

# CSE251: System Programming

## 15. Virtual Memory (1)

Seongil Wi

# HW4: Cache

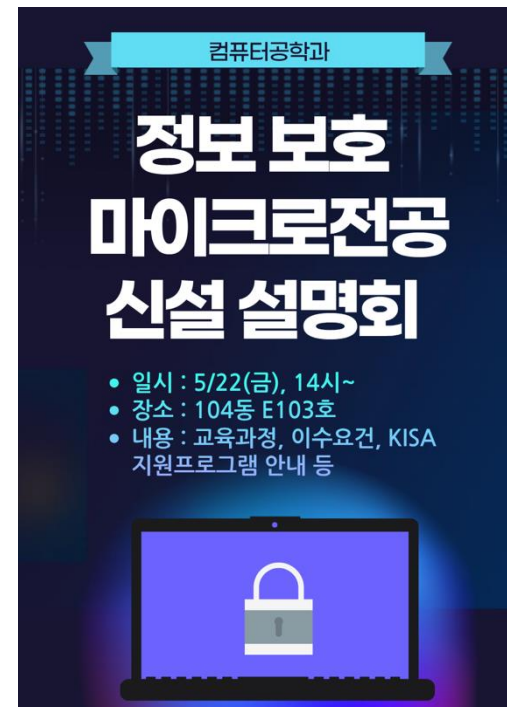
---



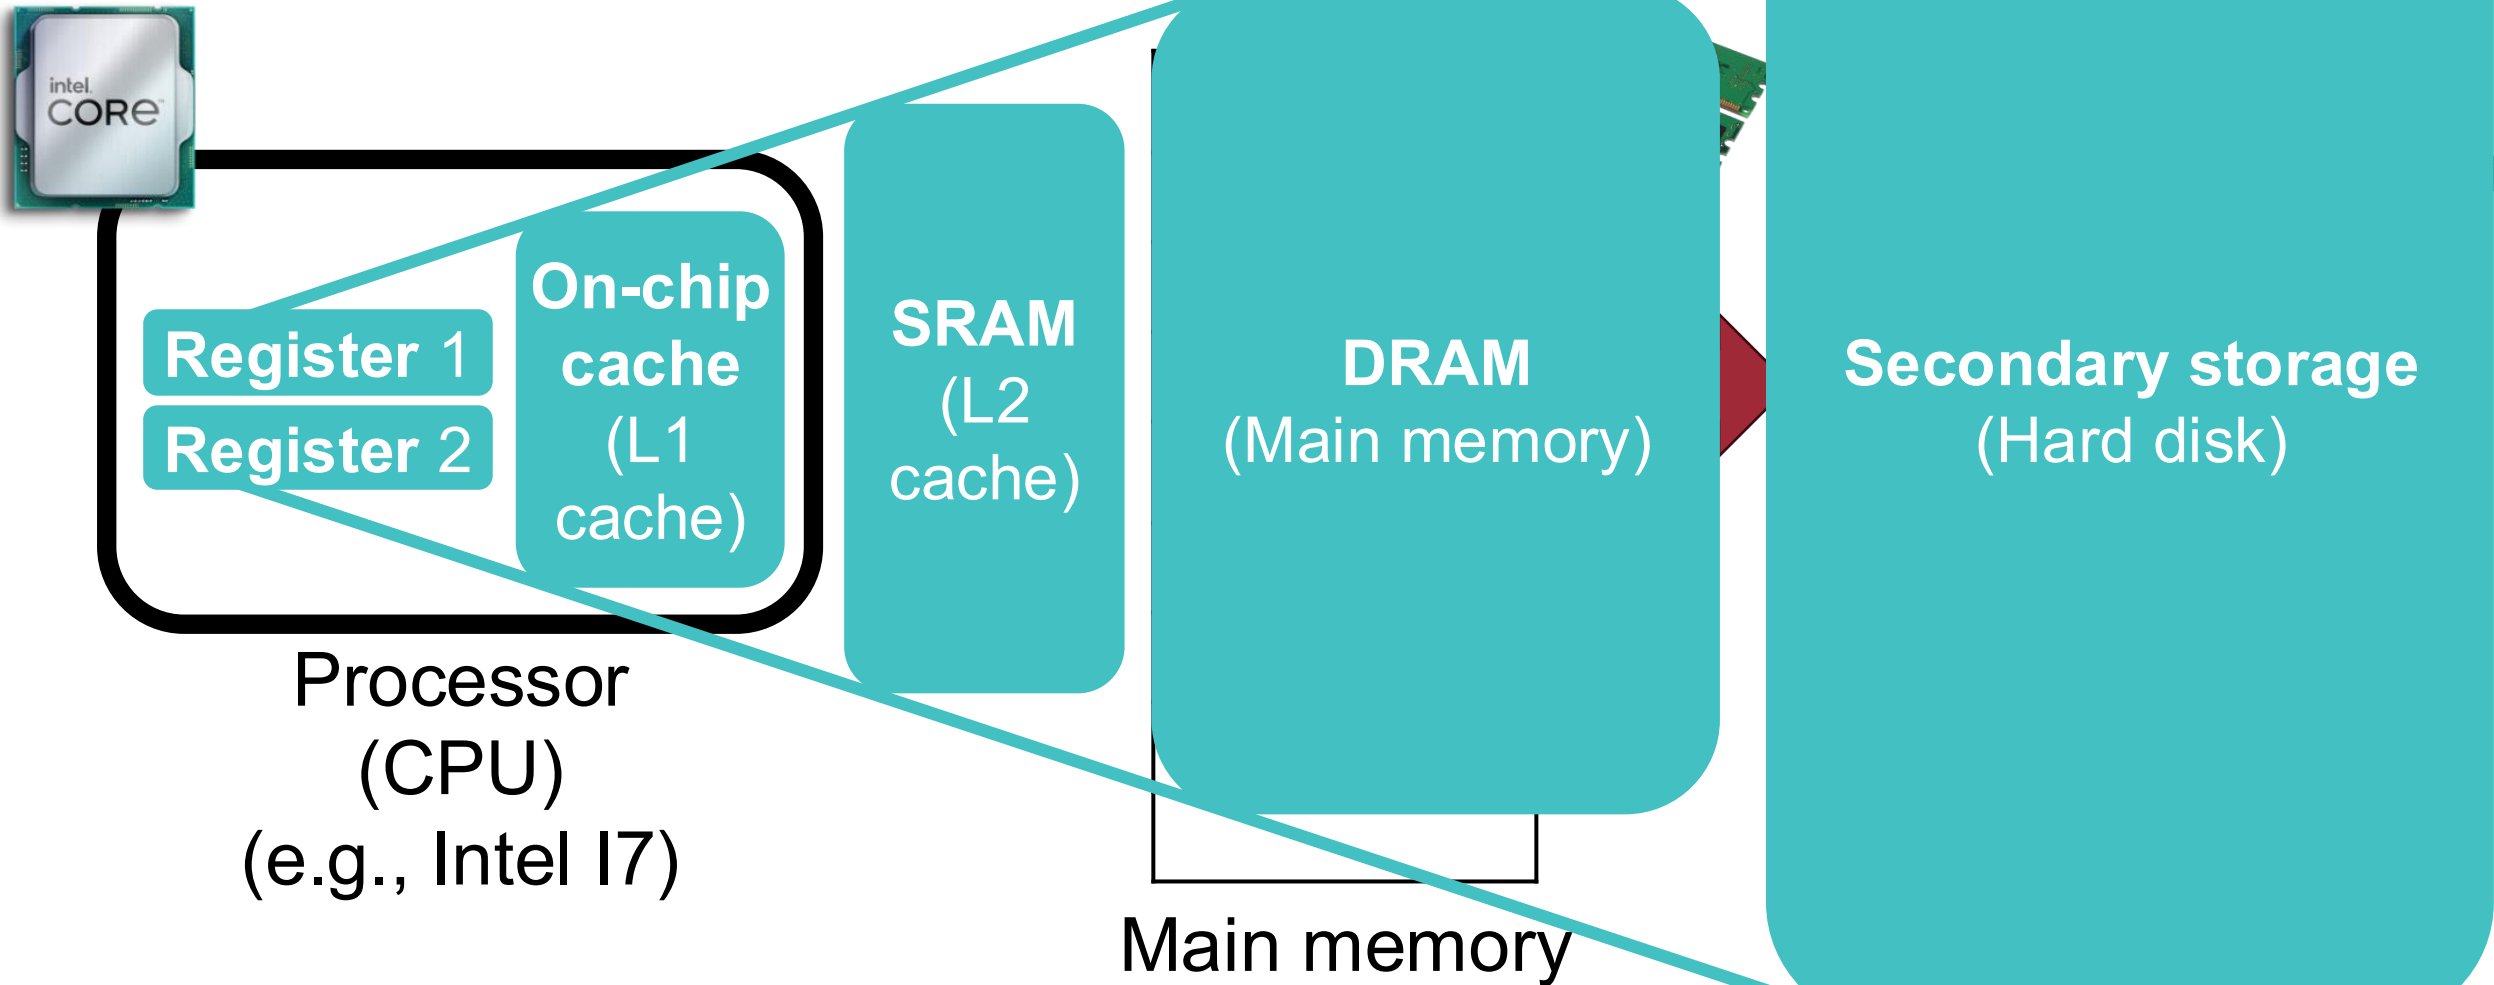
1. Write a small C program (about 200–300 lines) that simulates the behavior of a cache memory
  2. Optimize a matrix transpose function with the goal of minimizing the number of cache misses
  3. Answer the given questions
- 
- Due: May. 18, 11:59PM
  - Responsible TA: Jongyoung Kang
  - If you have any requests or questions, please ask via Blackboard Q&A session

# Information Session: Information Security Micro Minor <sup>3</sup>

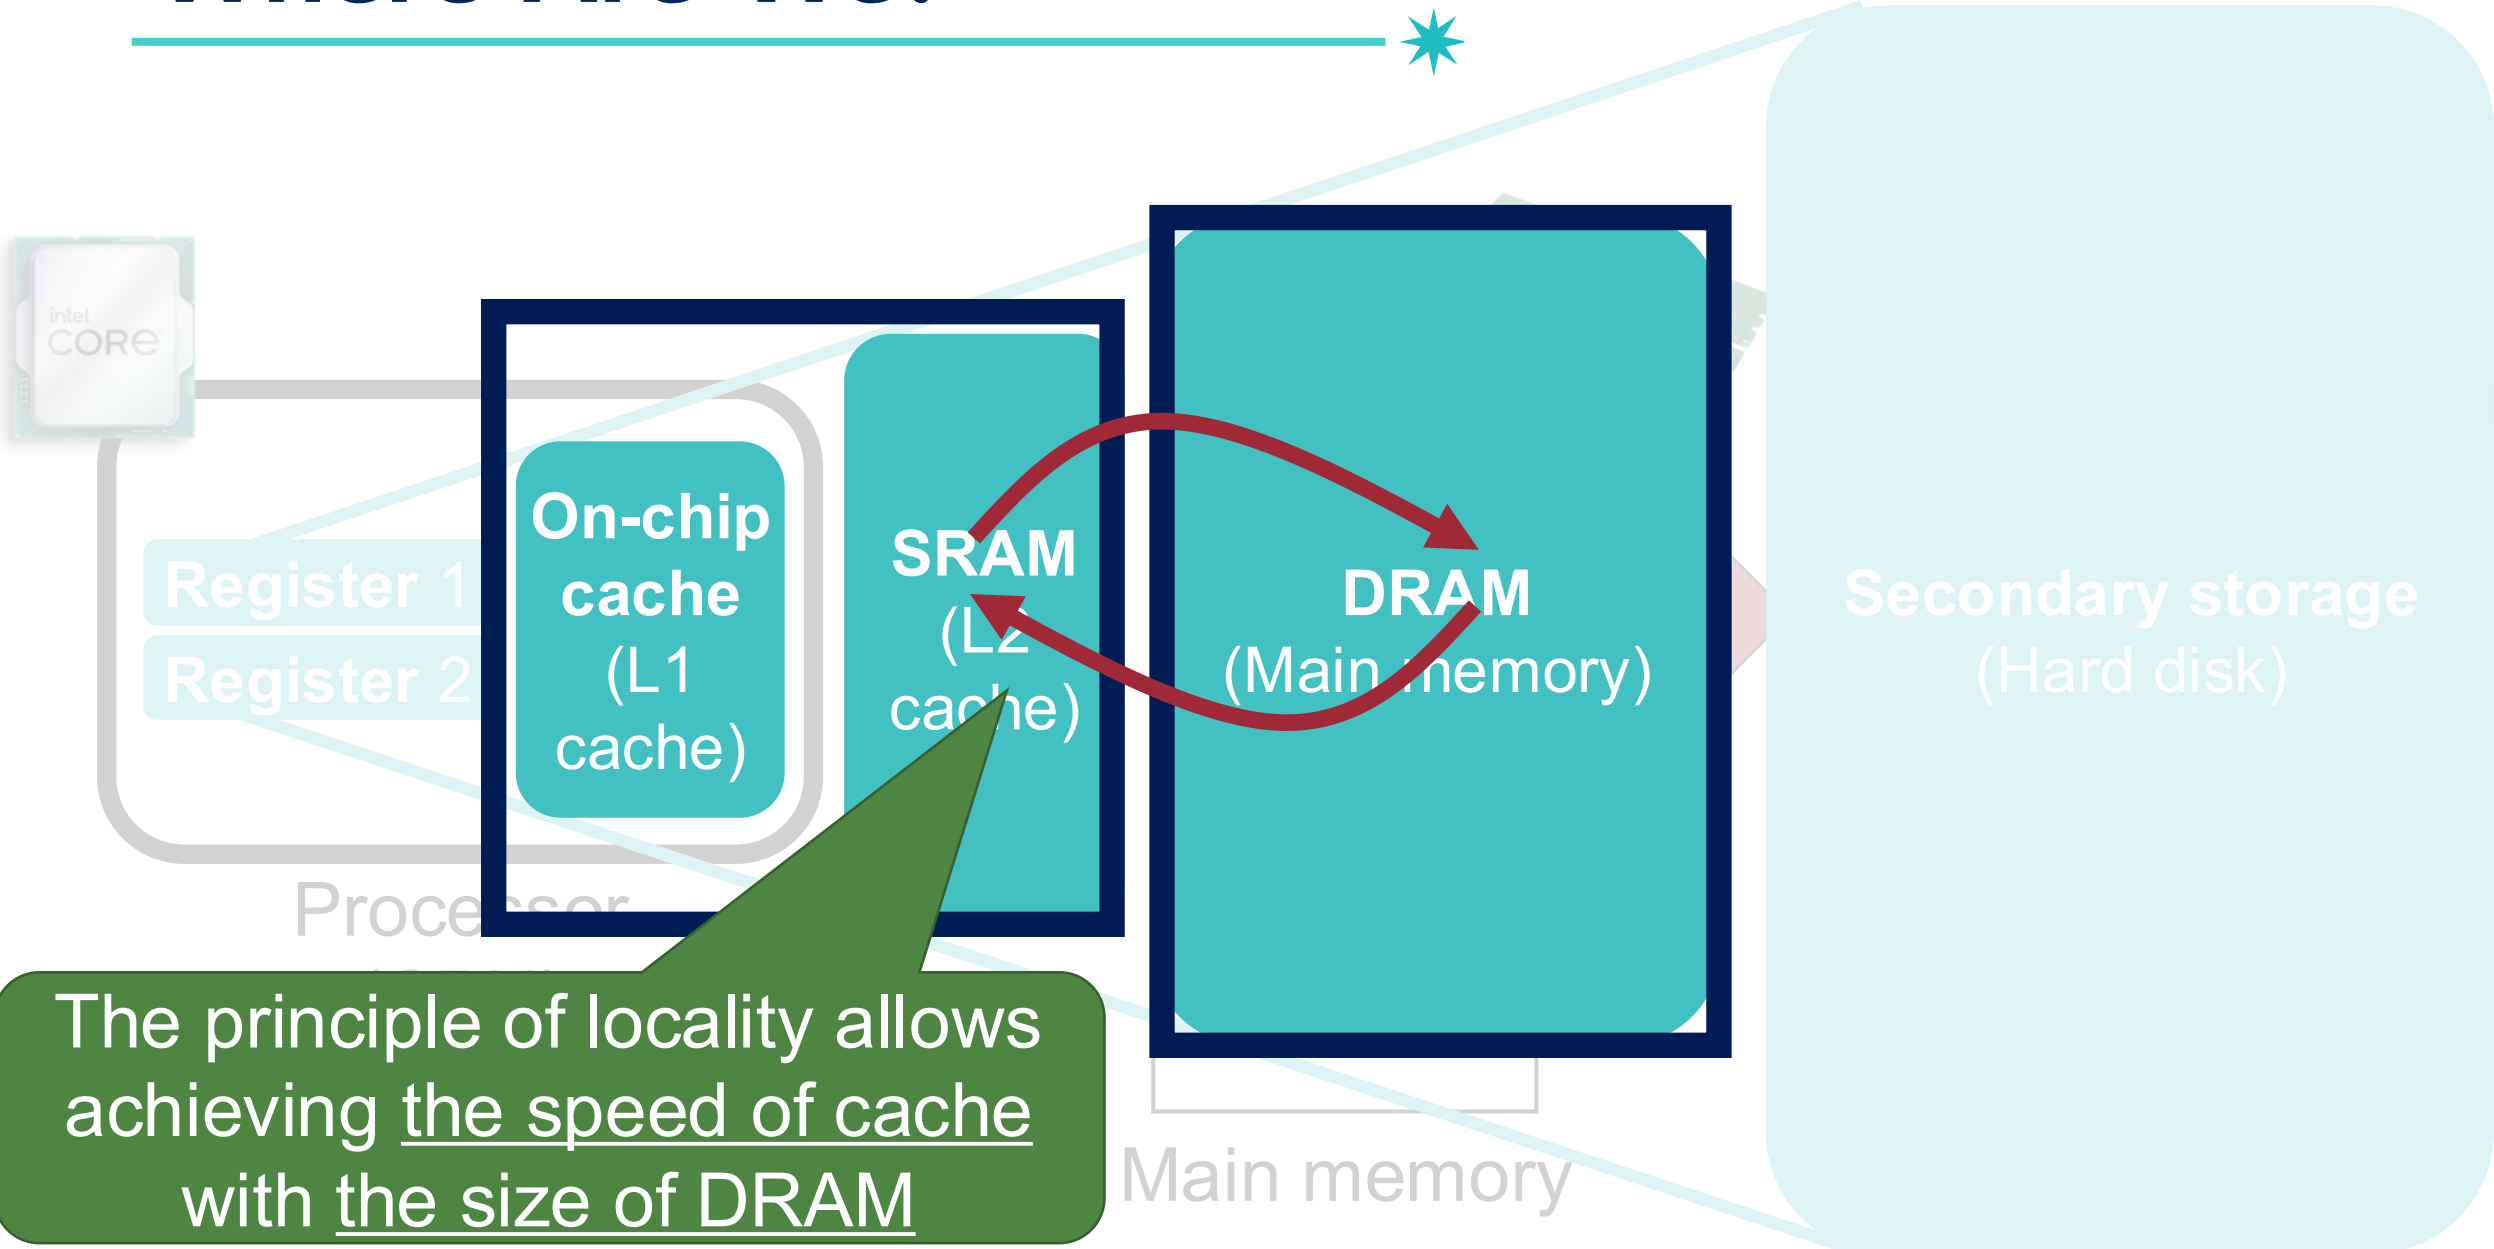
- Date and Time: May 22 (Fri.), 2PM~
- Venue: Room E103, Building 104
- Target: Interested Undergraduate Students
- Details: Information on the curriculum, completion requirements, KISA support programs, etc.
- Please register in advance!  
<https://forms.gle/KQMnCVmZXezAsgmj8>



