Today [we are starting! (volatile, not persistent)]

=> Virtualization ✓

=> Concurrency ✓

=> [Persistence]

=> power off
=> info lost
instructions
data

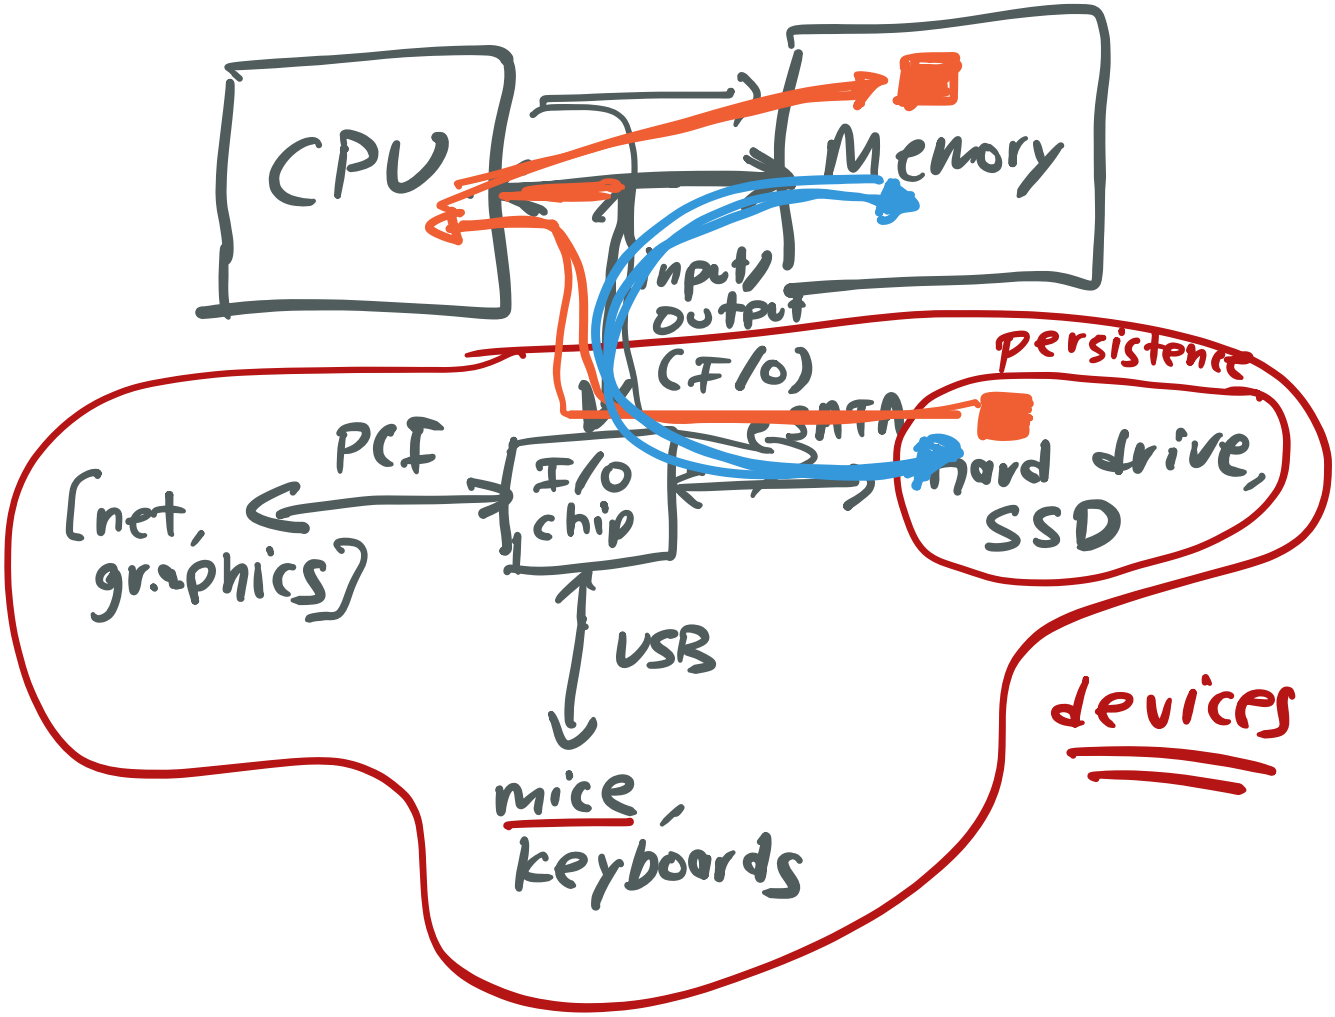

Today:

→ Device interaction (I/O)
  → how?
  → efficient?

