# CS 537 - Week 10

1. Zemaphores
2. Monitors

---

3. I/O Devices
4. Disks
5. Disk Scheduling

---

## Monitors

→ only one thread can be active in the monitor.

C ... (sleep) ..... empty buffer.

P ..... (full)!

C ..... consume ... sleep

$t_1$: allocate (100) → sleep

$t_2$: allocate (30) → sleep.

---

$t_3$: free (50)

# Persistence

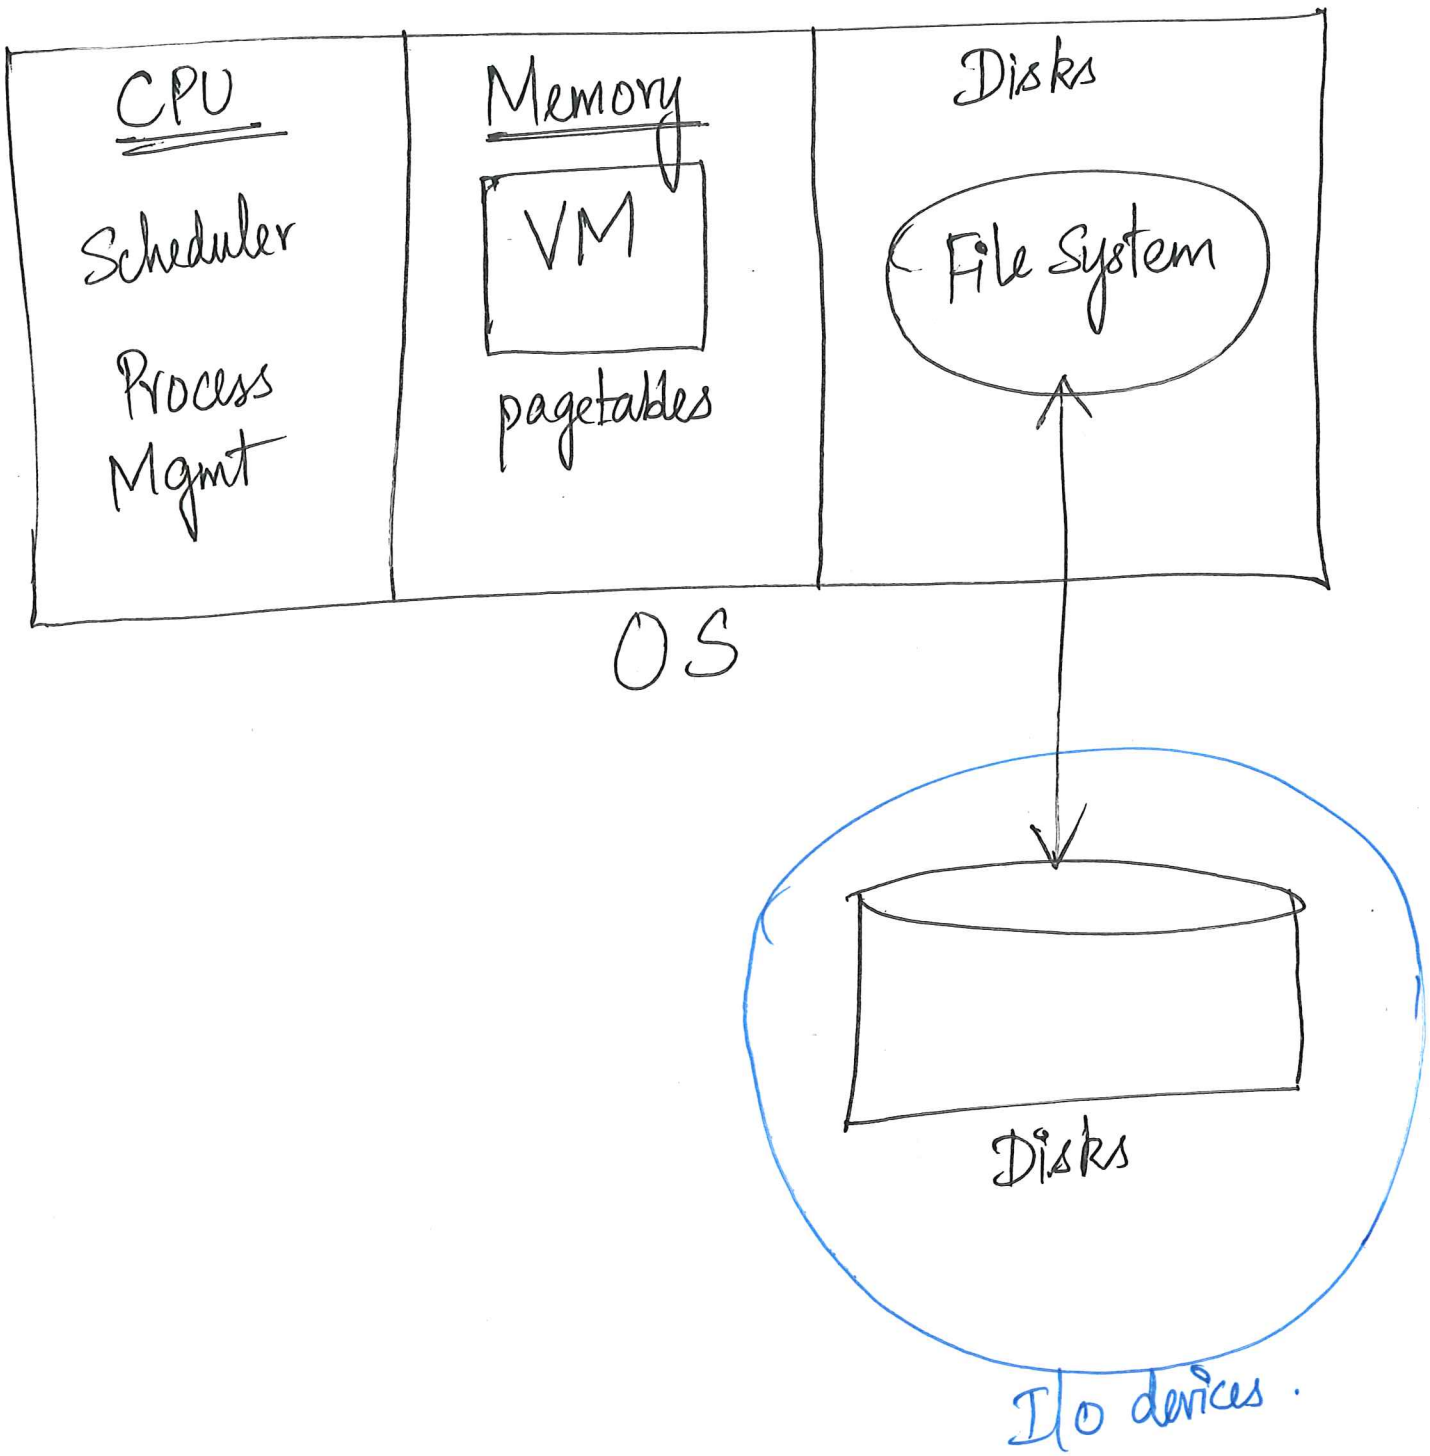

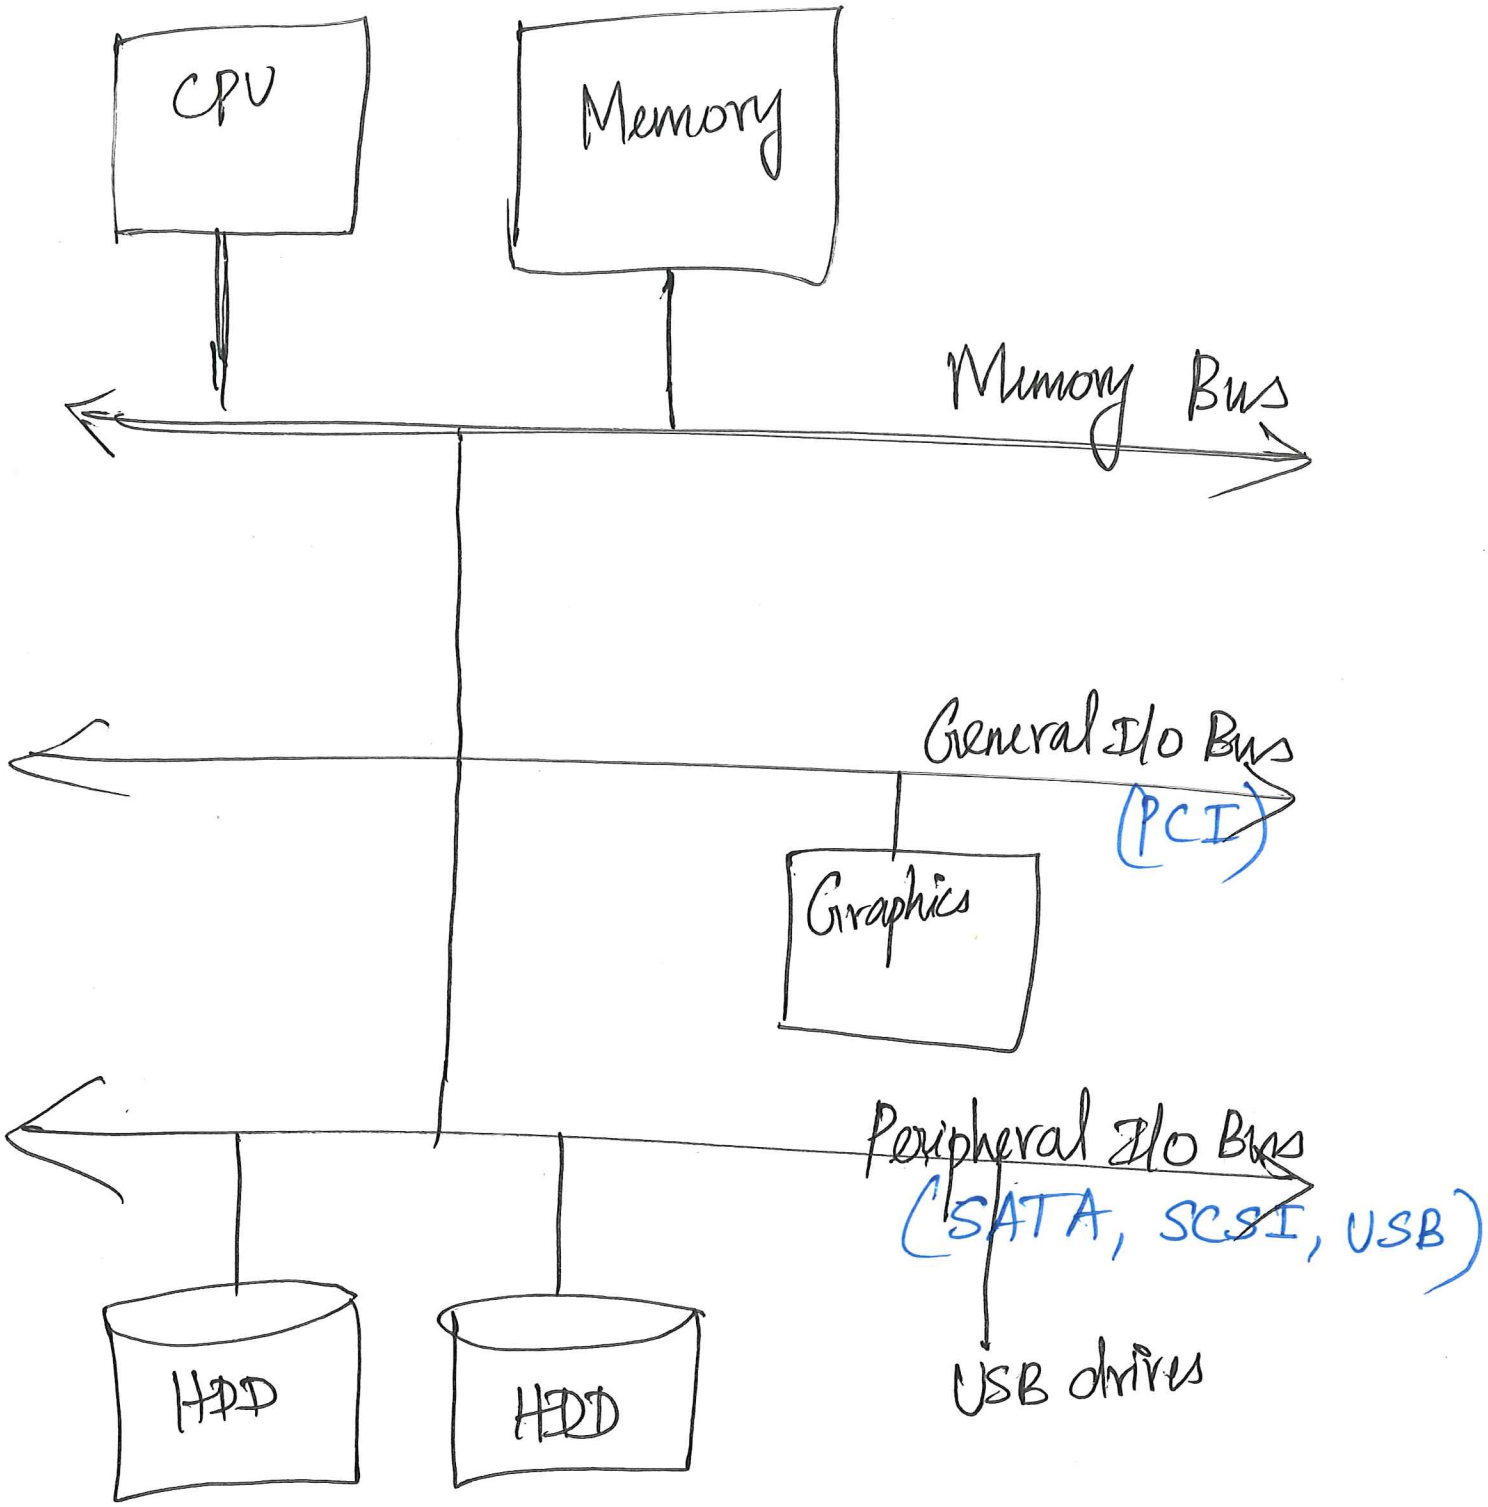
CPU
Memory
Memory Bus
General I/O Bus
(PCI)
Graphics
Peripheral I/O Bus
(SATA, SCSI, USB)
HDD
HDD
USB drives

