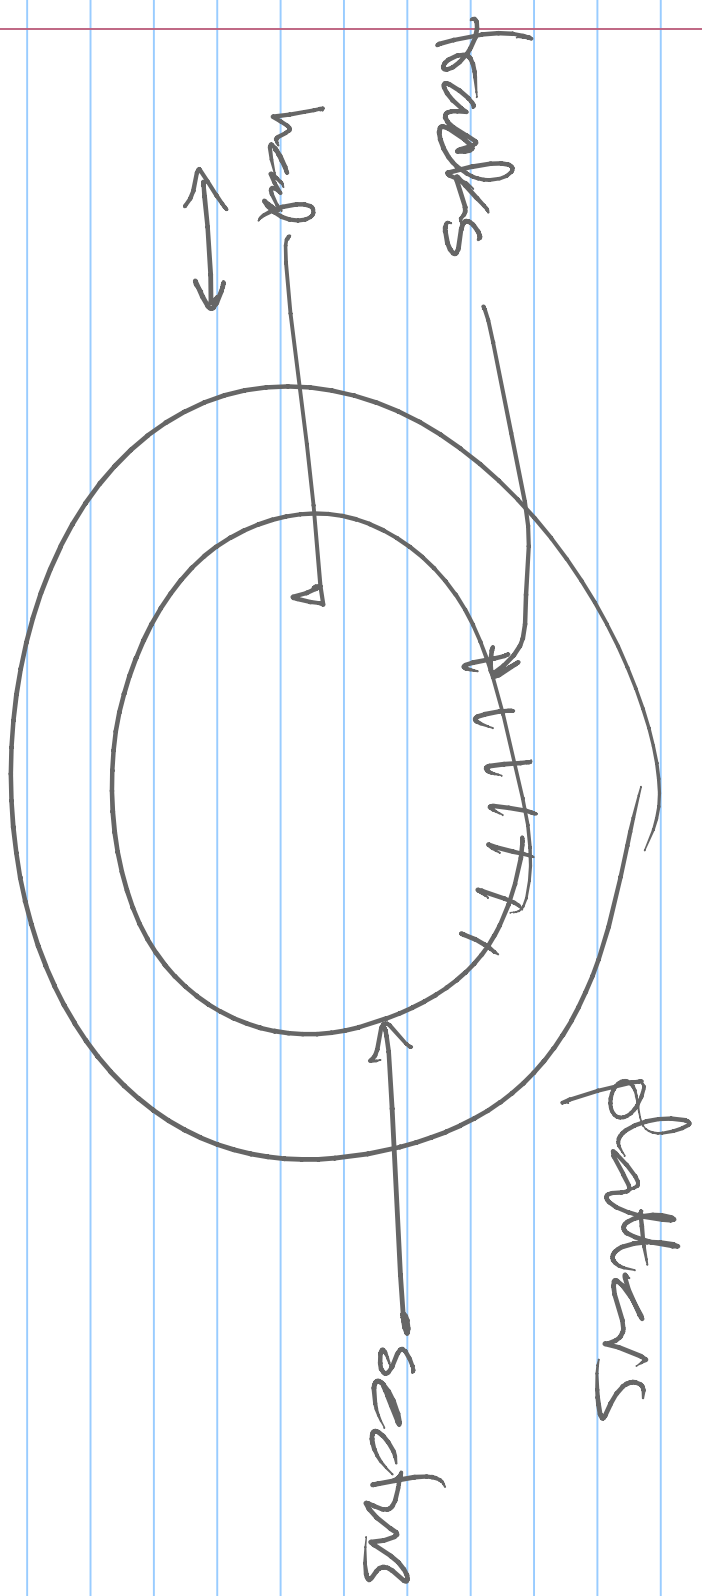


Input Output / Storage devices

↳ hard disks →

↳ SSD → solid state disks

hard drive



lots of sectors -  $N$  per head

lots of heads  $\rightarrow m$  per disk

Total sectors  $m \times N$

Sector contains a block of data  
physically



to write a block of data

① put data in mem

② Tell controller block # & mem address

③ controller does the rest  
(takes time)

④ " done "