# Where Are We?

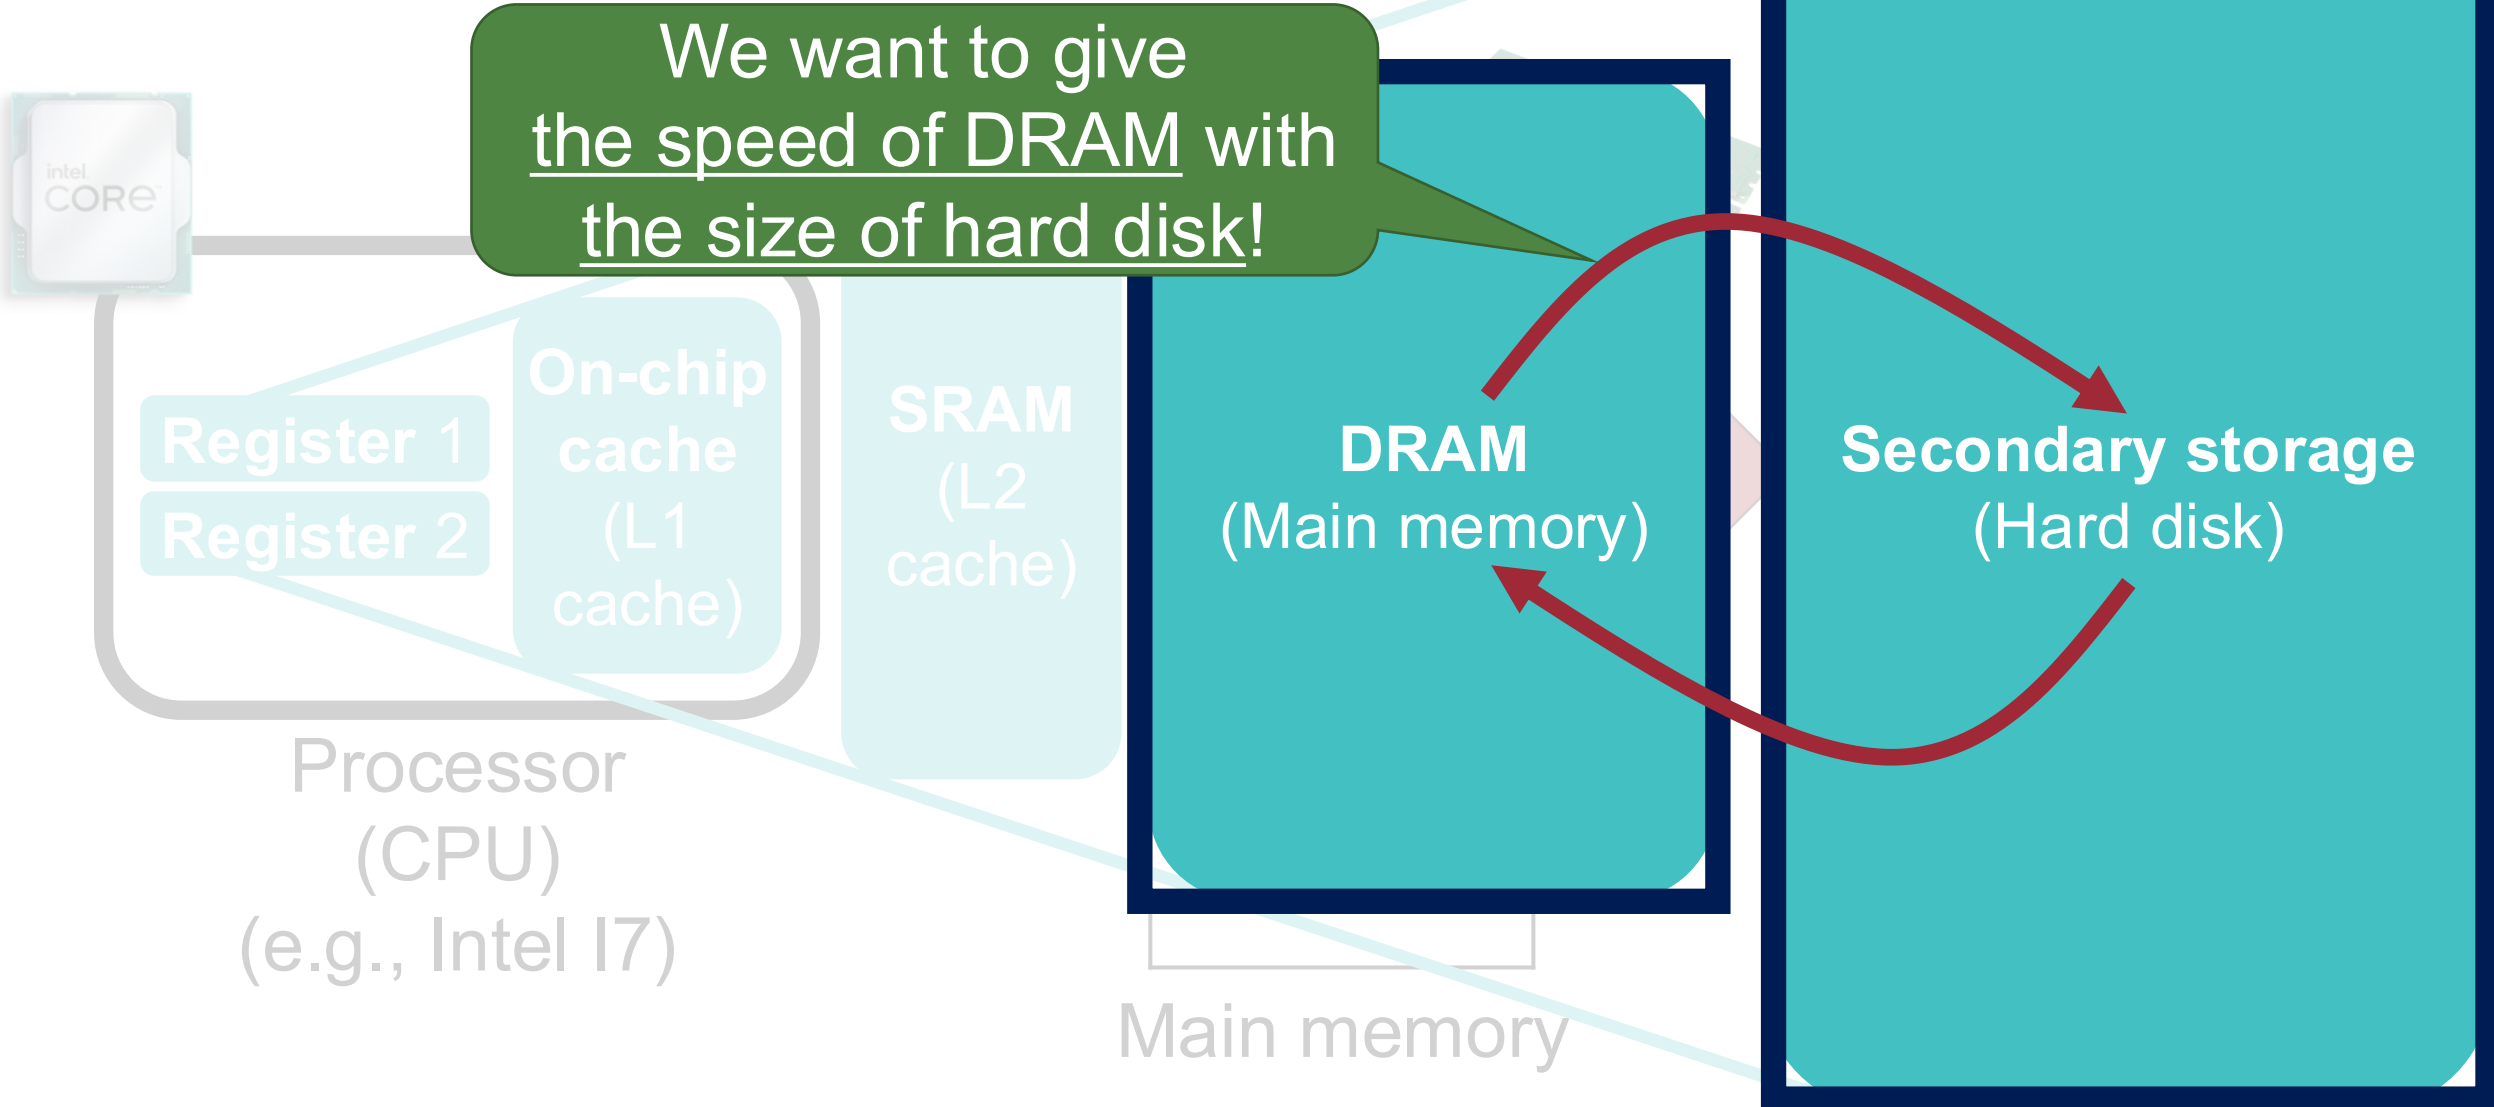


# Where Are We?



The principle of locality allows achieving the speed of cache with the size of DRAM

# Today's Topic: DRAM/HDD




# Introduction to Virtual Memory

# In the Physical World, We Have Limited-sized Memory <sup>8</sup>



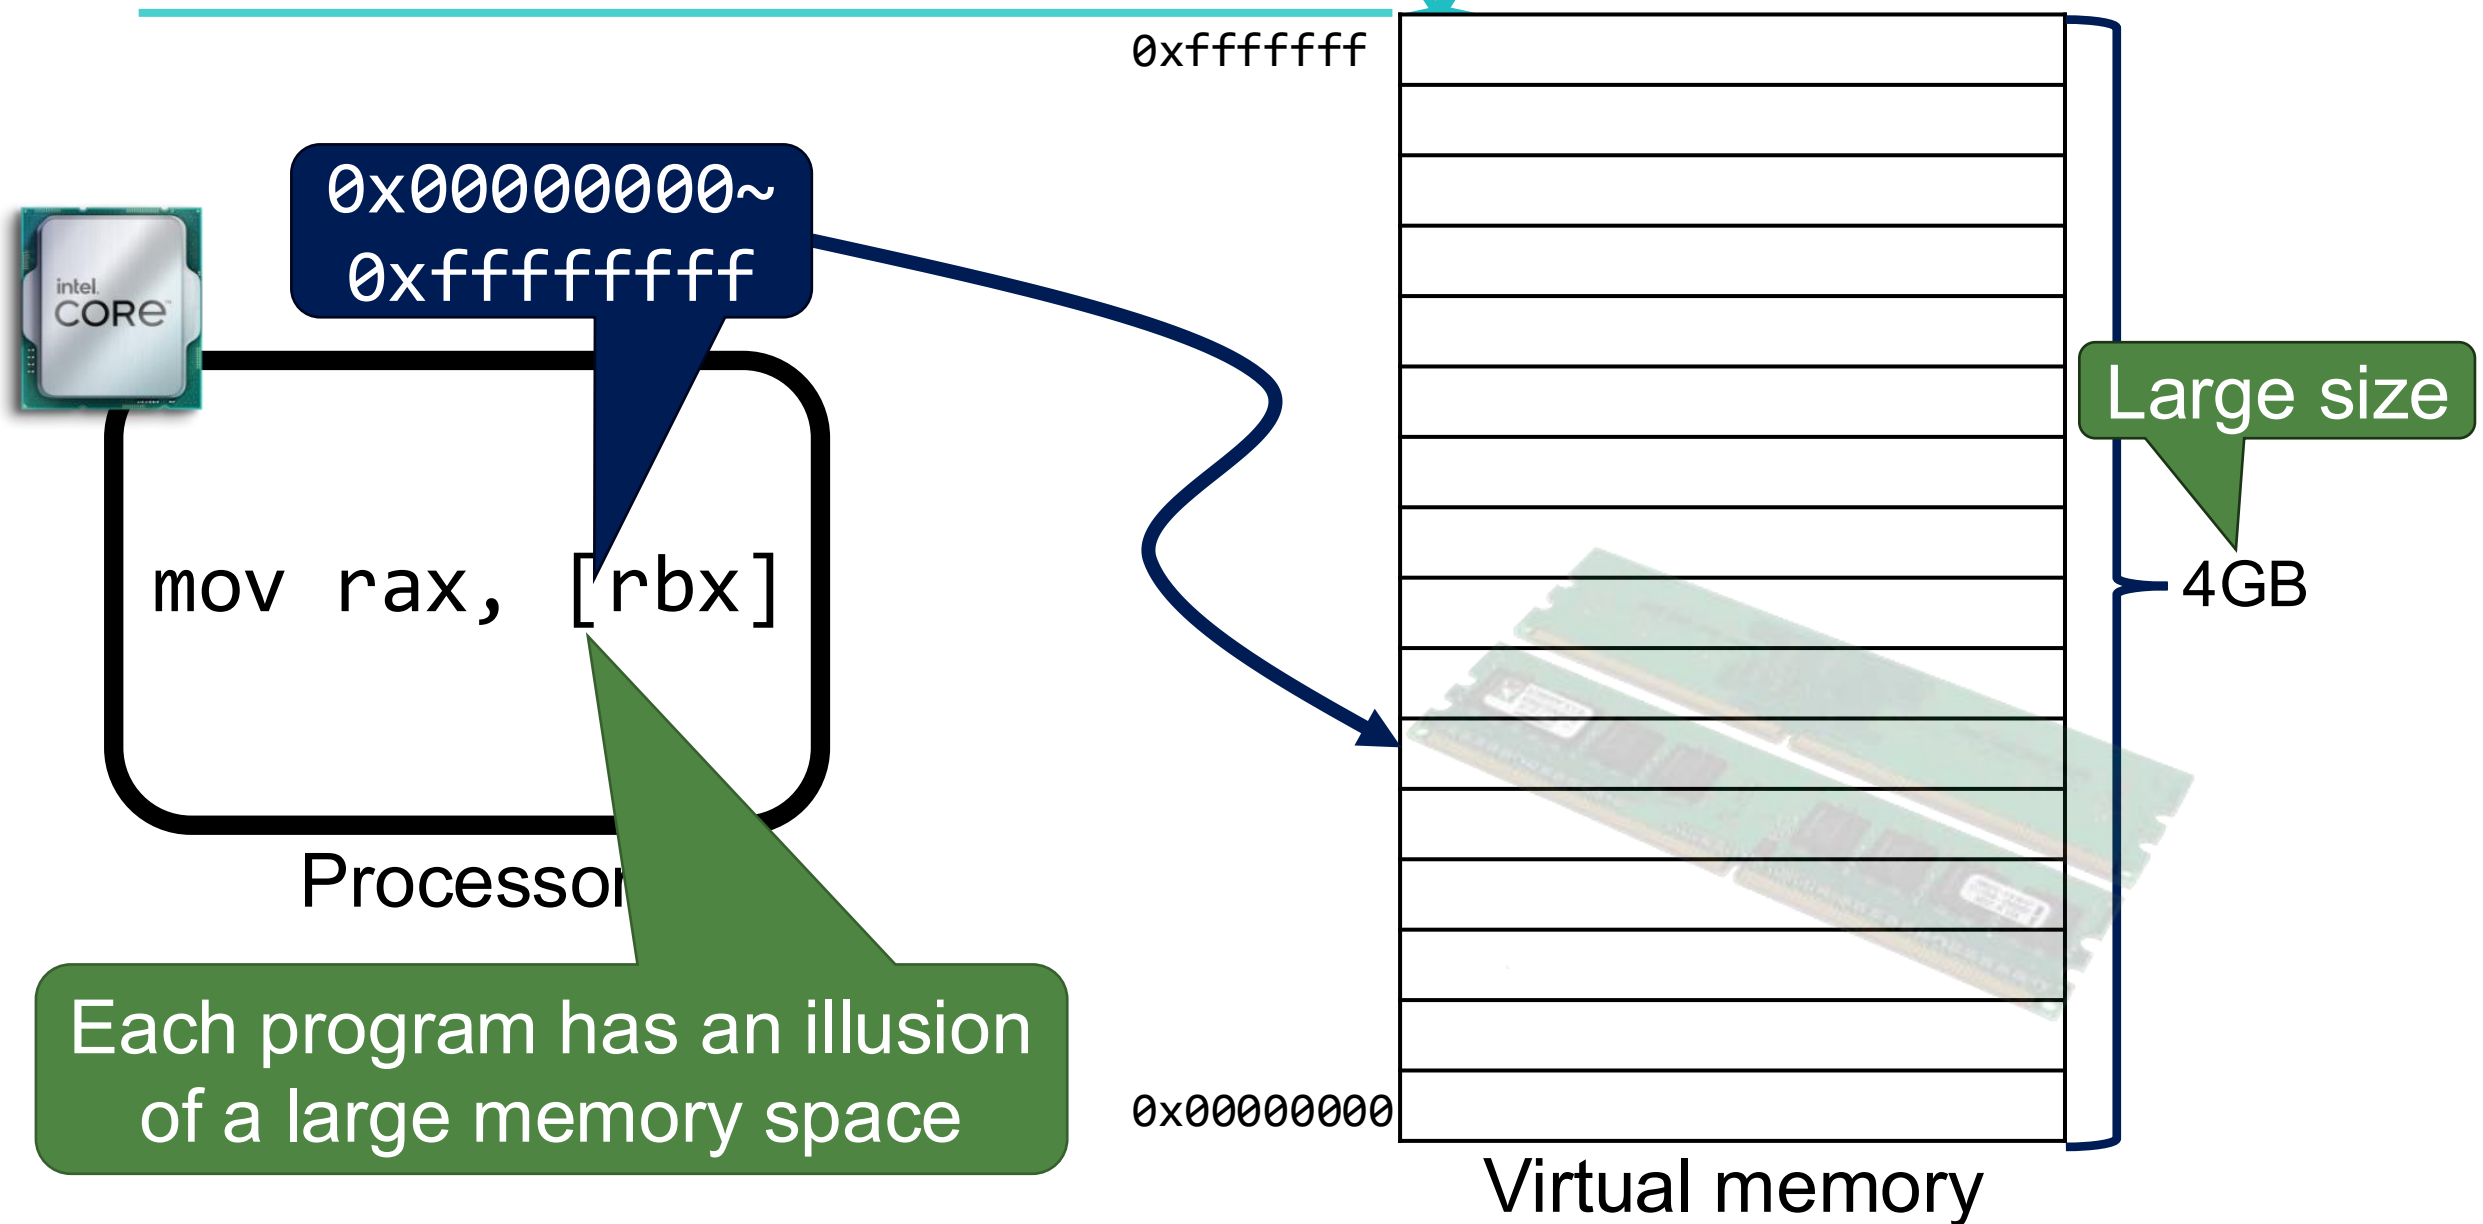
```
mov rax, [rbx]
```

Processor



Physical memory

# But, In the Virtual (Program) World, ...



# Virtual Memory

---



- Give programmers **an illusion of a large memory space** irrespective of actual capacity

# Virtual Memory

---



- Give programmers **an illusion of a large memory space** irrespective of actual capacity



*How do we give an illusion despite having small-sized physical memory?*

**Memory hierarchy between RAM and HDD**

# Memory Hierarchy: DRAM/HDD


12



0x00000000~  
0xffffffff

```
mov rax, [rbx]
```

Processor



DRAM  
(256MB)



Secondary storage  
(100GB)


Give the  
speed of the DRAM with  
the size of hard disk!

