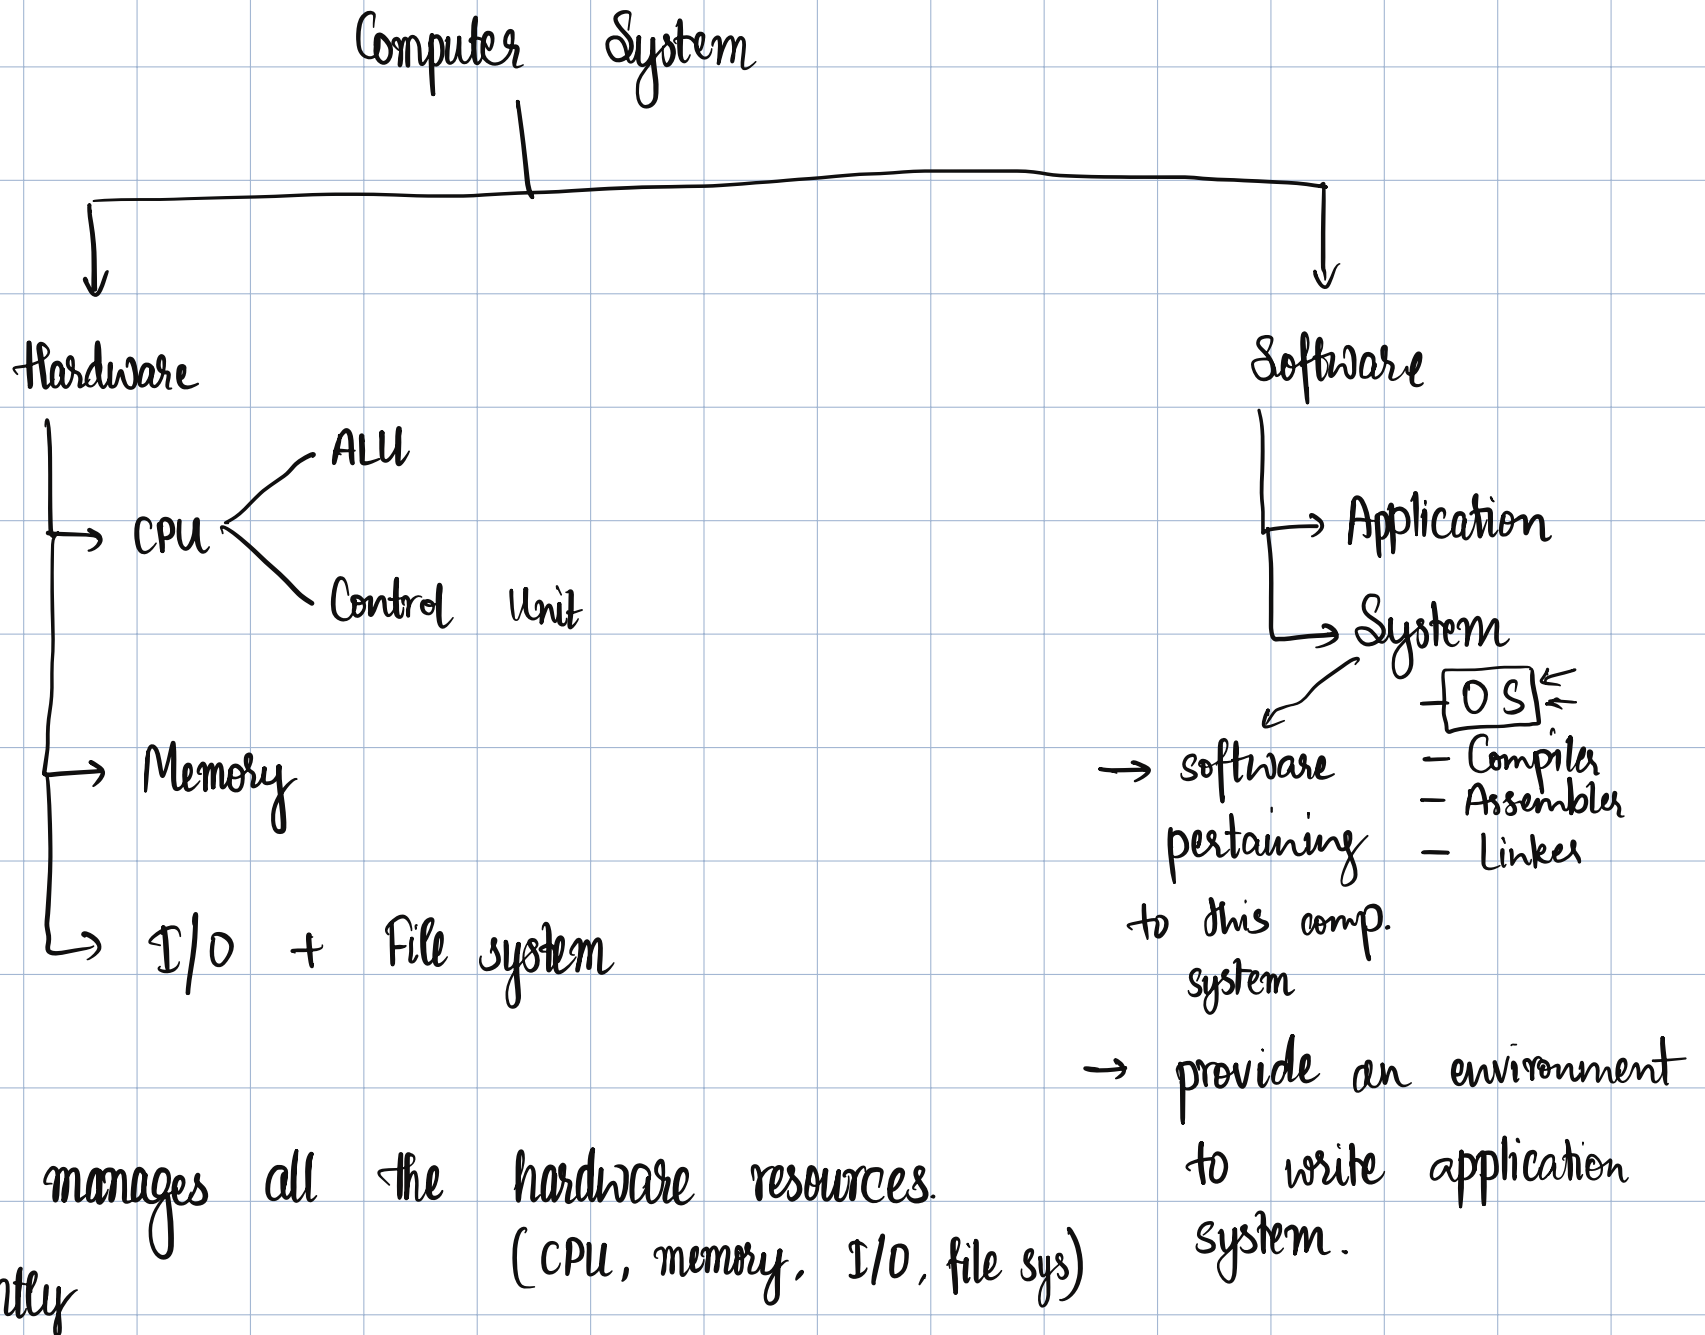


2024/10/22 - Operating Systems - Week 01

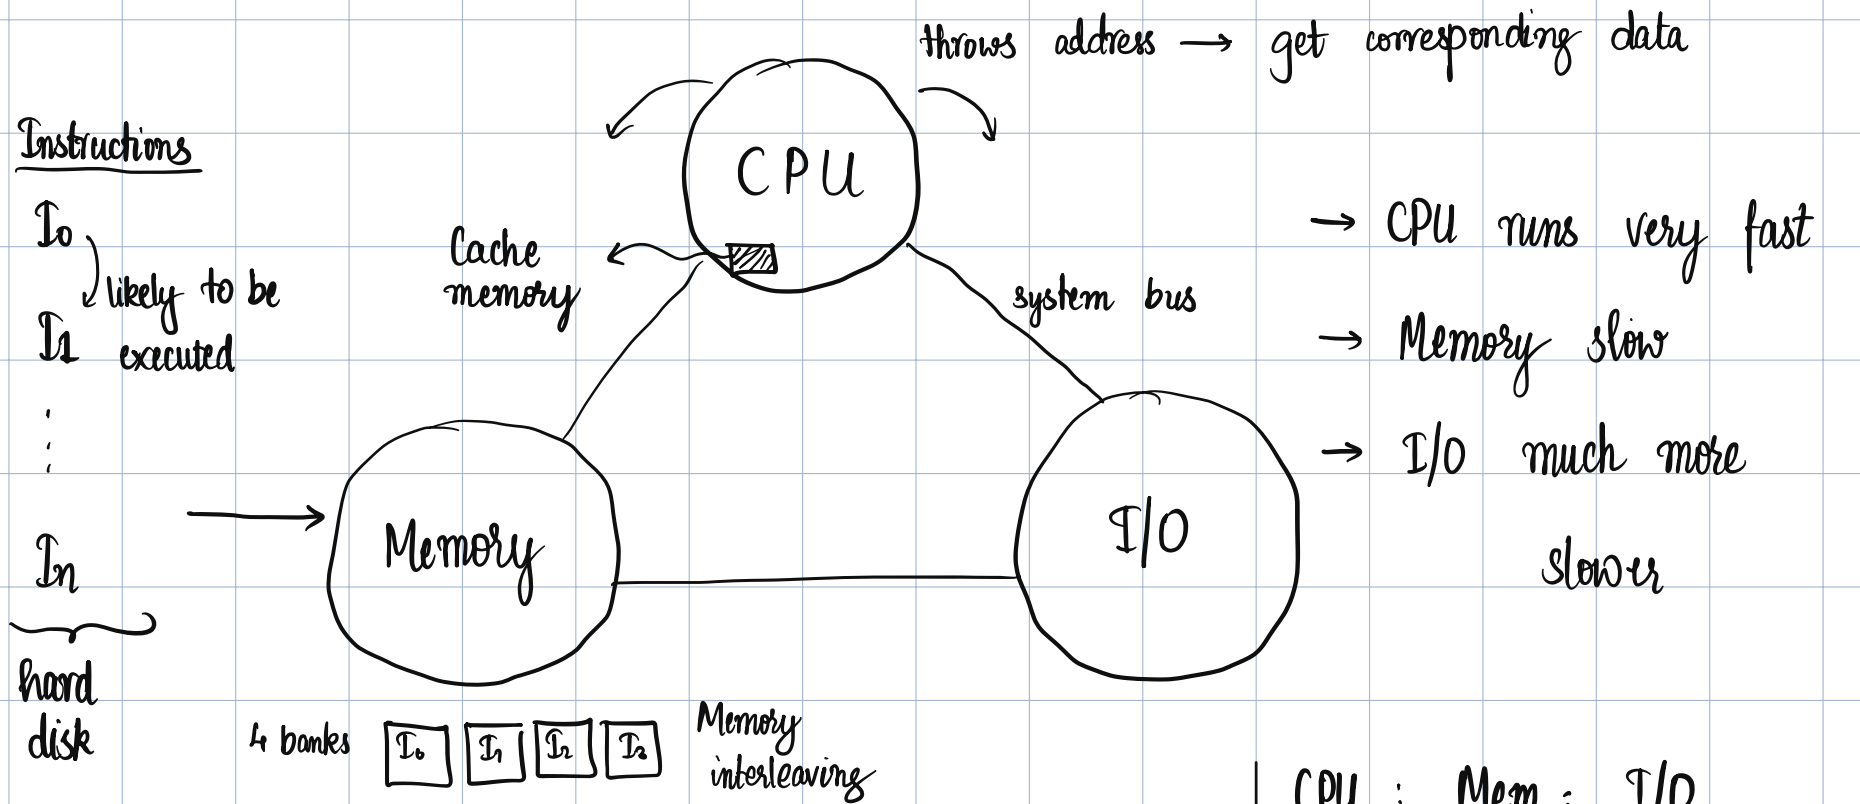
Textbooks : - "Operating Systems" → Andrew S Tenenbaum
→ Silber → OS-II
→ Stallings

1 final examination → Nov 24 sun (9 - 10:30)



* O.S manages all the hardware resources.
efficiently (CPU, memory, I/O, file sys)

- keyboard } → cannot be directly connected to a computer
- mouse } → need drivers
- display }



- Cache memory → cannot be large (cost of designing)
- locality of reference.

