
High Performance Computing

Lecture 34

Matthew Jacob

Indian Institute of Science

Understanding Program Behaviour

- Profiling: We have seen how to learn more about the important parts of our programs
 - to concentrate our optimization efforts there
- Earlier, we learned about the hardware and software that are involved in making our program execute on a computer system
- There are (at least) 2 problems with this that we had not considered

1. Problem with Silicon Memory

- Recall: Circuits that can remember things
 - Used in implementing registers, cache RAM, main memory
 - The circuits remember by
 - the state that a flip-flop is in or
 - the amount of charge stored in a capacitor
 - In both cases, the information is lost when the power source is turned off

2. Problem with How Programs Execute

% a.out

- Program runs as a process
- Memory image made up of text, data, stack and heap
- But that all disappears when the program finishes executing

Recall: Lifetime of Data

1. Lifetime = Execution time of program
2. Lifetime = Time between explicit creation of data & explicit deletion of data
3. Lifetime = During execution of a function (i.e., time between function call and return)
4. Lifetime = Beyond the execution time of the program
 - even if the power fails while the program is running

Agenda

1. Program execution: Compilation, Object files, Function call and return, Address space, Data & its representation (4)
2. Computer organization: Memory, Registers, Instruction set architecture, Instruction processing (6)
3. Virtual memory: Address translation, Paging (4)
4. Operating system: Processes, System calls, Process management (6)
5. Pipelined processors: Structural, data and control hazards, impact on programming (4)
6. Cache memory: Organization, impact on programming (5)
7. Program profiling (2)
8. **File systems: Disk management, Name management, Protection** (4)
9. **Parallel programming:** Inter-process communication, Synchronization, Mutual exclusion, Parallel architecture, Programming with message passing using MPI (5)

What is a File?

1. Storage for data that continues to exist beyond the lifetime of program
 - ❑ Persistent data
 - ❑ This will be possible only through the use of non-volatile, persistent secondary storage devices
 - ❑ like hard disk
2. A named sequence of data on a persistent storage device

Secondary Storage


Today, there are 3 main kinds

1. Magnetic: information stored on magnetic medium
 - e.g., hard disk drive, floppy disk, mag tape cartridge
2. Optical: information stored by optical properties
 - e.g., CD, DVD
3. Flash: information stored like in RAM cell, but with very slow rate of leakage
 - e.g., memory stick