# Memory Hierarchy: DRAM/HDD



0x00000000~  
0xffffffff

Effective address  
(virtual address)



Processor

0xffffffff



0x00000000

Virtual memory




4GB  
Secondary storage  
(100GB)




# Memory Hi

Address translation table (OS-managed table)  
Virtual address → Physical address





Processor

Virtual addresses




Address translation

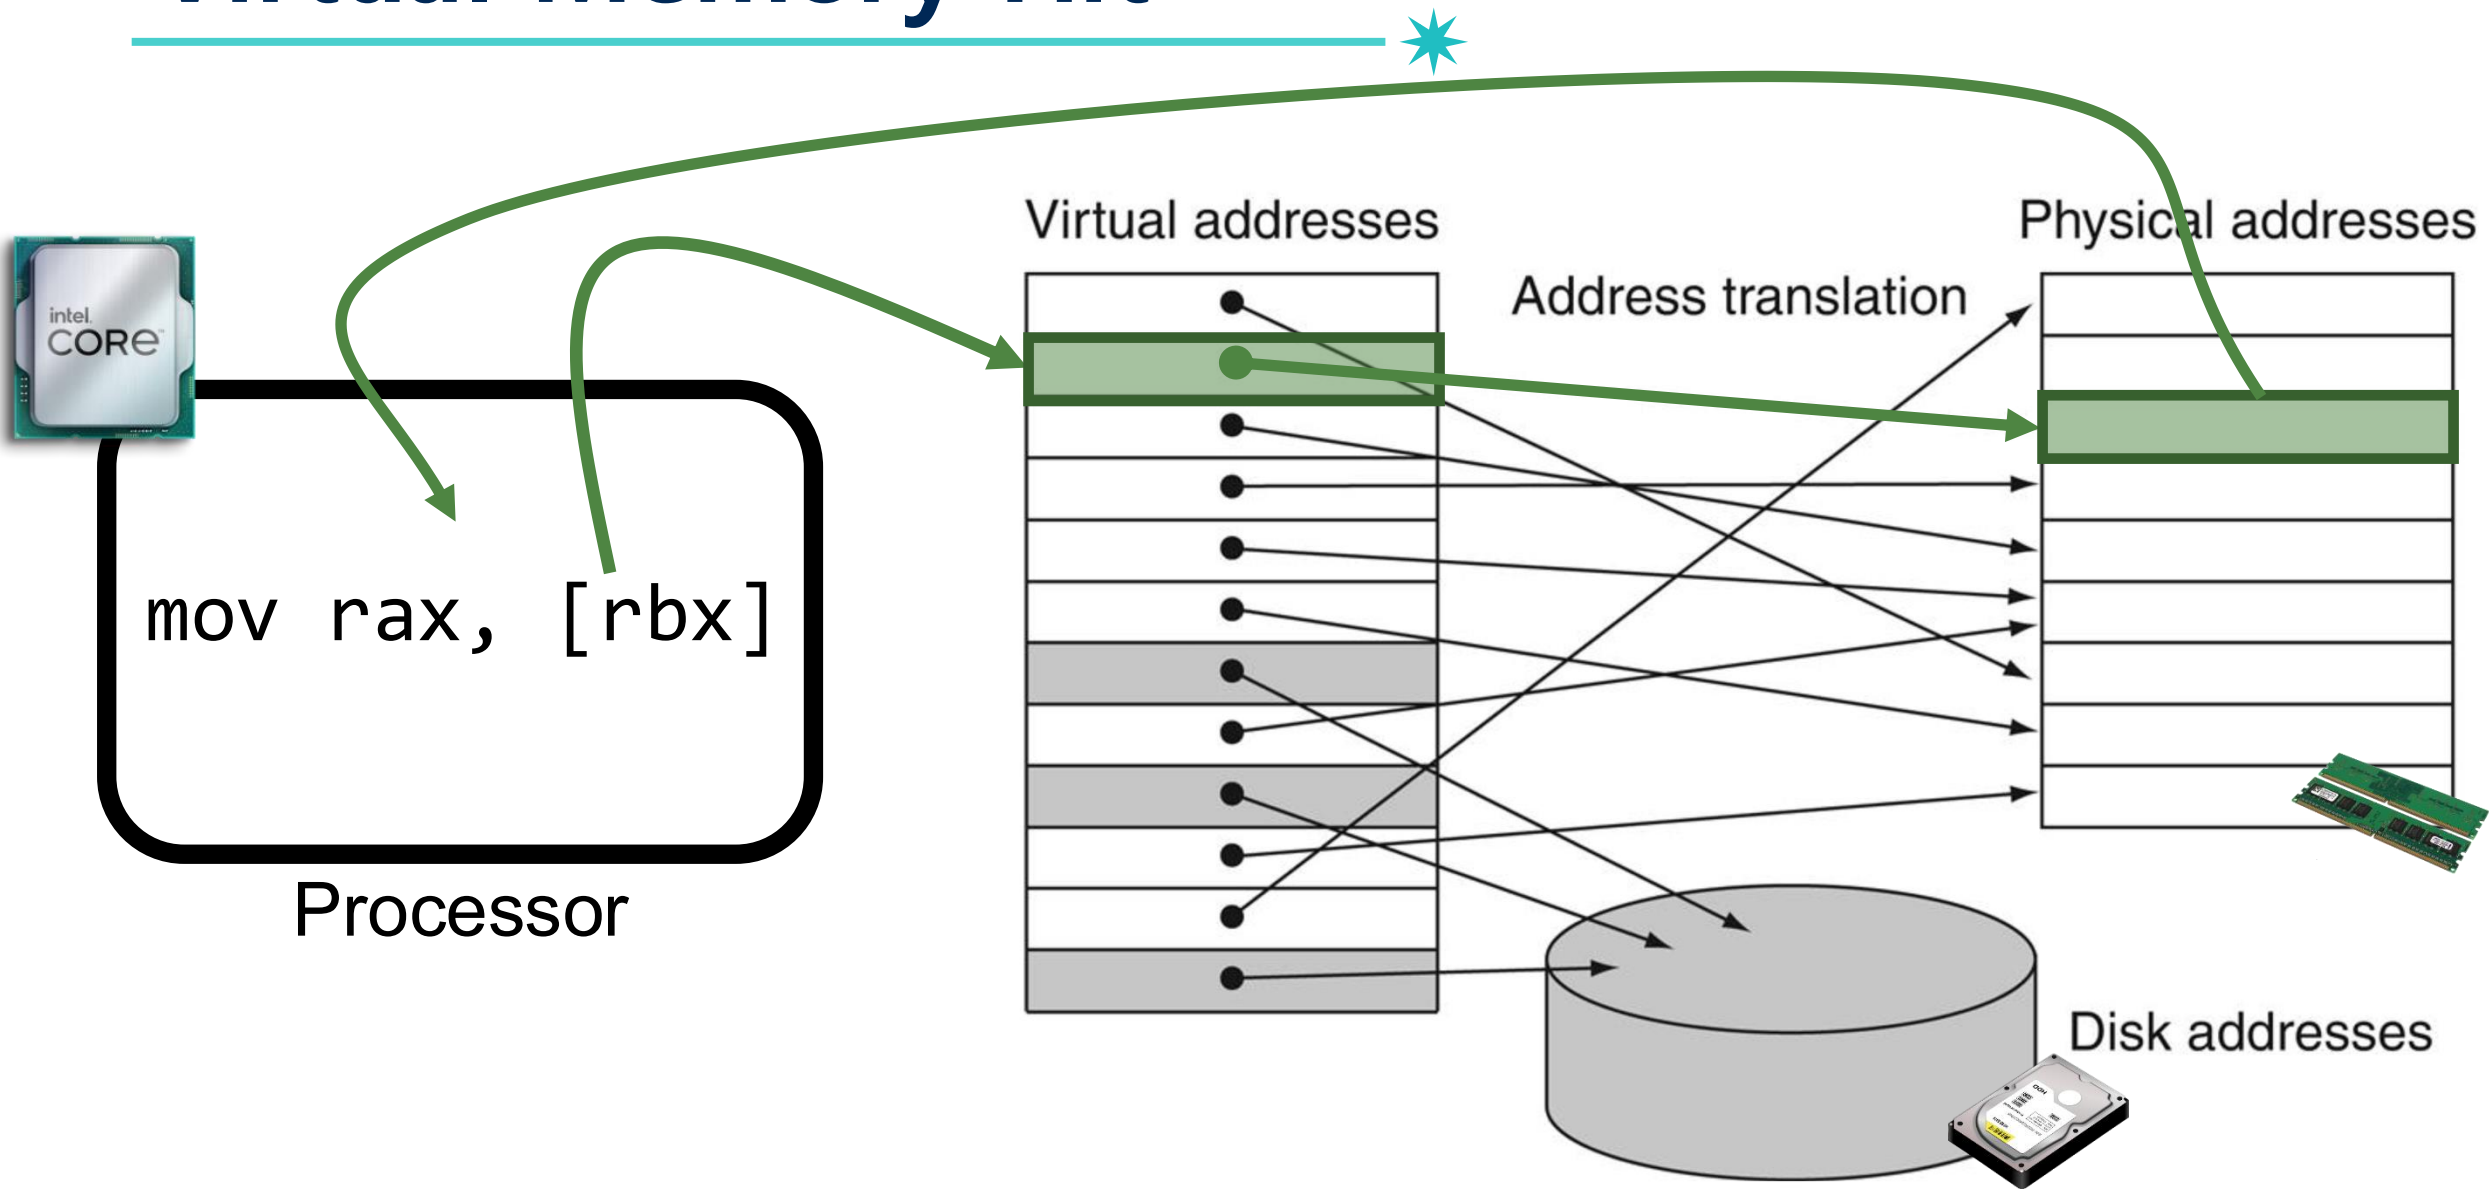
Physical addresses



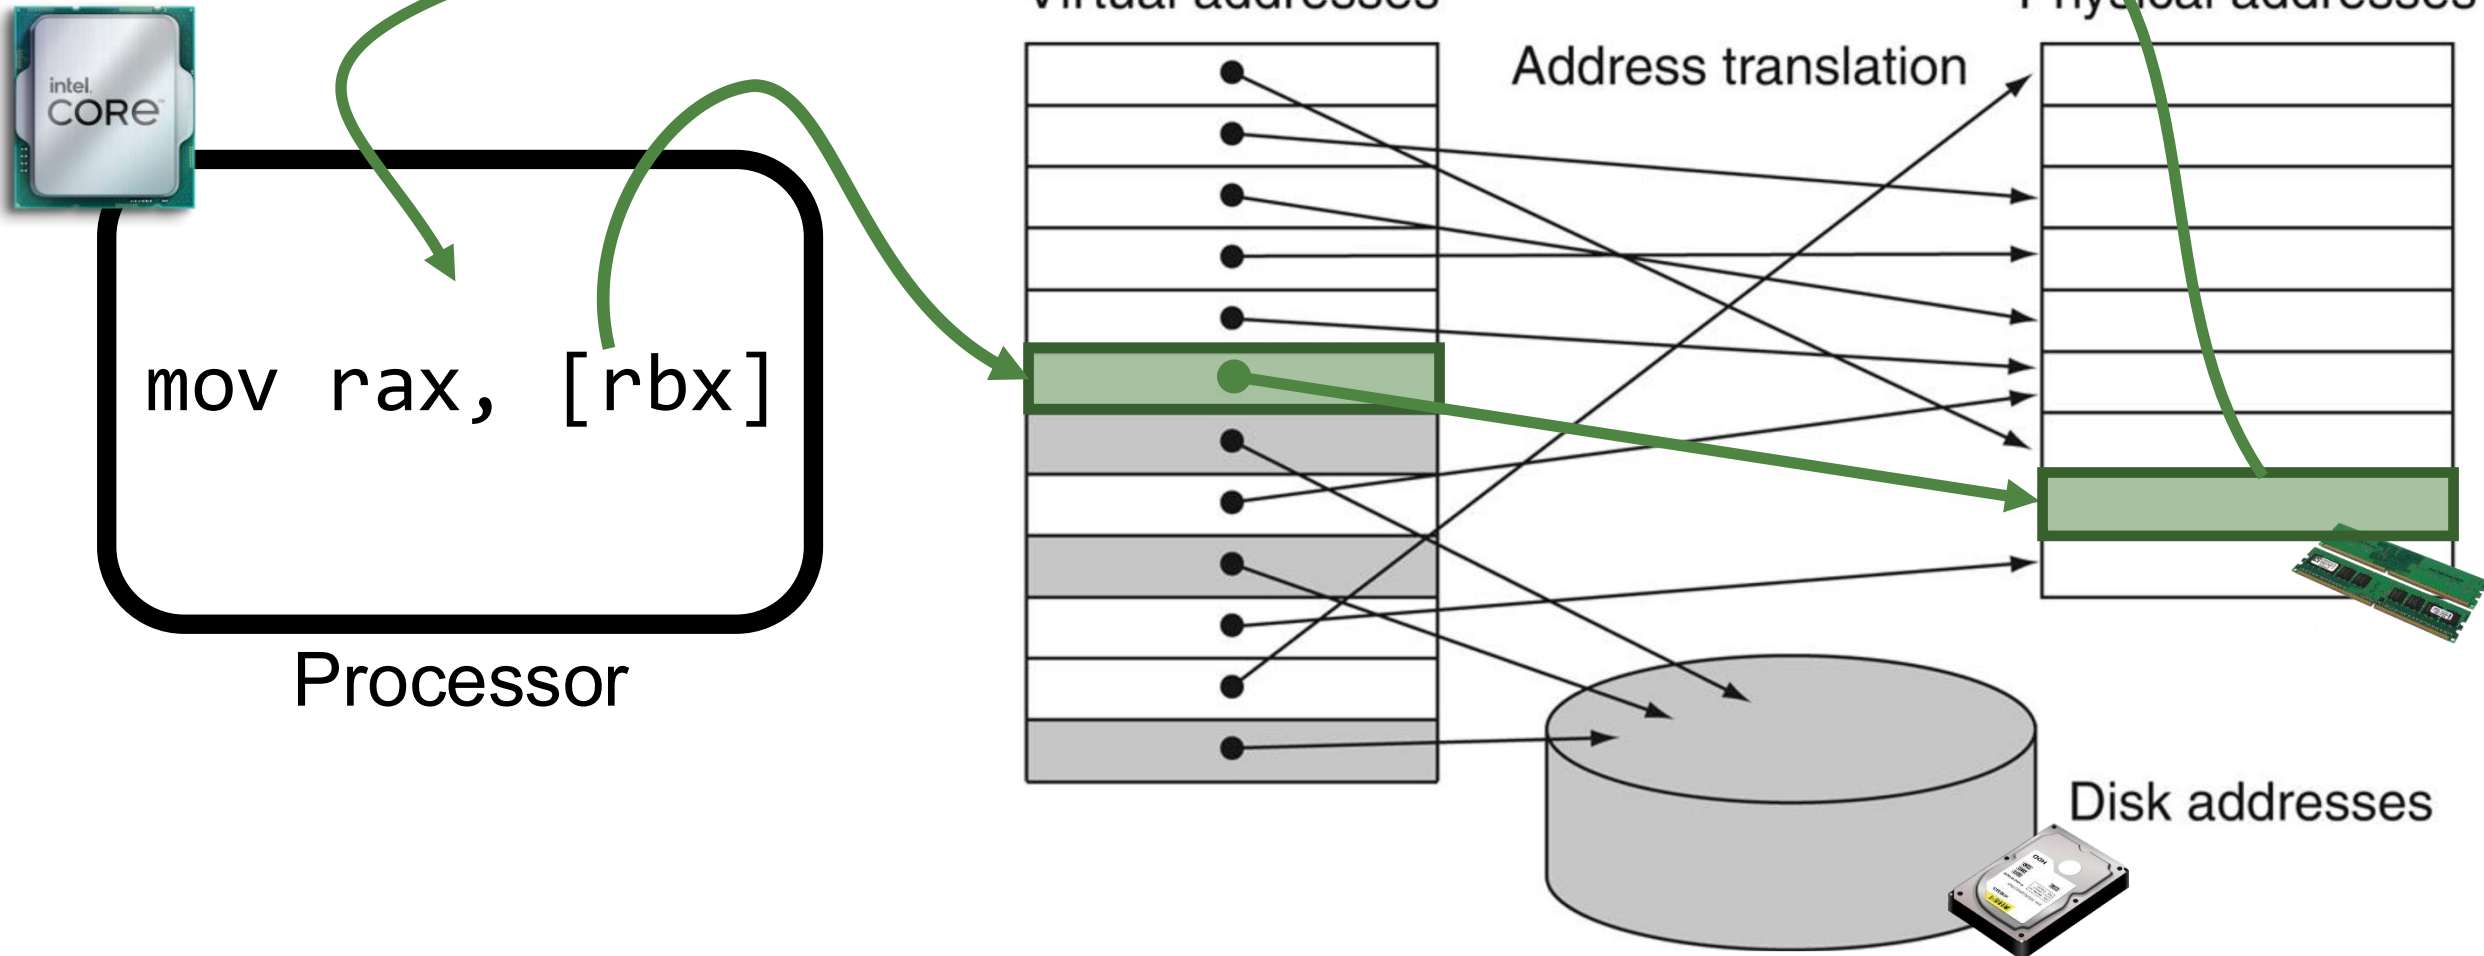
Disk addresses



# Virtual Memory Hit




# Virtual Memory Hit

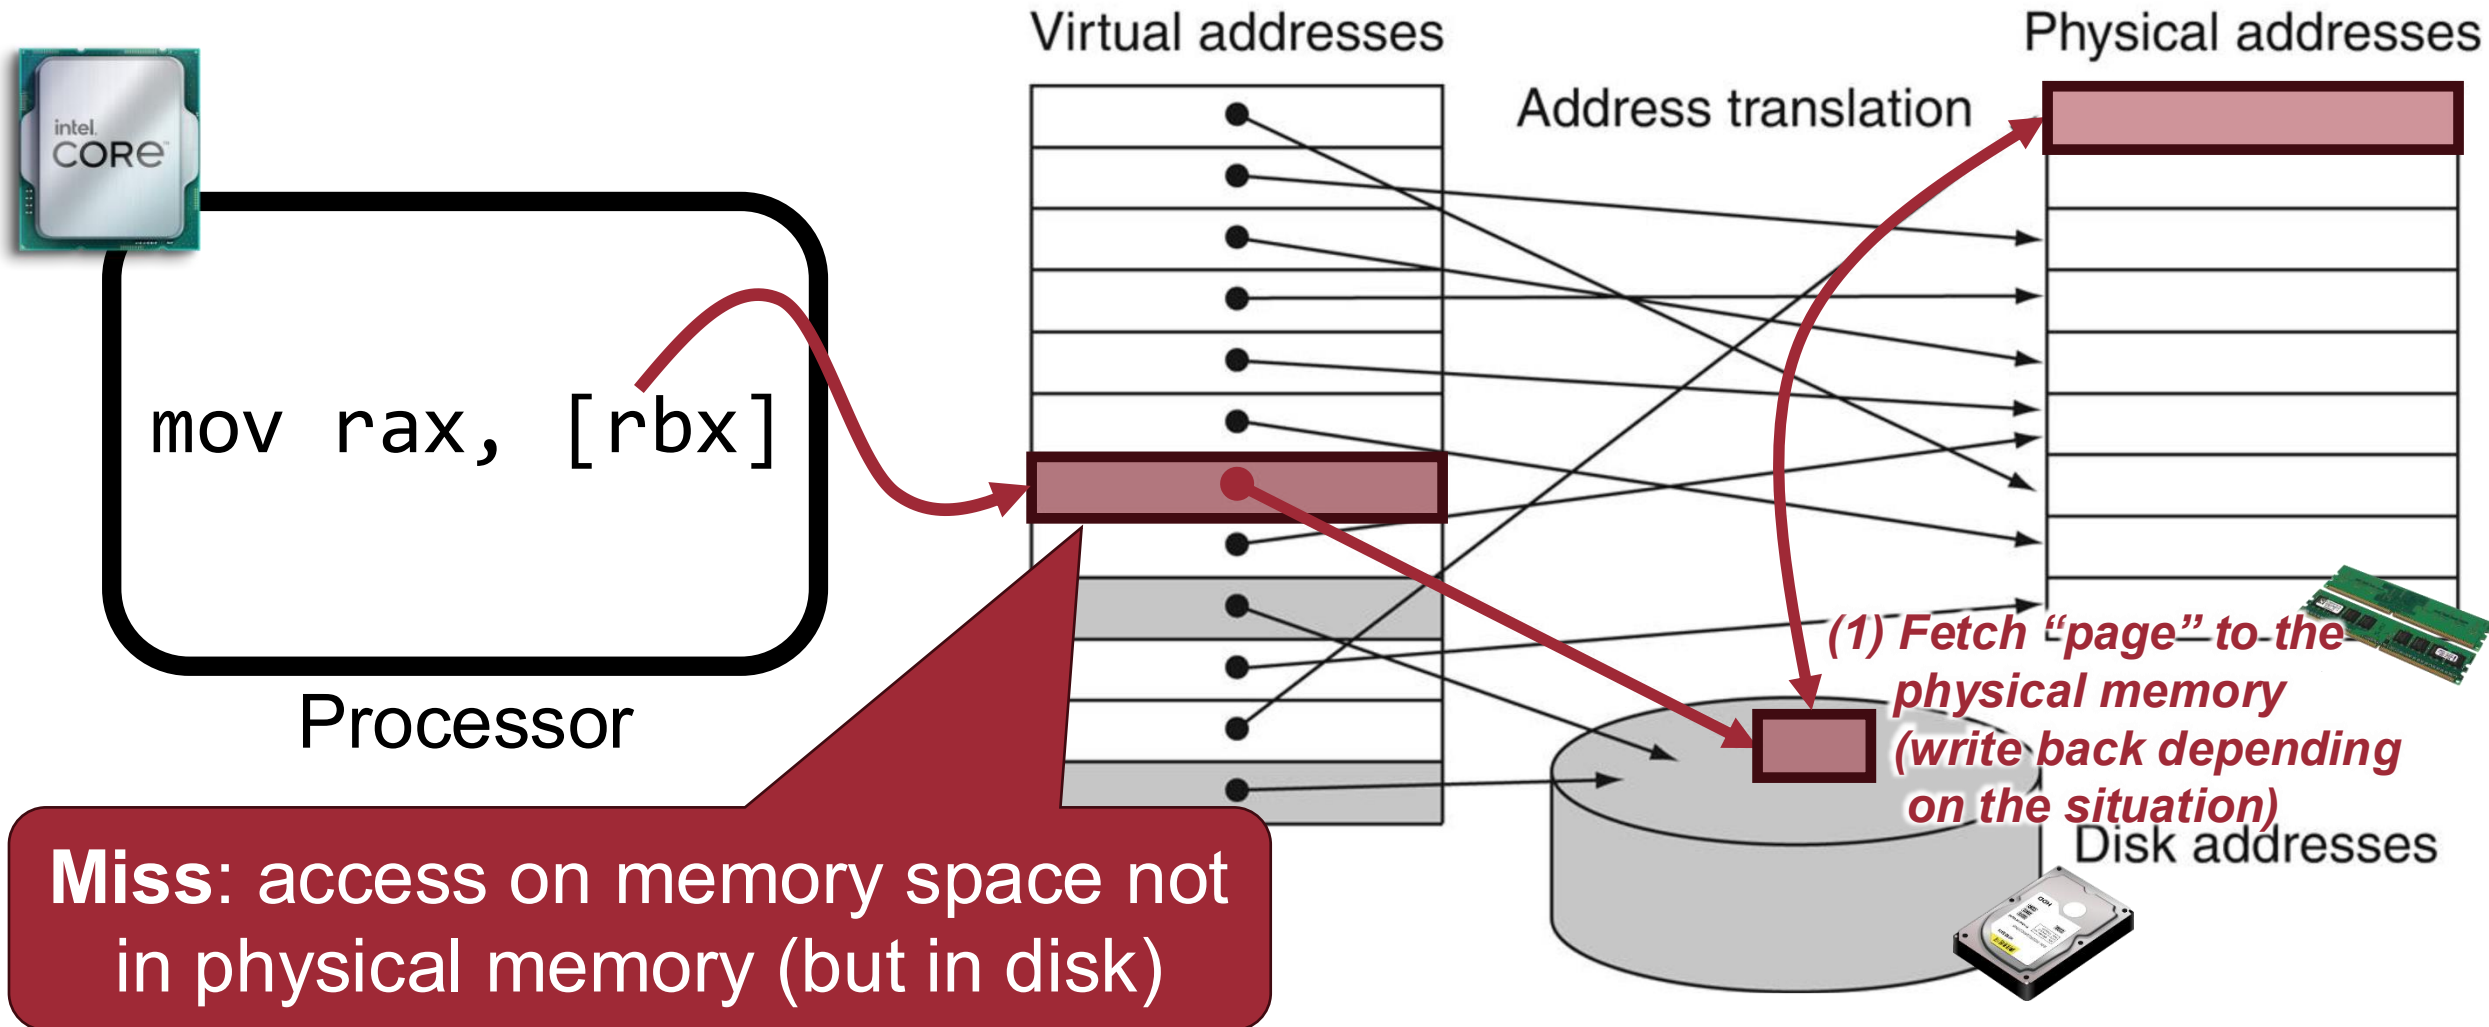




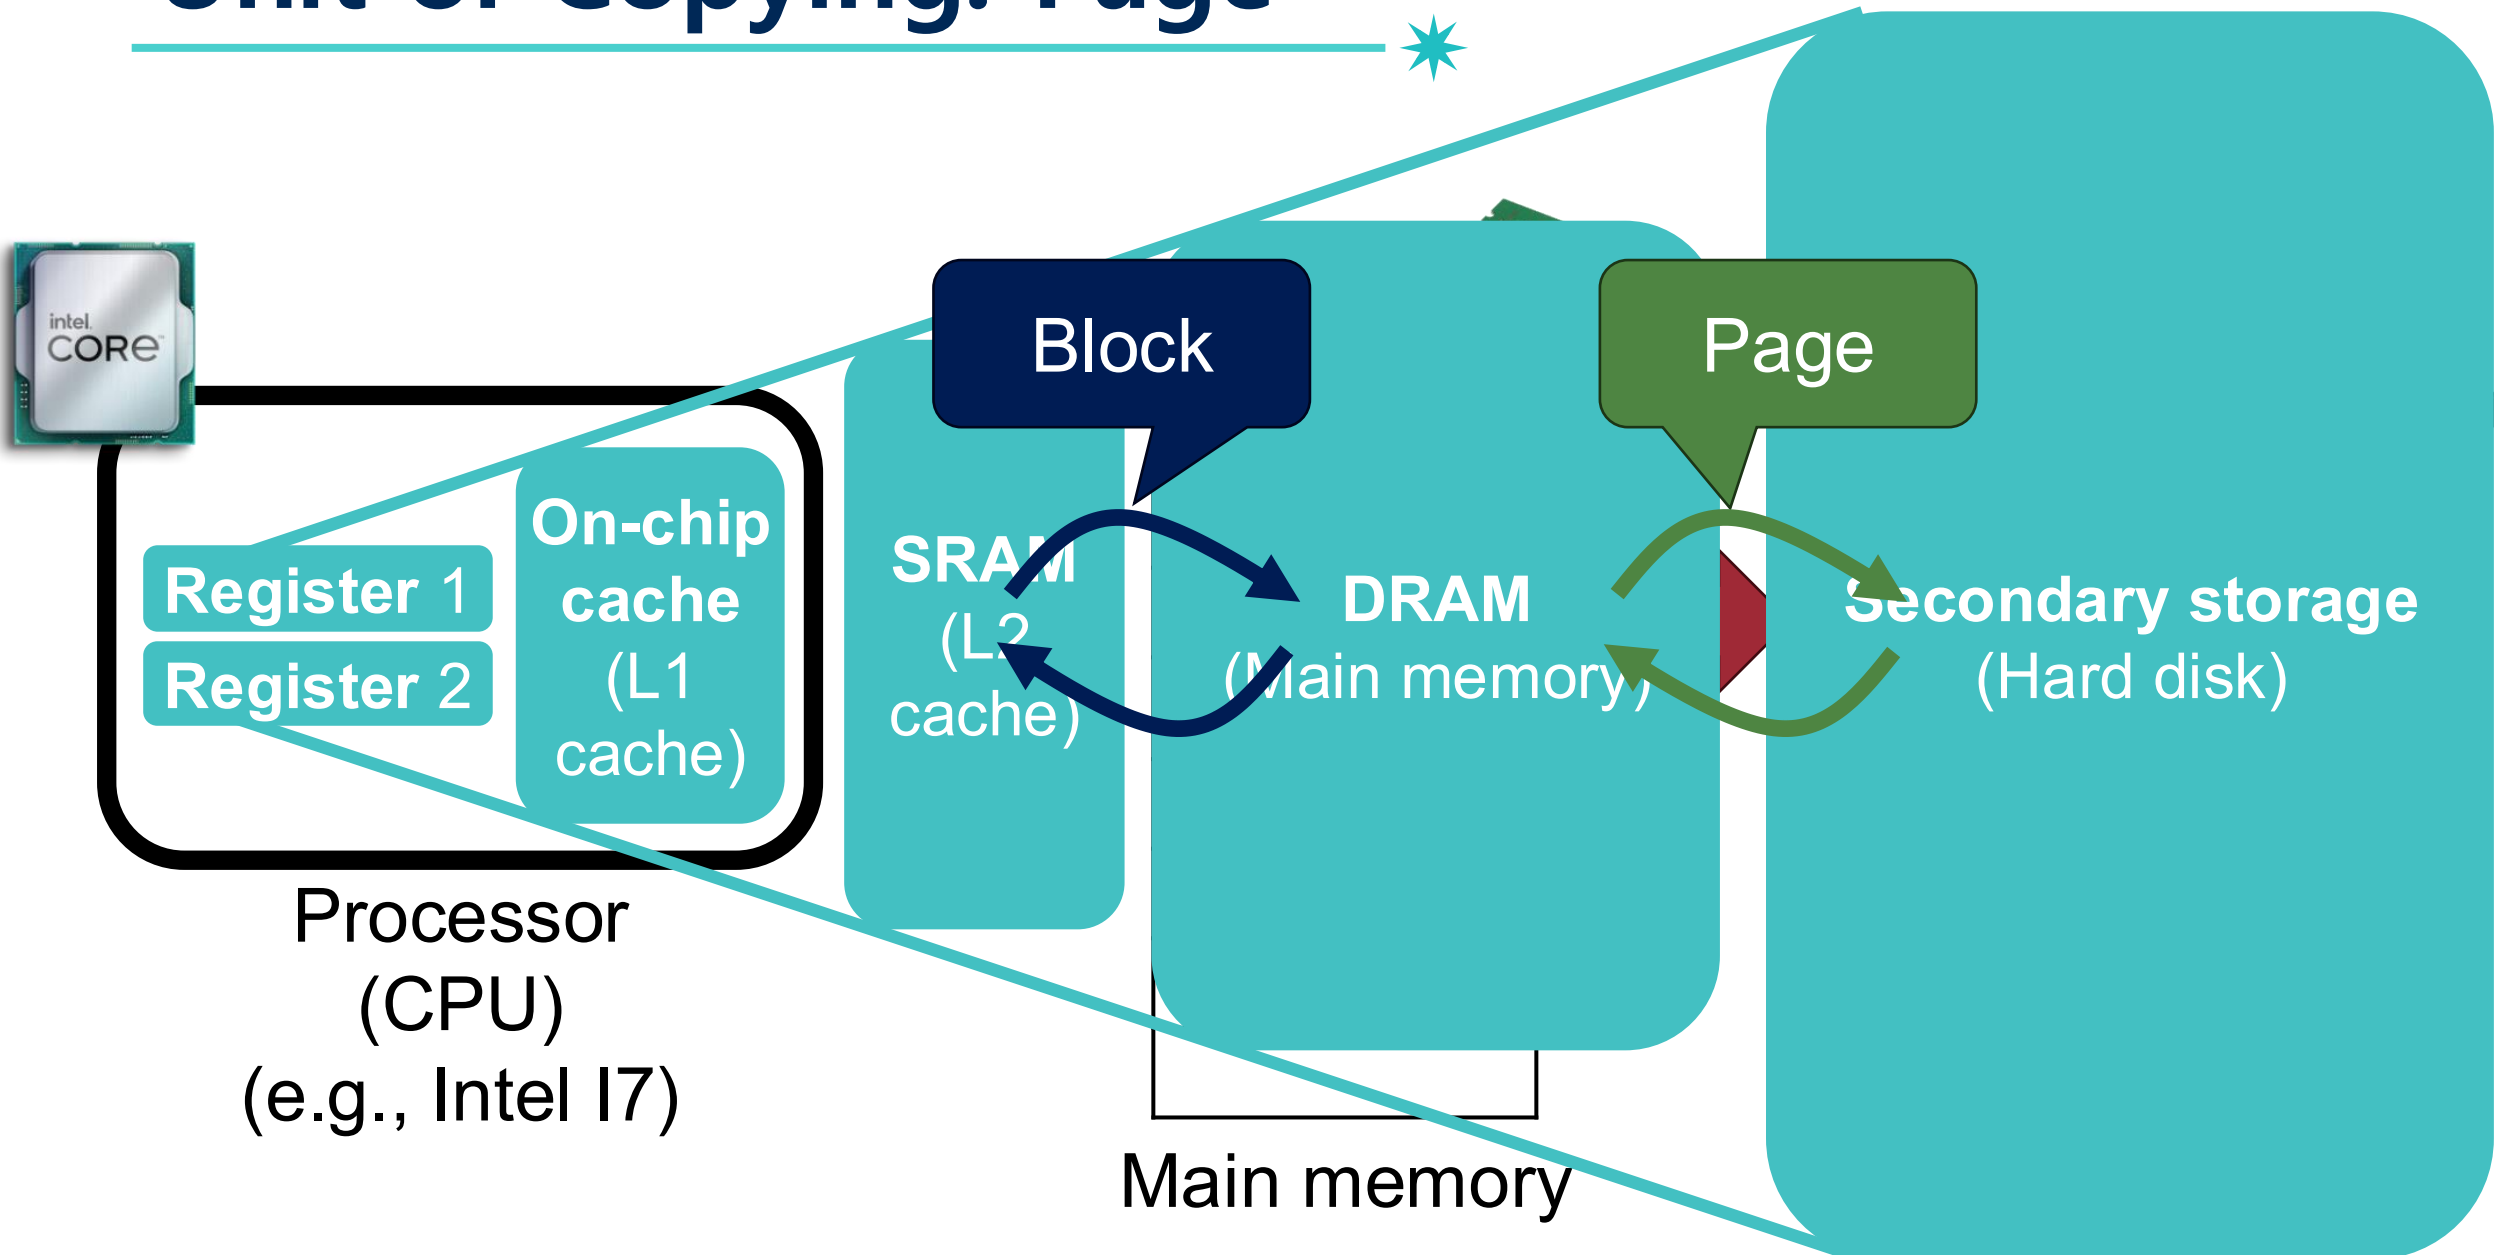
# Virtual Memory Miss (a.k.a., Page Fault)