Specific Device: Hard Drive

Future:

Device up => [OS : File Systems]

Generic Device:

hardware interface (generalization)

how does OS read/write?

- → command register (what to do)
- → data register
- → status register

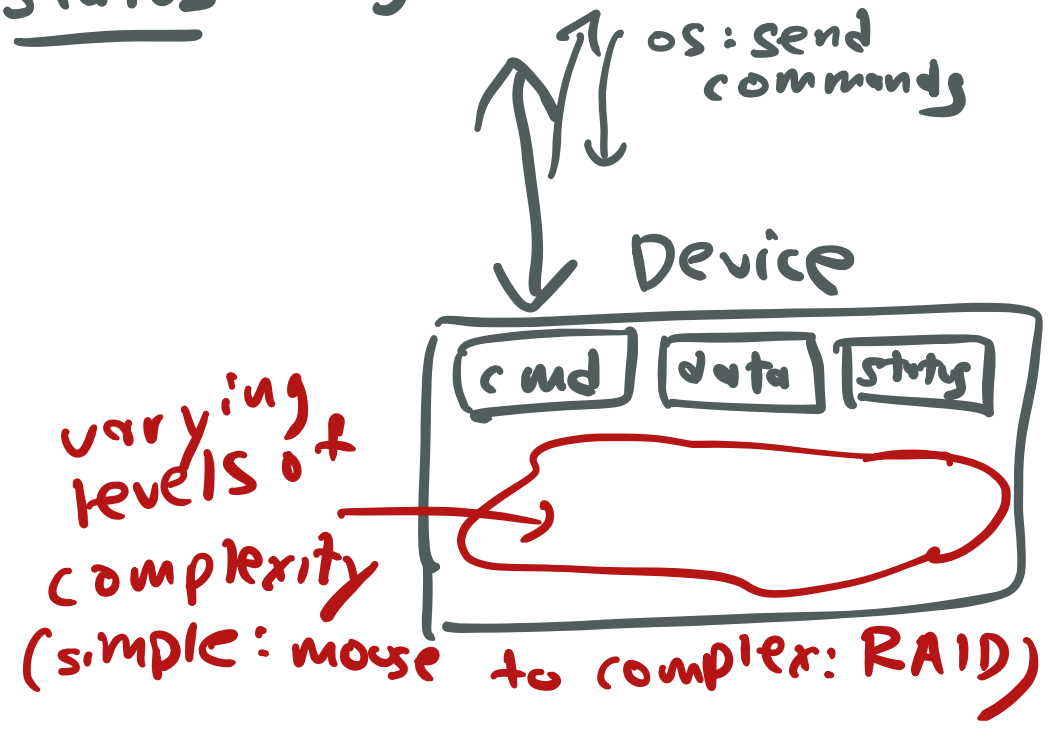

2 methods of access:

- => special I/O instructions
  OR
- => memory-mapped I/O

Protocol for access:

1) check if device is available
```
while (status == BUSY)
    ; // spin
```

2) move data
```
for (size of data)
    move data → data register
```

3) move cmd: write command → cmd register

4) wait:
```
while (status == BUSY)
    ; // spin
```

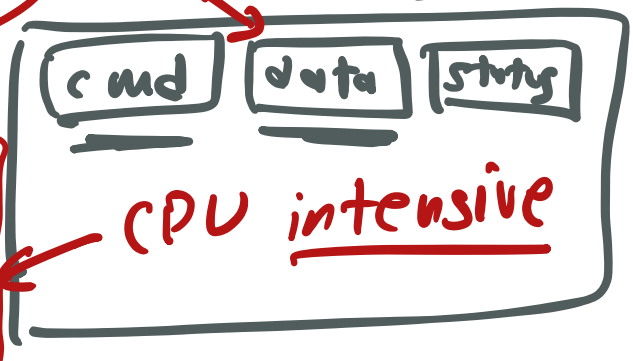

Problems:

#1 → lot of spinning (polling the device)

#2 → data movement: CPU intensive

Solutions:

#1: instead of spinning, go to sleep (block)

```
while (status == BUSY)
    sem_wait (& io);
```

what makes it up?

