
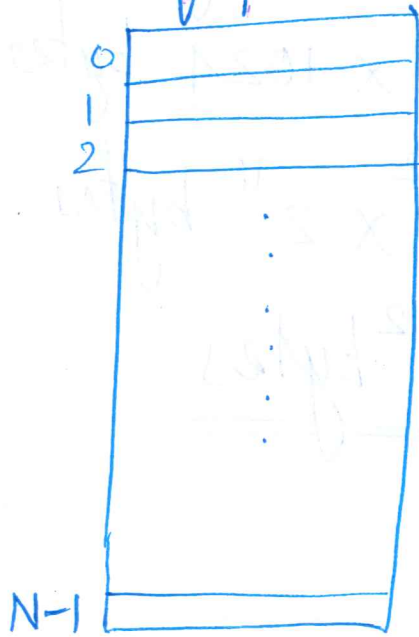


# Lecture - 35

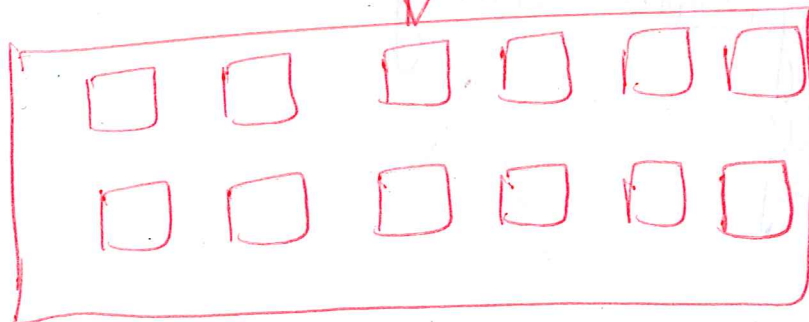
## Virtual Memory

VM of process A

①



Remember "block size" in cache memories?



Block  $\equiv$  Virtual Page in VM system.

Each Virtual Page is  $P = 2^p$  bytes in size.

eg. Virtual page size = 4 KB = 4096 bytes

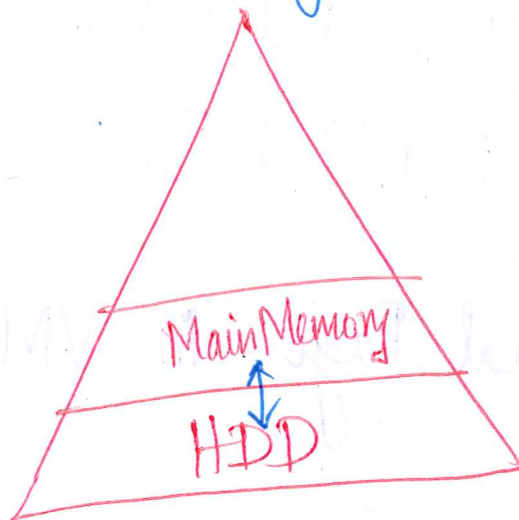
$$= 4 \times 1024 \text{ bytes}$$
$$= 2^2 \times 2^{10} \text{ bytes}$$
$$= \underline{\underline{2^{12} \text{ bytes}}}$$

Physical Page

Also  $P$  bytes in size.

ie.  $\boxed{\text{Virtual page size} = \text{Physical page size}}$

$\therefore$  Blocks should be of the same size between 2 adjacent storage devices in the memory hierarchy.



# 3 types of virtual pages



## Unallocated

- \* pages that are not yet created by the VM system.
- \* no data associated with them.
- \* do not occupy any space on disk.

## Cached

- \* Allocated pages that are currently cached in physical memory.


## Uncached

- \* Allocated pages that are NOT cached in physical memory.

(4)

# Virtual Memory

# Physical Memory




Virtual Pages (VPs) stored on disk

Physical Pages (PPs) cached in DRAM.