# A Simple Device

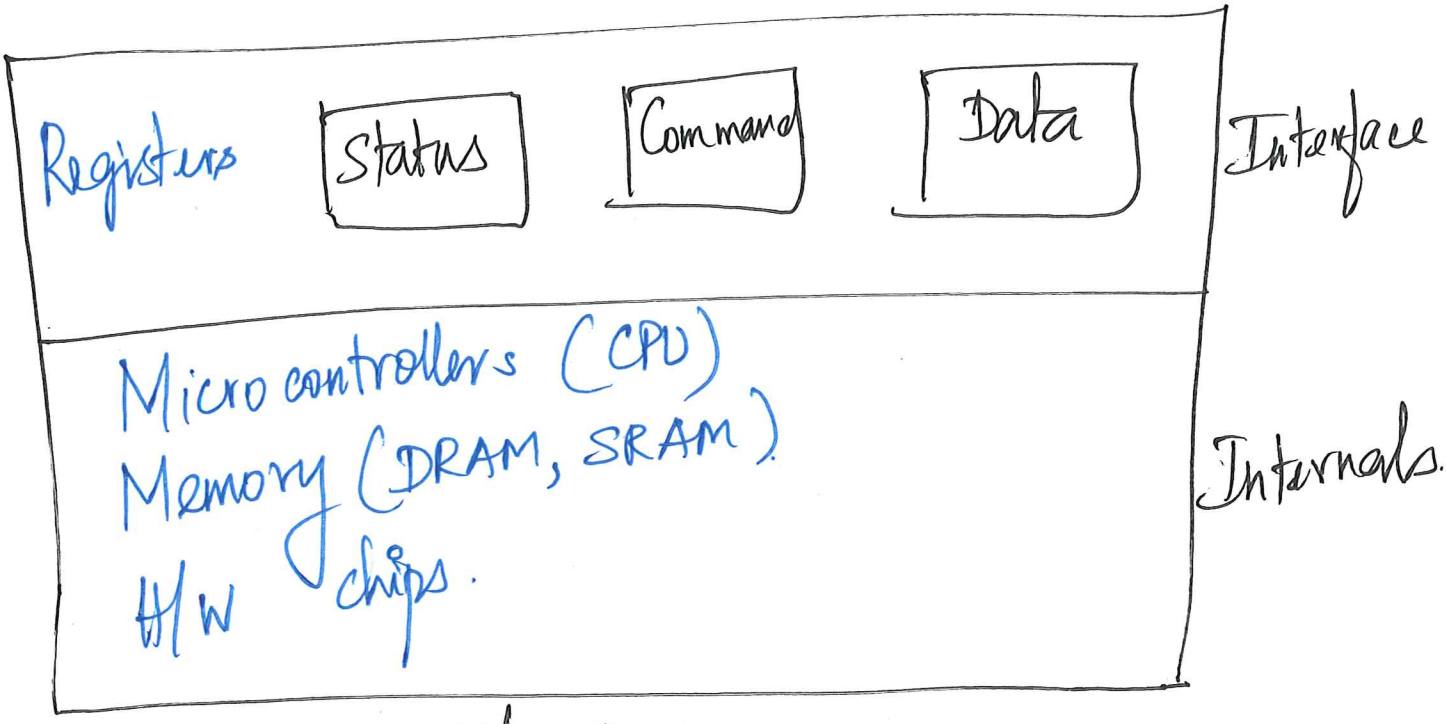


I/O Device

Protocol

```
while ( STATUS == BUSY )
    ;
Write data to DATA register.
Write ~~data~~ command to COMMAND register.
while ( STATUS == BUSY )
    ;
```

Polling

# Polling vs Interrupts


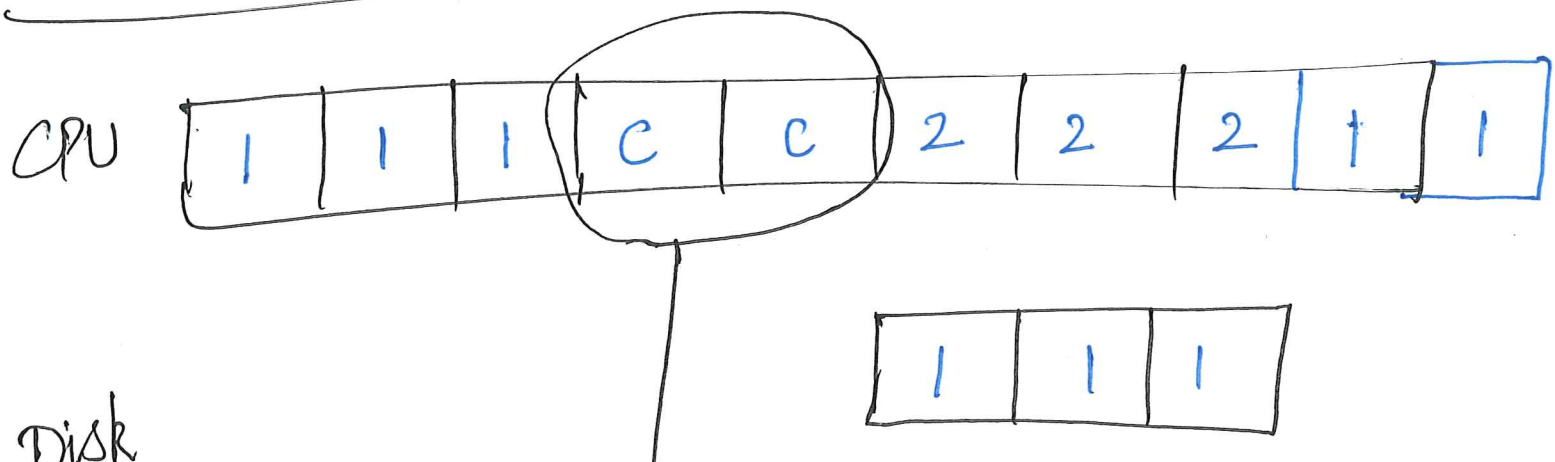


## DMA - Direct Memory Access.

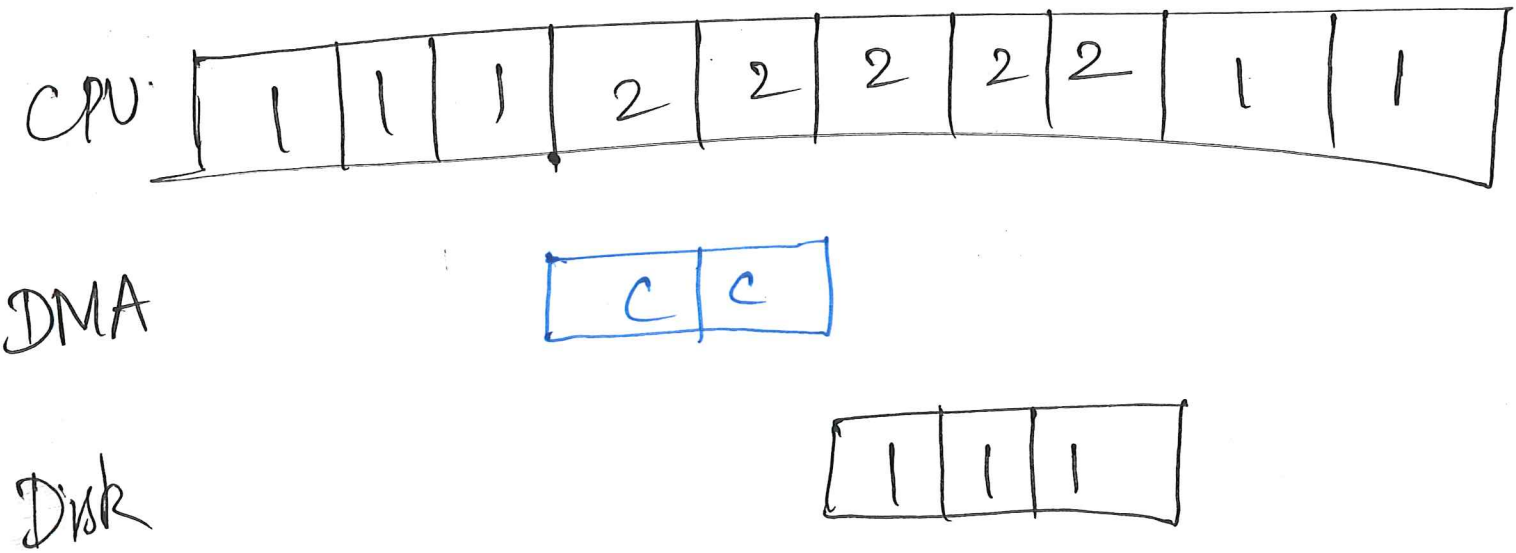

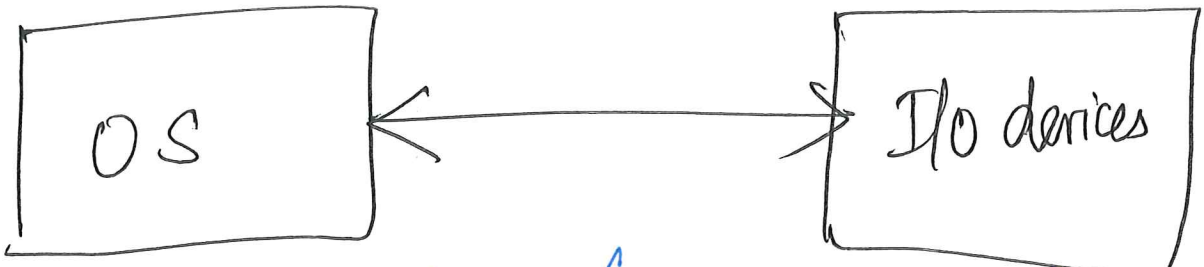


1. Special instruction (x86 in, out)
2. Memory mapped I/O.

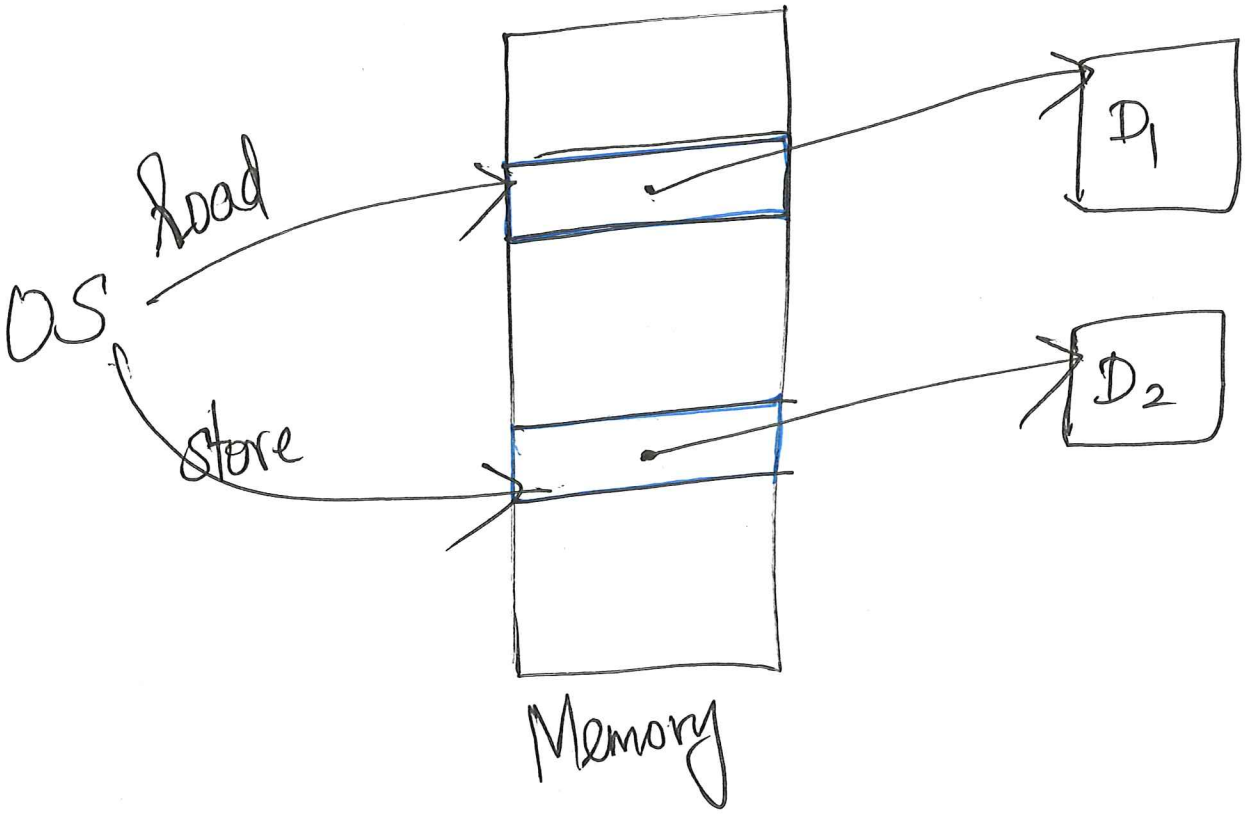

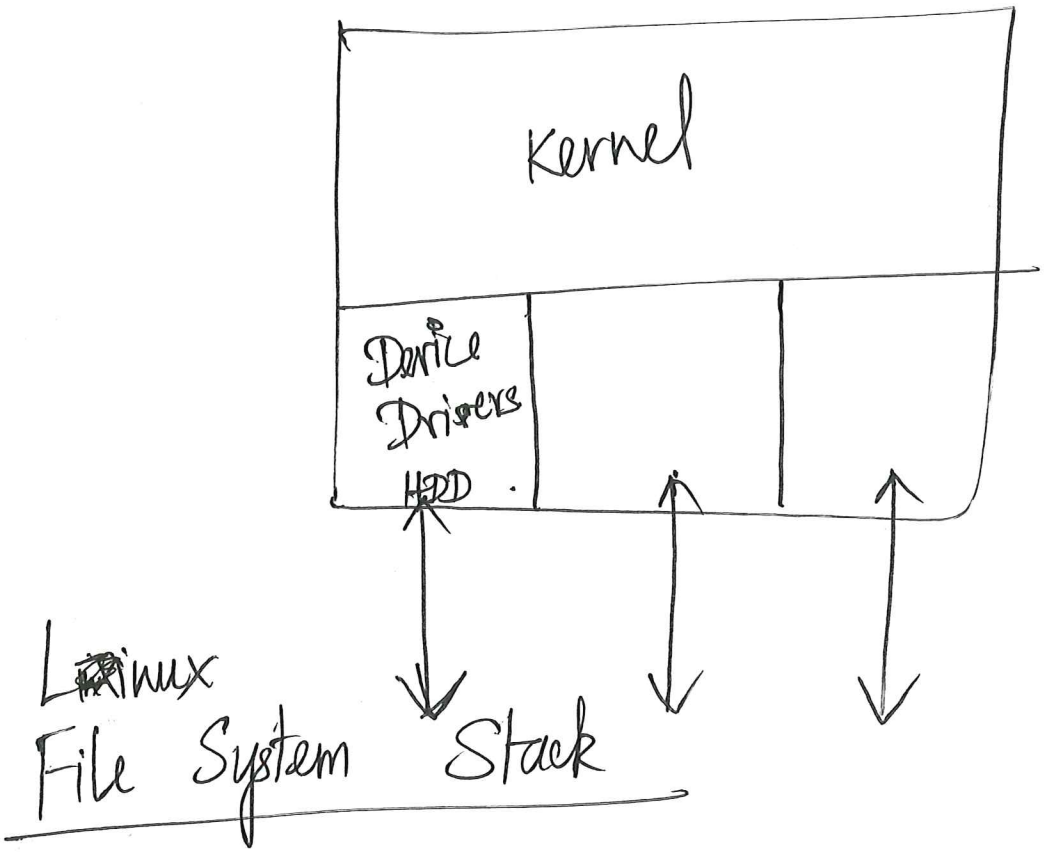
Kernel
OS
Device Drivers
HDD
Linux
File System Stack

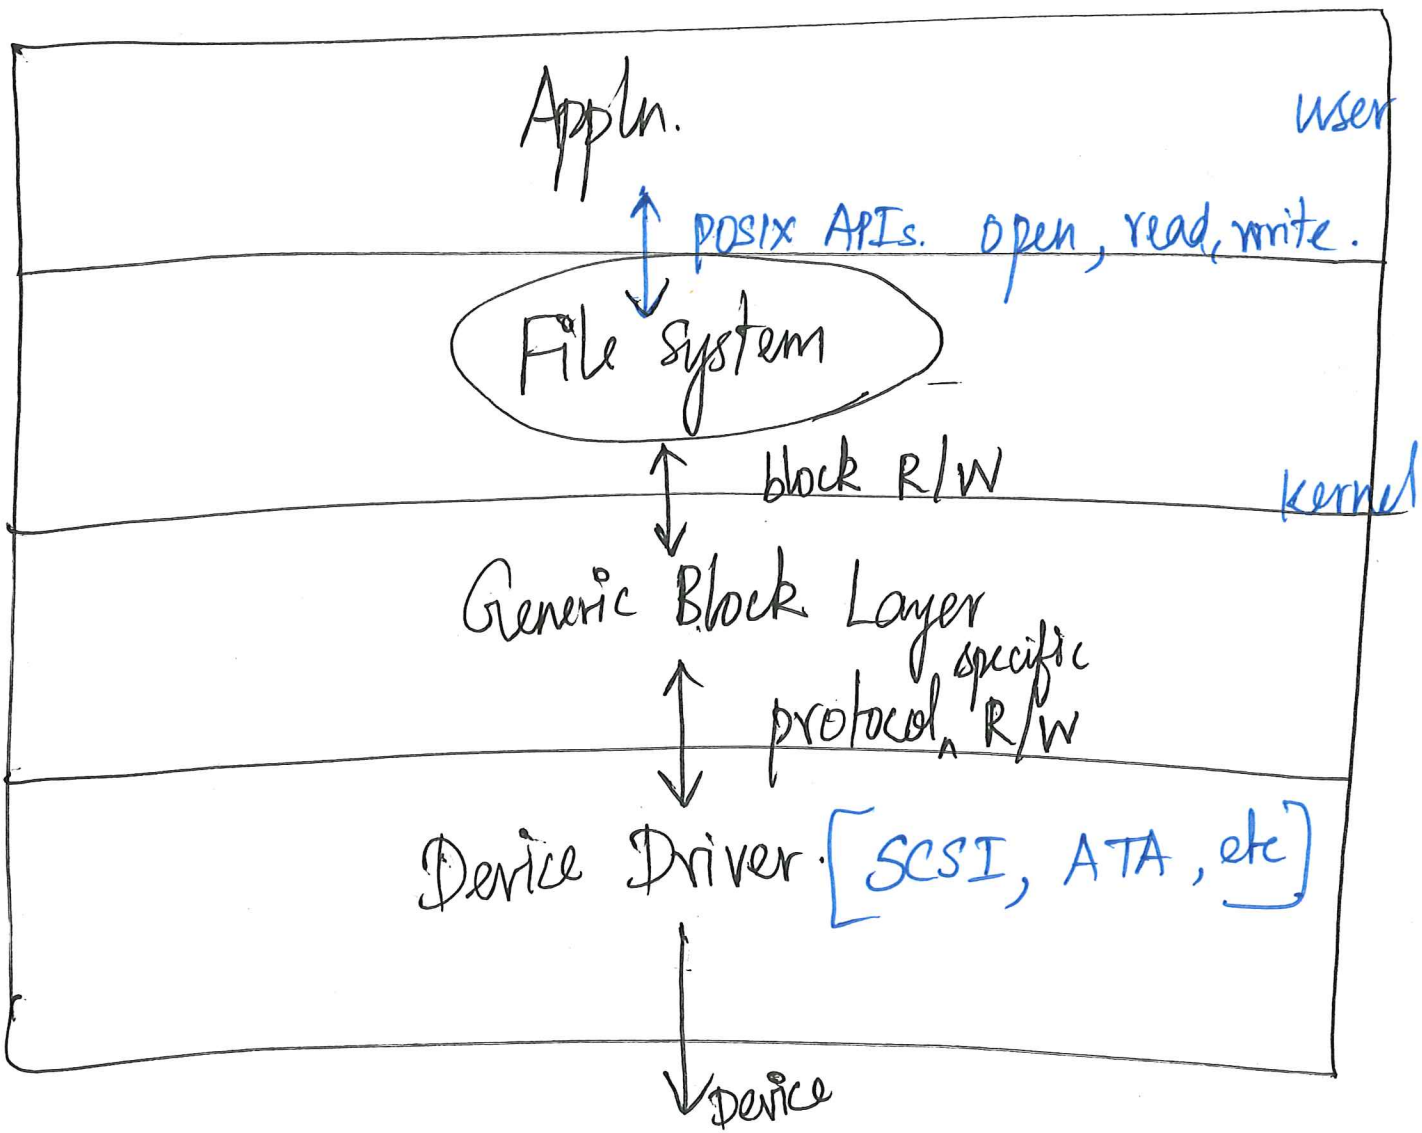
Appln.
user
POSIX APIs. open, read, write.
File system
block R/W
kernel
Generic Block Layer
specific protocol R/W
Device Driver [SCSI, ATA, etc]
Device

# Hark Disk Drives (HDD)

## Interface

sectors

0 — n-1 sectors.

1 sector = 512 bytes

## Geometry

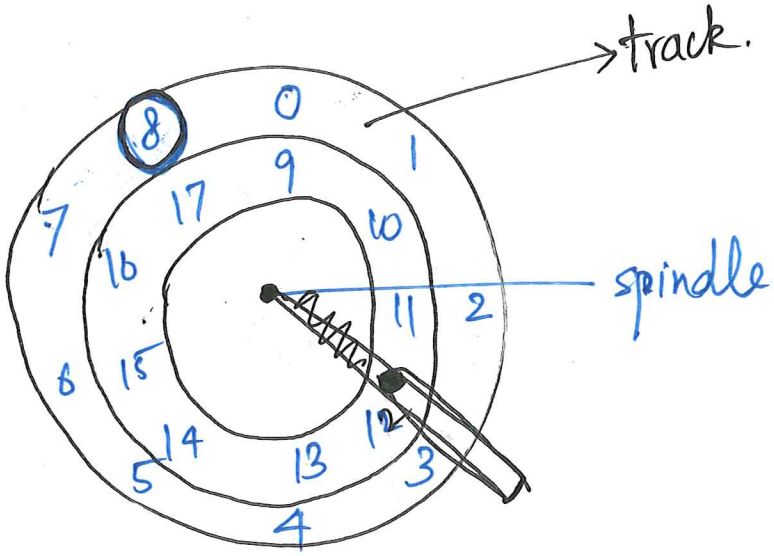


RPM

7200 OR 15000

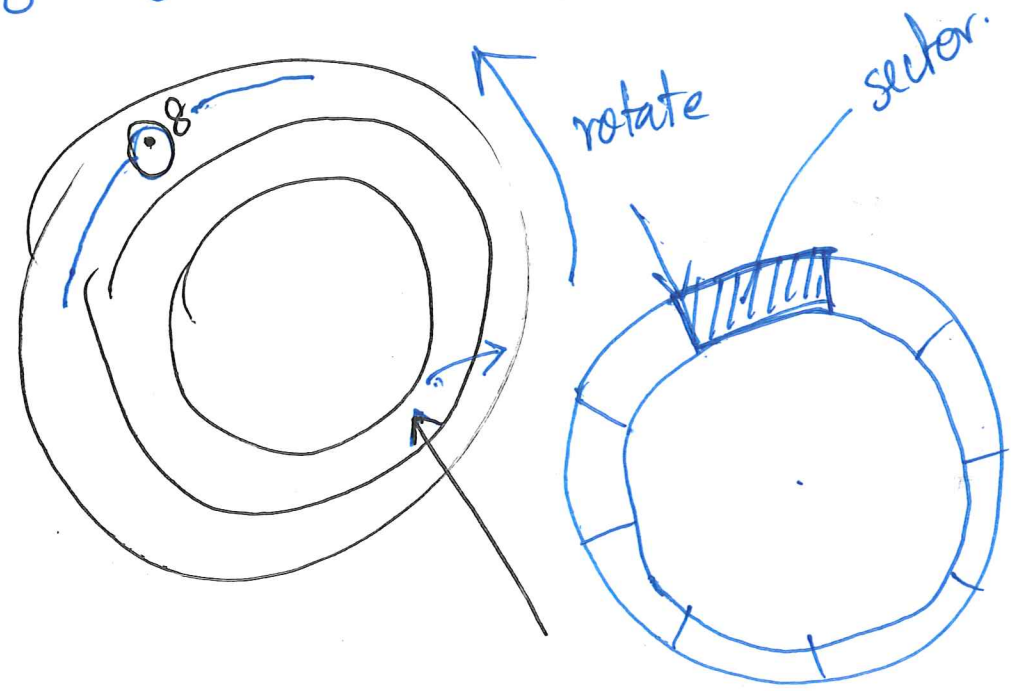


1. seek
2. rotate
3. data transfer.

# Workload

- Sequential Read
- Random Read.

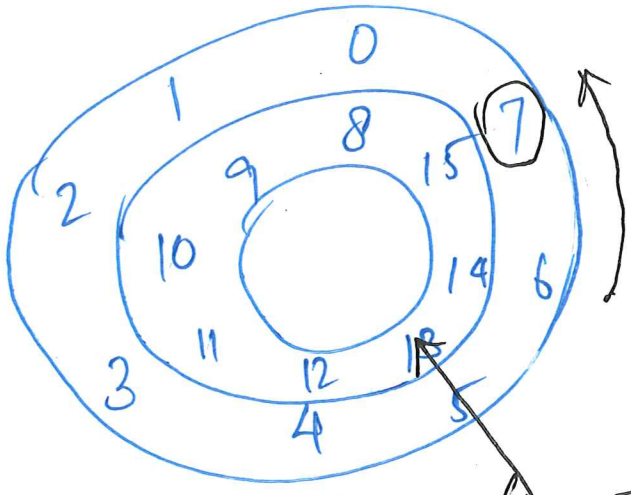


Sectors reads: 15, 0, 3, 9, 6 → Random Read.

4, 5, 6, 7, 8, 9, 10 → Seq. Read.

## Random Workload

- small (4KB) from random locations on disk.

8 sectors

# I/O Time

$$T_{I/O} = T_{seek} + T_{rotation} + T_{data\ transfer}.$$

$$Rate_{I/O} = \frac{Size_{transfer}}{T_{I/O}}$$

| | Cheetah 15K.5 | Barracuda |
|---|---|---|
| Capacity | 300 GB | 1 TB |
| RPM | 15,000 | 7,200 |
| Avg. Seek | 4 ms | 9 ms |
| Max Transfer | 125 MB/s | 105 MB/s. |

# Cheetah 15K.5

$$T_{I/O} = T_{seek} + T_{rotation} + T_{data\ transfer}$$

$T_{seek}$ → 4ms

$T_{rotation}$ (Rotational Delay)

**$T_{rotation}$**

RPM = 15,000

15,000 rots → 1 min (60 secs)

$$1\ rot \rightarrow \frac{\cancel{60}\ 4}{\cancel{15{,}000}}\ s = 4\ ms$$

avg. rot = 2 ms.

**$T_{transfer}$**

125 MB → 1 s

$$4\ KB \rightarrow \frac{4\ KB}{125\ MB}\ s \simeq 30\ \mu s$$

$$T_{I/O} = 4ms + 2ms + 30\mu s$$

$$= 6.03\ ms \simeq 6\ ms$$

$$R_{I/O} = \frac{4\ KB}{6\ ms} = 0.66\ MB/s$$

## Barracuda

$$T_{I/O} = T_{seek} + T_{rot} + T_{transfr.}$$

$$= 9ms + 4ms + 38\,\mu s$$

$$= 13.038ms \simeq 13\,ms$$

$$R_{I/O} = \frac{4KB}{13\,ms} = 0.31\,MB/s$$

| | Cheetah | Barracuda |
|---|---|---|
| Random | 0.66 MB/s | 0.31 MB/s |
| Seq. | 125 MB/s | 105 MB/s |

200

300

# Sequential Workload

100 MB

1 seek + 1 rotation.

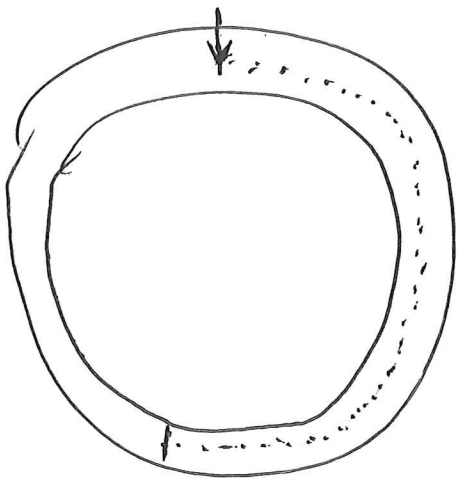

## Cheetah

$R_{I/O}$ = 125 MB/s

125 MB → 1 s

100 MB → $\frac{100}{125}$ s = 800 ms

$T_{I/O}$

## Barracuda

$R_{I/O}$ = 105 MB/s

105 MB → 1 s

100 MB → $\frac{100}{105}$ = 950 ms

$T_{I/O}$