About Magnetic Disks

- **Platter**: rotating metal disk covered with magnetic material


Top view of a platter



About Disks


- Platter: rotating metal disk covered with magnetic material

Side view of a platter



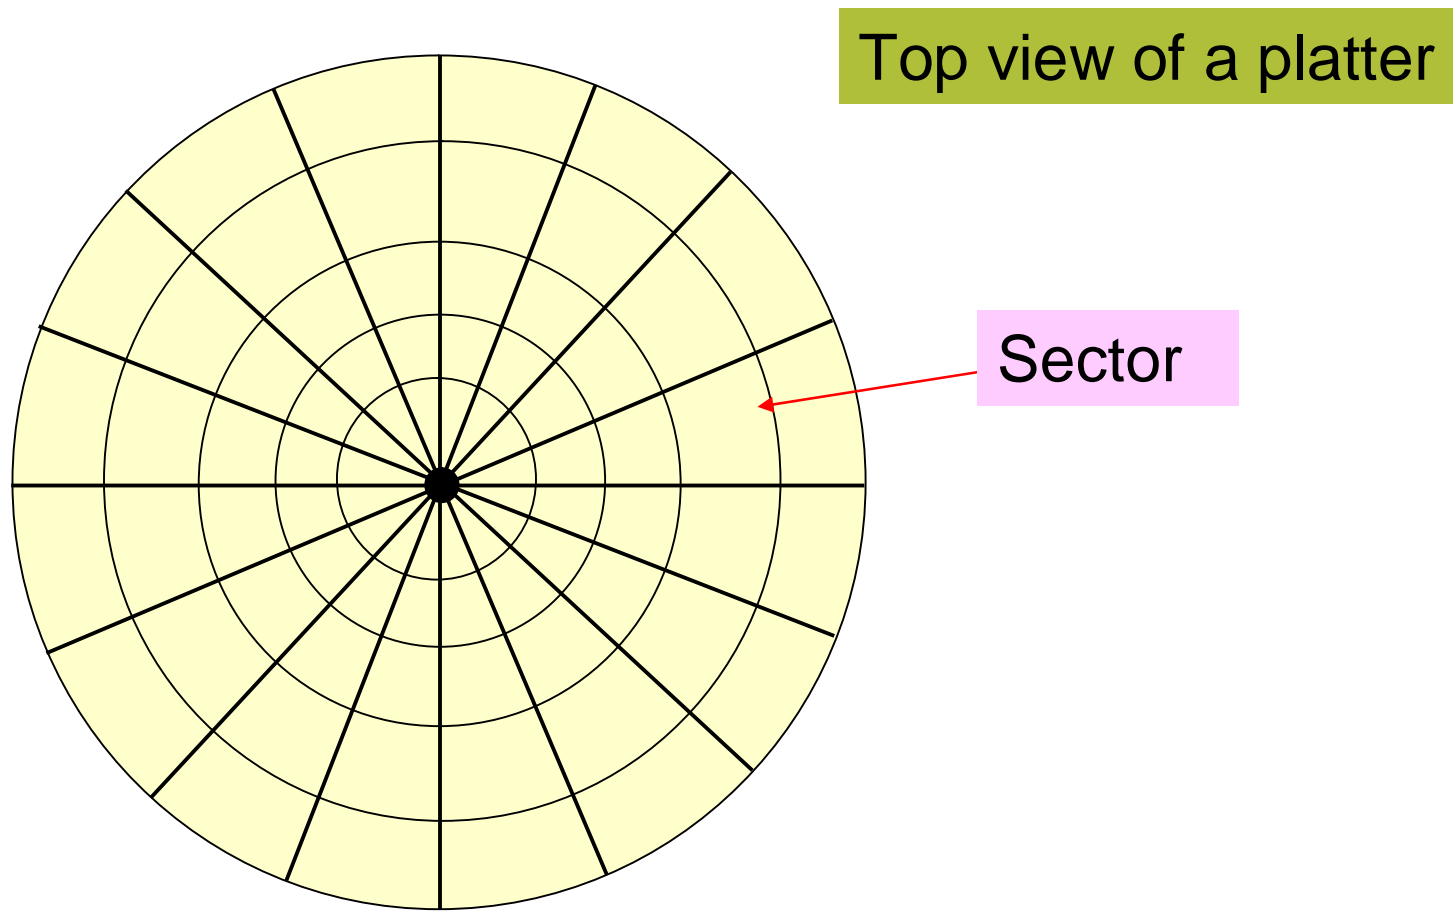
About Disks

- **Track**: concentric circular recording surface



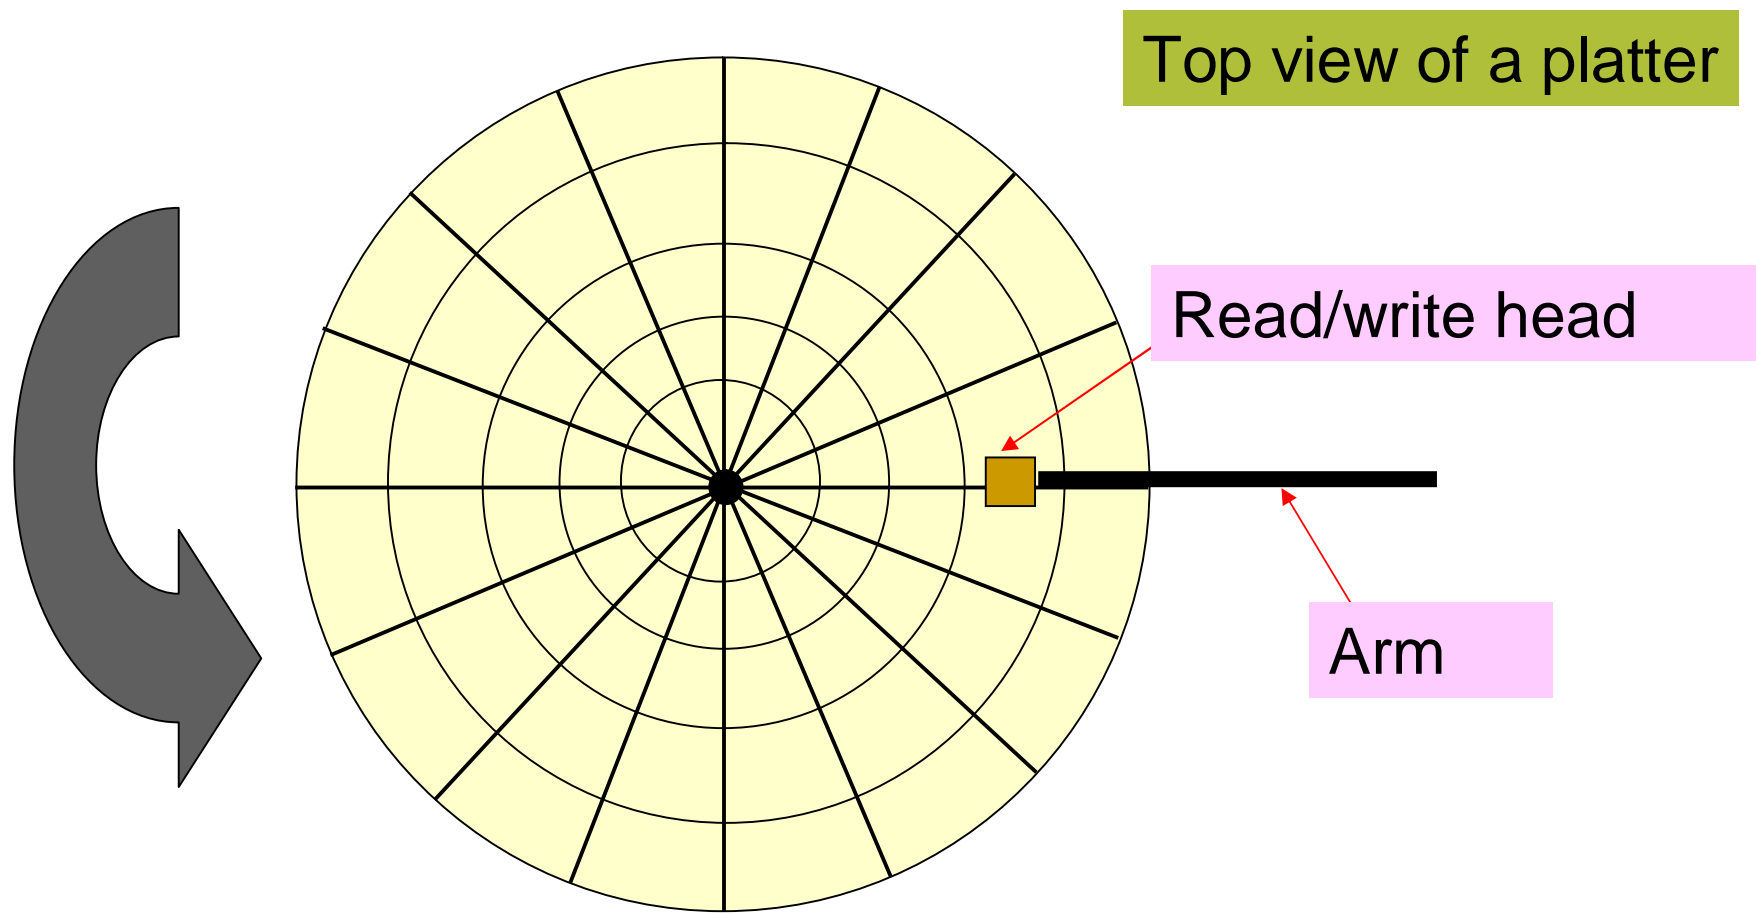
About Disks

- **Sector** or **block**: unit of track that is read/written at a time



About Disks


- Read/write head: An electromagnet that is used to read/write a sector



About Disks

- Multiple platters rotating together on a common spindle