- UA - Unallocated
- C - Cached
- UC - Uncached

# Main Memory (DRAM) as a cache

1. Block size (Virtual Page size) = 4KB to 2MB



- 2. Associativity: fully associative
- 3. Replacement policy: More sophisticated than L1, L2, L3 caches!
- 4. Write policy: write-back!

## PAGE TABLES!

	Valid	n-bit address field
PTE 0	1	0x8079A2C1
1	0	
2	0	
⋮	⋮	⋮
N-1	1	

Physical Memory

Virtual Memory (disk)

*Handwritten note: PTE = Page Table Entry*

(6)

Valid PPN (or) disk address  
disk address

Physical Memory (DRAM)

PTE 0	0	NULL
1	1	
2	1	
3	1	
4	0	NULL
5	0	NULL
6	0	
7	1	

VP1
VP2
VP7
VP4
VP3

Victim page  
Virtual Memory (disk)

VP0
VP1
VP2
VP3
VP4
VP5
VP6
VP7

PAGE TABLE!

Page Hit

Read VP2!

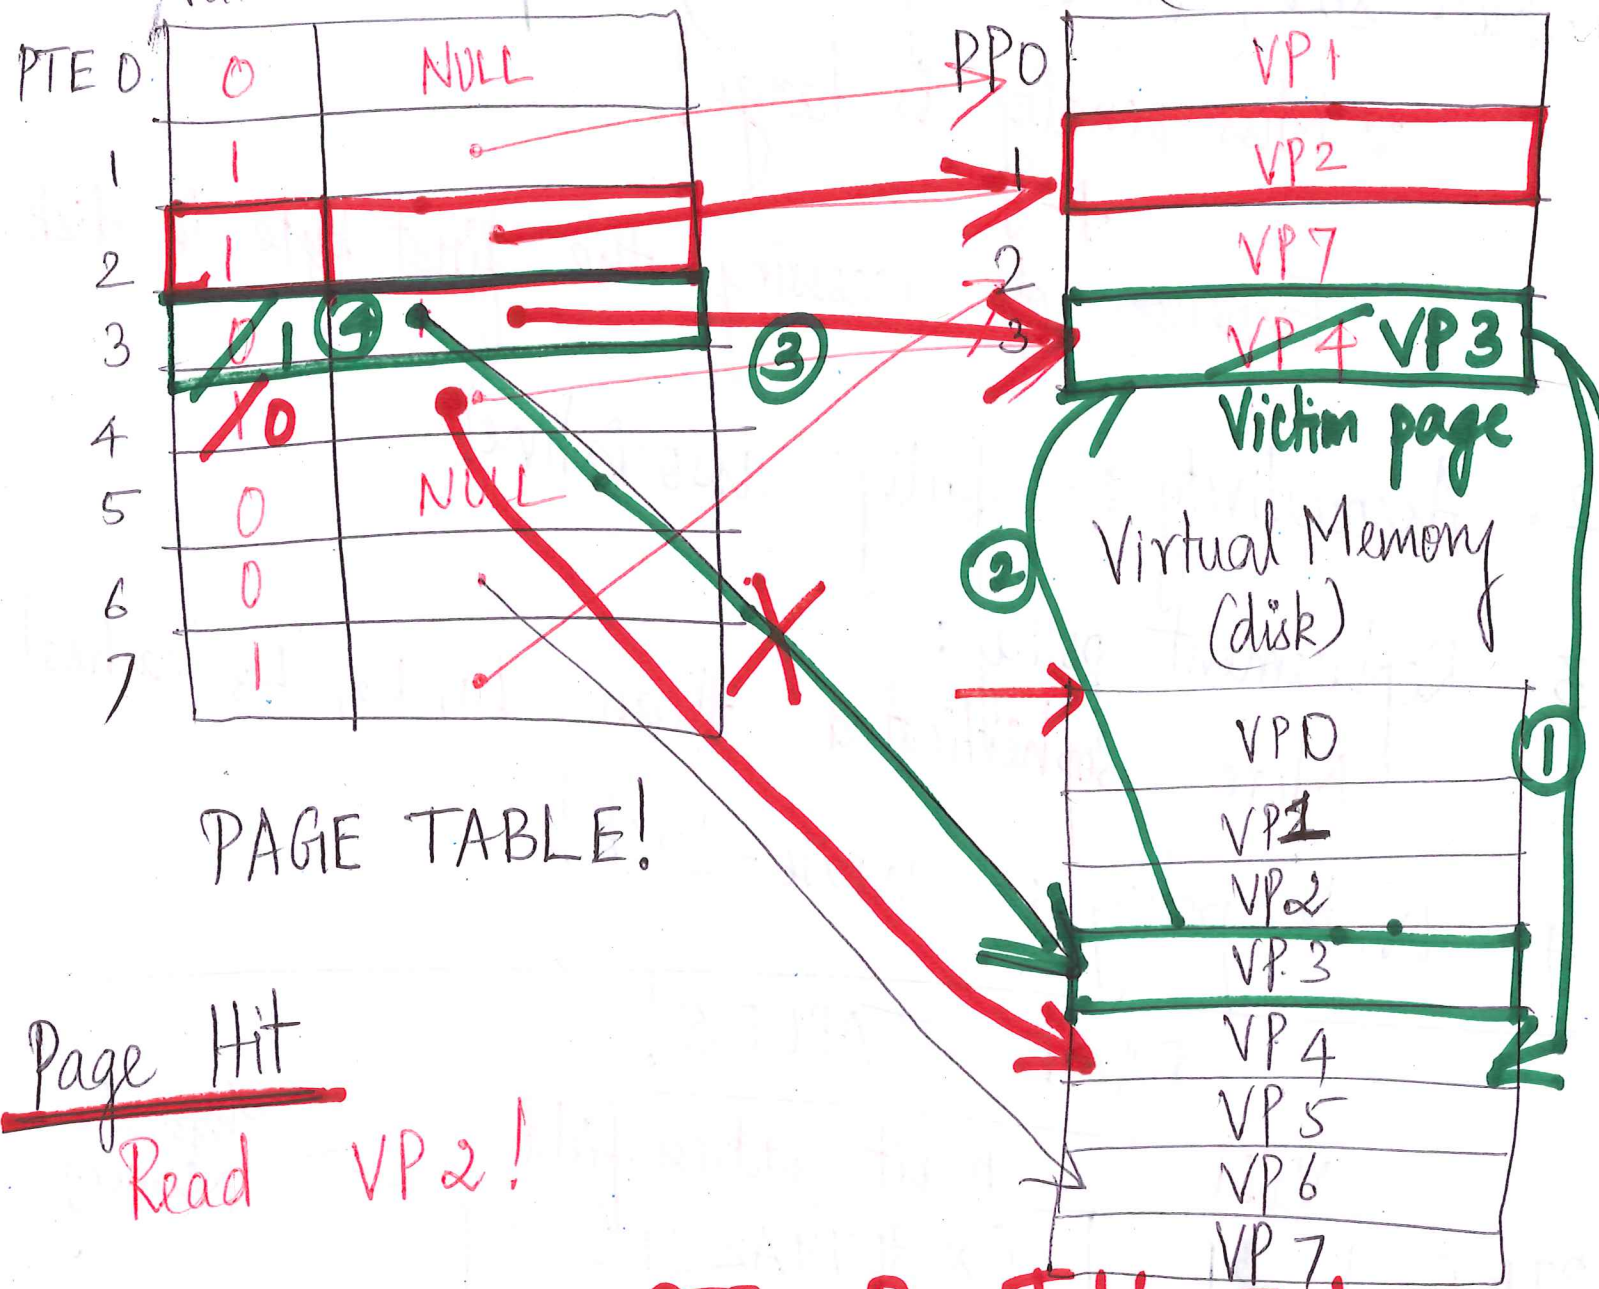
Page Fault



Read VP3!