=> interrupt : signal to system that device is done

run interrupt handler

=> sem_post(&io);

#2: DMA engine : hardware

[OS can program the h/w to move data directly between memory/device]

old:

CPU: spinning | data transfer | spinning

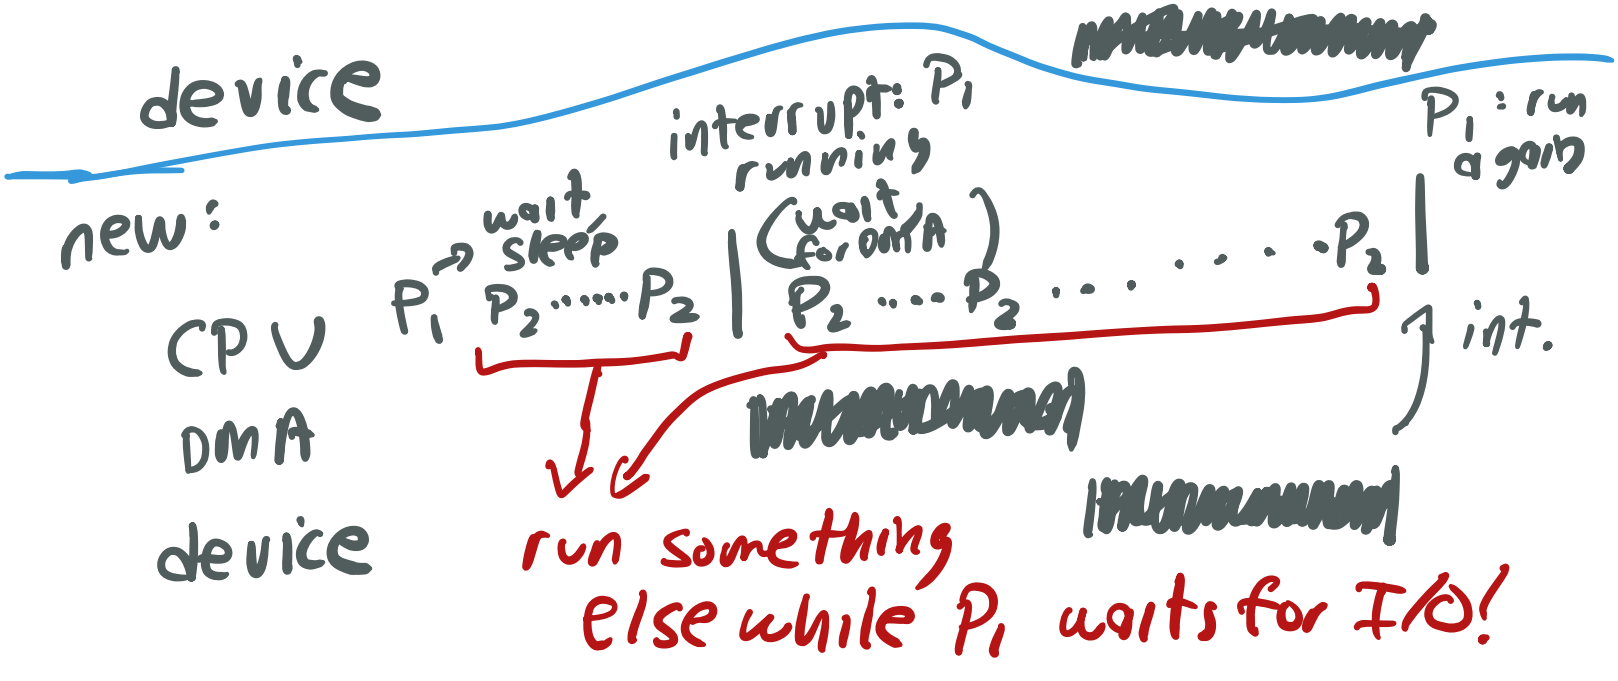

Hard Drive: How?

