

Today (4/3) April (snow)

3 Parts:

~~Virtualization~~

~~Concurrency~~

Persistence (Input/Output)

Focus: Disk, SSDs +
software in OS
(file system)

Today:

→ Devices

→ Hard Drives (Hard Disk Drives)

HDDs

→ I/O scheduling (Cost)



Break:

→ Project

→ Midterm

Devices :

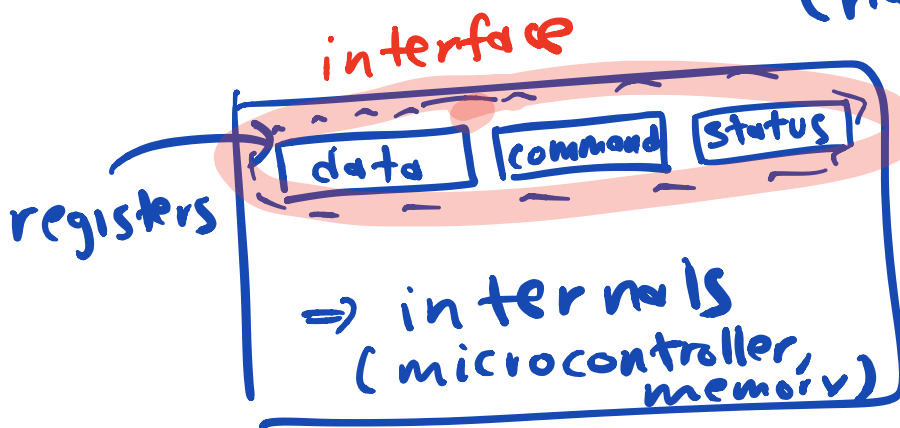
- => many types
- => hard drive
- => SSD
- => other storage
- => keyboard/mouse
- => graphics
- => sound
- => RGB light ⇐ FOCUS
- => Network





Example: (Intel Z270 chipset)
hot hardware.com

Device interface: device controller



(hardware that presents interface + implements device functions)

protocol:

```

while (status == BUSY)
{
    write ; // spin (wait)
    write data => data register
    write cmd => cmd register
}
while (status == BUSY) // device begins

```

`;/spin (wait)` operation for cmd

Questions:

Inefficiencies

⇒ spin waiting

⇒ data movement

How to do I/O?



⇒ how to read/write registers on device?

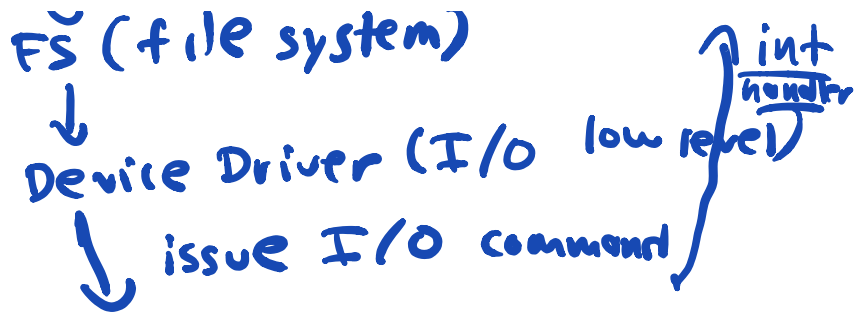
ignoring: error handling

Spin waiting: Avoid CPU waste (esp. for slow devices)
(polling)

⇒ Interrupts

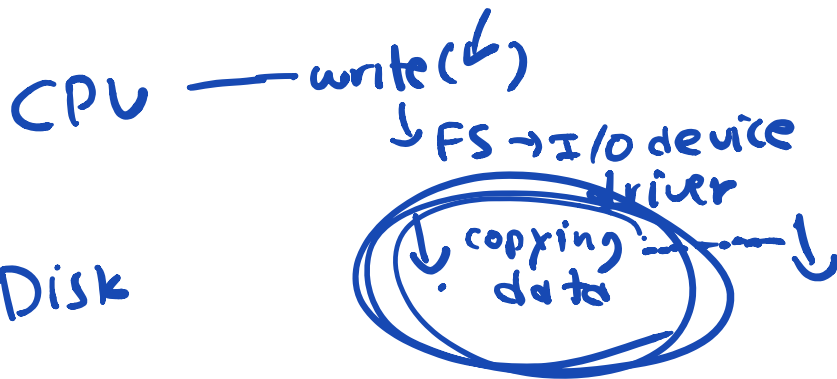
overlap






Disk

time



⇒ reduce CPU burden
w/ more specialized
hardware:

Direct Memory Access
(DMA)



How to interact w/ device?

{ \Rightarrow I/O instructions (x86) }
 { \Rightarrow memory mapped I/O in/out }

Break:


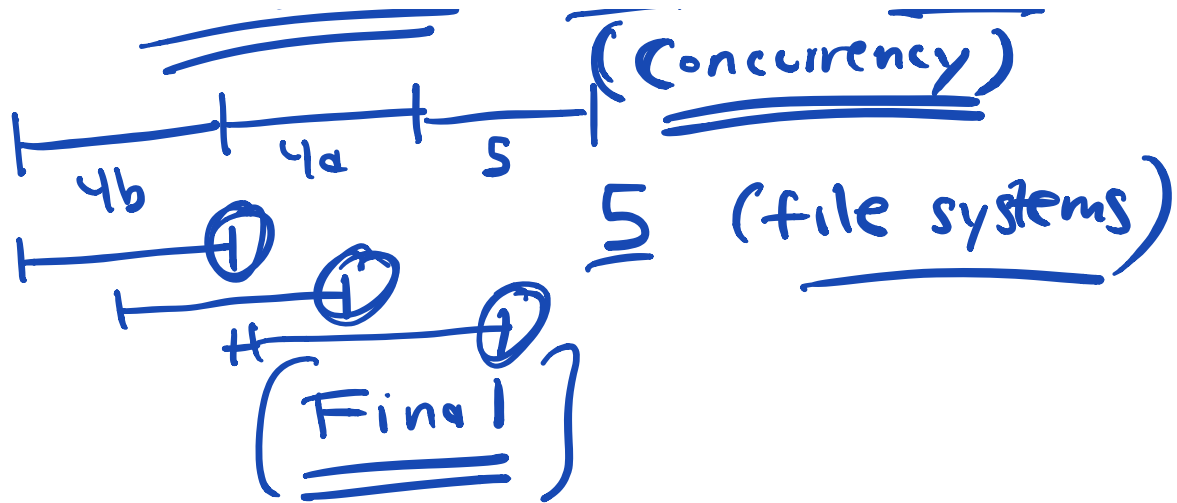
Midterm:

info : coming
 (answer key + your answers)

Projects:

4a

4b



Hard Drives : { Interface
Internals }
 => Persistent Storage Device
 "CS" perspective:
 => Controller ← interface
 ↓
control operation

⇒ Mechanics:
Platters, arm, etc.

Interface: Hard Drive

Read sector, len ↗ return data

Write sector, len, data → return error code

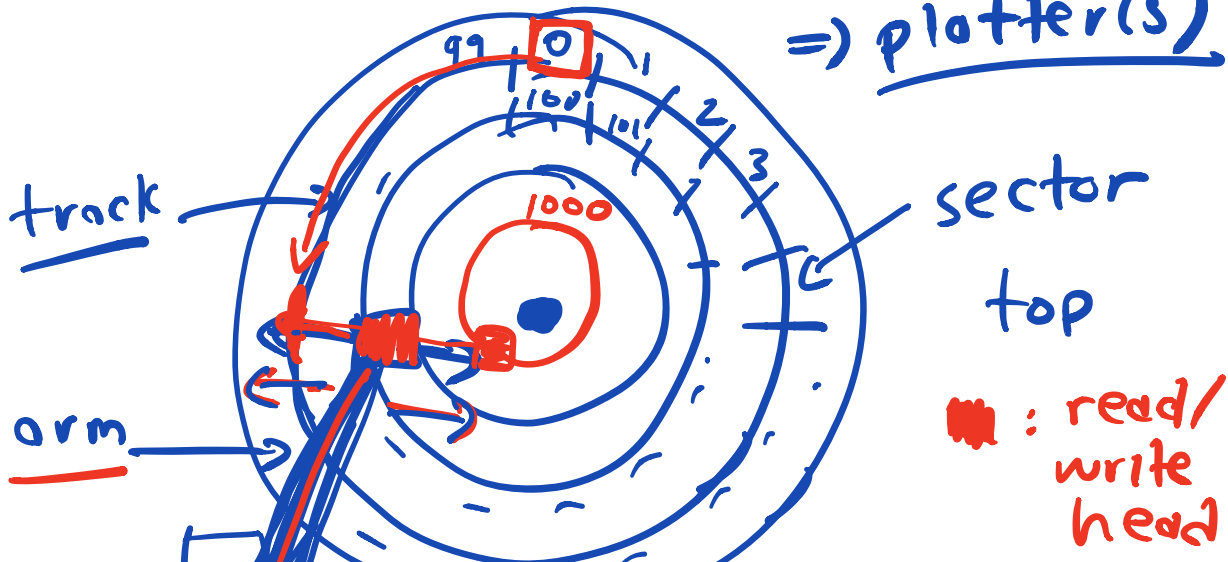


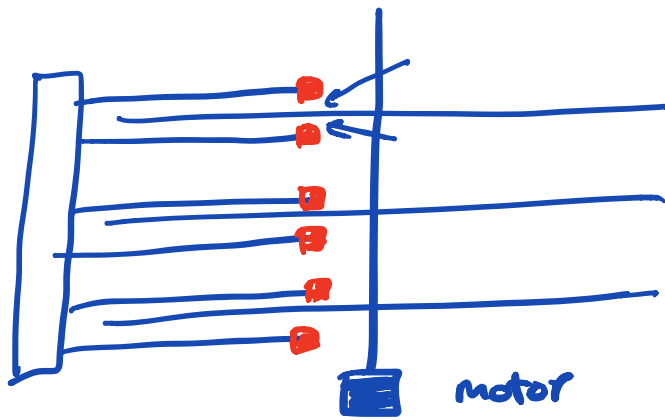

sectors (512 bytes)

Internals:

mapping: interface ⇒
platter, surface,
track

⇒ platter(s)





Read/Write:

Seek:

positioning read/write head over correct track

over head

next

rotation
(rot. delay):

waiting for
desired sector
to rotate under
head

transfer:

do read/write

seek rotate transfer

$$T_{\pm IO} = T_{\text{seek}} + T_{\text{rotate}} + T_{\text{transfer}}$$

large I/Os vs. small I/Os

transfer: 100 MB/s
seek avg: 7 ms
rot avg: 3 ms

1 MB I/Os vs. 1 KB I/Os


Perf? seek: 7ms } 10ms }
 rot: 3ms }

transfer: 10 ms → transfer: ~0
rate

$$1 \text{ MB} \left(\frac{1 \text{ sec}}{100 \text{ ms}} \right) \left(\frac{1000 \text{ ms}}{1 \text{ sec}} \right) =$$

transfer
(1 MB)

(10 ms)




50%

$$\frac{1 \text{ MB}}{20 \text{ ms}} \Rightarrow \text{MB/s}$$

$$\left(\frac{1 \text{ MB}}{20 \text{ ms}} \right) \left(\frac{1000 \text{ ms}}{1 \text{ sec}} \right) = 50 \text{ MB/s}$$

$$\uparrow \frac{1 \text{ KB}}{10 \text{ ms}} \left(\frac{1000 \text{ ms}}{1 \text{ sec}} \right) \Rightarrow \underline{100 \text{ KB/s}}$$



Break: 3 mins

Scheduling (I/O)


⇒ Given set of requests:
have to decide which
to do first

Goals:

Fairness

Performance

queue




"Job": I/O request

desire: shortest job first
need to know: how long
it runs for

here: (kind of) possible

OS: but hard drive internals
hidden

inside drive: have much more
knowledge





Sched:

FIFO: X

Shortest Seek Time

First (SSTF)

→ (Starvation)

→ (no acct'ing of rotate)


→ Shortest Access Time

First (SATF)

(rot + seek)


→ Solving starvation:

"SCAN" (elevator)



→ Bounded Queue





Is SATF best?
(optimal)