PTE - Page Table Entry

demand paging




# VM as a tool for Memory Management

- 1. Simplifying linking
- 2. " loading
- 3. " sharing
- 4. " memory allocation.


# VM as a tool for memory protection.

Virtual Address Spaces


Process i:



Process j:



Physical Memory



Per process page table.

# Page-level memory protection

Page tables with "permission bits"

## Physical Memory

	SUP	READ	WRITE	Address
VP0	No	Yes	No	PP6
VP1	No	Yes	Yes	PP2
VP2	Yes	Yes	Yes	PP10

Process i

	SUP	READ	WRITE	Address
VP0	No	Yes	No	PP9
VP1	Yes	Yes	Yes	PP6
VP2	No	Yes	Yes	PP7


Process j





# ⑨ Address Translation

N-element  
Virtual Address Space  
(VAS)

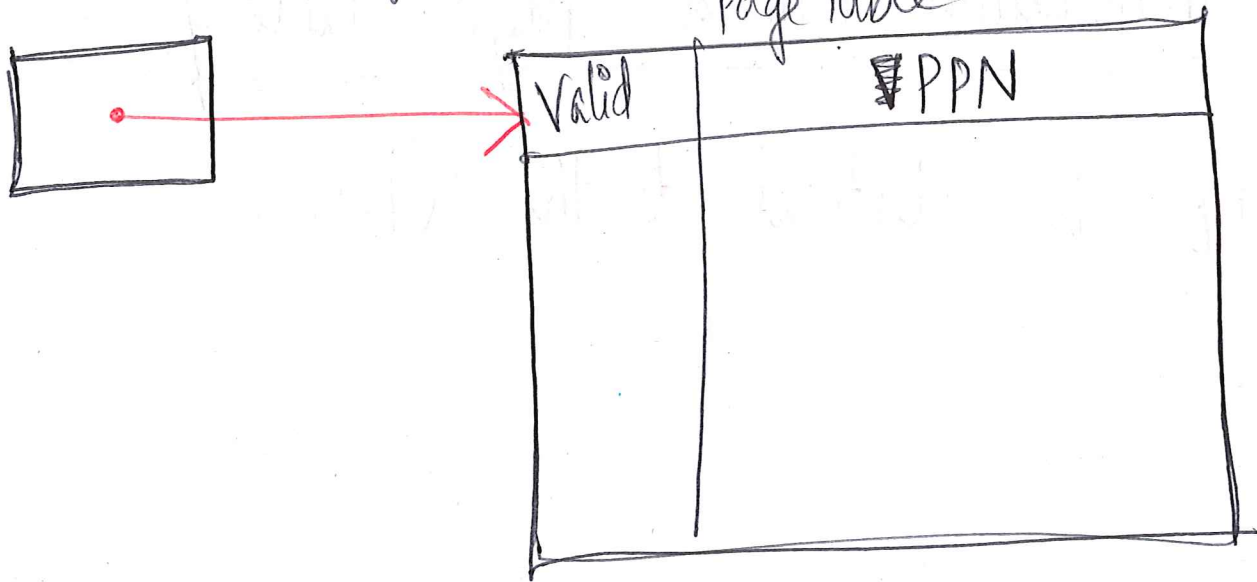


M-element  
Physical Address Space  
(PAS)

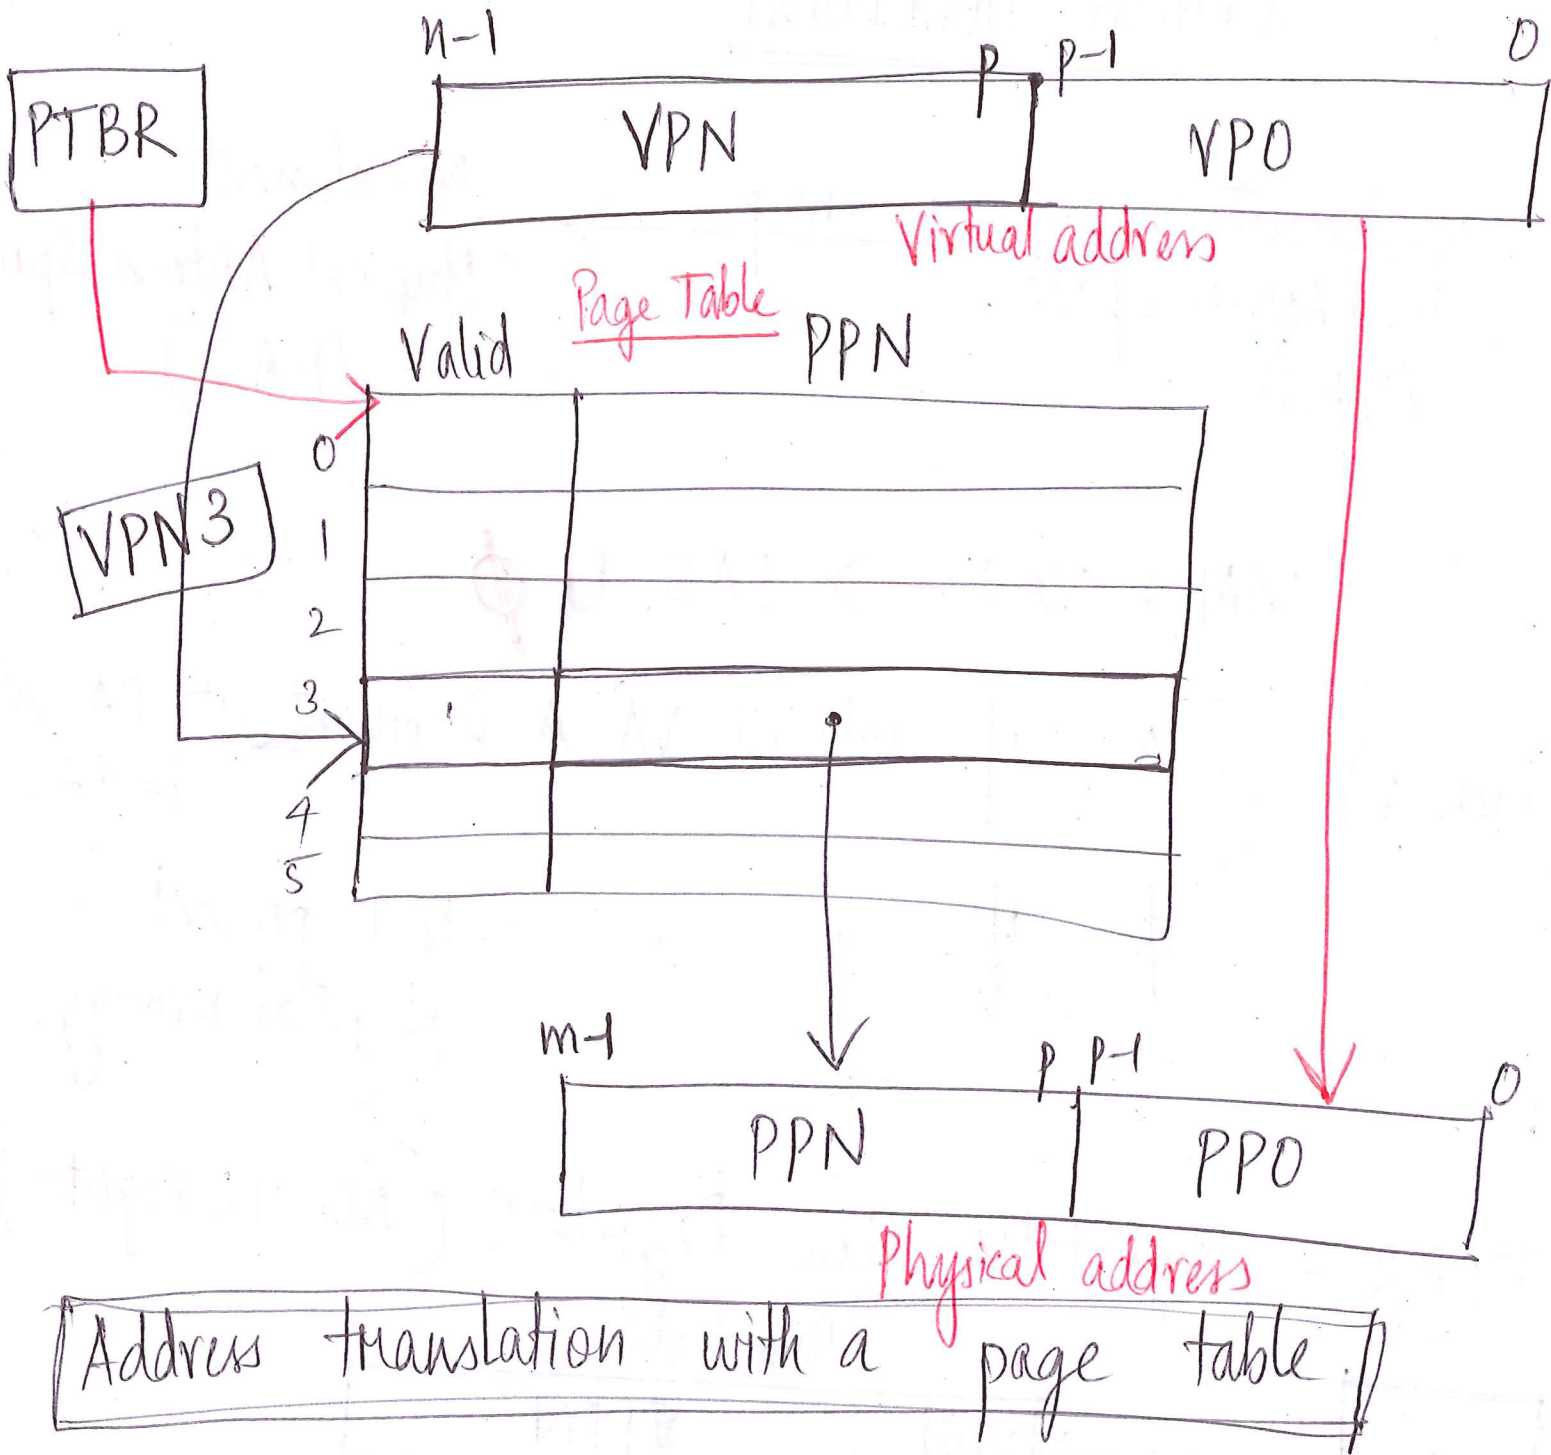
$$\text{MAP} : \text{VAS} \rightarrow \text{PAS} \cup \emptyset$$

$$\text{MAP}(A) = \begin{cases} A' & \text{if data at VA } A \text{ is present at PA } A' \text{ in PAS.} \\ \emptyset & \text{if " " " NOT present in physical memory.} \end{cases}$$

PTBR - Page table Base Register. (CPU Register)

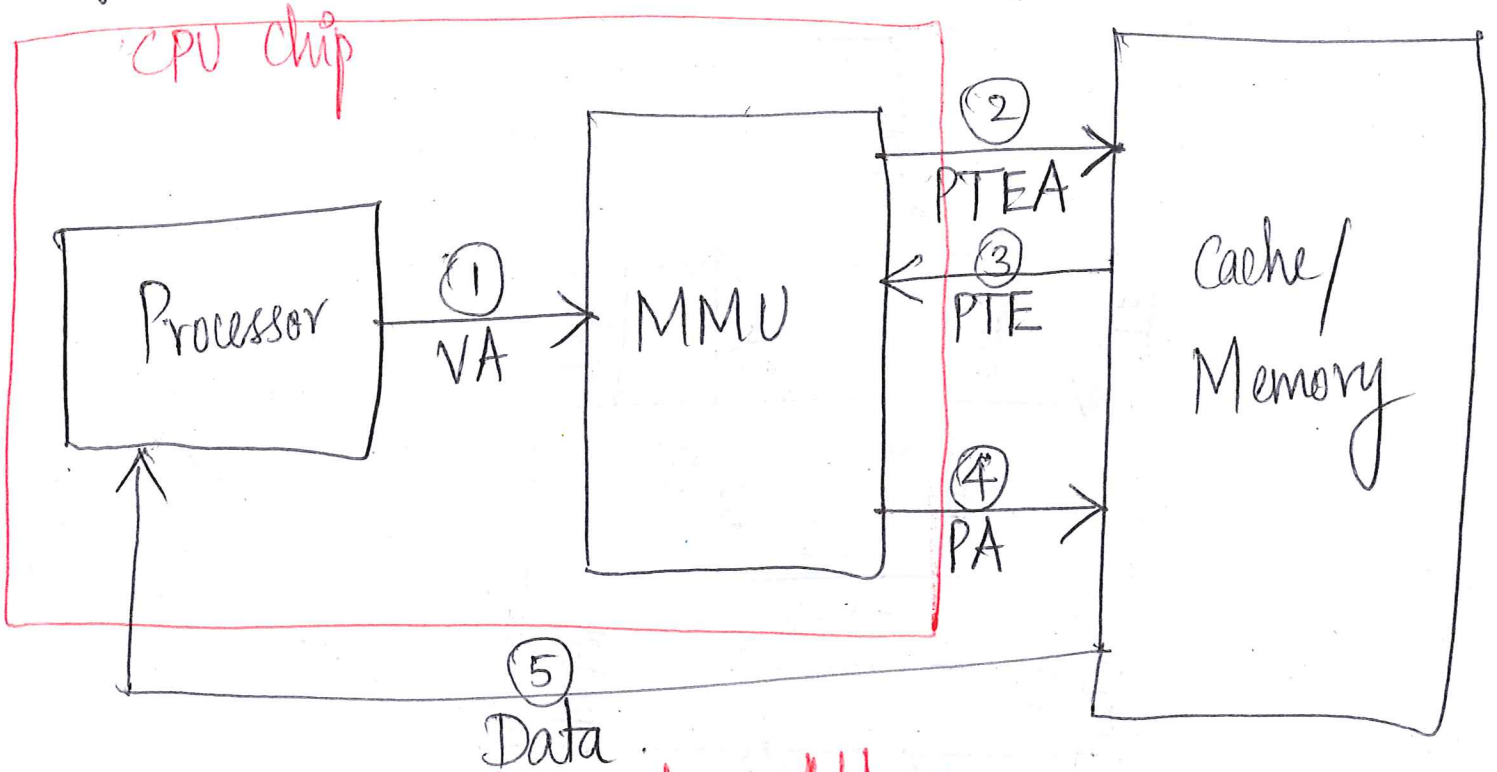


(10)