Side view of a disk with multiple platters



Magnetic coating possibly on both upper and lower surfaces of platters

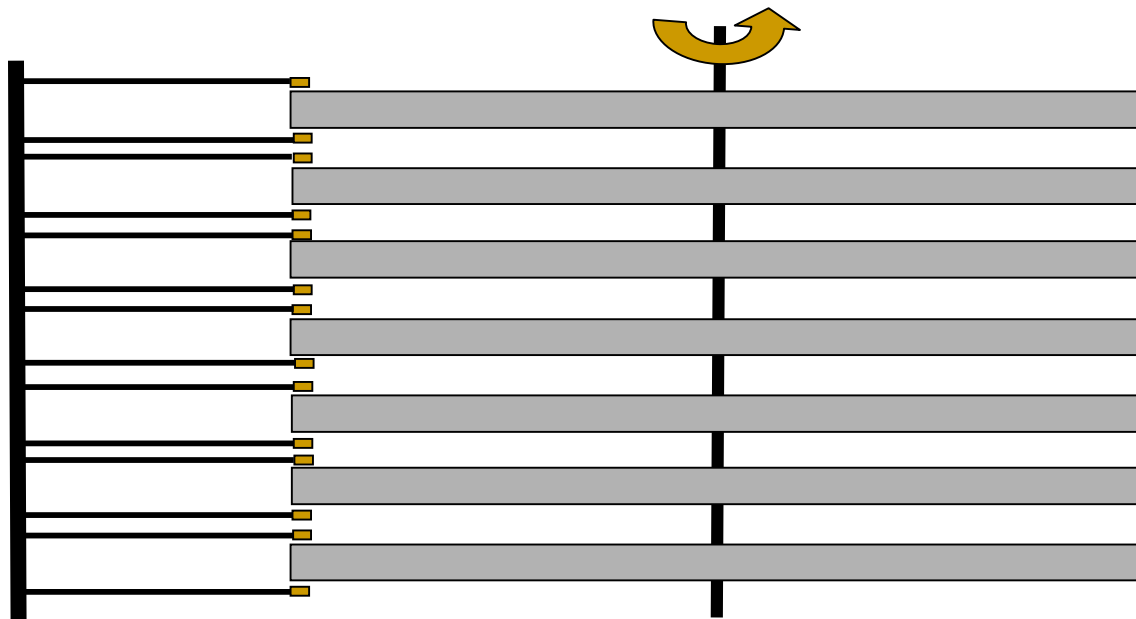
About Disks

- Multiple platters rotating together on a common spindle

Side view of a disk with multiple platters

With a separate Read/write head for each surface



All heads being connected to a single actuator (by which they move in or out)



About Disks

- **Cylinder**: all the tracks associated with a given actuator position

Current cylinder: Outermost track of each surface

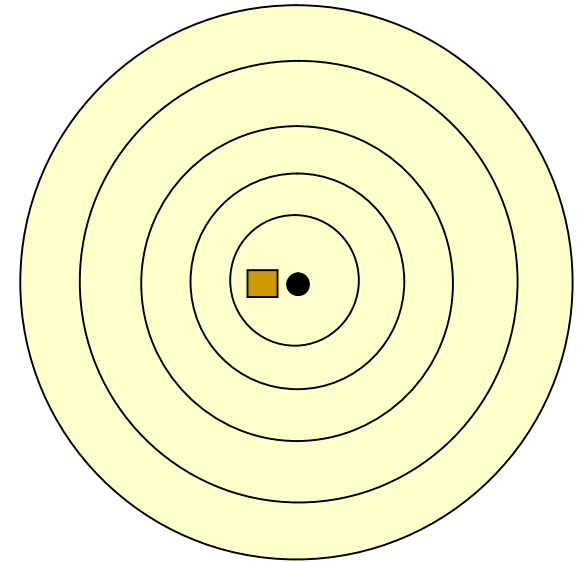



If the actuator then moves the heads in completely...

About Disks

- **Cylinder**: all the tracks associated with a given actuator position

Current cylinder: Innermost track of each surface



About Disks

- How long does it take to read/write a disk sector?
 - **Seek latency:** Time for actuator/disk arms to move to the correct cylinder
5-10 msec
 - **Rotational latency:** Time for correct sector to rotate to under the read/write head
2-3 msec for a 15,000 rpm disk
 - **Transfer time:** Time for the data to be transferred from the disk to the main memory
at 30MB/sec
 - Disk may currently be in a low power consumption mode (not spinning)
100s of msec