SSD → non volatile RAM

→ internally uses RAM/physical addresses

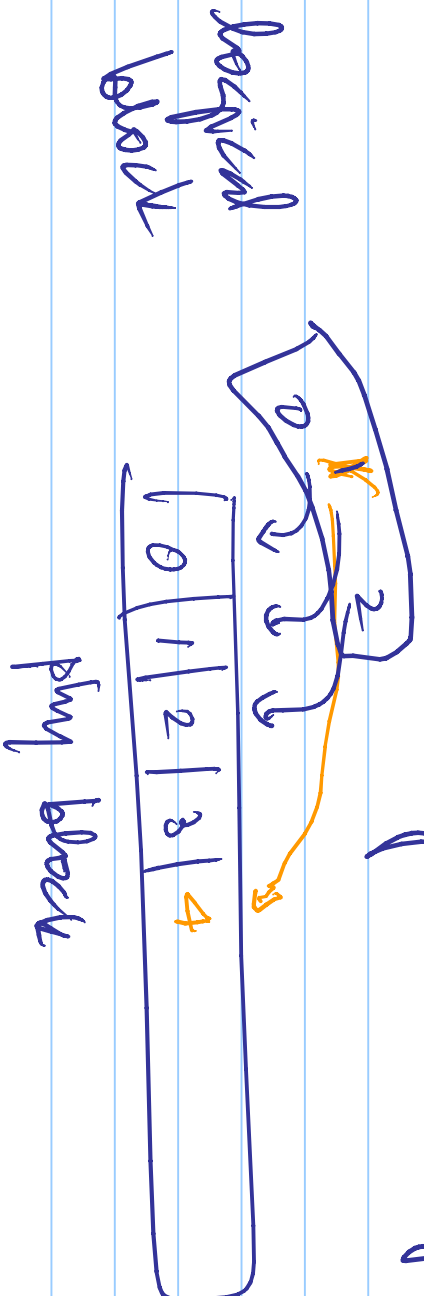
→ controller (controller) makes the RAM look like a hard drive (read/write blocks)

write to SSD

① slow → erase then write

② destructive

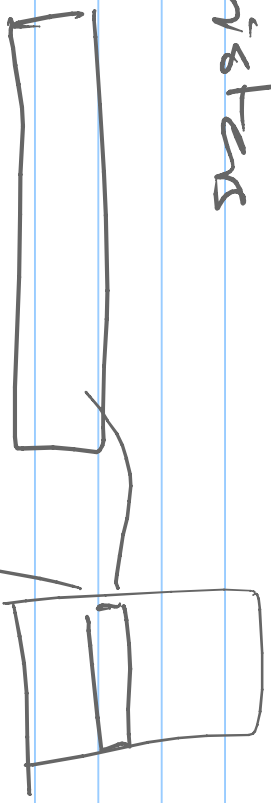
load leveling (done by hardware)



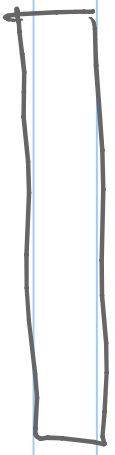
How to do I/O (talk to controllers)

controllers have registers

Status Register  
(& control)

