Physical Data Organization

CPS 116

Introduction to Database Systems

Announcements (November 3)

- Homework #3 due today
- Project milestone #2 due in a week

Outline

- It's all about disks!
  - That's why we always draw databases as 
  - And why the single most important metric in database processing is the number of disk I/O's performed
- Storing data on a disk
  - Record layout
  - Block layout

Storage hierarchy

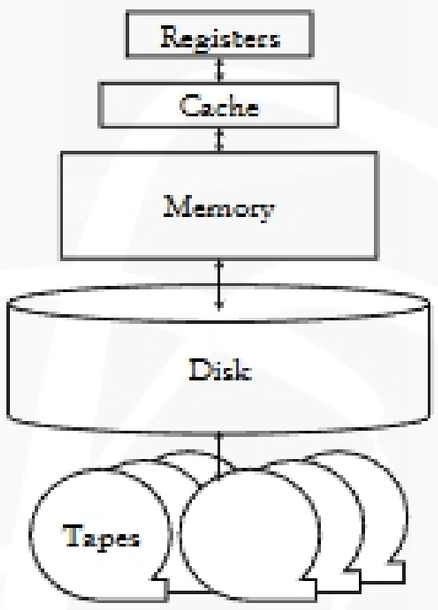

Why a hierarchy?

How far away is data?

Location	Cycles	Location	Time
Registers	1		
On-chip cache	2		
On-board cache	10		
Memory	100		
Disk	10^6		
Tape	10^9		

(Source: AlphaSort paper, 1995)

☞ I/O dominates—design your algorithms to reduce I/O!

A typical disk

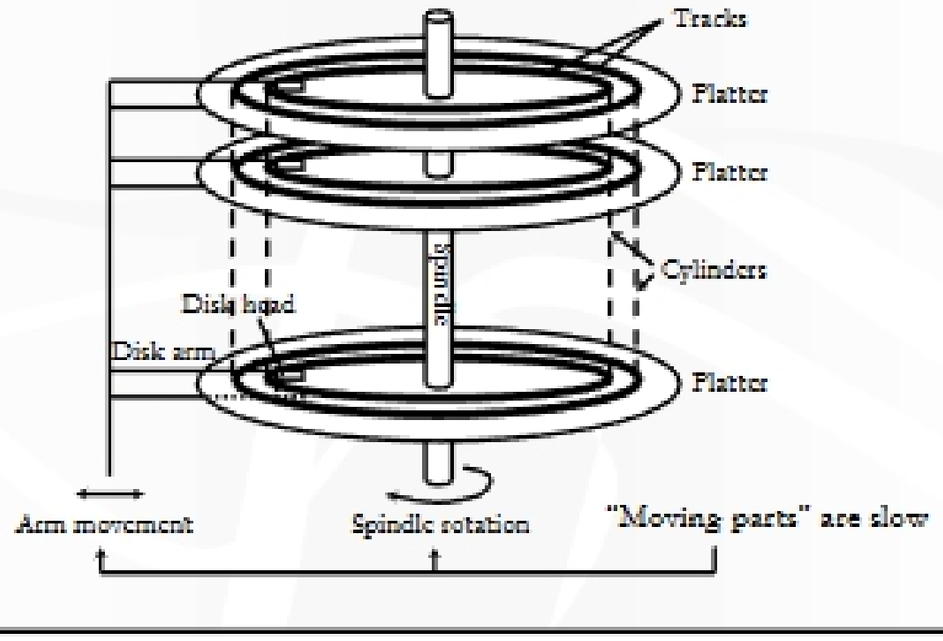

Top view

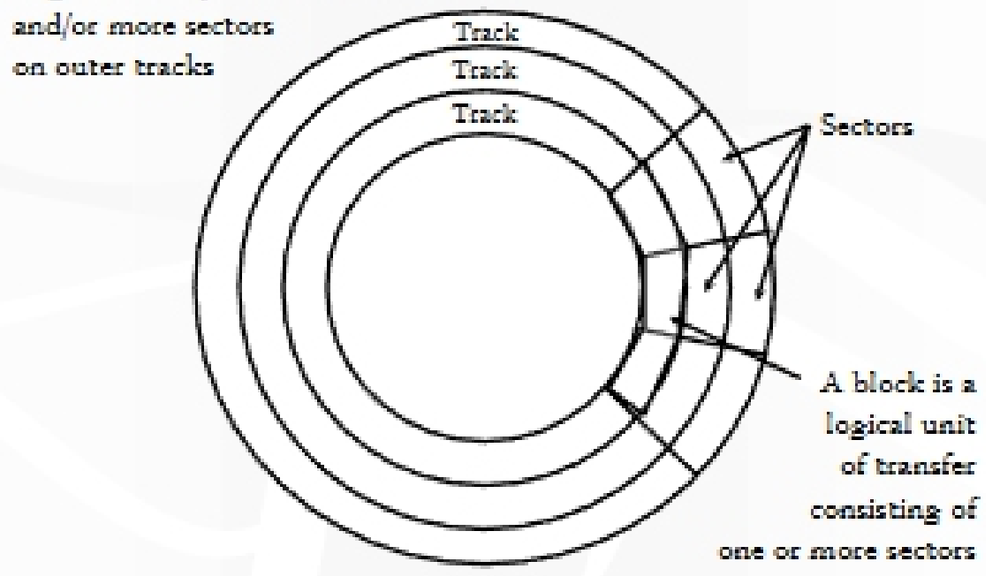

Disk access time

Sum of:

- Seek time: time for disk heads to move to the correct cylinder
- Rotational delay: time for the desired block to rotate under the disk head
- Transfer time: time to read/write data in the block (= time for disk to rotate over the block)

Random disk access

Seek time + rotational delay + transfer time

- Average seek time
  - Time to skip one half of the cylinders?
  - "Typical" value: 5 ms
- Average rotational delay
  - Time for a half rotation (a function of RPM)
  - "Typical" value: 4.2 ms (7200 RPM)