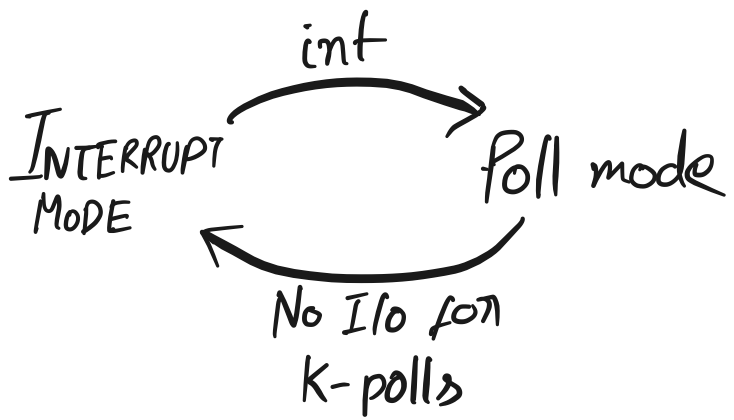


CS202: I/O, Disks

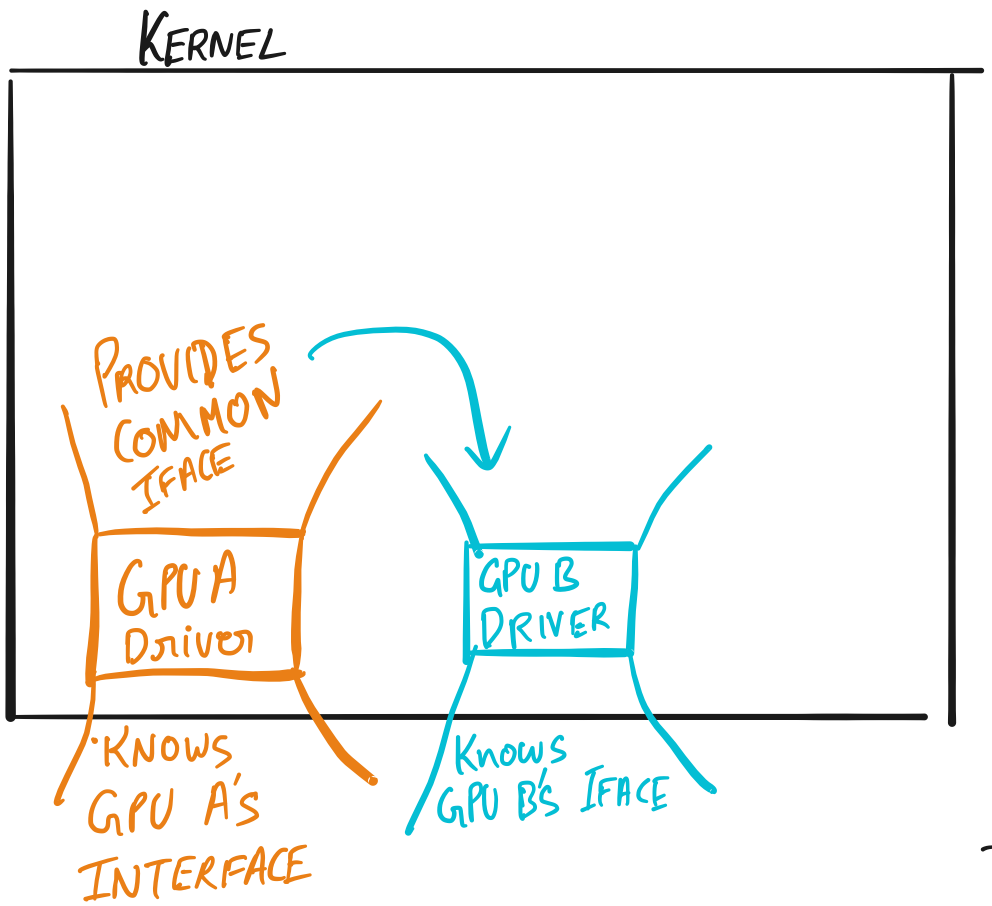
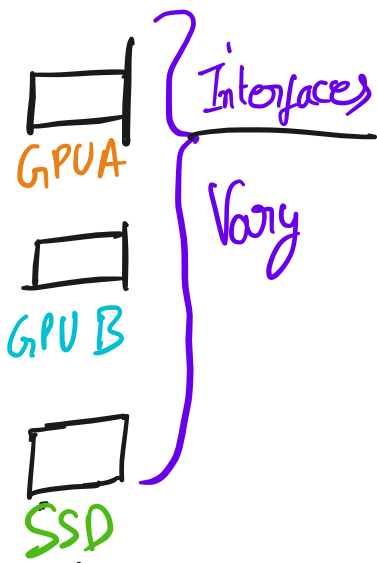
Where we were