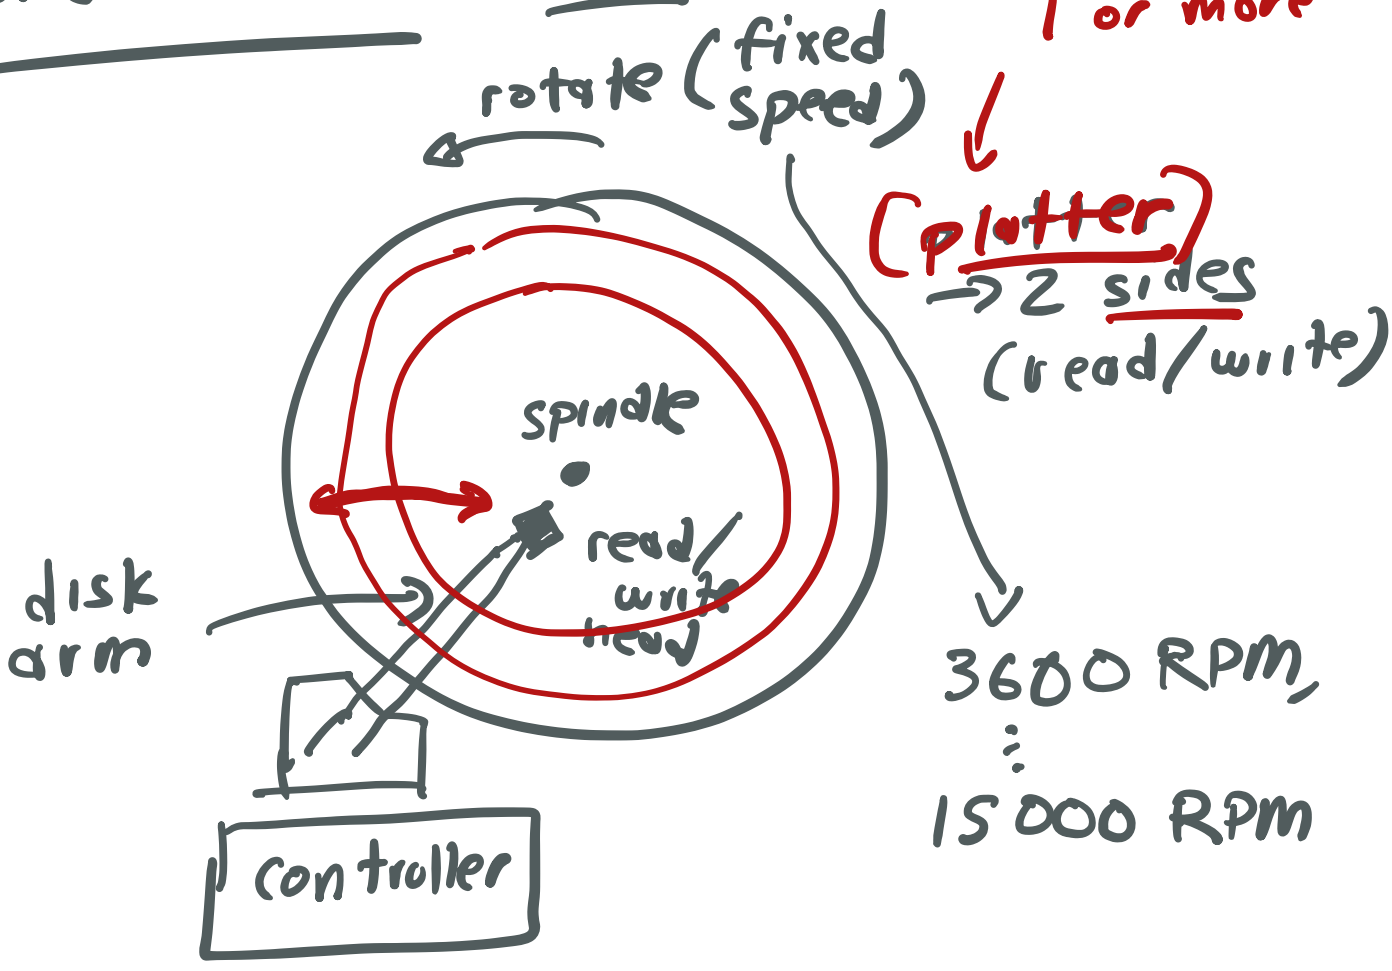

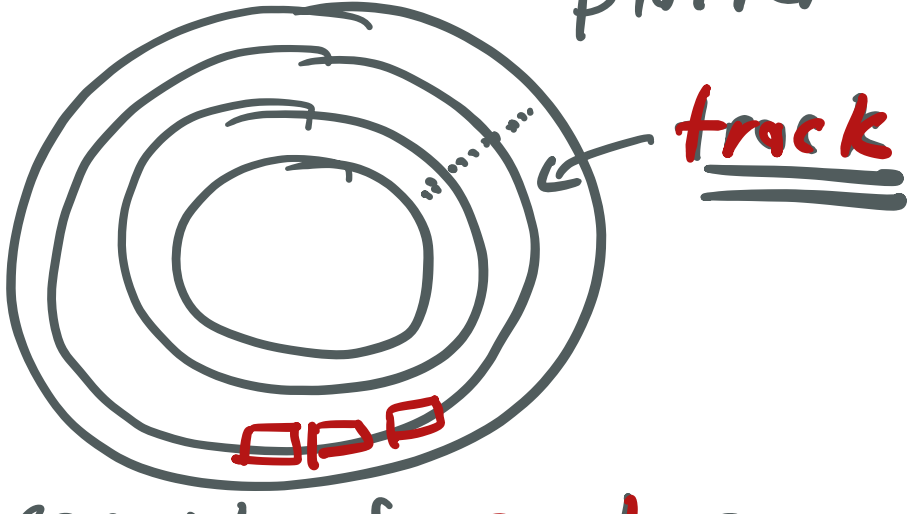

track: consists of **sectors** (512 byte)

to **read** or **write**:

e.g. read a sector

how many sectors

address, (size)

