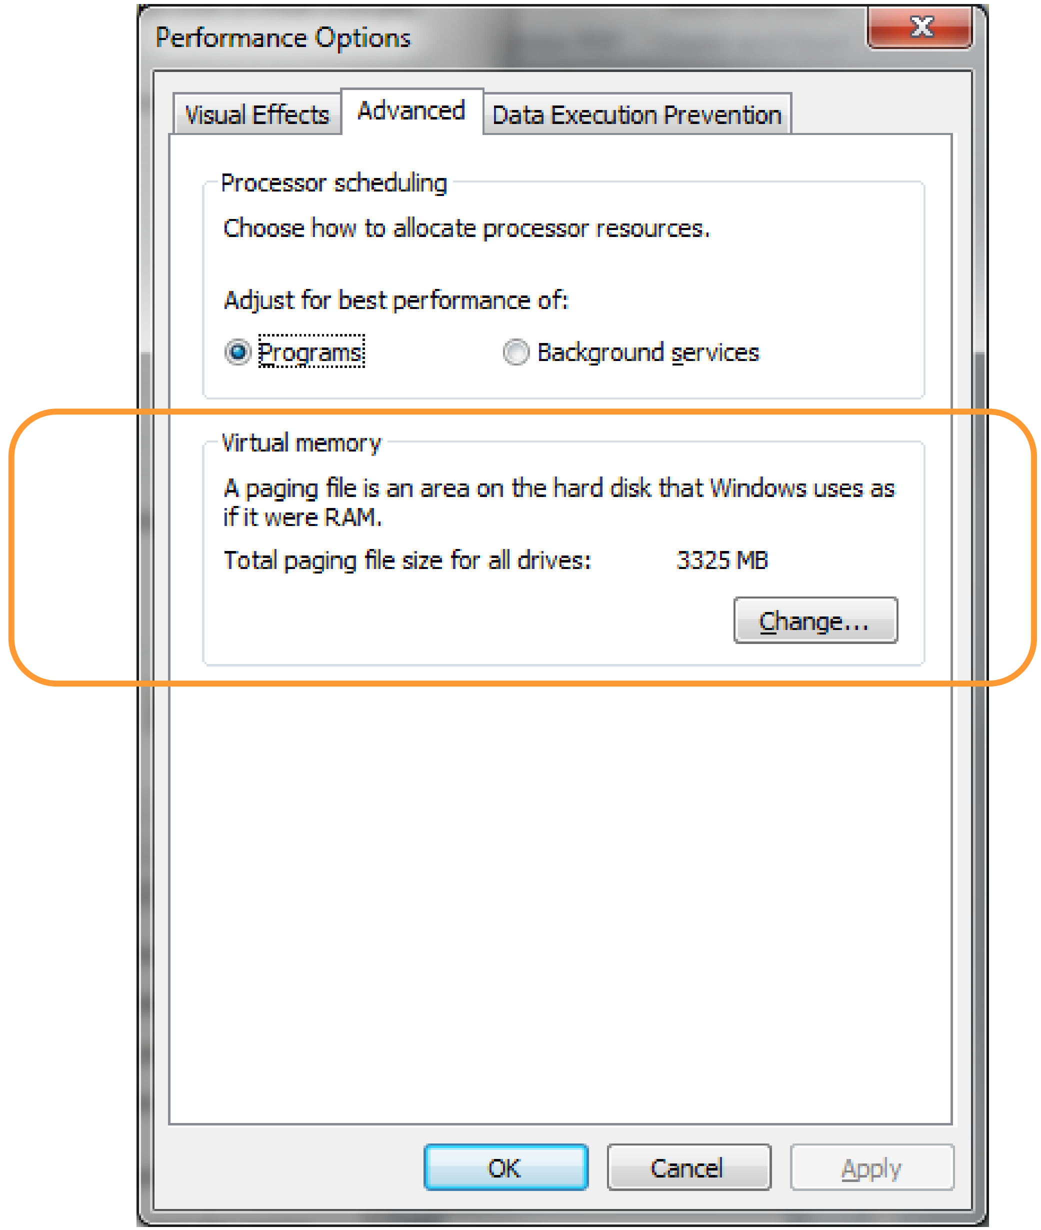

The write cache

A disk OS will usually provide a "write-behind cache".

