- [x] 1. Last time
- [x] 2. Disks (HDDs)

Next time: intro to file systems

---

2. Disks

- [x] A. What is a disk?
- [x] B. Geometry
- [x] C. Performance
- [x] D. Common numbers
- [x] E. Interface to disk
- [x] F. Performance II
- [ ] G. Disk scheduling (performance III)
- [x] H. Technology + systems trends

$/GB

- stack of magnetic platters
- rotate together on central spindle

3600 - 15,000 RPM (=60 - 250 rot/sec)

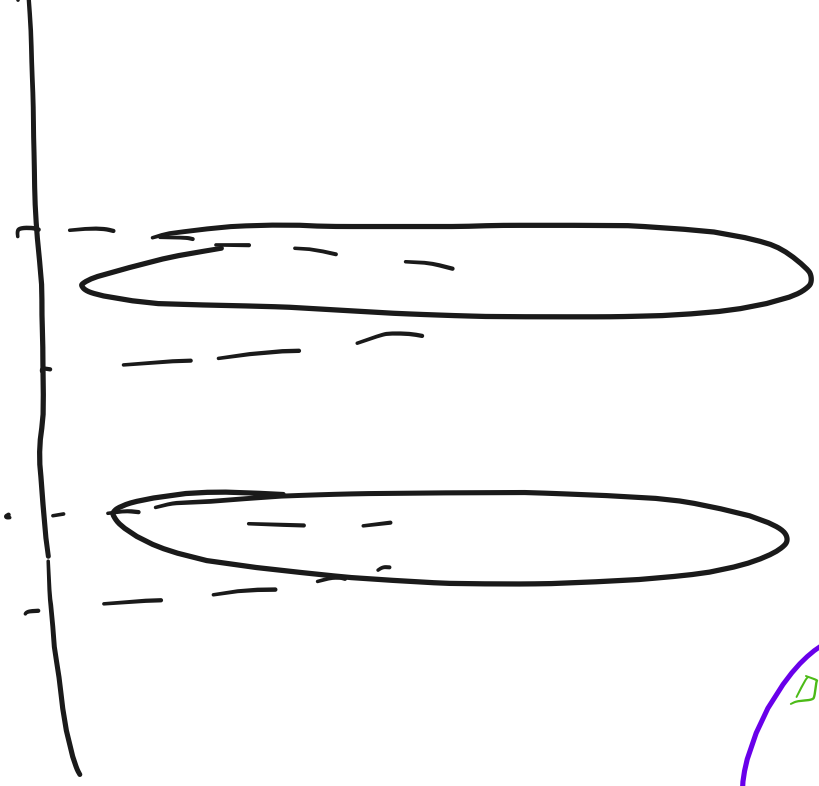

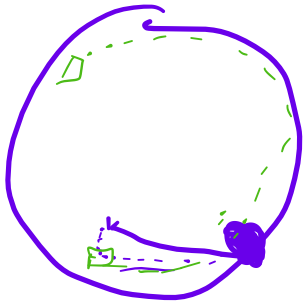

Arms move together

Arms contain disk heads; heads read + write to platters

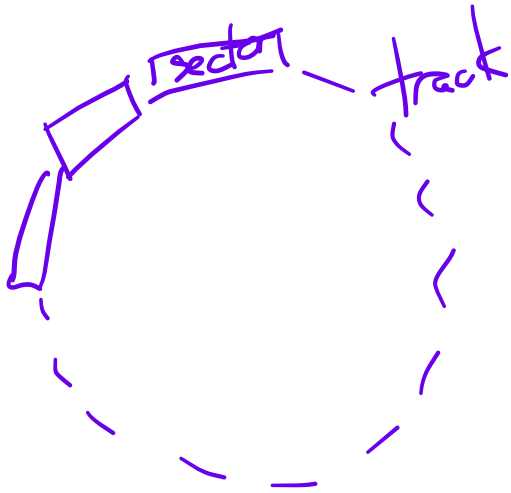

Geometry

- track: circle on a platter
- sector: chunk of a track
- cylinder: all tracks of fixed radius on all platters
- generally only one head active at a time
- disk positioning system: moves head to a track and keeps it there

- seek: 4 phases: speedup, coast, slowdown, settle

Performance

Components of total transfer time:
- rotational delay
- seek delay
- transfer time

"Avg seek time"
- "time to seek 1/3 of the disk", but manufacturers might report:

1/3 of "time to seek the whole disk"

Common #s

Capacity: TBs common (10^{12} bytes vs. 2^{40} bytes)

Platters: 8

Number of cylinders: tens of thousands

Sectors per track: ~1000

RPM: 10,000

Transfer rate: 50-150 MB/s

MTBF: ~1 million hours

How driver interfaces to disk

Sectors

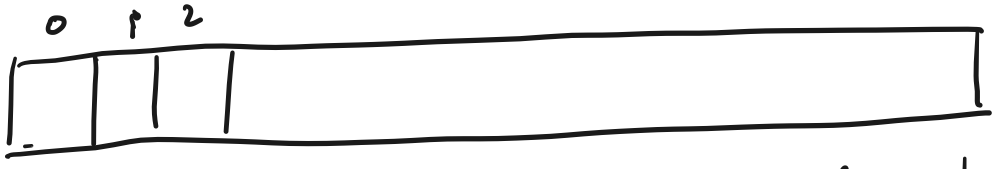

Interface to disk is a linear array of sectors

Sector: 512 bytes, moving to 4KB written atomically

Disk does some cool things under the hood (invisibly to OS)

Zoning

Skewing

Sparing

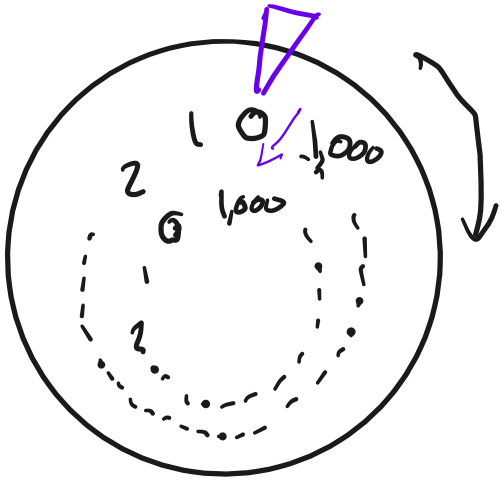

Disk performance example

Spindle speed: 12,000 RPM

Avg. seek time: 12 ms

Transfer rate: 128 MB/s

Sector: : 512 bytes

(a) What is the throughput if doing 500 sector reads, spread randomly over the disk and serviced in FIFO order?

(b) Same question, but now the reads are sequential

(a)

$$TPUT = \frac{\text{bytes}}{\text{time}} = \frac{500 \times 512 \text{ bytes}}{500 \text{ random reads}}$$

$$1 \text{ read} = \underbrace{\text{rot. delay}}_{2.5\text{ ms}} + \underbrace{\text{seek time}}_{12\text{ ms}} + \underbrace{\text{transfer time}}_{4\,\mu\text{s}}$$

$$\frac{\text{minutes}}{12{,}000 \text{ rot}} \times \frac{60\text{s}}{\text{min}} \times \frac{1}{2} = \frac{60\text{s}}{24{,}000 \text{ } \frac{1}{2}\text{ rot}} = \frac{0.005\text{ s}}{400 \text{ } \frac{1}{2}\text{ rot}} = \frac{2.5\text{ ms}}{(\frac{1}{2})\text{ rot}}$$

$\frac{s}{128MB} \times 512B = \frac{s \cdot 2^9 B}{2^{27} B} = 2^{-18} \cdot s \approx 4\mu s$

$2^{-18} = \frac{1}{2^{20}} \cdot 4$

$t_{put} = \frac{500 * 512}{(14.5ms + 4\mu s) \cdot 500} \approx \frac{512\,B}{14.5ms}$ $\quad 14.5 \cdot 70 \approx 1000$

$\frac{1}{14.5ms} = \frac{1}{.0145s} = 70 = \frac{70 \cdot 512\,B}{s}$

$\approx \boxed{\frac{35kB}{s}}$

$t_{put} = \frac{bytes}{time} = \frac{500 \times 512}{rot + seek + transfer}$

rot delay: 2.5 ms
seek delay: 12 ms
transfer: ? 2 ms

$\frac{s}{128MB} \times 500 \times 512B \approx \frac{2^9 \cdot 2^9 \cdot B \cdot s}{2^{27} MB} = \frac{2^{18} s}{2^{27}}$

$= 2^{-9} s \approx 2 ms$

$t_{put} = \frac{500 \times 512\,B}{16.5 ms} \approx 60 \times 512 \times 512 \frac{B}{s}$

$\approx 64 \times 512 \times 512\,B/s$

$\approx 2^6 \cdot 2^9 \cdot 2^9\,B/s$

$= 2^{24}$ B/s $=$ **16 MB/s**

(a)

$$\text{TPUT} = \frac{\text{bytes}}{\text{time}} = \frac{500 \times 512 \text{ bytes}}{500 \text{ random reads}}$$

1 read = rot. delay + seek time + transfer time

$$\text{rot. delay} = \frac{60\text{s}}{7200} \cdot \frac{\text{min}}{\text{rot}} = \frac{6}{7200} \approx 5\text{ms}$$

1 rot. $= \frac{60}{\text{min}} \times \frac{}{12{,}000 \text{ rot}} = \frac{}{1200} \approx$

avg. rot. delay = 2.5 ms

avg. seek time = 12 ms

transfer time $= \frac{1 \text{ sec}}{128 \text{ MB}} \times 512 \text{ bytes} = \frac{2^9}{2^{27}} = 2^{-18}$

$\approx 4 \mu s$

$$500 \times \frac{512 \text{ bytes}}{500 \times (2.5 \text{ ms} + 12 \text{ ms} + 4 \mu s)} \approx \frac{512 \text{ bytes}}{14.5 \text{ ms}} \approx \frac{512 \times 70 \text{ bytes}}{s}$$

$$\left\{ \frac{1}{14.5 \text{ ms}} \approx \frac{70}{70 \times 14.5 \text{ ms}} = \frac{70}{1000 \text{ ms}} = \frac{70}{s} \right\}$$

$$\approx 500 \times 70 \frac{\text{bytes}}{\text{sec}} = \frac{35000 \text{ bytes}}{\text{sec}} = \boxed{\frac{35 \text{ KB}}{\text{sec}}}$$

(b)

$$\text{TPUT} = \frac{\text{bytes}}{\text{time}} = \frac{500 \times 512 \text{ bytes}}{1 \text{ seq. read}}$$

avg. rot. delay = 2.5 ms

avg. seek time = 12 ms

transfer time $= \frac{1 \text{ sec}}{128 \text{ MB}} \times 500 \times 512 \text{ bytes} \approx \frac{2^9 \times 2^9}{2^{27}} = 2^{-9}$

$= 2 \text{ ms}$

$$\frac{500 \times 512 \text{ bytes}}{2.5\text{ms} + 12\text{ms} + 2\text{ms}} = \frac{500 \times 512 \text{ bytes}}{16.5\text{ ms}}$$

$$\approx \frac{30 \times 512 \text{ bytes}}{1\text{ ms}}$$

$$= \frac{3 \times 5\text{ KB}}{1\text{ms}} = \frac{15\text{ MB}}{s}$$