CPU	Mem	I/O
1	5	100
↓ CPU idle most of the time		

* Temporal and spacial locality

→ Cache

→ instruction cache

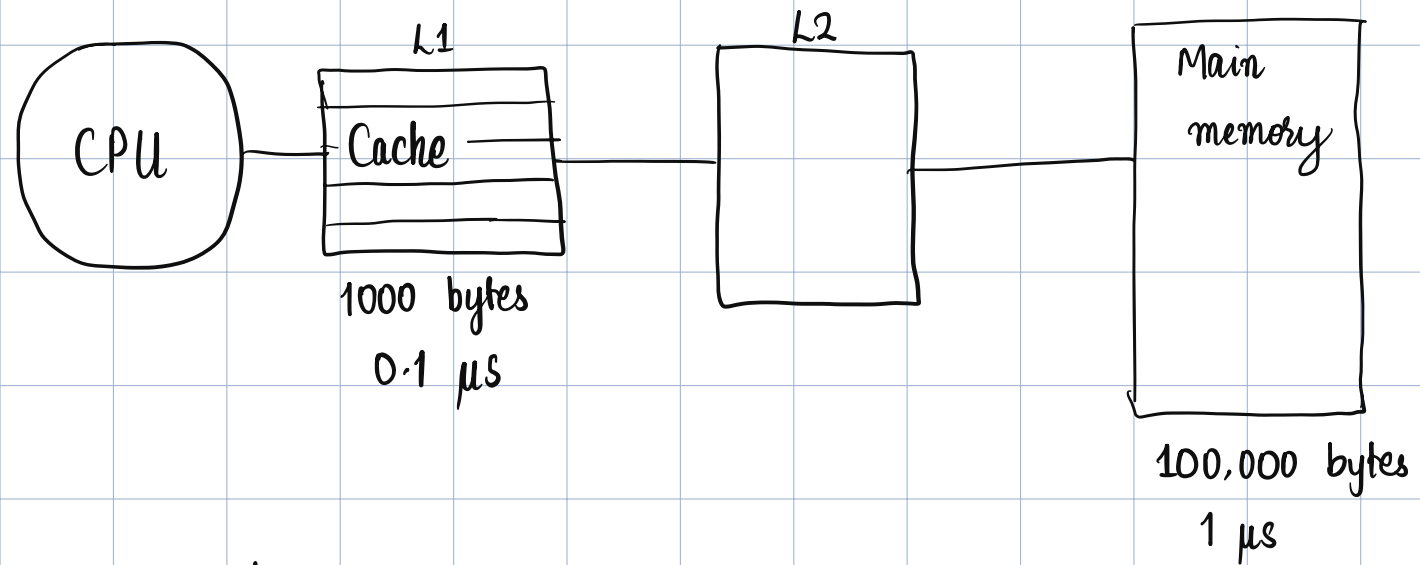
→ data cache

Should not
be mixed

CPU :	Mem :	I/O
1	5	1000

2024/10/24

→ Mismatch in speeds : CPU vs Memory → Cache, Memory interleaving
Cache circumvents this problem to some extent.



→ bigger cache \Rightarrow cost increases

$$\begin{aligned}\langle \text{access time} \rangle &= h \cdot t_c + (1-h) \cdot (t_c + t_m) \\ &= 80\% (0.1 \mu\text{s}) + 20\% (0.1 + 1) \mu\text{s}\end{aligned}$$

assume

$$\begin{aligned}h &= \text{hit rate} \\ &= 80\%.\end{aligned}$$

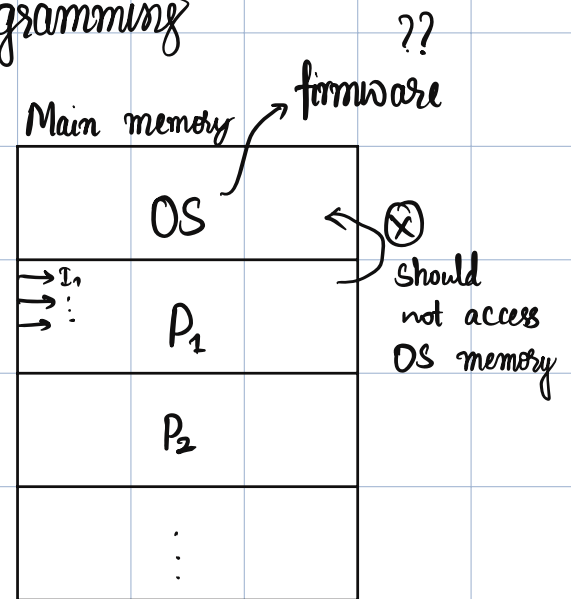
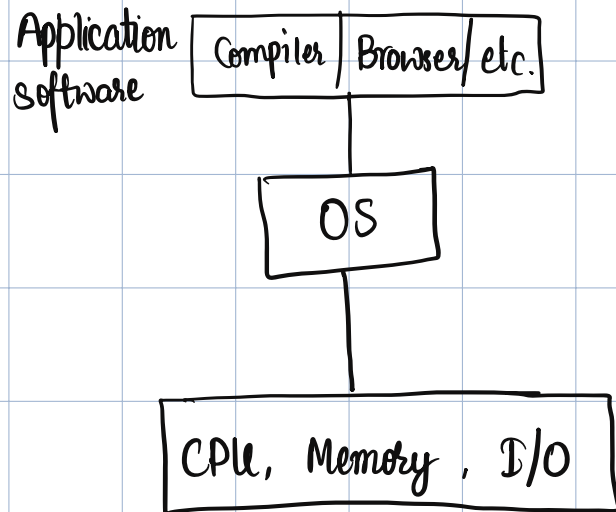
→ Mismatch in speeds : CPU vs. I/O → Multiprogramming

→ Problems

→ DMA (direct memory access)
→ keeps CPU & I/O busy

→ Mismatch in speeds : memory & I/O


↓
Multiport memories



P₁: only a part of the program (remaining part in disk) → because we want to support multiprogramming

Batch O.S

→ Cache is transparent to the CPU
in CS:
what is not seen



job = program + data
(Matrix multiply) (input matrices)

job takes a lot of
time \nRightarrow program is large
 \Rightarrow computational complexity
of program \uparrow

→ throughput : number of jobs per hour completed by the CPU.

→ turnaround : time difference : $t_{\text{completion}} - t_{\text{submission}}$

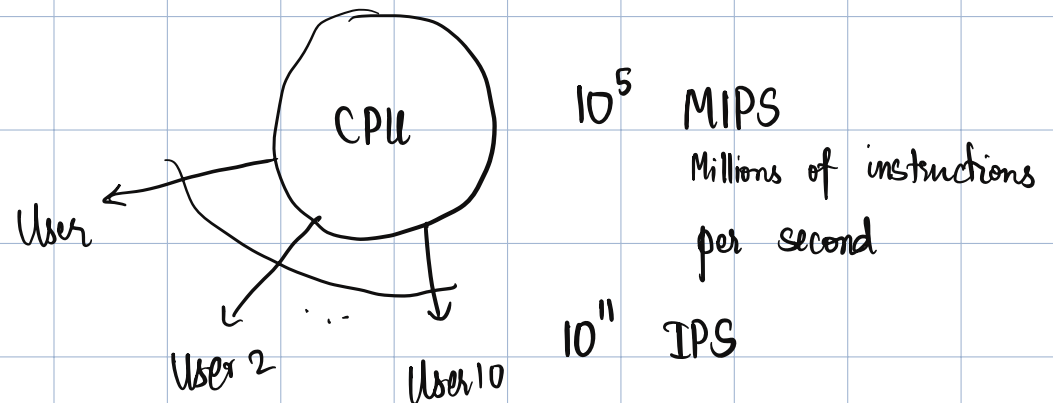
↑ throughput , ↓ turnaround

Time Sharing O.S [Interactive O.S]

→ Virtualization

Gives an illusion that
there are 10 CPUs

→ minimize response time,
maximize throughput.



→ O.S requires hardware support

→ Time sharing + batch O.S
day night

2024/10/25


→ Batch O.S

→ Time Sharing O.S

→ Real-time O.S.

Batch OS : example

<u>Job</u>	<u>Arrival time</u>	<u>processing time</u>	<u>Completion time</u>	<u>Turnaround time</u>
1	0	10	10	10
2	2	5	15	13
3	4	8	23	19



→ OS schedules jobs to the CPU

→ Turnaround time ↓ ⇒ better

Time sharing OS:

→ Response time (time taken to start responding)
virtual CPU

Batch OS
turnaround
time

Time sharing
Response time

Real time OS

→ programs have "deadlines"; must be completed by deadlines.

→ e.g.: satellite systems

If p_2 is done first
 $t_{\text{completion}} = 35$

$t_{\text{completion}} = 55$

Process	deadline	Processing time
p_1	100	35
p_2	50	20

If program doesn't complete
deadline

Hard RTOS

Missile launch
at approaching
target

Soft RTOS


Multimedia
streaming

Batch OS Time Sharing Real time
→ turnaround, response time, meeting deadline

Process Program in execution

requires
ID, state, priority, program, data, stack, heap, context
instructions function calls dynamic memory

→ stored in process control block (PCB)
or process table → array of records (structs)



resume at what instruction?

→ not required in Batch OS
(each program executes completely, doesn't need to switch)

→ only context, right?

→ Process table → in memory
→ cannot be accessed by the user.
→ only by the O.S.

context to be stored and loaded

- store program counter (PC), registers
- list of files in use
- accounting info
 - no. of CPU cycles spent

→ switching is not free of cost.

context switching overhead

→ Who creates the processes? OS.

UNIX

fork()

exit()

kill()

Windows

create process()

exit process()

terminate process()

→ only 1 state at a time.

Process state diagram