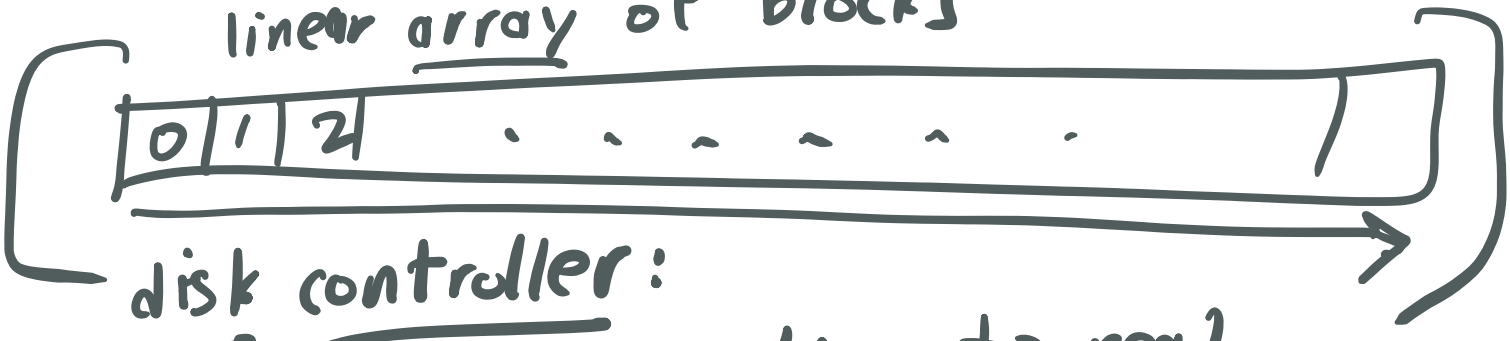

disk controller:
=) internal action to read or write

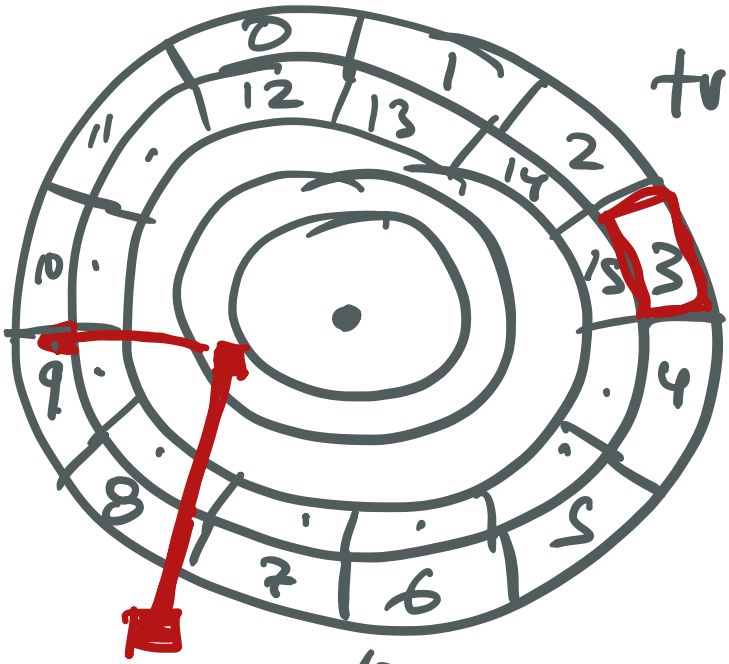

track (many sectors)

to do an I/O:

=) seek : moves disk arm to correct track

→ wait (rotation) : wait for sector to rotate under head

→ transfer : read/write

Example:

avg seek: ~ 7ms] 10ms

avg rotate: ~ 3 ms }

transfer: ~ 0 ms

~ 200 MB/s

512 b (½ K)

½ KB

[random I/Os: dominated by seek/rotation]

throughput: $\frac{512\ b}{10\ ms} \cdot \frac{1000\ ms}{seconds} = \frac{½\ MB}{s}$

random I/O is slow,

[sequential I/O (or large I/Os) are "fast"]

disk scheduling: