- [x] 1. Last time
- [x] 2. Disks (HDDs)
- [ ] 3. Intro to file systems (next time)

SSD $/GB VHS VCR

---

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) (see notes)
- [x] H. Technology + systems trends

$/GB

- stack of magnetic platters
- rotate together on central spindle

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

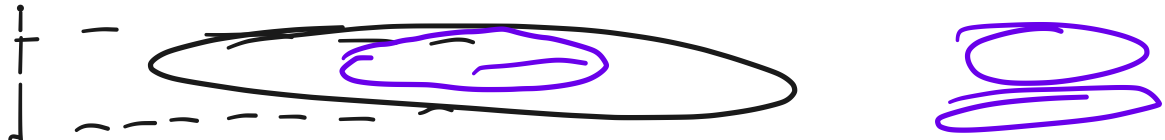

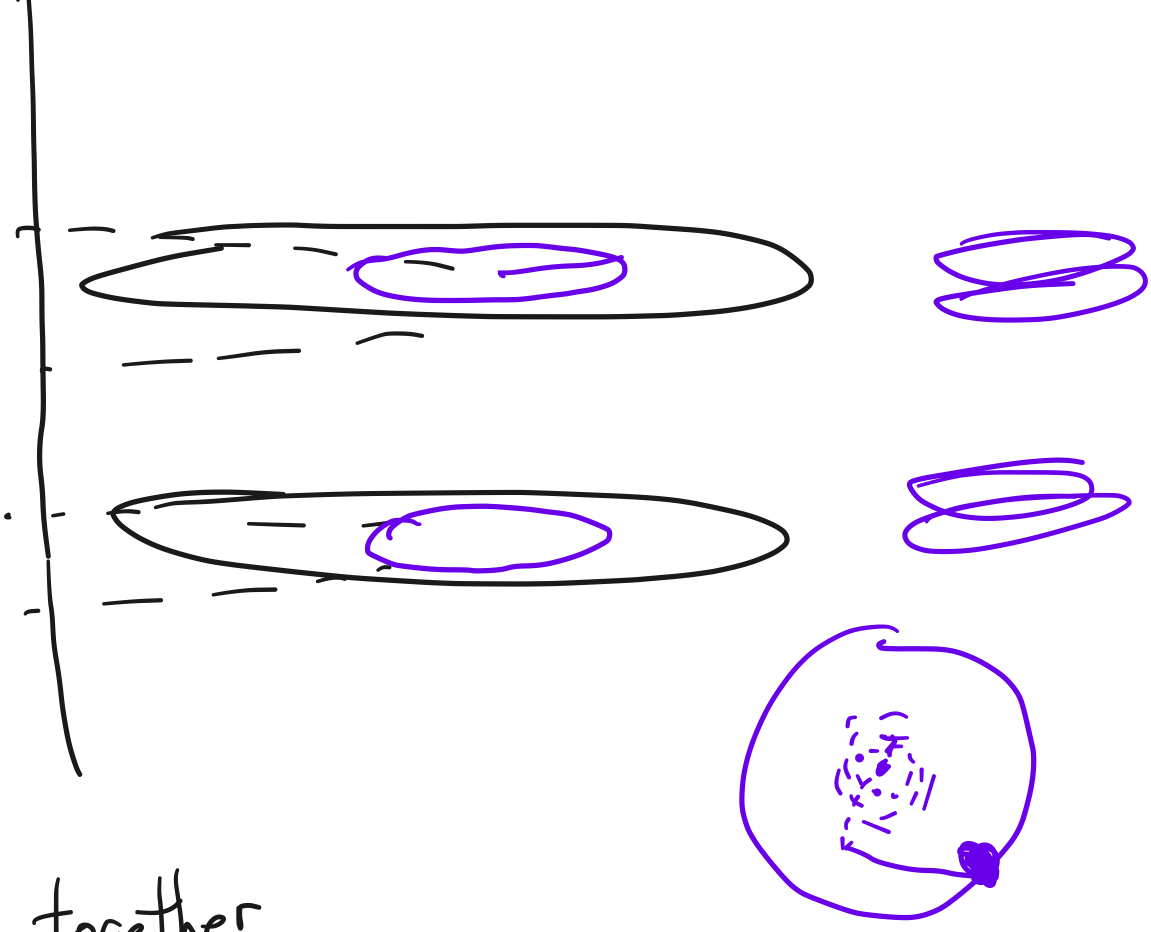

Arms move together

Arms contain disk heads; heads read + write to platters

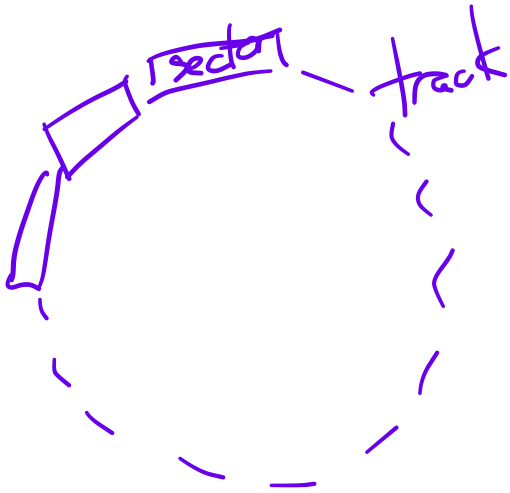

Geometry

- 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

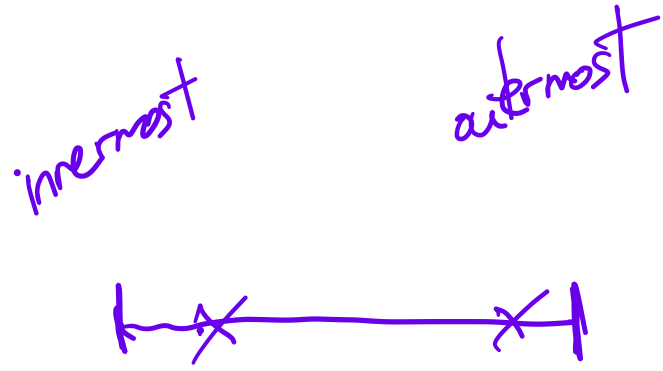

"Avg seek time"
- "time to seek 1/3 of the disk", but manufacturers might report:

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

Common #s

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

Platters: 8

# of thousands or more

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

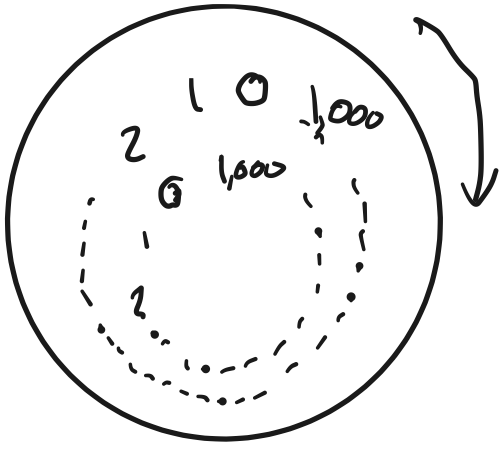

Disk performance example

Spindle speed: 12,000 RPM

Avg. seek time: 12 ms

Transfer rate: 128 MB/s

Sector: : 512 bytes

$512\,B \times \frac{1\,s}{128\,MB} =$

$2^9 B / 2^{27} B \cdot s = 2^{-18} s = \frac{1}{256k\,s} \approx 4\,\mu s$