- Transfer data to/from devices
 - o I/O instructions inb/inw/outb/outw
 - o Memory mapped I/O
 - o DMA
- Coordination w/devices
 - Interrupts
 - Don't waste cycles polling
 - High interrupt rates lead to
NO FORWARD PROGRESS
 - Polling
 - Might waste cycles
 - OS decides when to check
↳ FORWARD PROGRESS
- IN PRACTICE

ADAPTIVE



DEVICE DRIVERS



SYNCHRONOUS Vs ASYNCHRONOUS I/O

↳ FROM USERSPACE

So far: read/write/... block the calling thread

... Thread has nothing to

Assumption: Instead has nothing to do while waiting

But, what if thread could do other things
(e.g., switch to & run different usermode thread)

Most operating systems offer non-blocking APIs
Not standard

Do not need to know

- Linux: epoll, io-uring
- OS X: kqueue
- Windows: Completion ports

Similarity

↳ Poll for completion

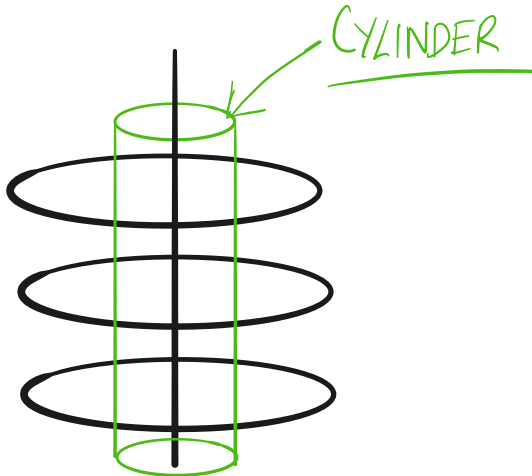
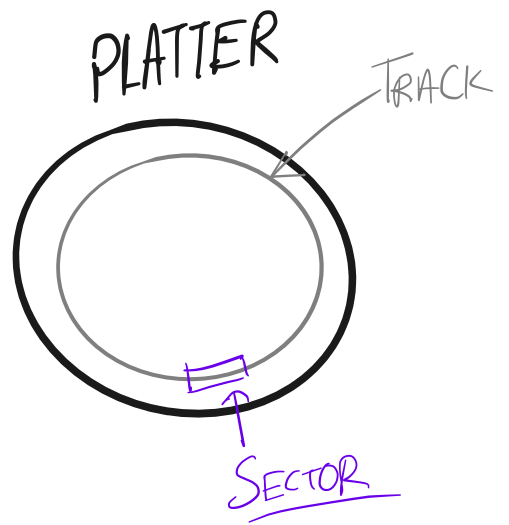

→ Mechanism to wait/block if necessary

DISKS: SPINNING DISKS

- Why?
- Still widely used
 - Dictate the design of most/many file systems



Platter
(1-8)



Interface: Linear array of sectors (generally 512 bits)



READ/WRITE 4K starting from Read(25)


How:

DISK CONTROLLER
(Implemented by
disk hardware)

translates sector →

< platter, track, sector >

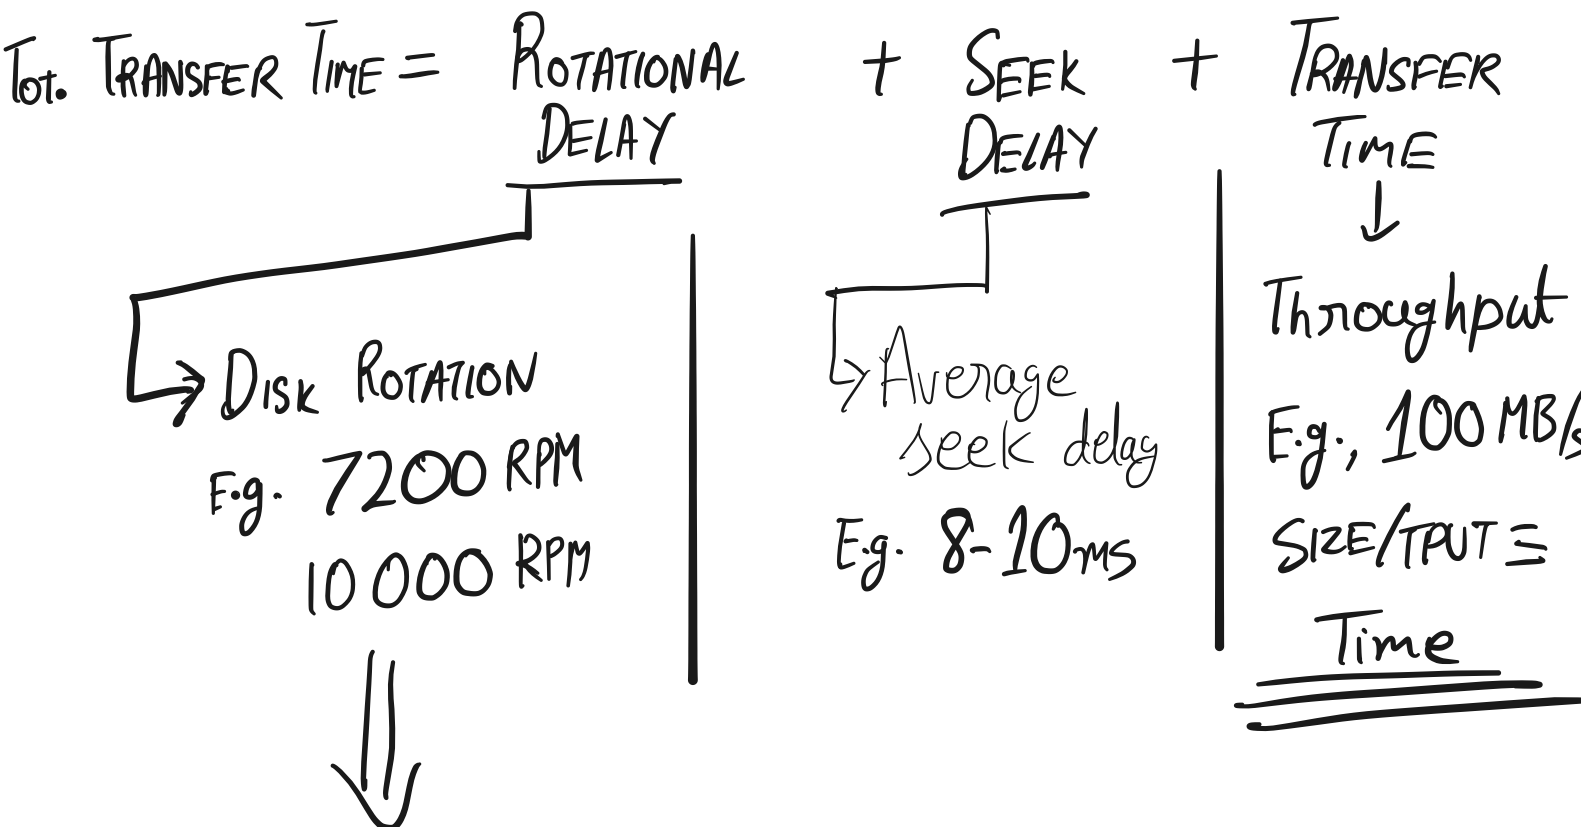

Read(25)