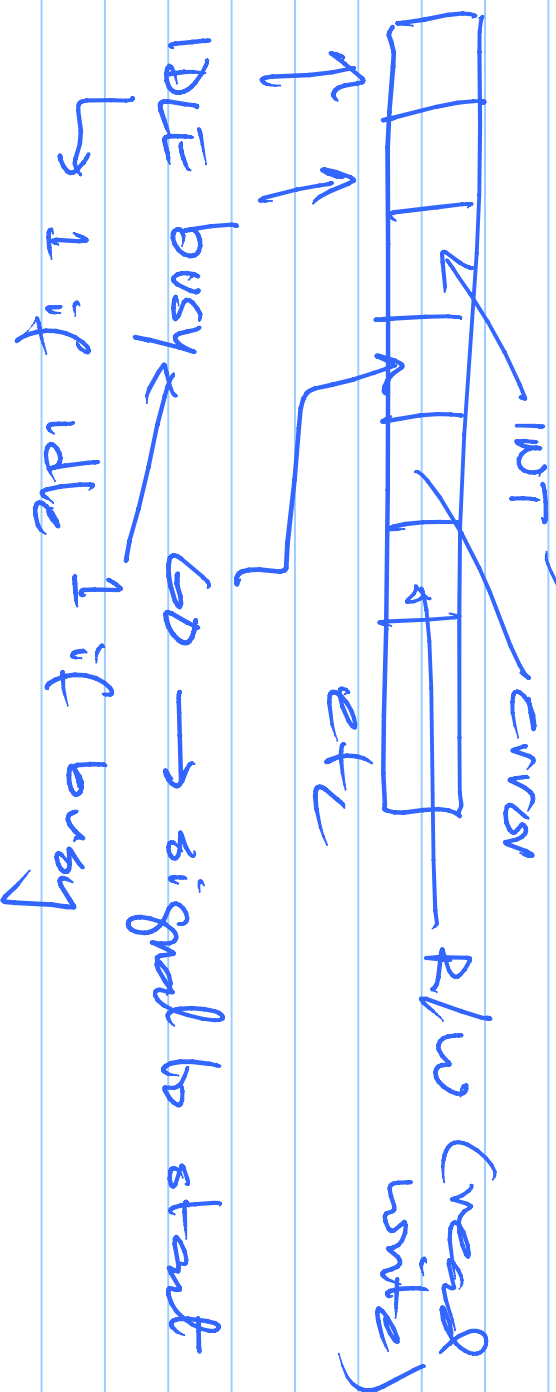


data Register

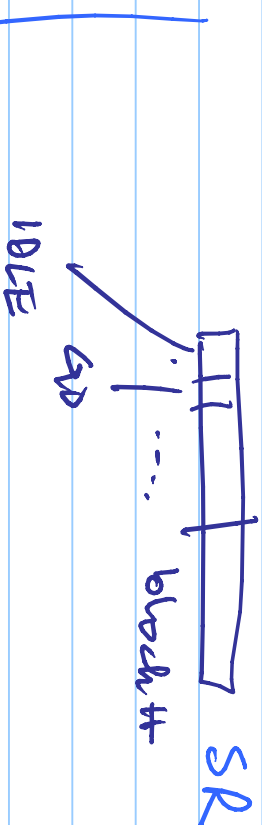


RAM  
address  
in mem

STATUS REGISTER for disk → int enable



BS wants to write block to from mem location  $\leftarrow$  disk addr



$\rightarrow$  wait till idle bit is set

put block # in SR

put  $n$  in data reg

Set R/W to W

Set 'go' bit to 1

# Parallel I/O

- wait for idle ] → busy wait  
- wait for done ] → busy wait

$P_1$

$P_2$

Busy Loop

→ wait for idle

→ wait for idle

set GO

→ wait for done

→ wait for done

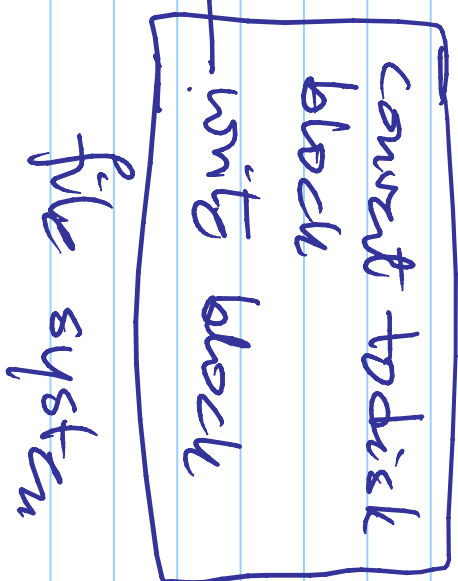
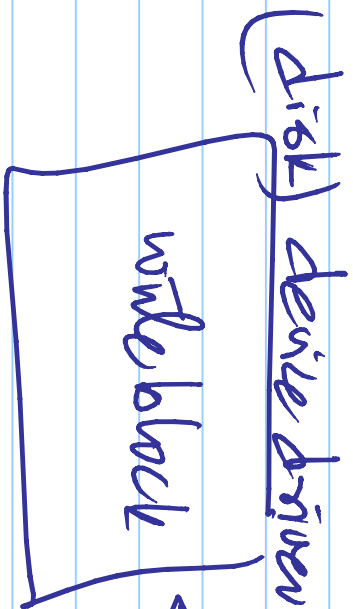
# Application

=====  
printf ( )

→ library

=====  
write

→ syscall



## interrupt driven I/O

- every time the disk controller finishes an operation, it signals
- sends an interrupt to the CPU
- device driver
- has 2 halves
  - user half
  - interrupt half

write a block (block#, addr)

{ P(disk\_mutex) ← I/O q

= put block# → SR  
= put addr → DR

set go bit

P(done) ← I/O burst

V(disk\_mutex)

}

INT

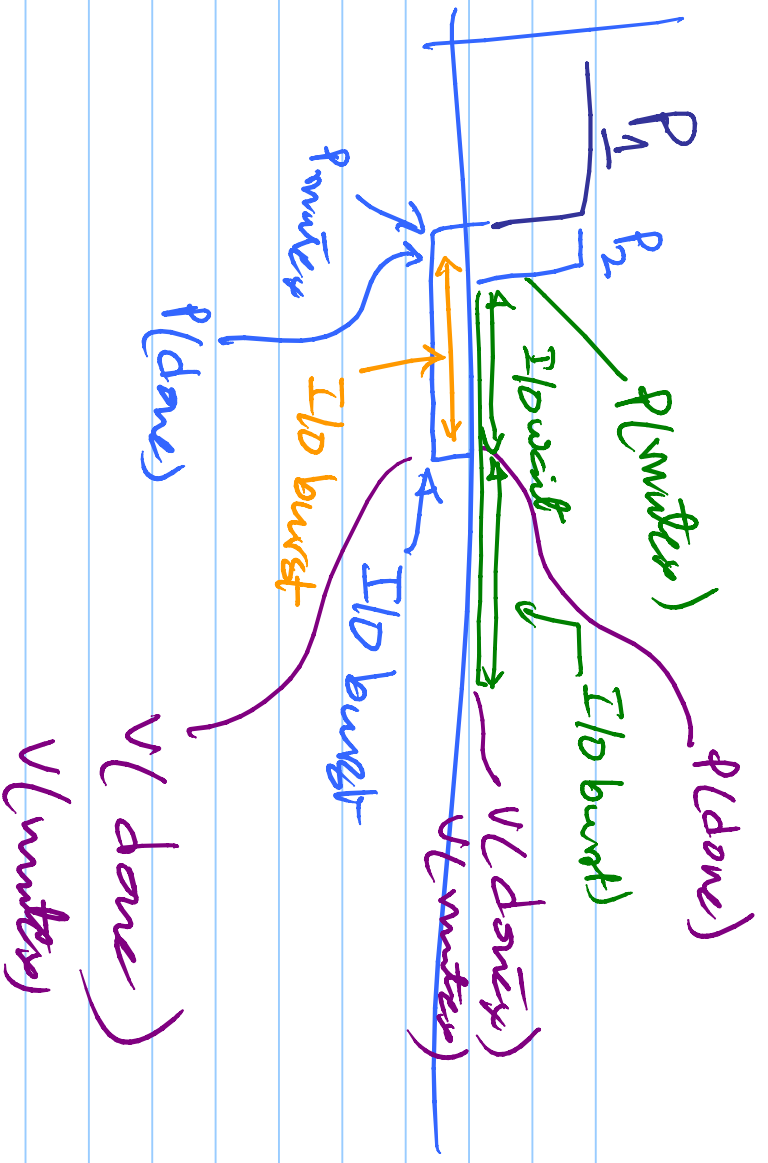
disk\_mutex = 1  
done = 0

INT

{

V(done)

}



# I/O scheduling

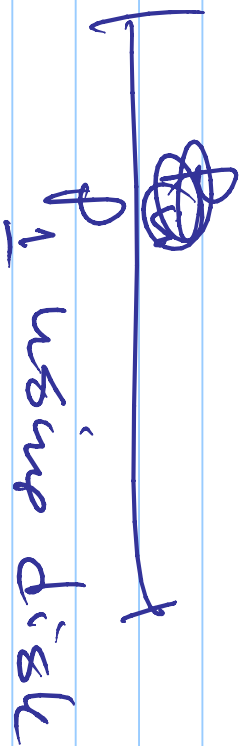
$P_1, P_2, P_3 \rightarrow$  all want to write to disk

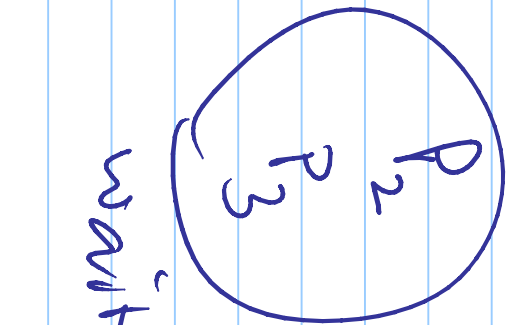
$\rightarrow$  what order (assume  $P_1$  came first & then  $P_2$  & then  $P_3$ )

$\rightarrow$  FCFS if  $P_3$

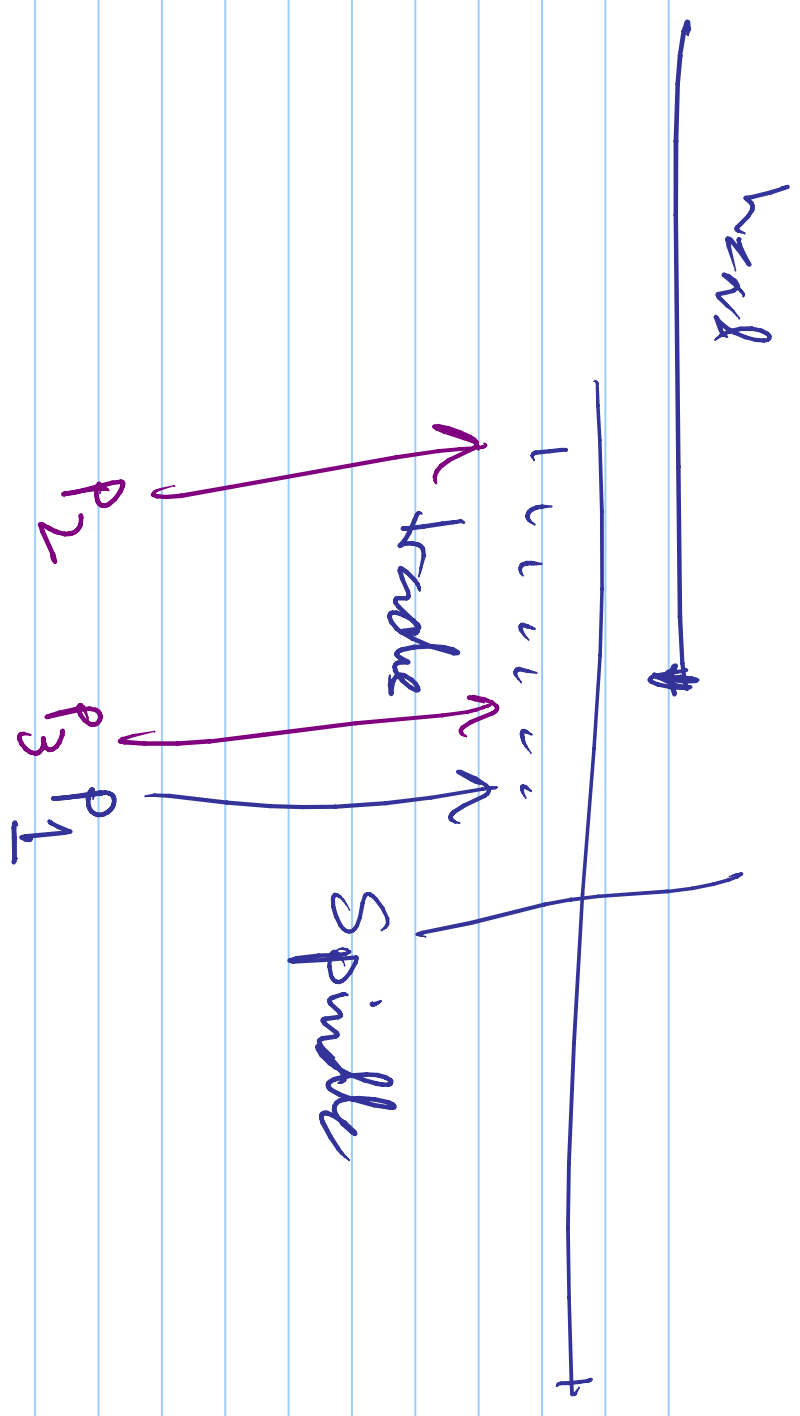
fair semaphores are used  $\rightarrow$

$P_1 \rightarrow$  gets  $\cancel{P}$  (mutex<sup>disk</sup>)

  
 $P_1$  using disk



$\rightarrow$  when  $P_1$  is done should we pick  $P_2$ ?  
waiting



# Disk I/O scheduling

1) Disk head scheduling

SEEK time

2) Disk <sup>rotational</sup> latency scheduling

latency

time taken to wait

time taken

to move head from  $t_1$  to  $t_2$  | head

# Seek optimization

① FCFS → not good

② SSTF → shortest seek time  
first