This can cause problems when recording audio.

If the cache is large, then once it becomes full, it will take a long time to commit to disk.

Some OS's can disable write caching for single files, but the programmer must know how.

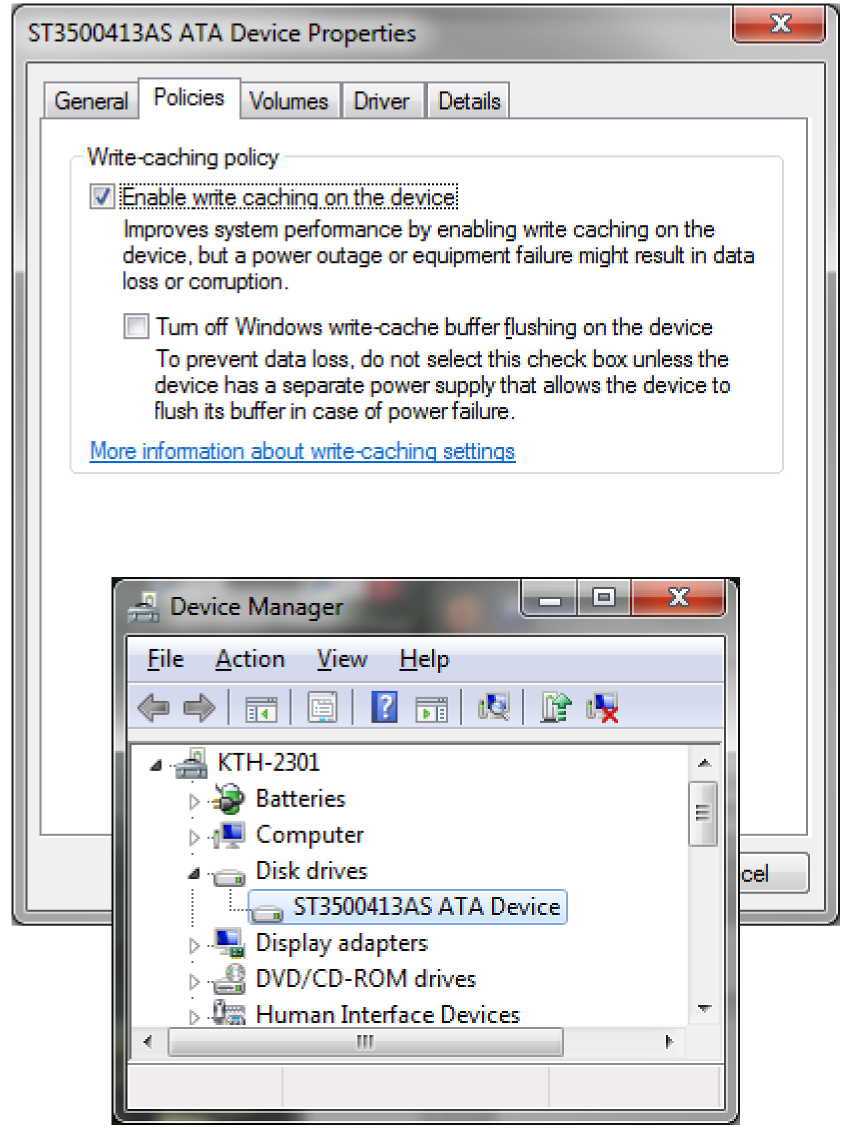

Summary

- Acquisition/playback is synchronous in real time,
  while storage is bursted for efficiency. This difference is resolved by memory buffering, usually in many different stages.
- Modern audio systems rely on very many subsystems, including CPU’s, buses, memories, memory-paged file systems and real-time converters.
- With modern hard disks, disk performance is no longer a bottleneck. Rather, it is the **complexity** of interfacing the subsystems, and of making the operating system attend to the **real-time data** concurrently with **other processes** that is difficult. Achieving smooth real-time performance with small delays, no audible gaps, and hardware flexibility is a big challenge.