# Virtual Memory Miss (a.k.a., Page Fault)



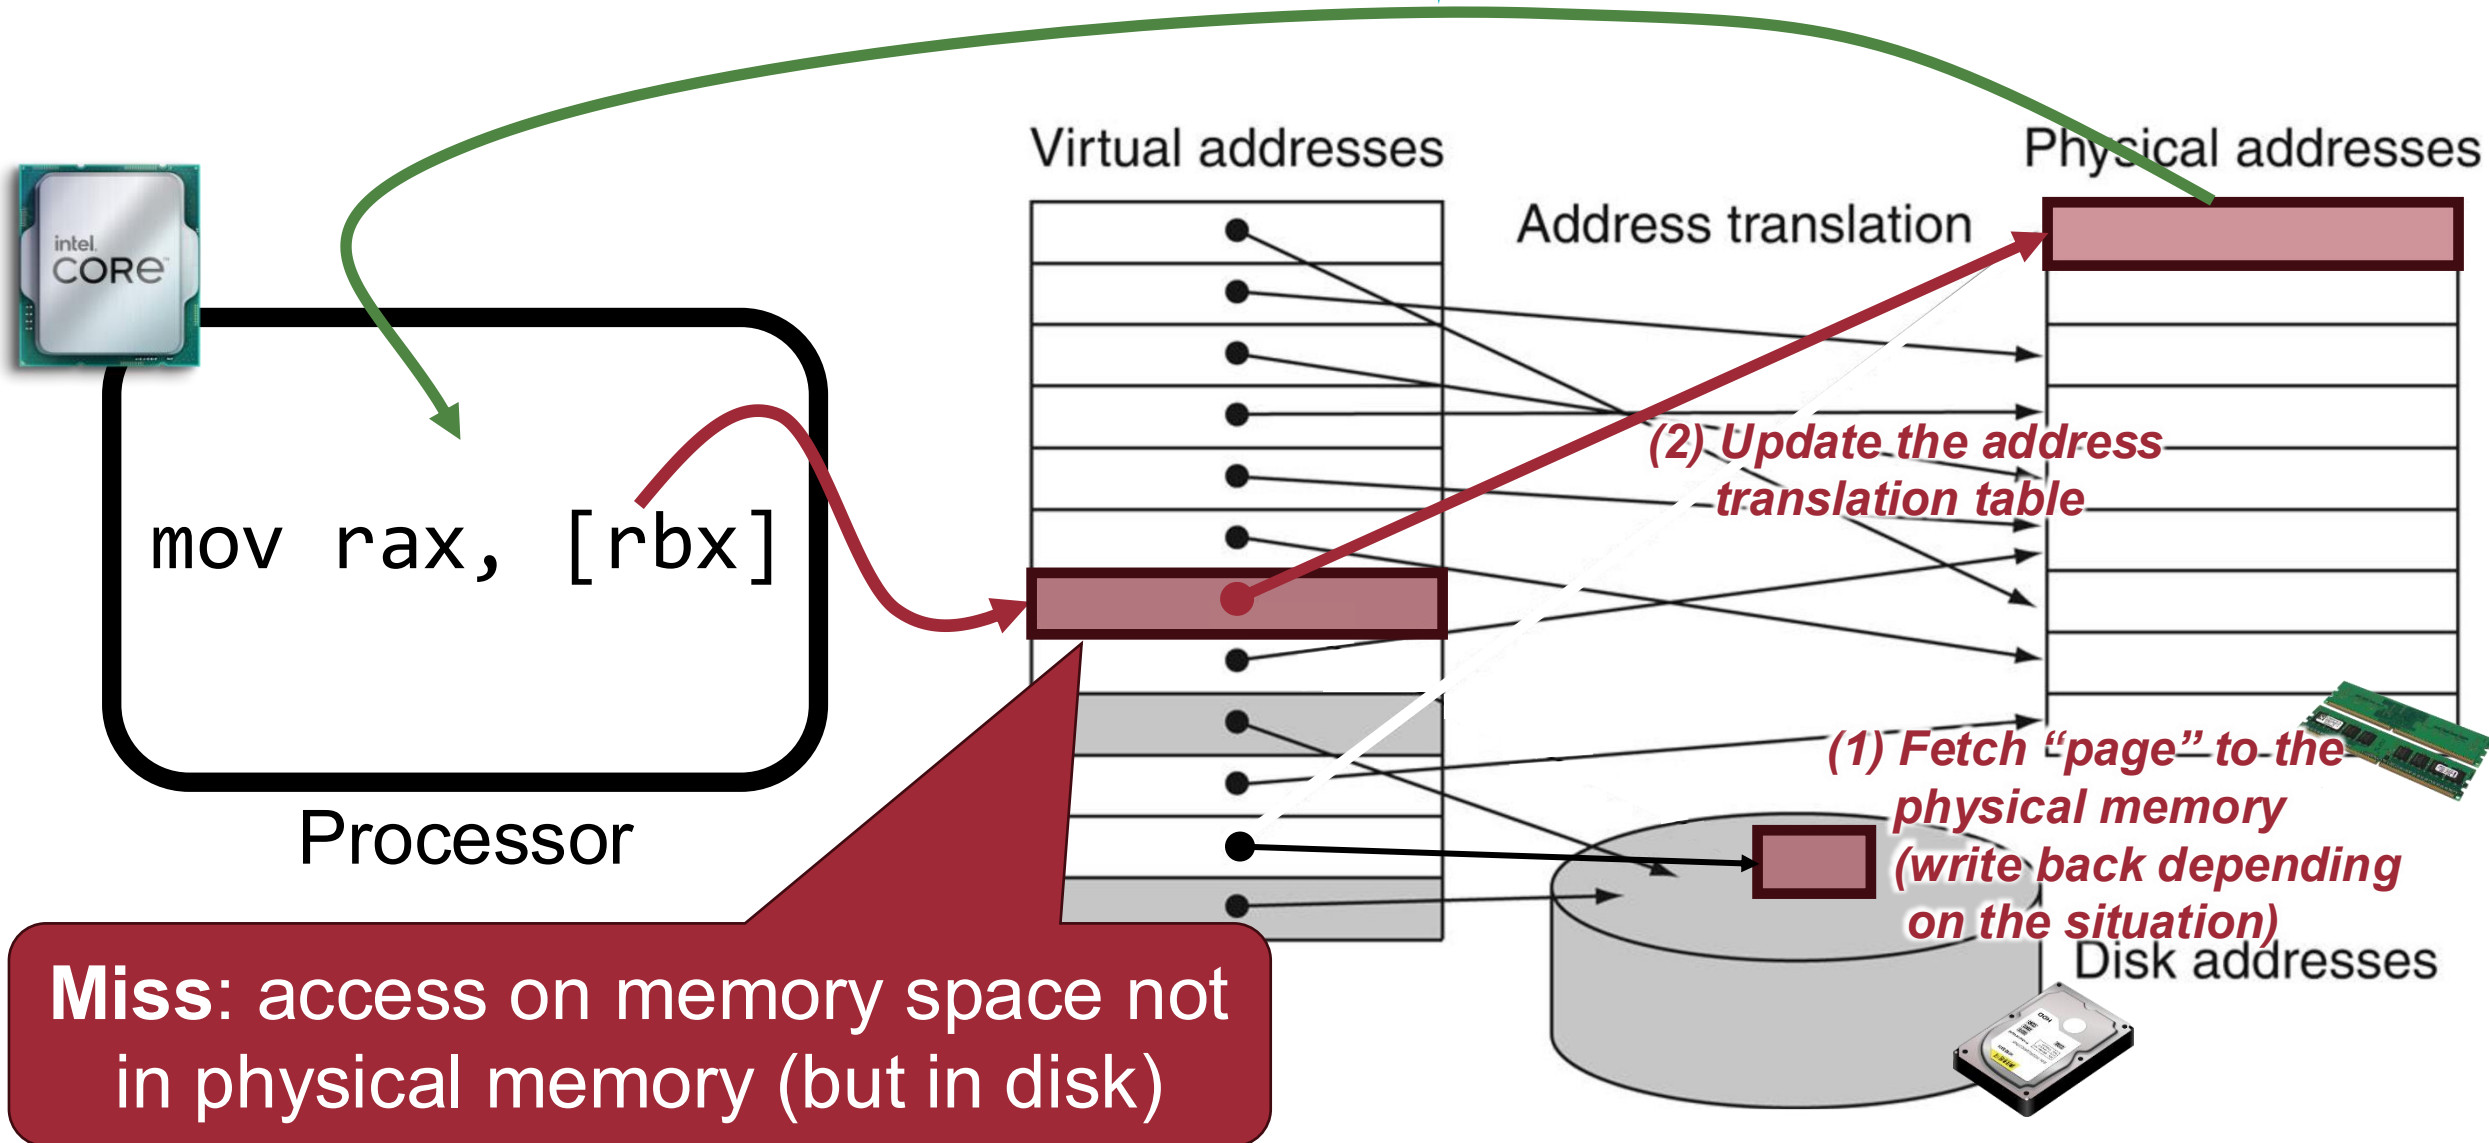
# Unit of Copying: Page



# Virtual Memory Miss (a.k.a., Page Fault)



# Virtual Memory Miss (a.k.a., Page Fault)



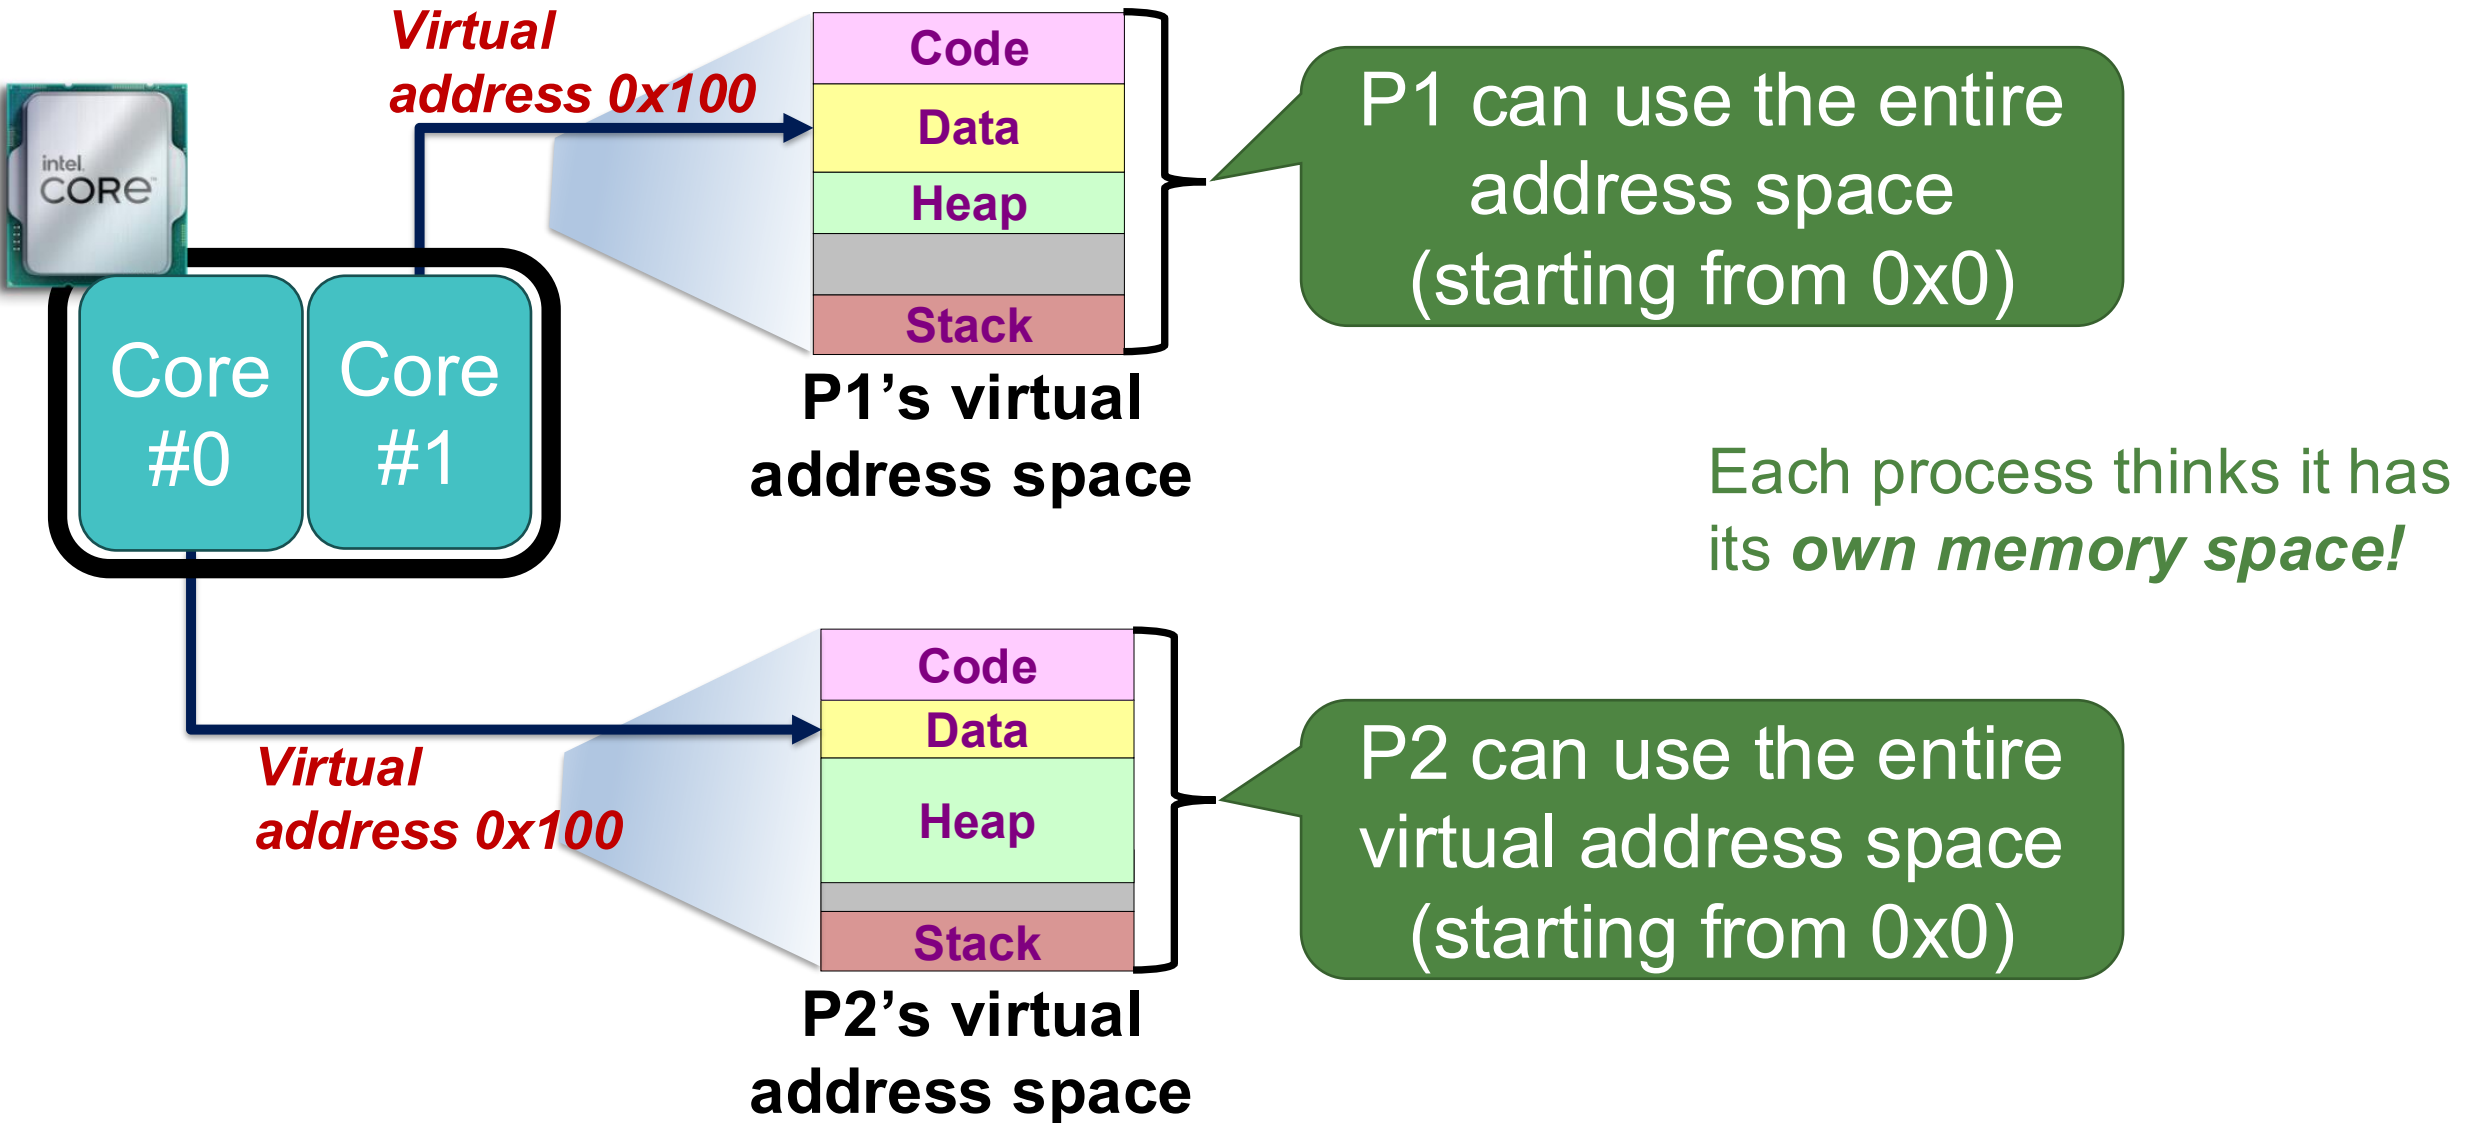
# Virtual Memory

---

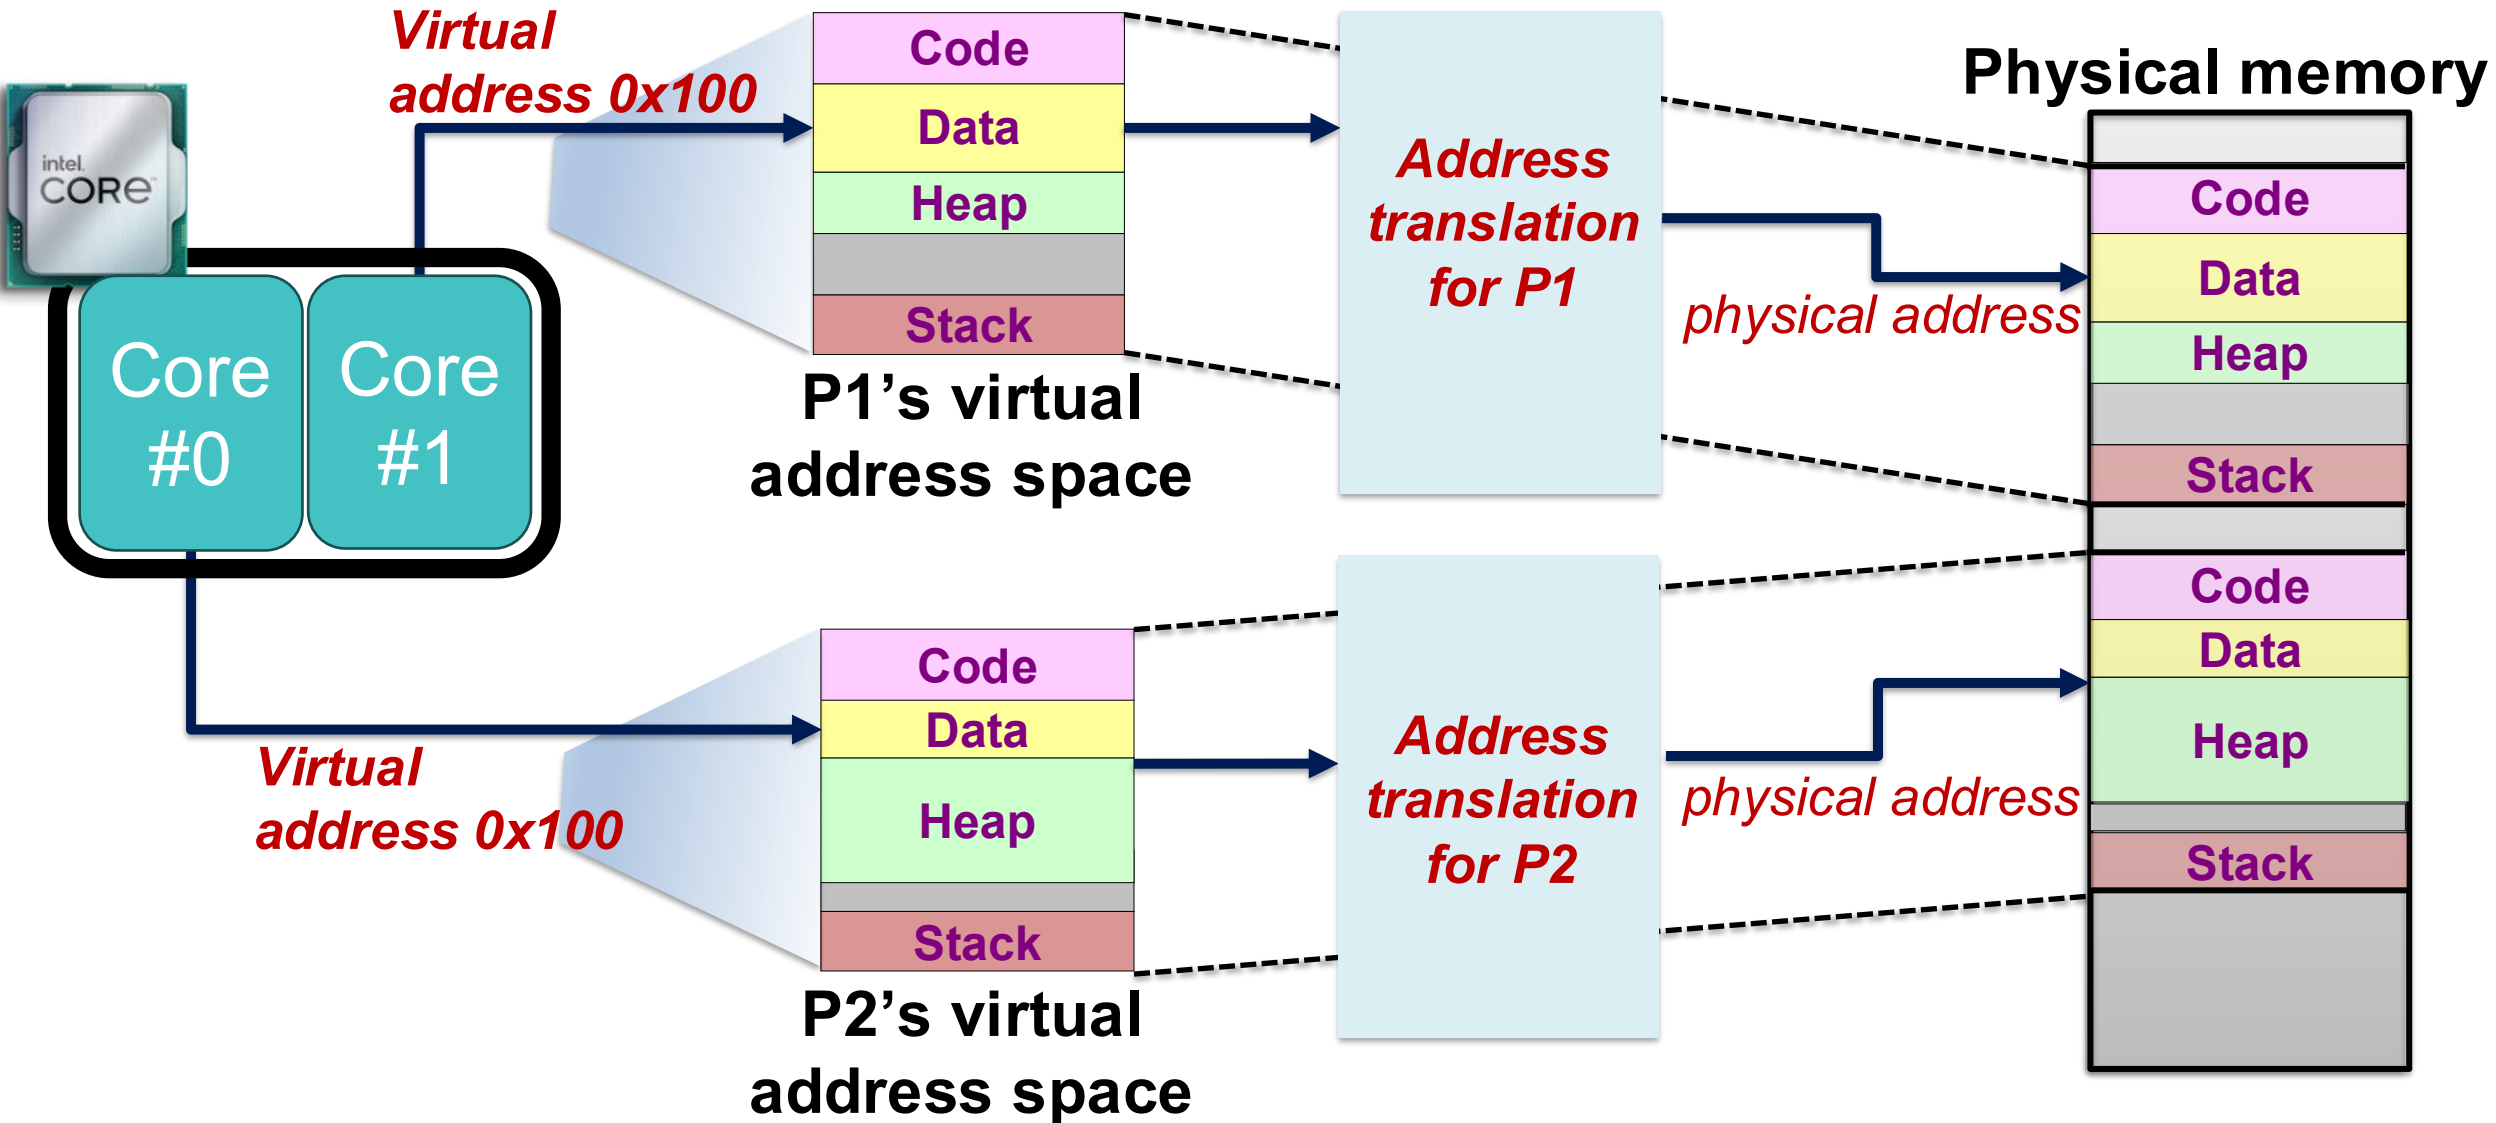


- Give programmers **an illusion of a large memory space** irrespective of actual capacity
- Uses main memory as a “cache” for the hard disk
- Other benefits?
  - **Program simplification**: give each program the illusion that it has its *own private memory* (e.g., 0x00000000~0xffffffff)