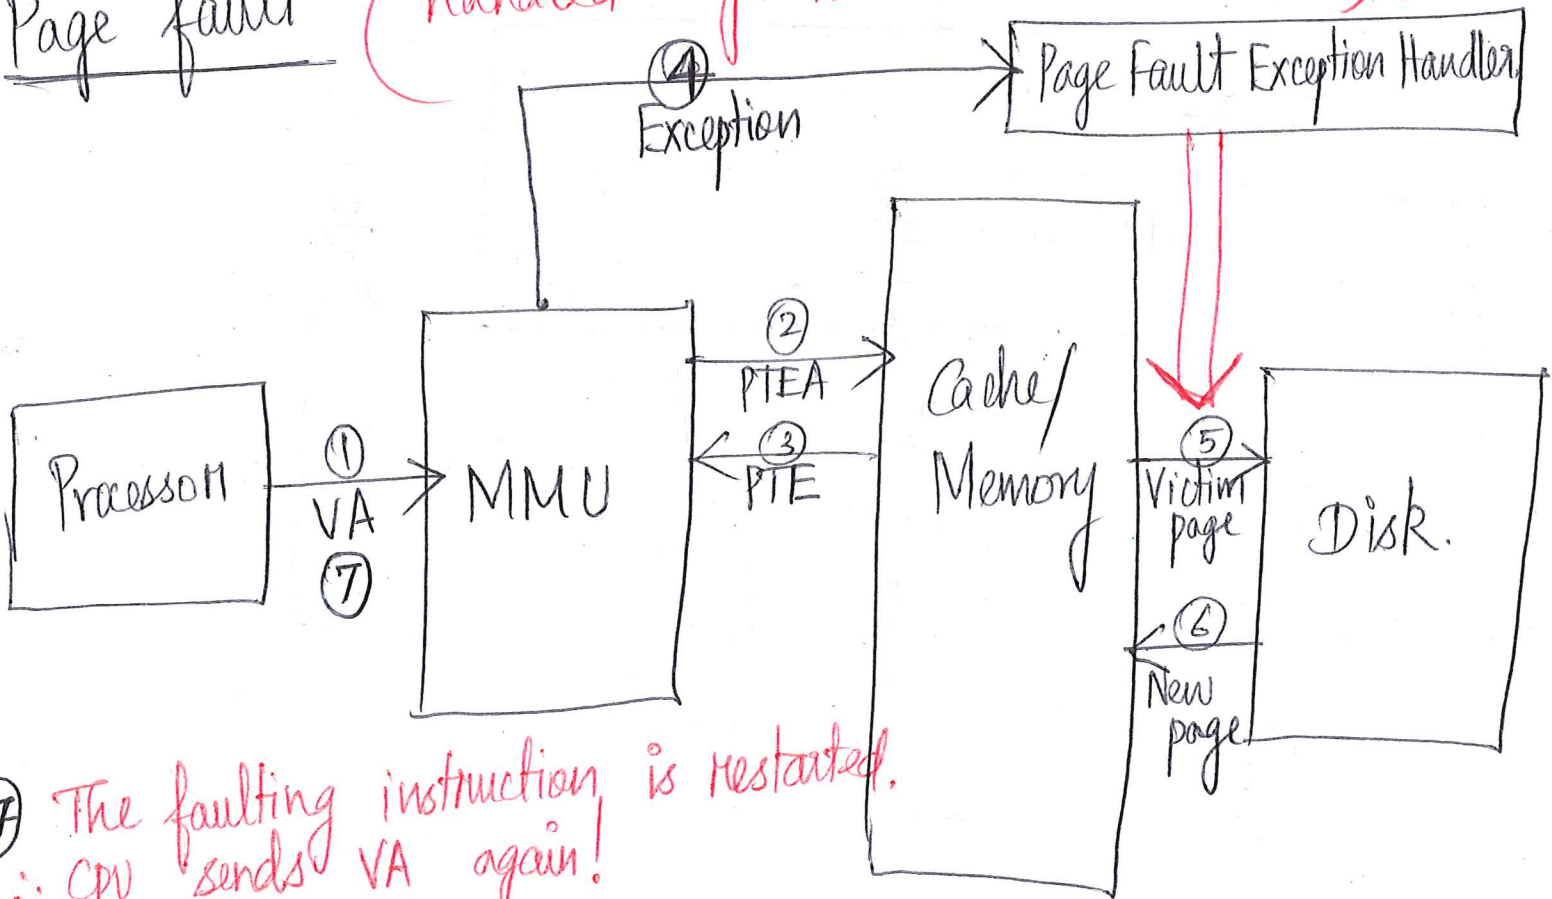
NOTE: PPO is identical to the VPO.

Page hit (handled entirely by the hardware) (11)



PTEA - Page Table Entry Address.

Page fault (handled by HW and OS Kernel)



7 The faulting instruction is restarted.  
 ∴ CPU sends VA again!

# Caches + VM