(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)

$$TPUT = \frac{\text{bytes}}{\text{time}} = \frac{500 \times 512}{500(\text{seek delay} + \text{rot. delay} + \text{transfr delay})}$$

seek delay: 12 ms

avg rot delay: $\frac{1\,rot}{2} \times \frac{1\,minute}{12{,}000\,rot} \times \frac{60\,s}{min} = \frac{6}{2400} = \frac{1}{400}$

transfer delay

(4) $TPUT = \frac{bytes}{time} =$

$\frac{512\,B}{12ms + 2.5ms} = \frac{512B}{14.5ms} \approx .0B$ $= 2.5ms$

$14.5ms \cdot \frac{1}{70}s$

$\frac{512B}{\frac{1}{70}s} = \frac{70 \times 512B}{s}$

$\approx \frac{35,000B}{s}$

$\approx 35\,KB/s$

$TPUT = \frac{bytes}{time} = \frac{500 \times 512}{\text{lat seek} + \text{lat. rot delay} + \text{transfr}}$

$= \frac{500 \times 512}{14.5ms +}$

$500 \times 512\,B \times \frac{1s}{128\,MB} \approx \frac{2^9 \times 2^9\,B \cdot s}{2^{27}\,B}$

$\approx 2^{-9}\,s = \frac{1}{512}\,s \approx 2\,ms$

$$TPUT = \frac{500 \times 512\,B}{14.5ms + 2ms} \approx \frac{(512)^2}{16.5ms} = \frac{(512)^2}{\frac{1}{60}\,s} =$$

$$= 60 \times (512)^2 \frac{B}{s}$$

$$= 60 \times 2^{18} \frac{B}{s}$$

$$= 64 \times 2^{18} \frac{B}{s}$$

$$= 2^6 \times 2^{18} \frac{B}{s}$$

$$= 2^{24} \frac{B}{s}$$

$$= 16\,MB/s$$