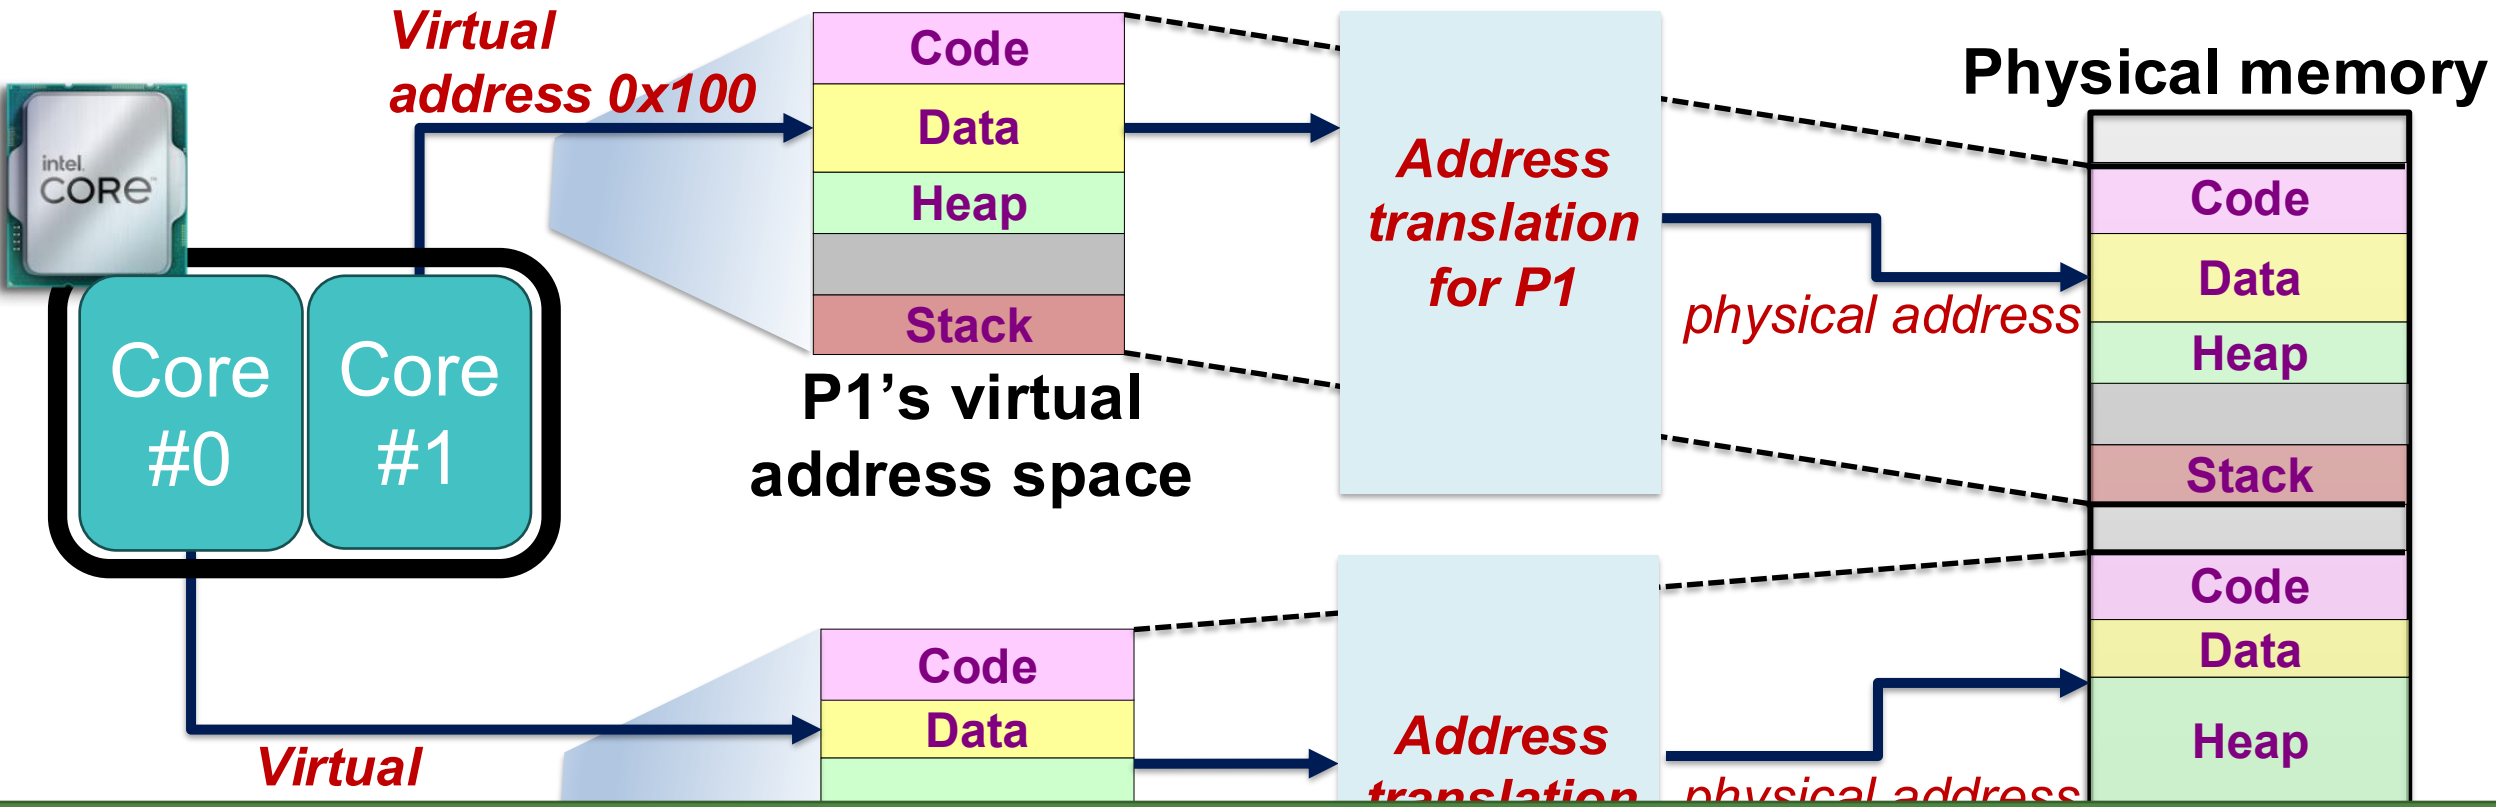
# Program Simplification



# Program Simplification



# Program Simplification



Programming and storage management ease

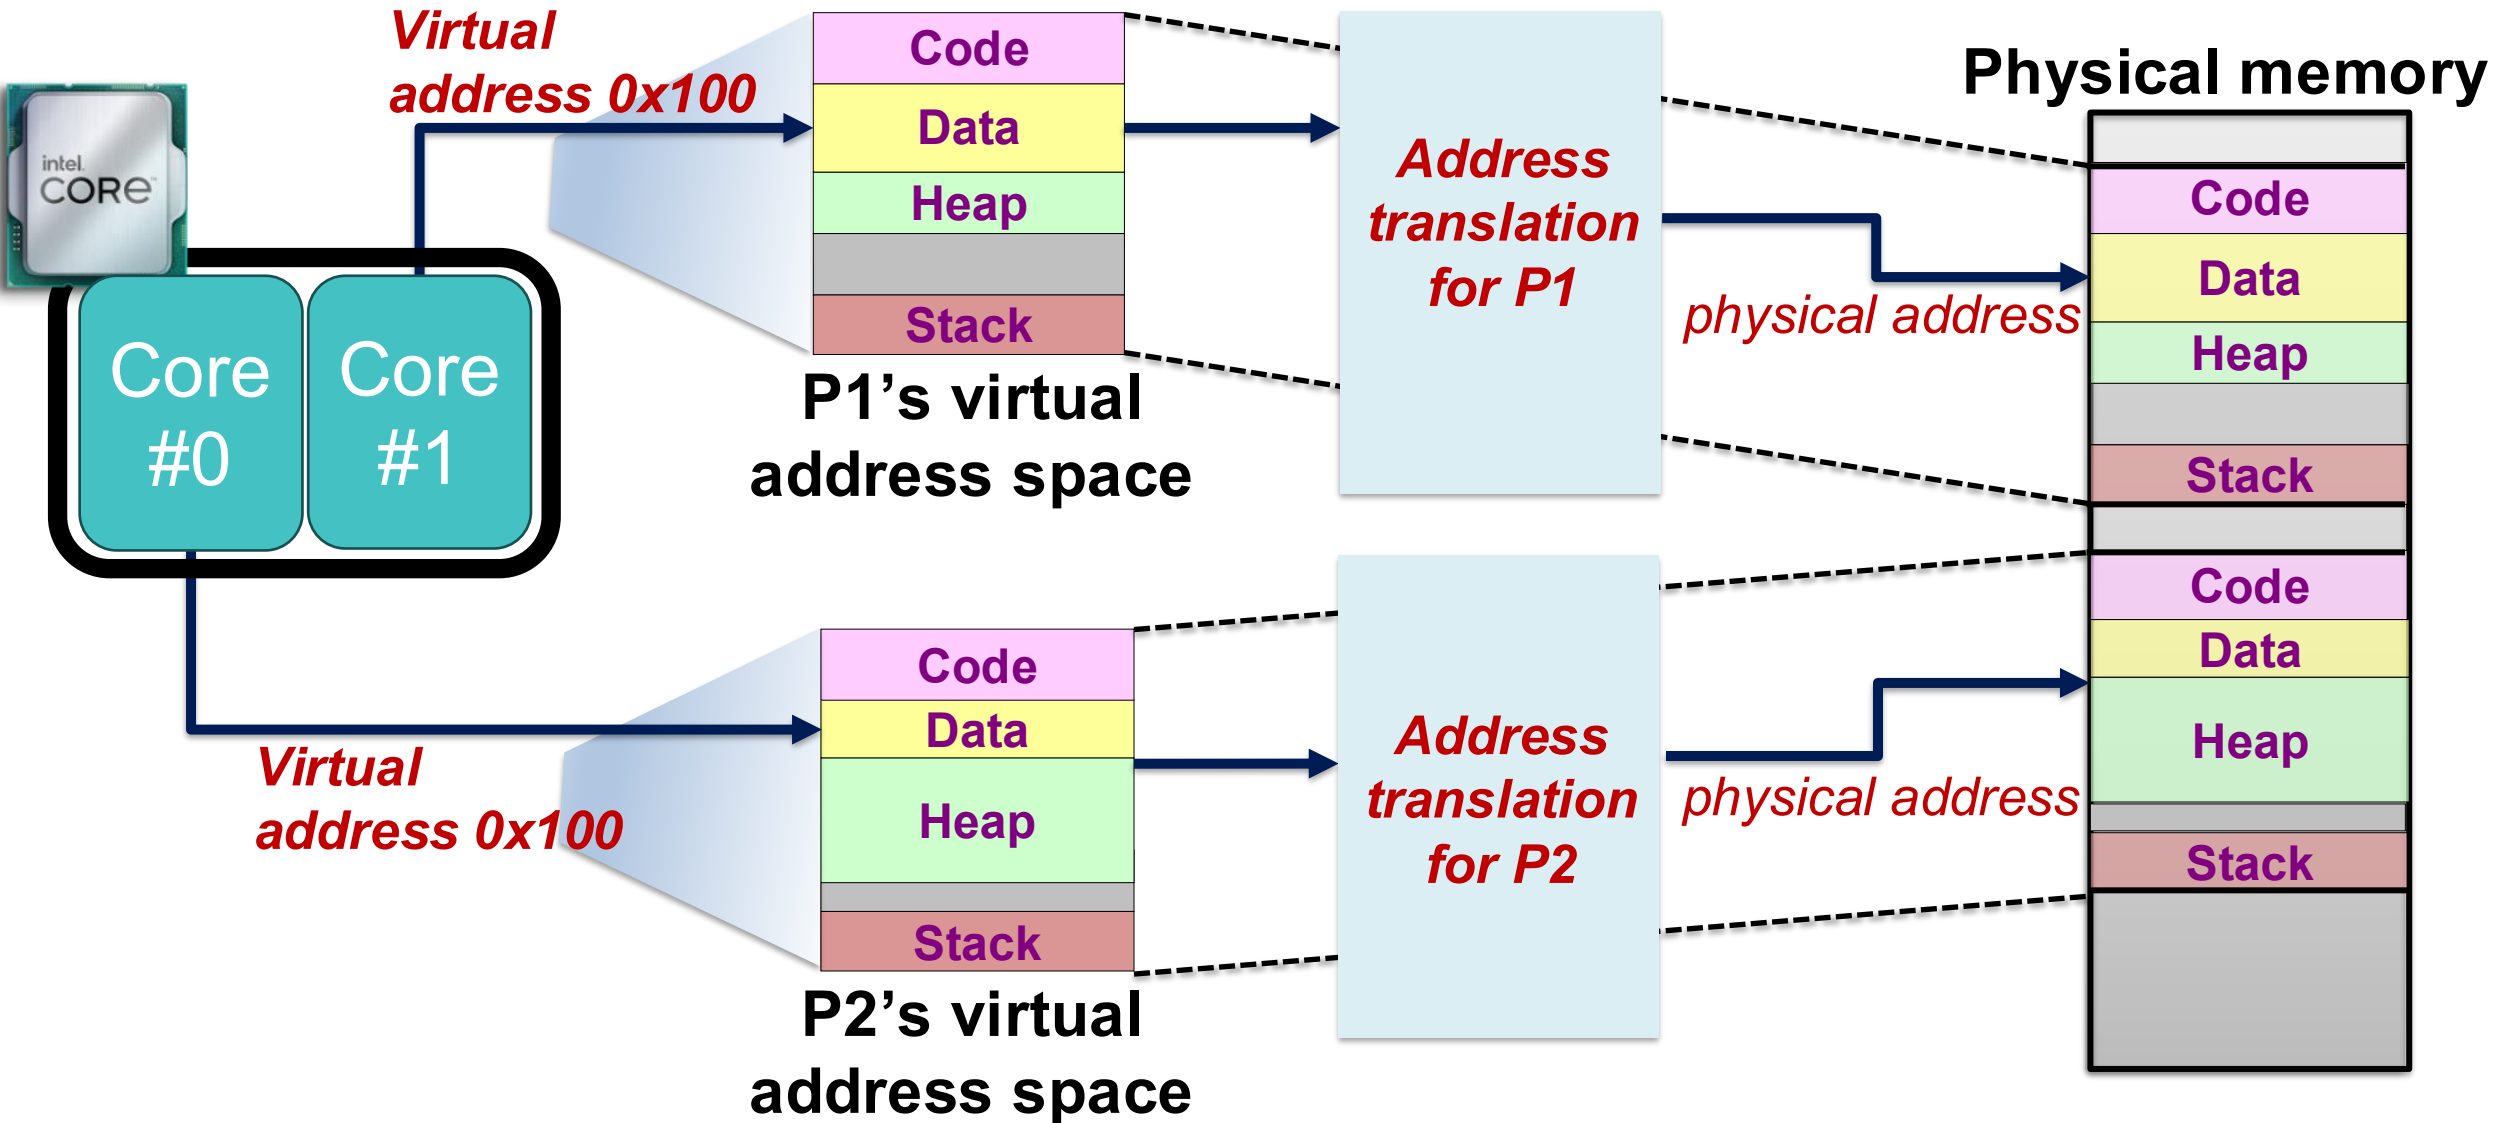
# Virtual Memory

---



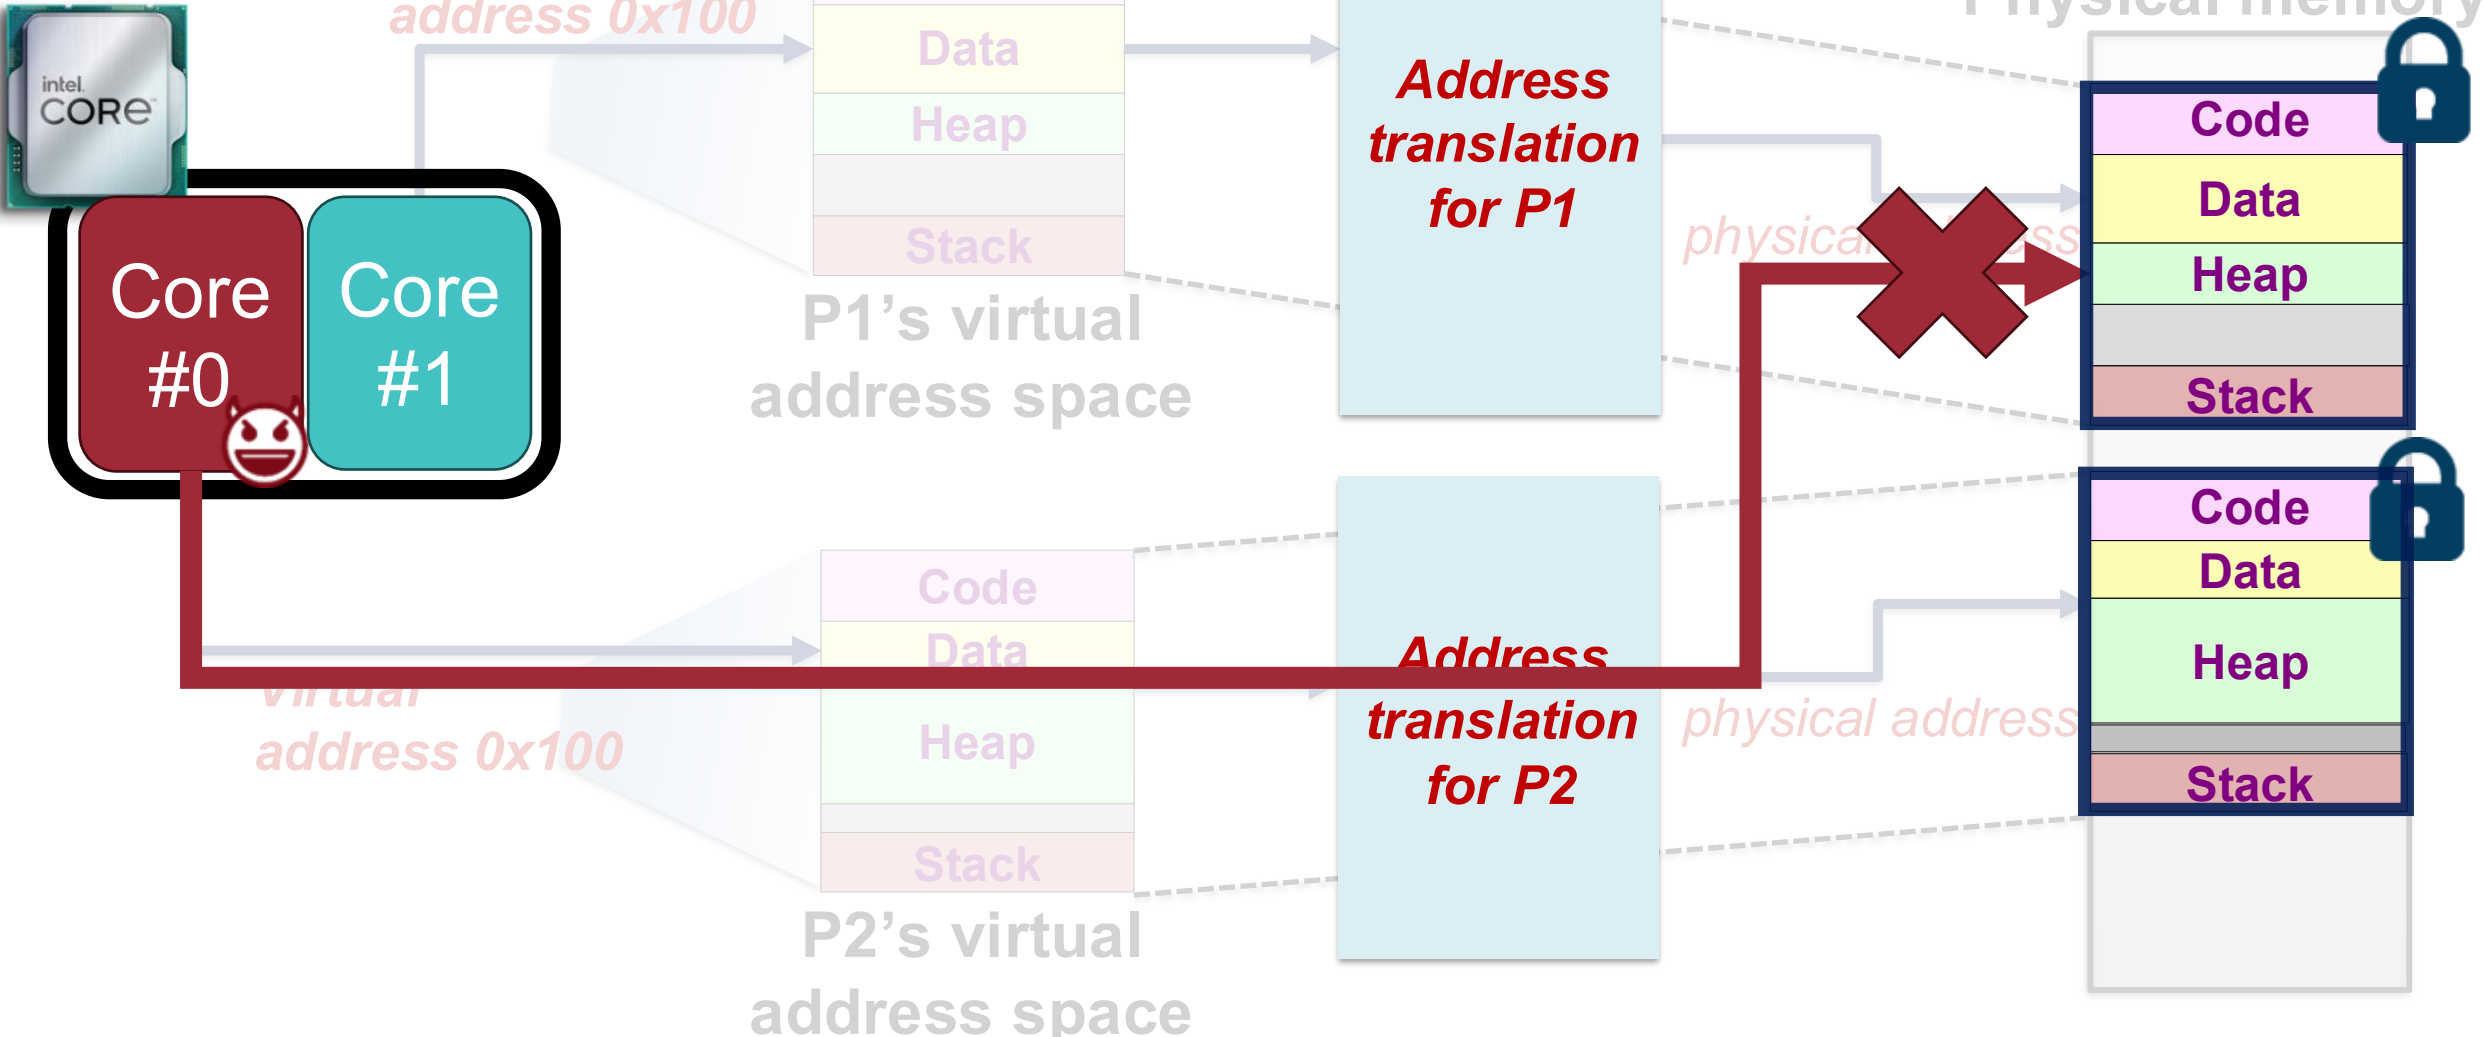
- Give programmers **an illusion of a large memory space** irrespective of actual capacity
- Uses main memory as a “cache” for the hard disk
- Other benefits?
  - **Program simplification**: give each program the illusion that it has its *own private memory* (e.g., `0x00000000~0xffffffff`)
  - **Security**: memory protection between processes