① SEEK DELAY

② ROTATION DELAY

③ TRANSFER TIME



Avg. Rotation time from RPM
7200 RPM \Rightarrow 120 RPS

\Rightarrow 1 Rotation every $\frac{1}{120}$ s \approx 8.3ms per rotation

\Rightarrow Avg Rotation time \approx 4.15ms



10000 RPM $\Rightarrow \sim 167$ RPS

$\Rightarrow \sim 6$ ms per rotation

\Rightarrow Avg Rotation time = 3ms

Let us use this:

Disk: Rotation 12,000 RPM

Avg. Seek Time 12ms

Transfer rate 128 MB/s

Assume 512 Byte reads \rightarrow 512 B sectors

(a) Throughput (bytes/second) to read 500 sectors spread randomly across disk, served in FIFO order?

Read 27, 12, 100, 17, ...



~~500 * 512 B~~
500 * T

Might need to seek for each sector

$T_{\text{put}} = \text{Data Read} / \text{Time taken}$.

Data read = 500 * 512 B

Time taken for 1 read (on avg)

$$\text{Seek time} = 12 \text{ ms}$$

$$\text{Rotation time} = 2.5 \text{ ms}$$

$$12000 \text{ RPM} = 200 \text{ RPS}$$

$$\Rightarrow 1 \text{ rotation every } \frac{1}{200} \text{ s} = 5 \text{ ms}$$

$$\Rightarrow \text{Avg Rot time} = 2.5 \text{ ms}$$

$$\begin{aligned} 1 \text{ MB} &= \\ 2^{10} \text{ KB} &= \\ &= 2^{10} \cdot 2^{10} \text{ B} \end{aligned}$$

$$\begin{aligned} T_{\text{TRANSFER}} &= \frac{512 \text{ B}}{128 \text{ MB/s}} = \frac{2^9}{2^7 \cdot 2^{10} \cdot 2^{10}} \text{ s} \\ &= \frac{1}{2^{18}} \text{ s} \approx 4 \mu\text{s} \end{aligned}$$

$$\begin{aligned} \text{Time for 1} &= 12 \text{ ms} + 2.5 \text{ ms} + 0.004 \text{ ms} \\ \text{TRANSFER} &\approx 14.5 \text{ ms} \end{aligned}$$

$$T_{\text{PUT}} = \frac{500 \times 512 \text{ B}}{500 \times 14.5 \text{ ms}}$$

$$= \frac{512 \text{ B}}{14.5 \text{ ms}} = \frac{512000 \text{ B}}{14.5 \text{ s}}$$

$$\approx 34.5 \text{ KB/s}$$

⑥ Throughput for 500 sequential sectors?

$$\frac{5, 6, 7, \dots \quad 50 \text{ s}}$$

What changes :- seek once to track, wait once for head to arrive at 5.

$$\text{Seek time} = 12 \text{ ms}$$

$$\text{Rotation time} = 2.5 \text{ ms}$$

$$\text{Transfer time} = \frac{500 \cdot 2^9 \text{ B}}{2^7 \cdot 2^{20} \text{ B/s}} = \frac{500 \text{ s}}{2^{18}}$$

$$\approx \frac{2^9}{2^{18}} \text{ s} = \frac{1}{2^9} \text{ s} = \frac{1}{512} \text{ s}$$

$$\approx 2 \text{ ms}$$

$$T_{\text{put}} = \frac{500 \cdot 2^9 \text{ B}}{14.5 \text{ ms}} = \frac{500 \cdot 10^3 \cdot 2^9 \text{ B/s}}{14.5}$$


~ 18 MB/s

OBSERVATION: SEQUENTIAL ACCESS IS MUCH FASTER

↳ CONSTRAINT FOR A LOT OF
FILE SYSTEM DESIGN.



write (



Read

write-