# Memory Protection



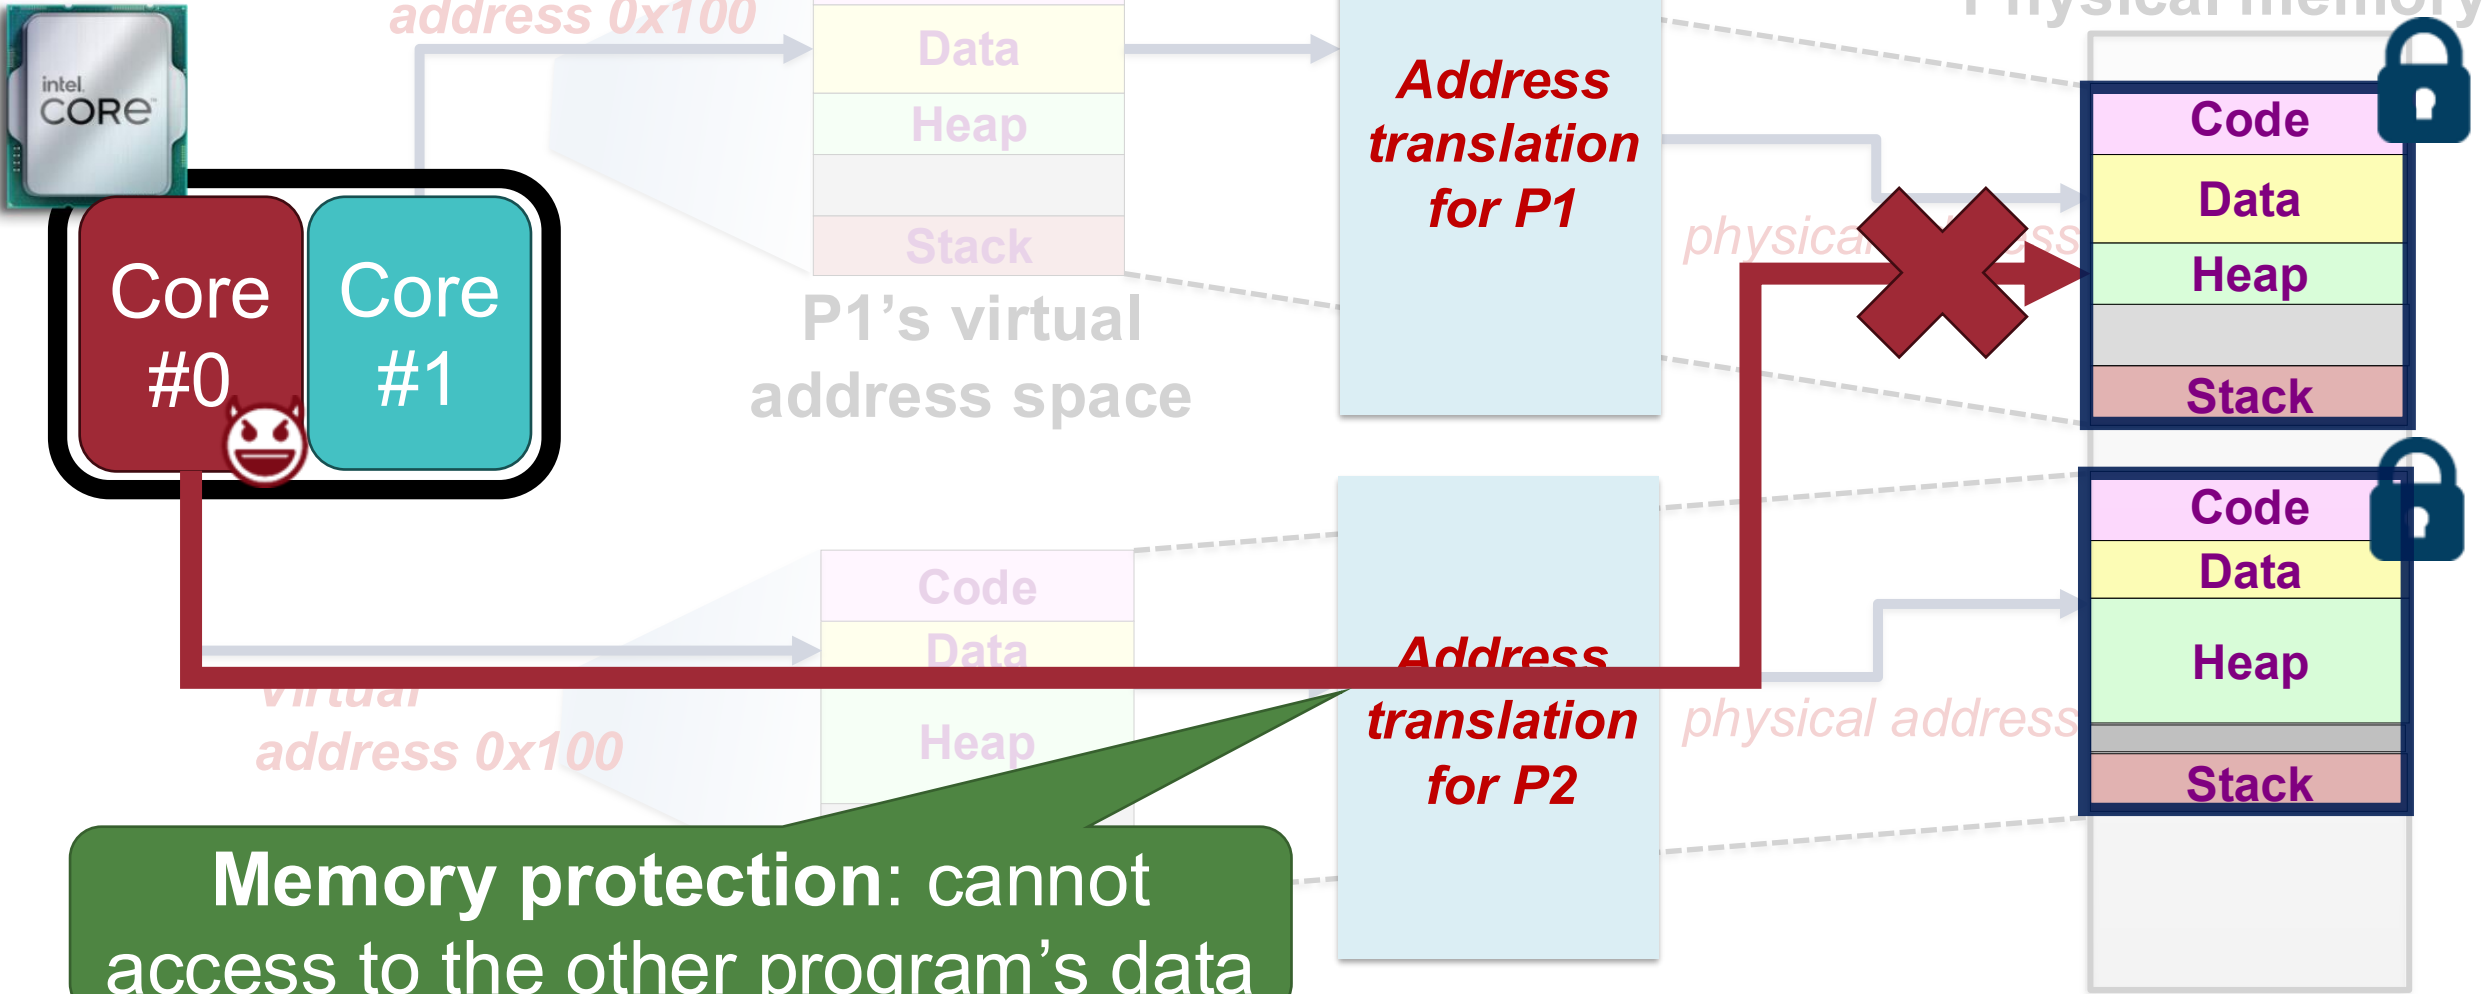
# Memory Protection

OS-managed table



# Memory Protection

OS-managed table



Memory protection: cannot access to the other program's data

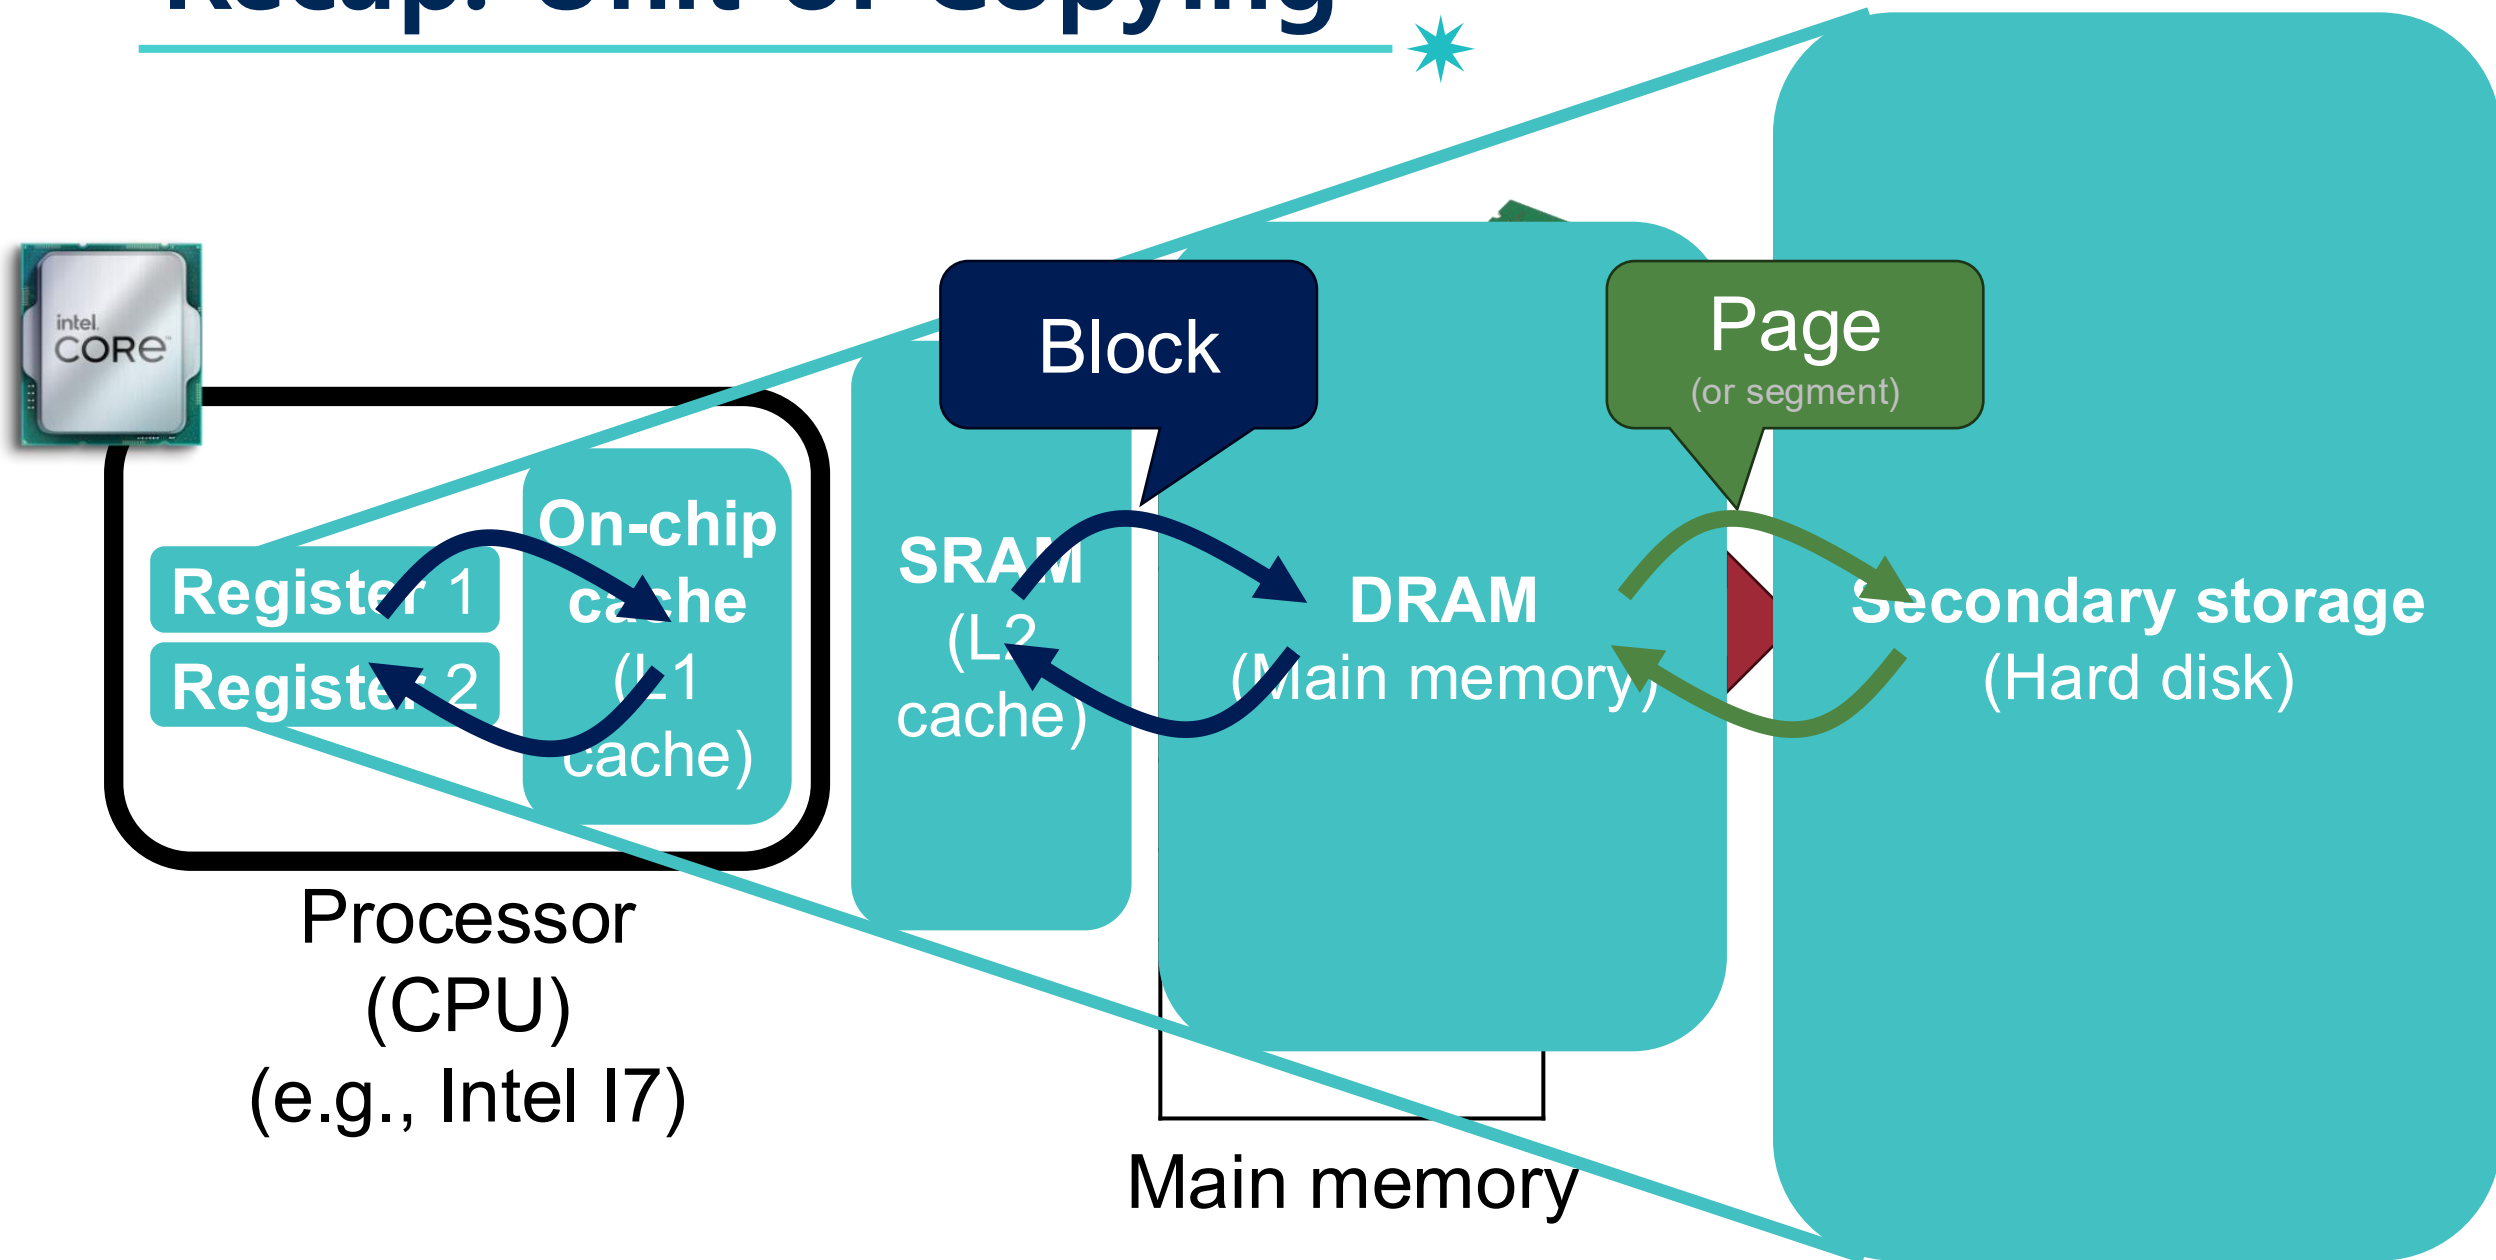
# Virtual Memory Summary

---

- Idea: Give programmers **an illusion of a large memory space** irrespective of actual capacity
  - Programmers assume they have “infinite” amount of physical memory
  - Programmers do not need to worry about how to manage it (simplifies memory management for multi-processing system)
  
- CPU and OS cooperatively manage physical memory space to provide the illusion on behalf of users
  - Illusion is maintained for each independent process
  - Let OS to protect programs or share memory (e.g., library code) from each other

**Unit of Copying: Page**

# Recap: Unit of Copying



# Virtual Memory Block Type #1: Page

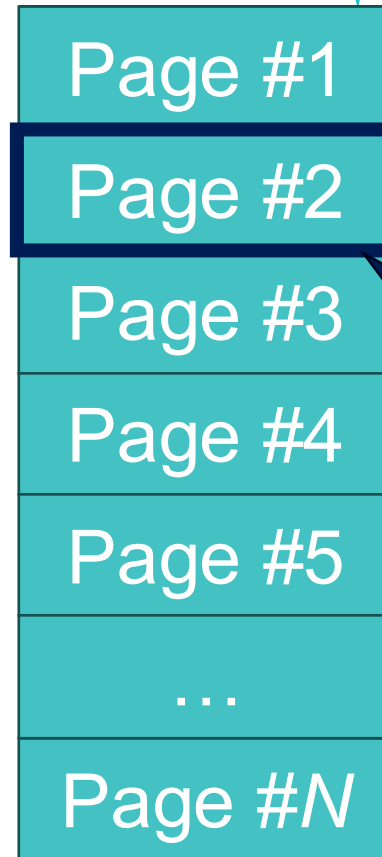
- A block of virtual memory
- Unit of copying
  - between DRAM and Disk
- Fixed length
  - e.g., 4KB
  - Generally, 4KB ~ 4MB



Process A  
(Virtual memory)

# Virtual Memory Block Type #1: Page

- A block of virtual memory
- Unit of copying
  - between DRAM and Disk
- Fixed length
  - e.g., 4KB
  - Generally, 4KB ~ 4MB

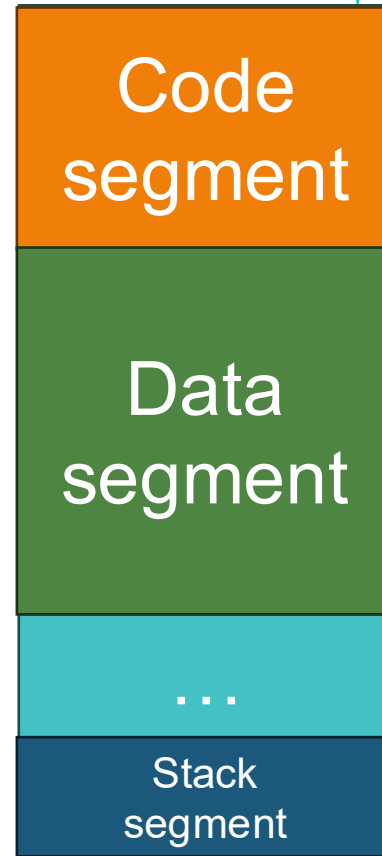


Process A  
(Virtual memory)

Fixed-length  
virtual memory block

# Virtual Memory Block Type #2: Segment

Segment length is variable!



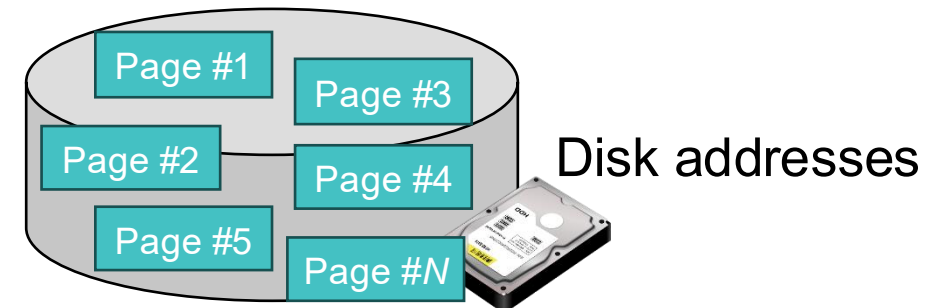
Process A  
(Virtual memory)

# Page

- A block of virtual memory
- Unit of copying
  - between DRAM and Disk
- Fixed length
  - e.g., 4KB
  - Generally, 4KB ~ 4MB



Process A  
(Virtual memory)

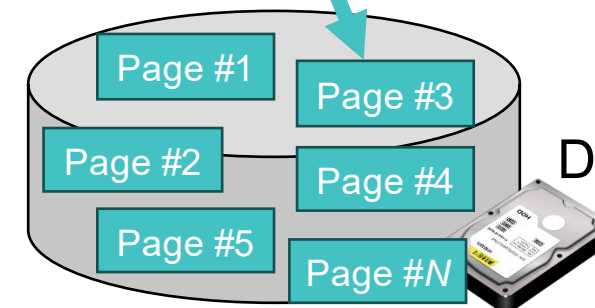


# Memory Paging (a.k.a., Swapping)

Memory management scheme by which a computer **stores and retrieves data** from secondary storage for use in main memory

Page: unit of copying!

*Paging  
(Swapping)*



Physical addresses

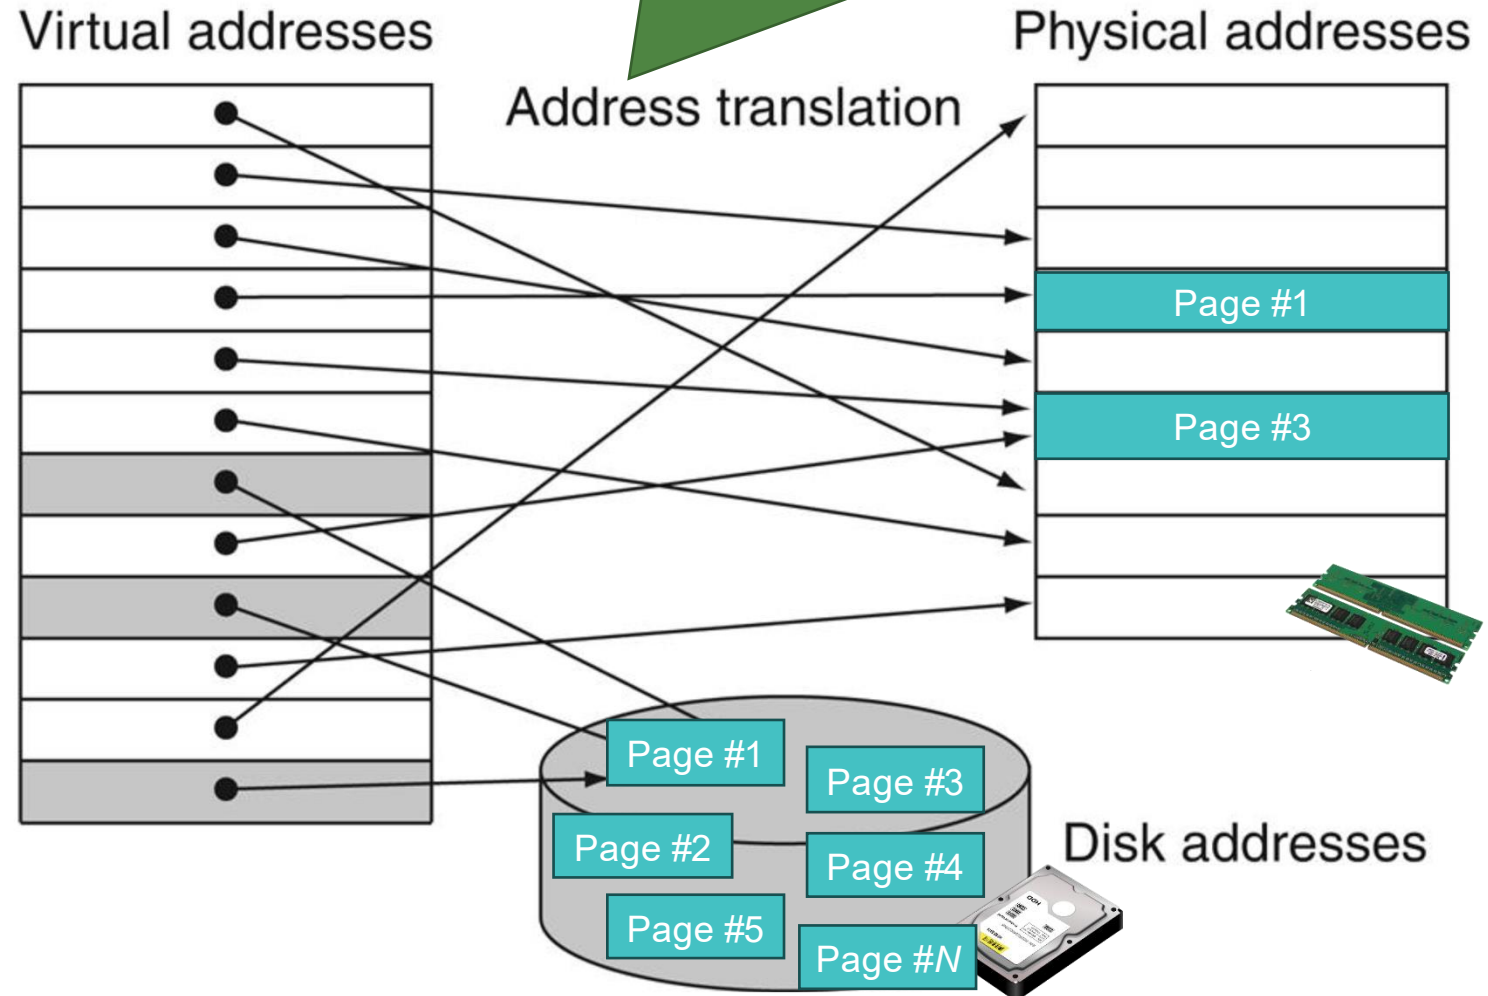


Disk addresses

# Address Translation

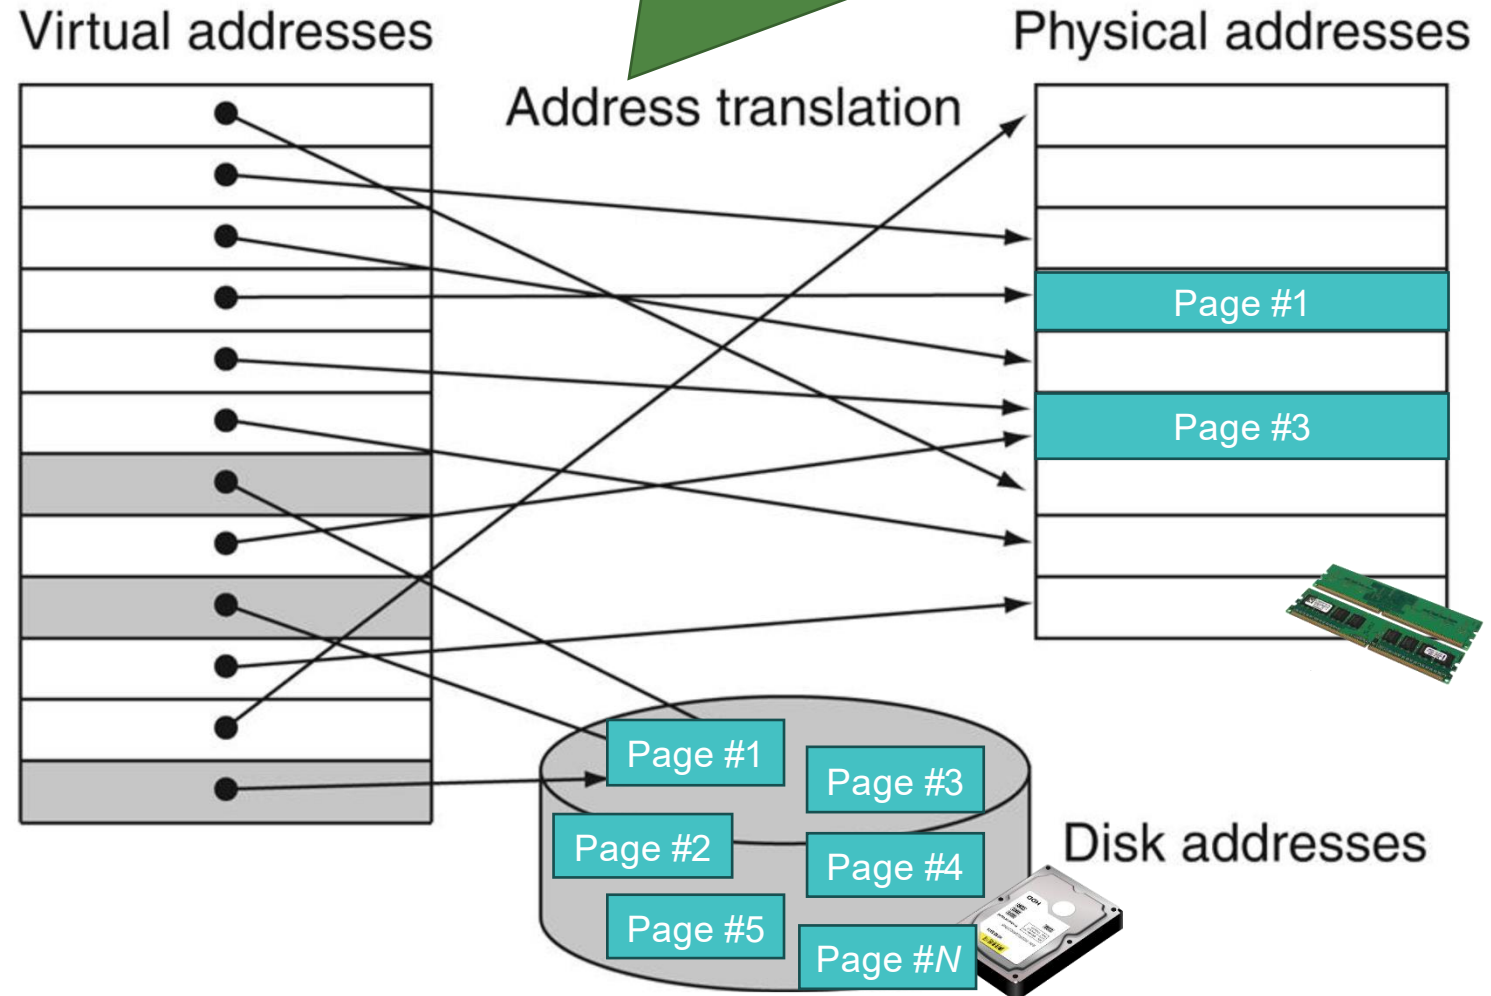
# Memory Hi

Address translation table (OS-managed table)  
Virtual address → Physical address



# Memory Hi

## “Page Table”



# Page Table: Address Translation

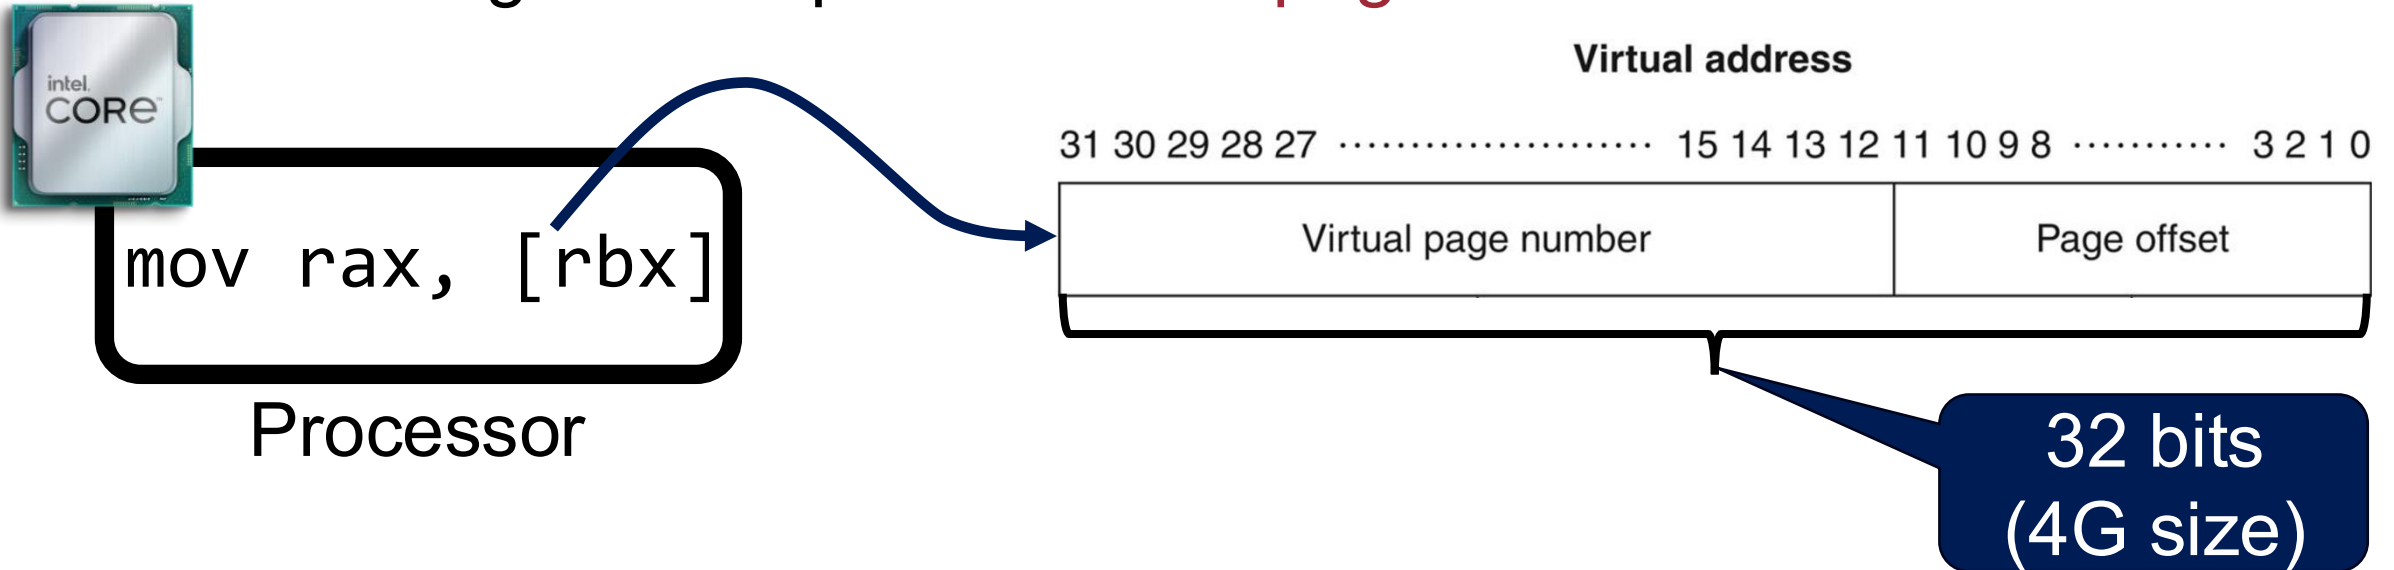
---



CPU converts virtual addresses into physical addresses via an OS-managed lookup table called **page table**


# Page Table: Address Translation

CPU converts virtual addresses into physical addresses via an OS-managed lookup table called **page table**




# Programmer's Point of View

CPU converts virtual addresses into physical addresses via an OS-managed lookup table called **page table**



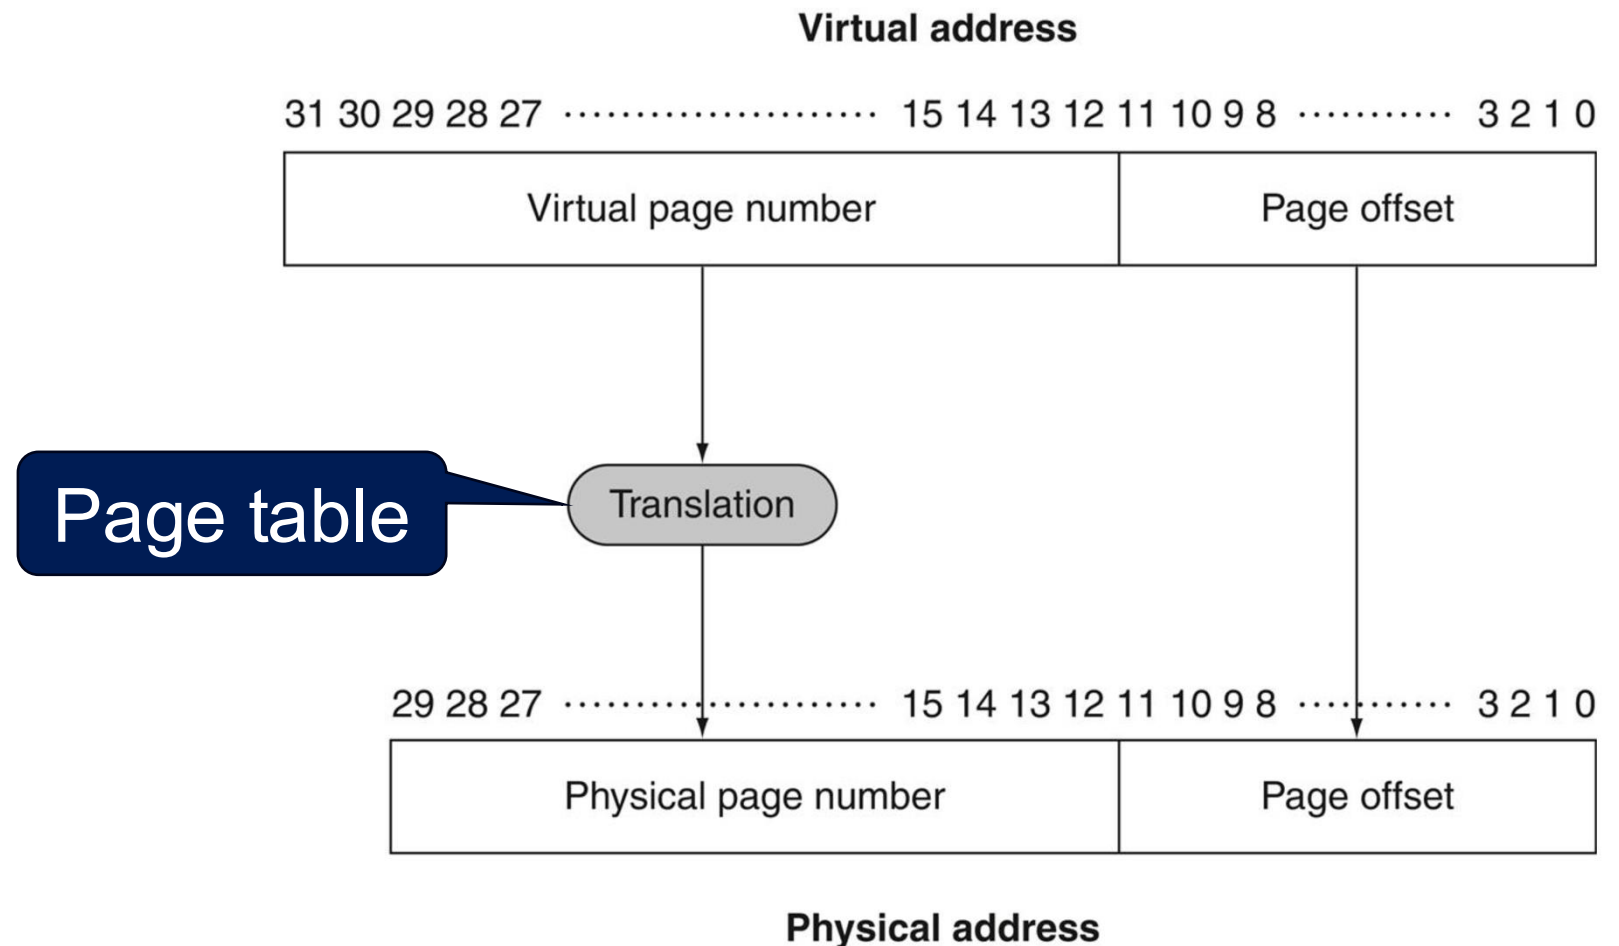
# Programmer's Point of View

CPU converts virtual addresses into physical addresses via an OS-managed lookup table called **page table**



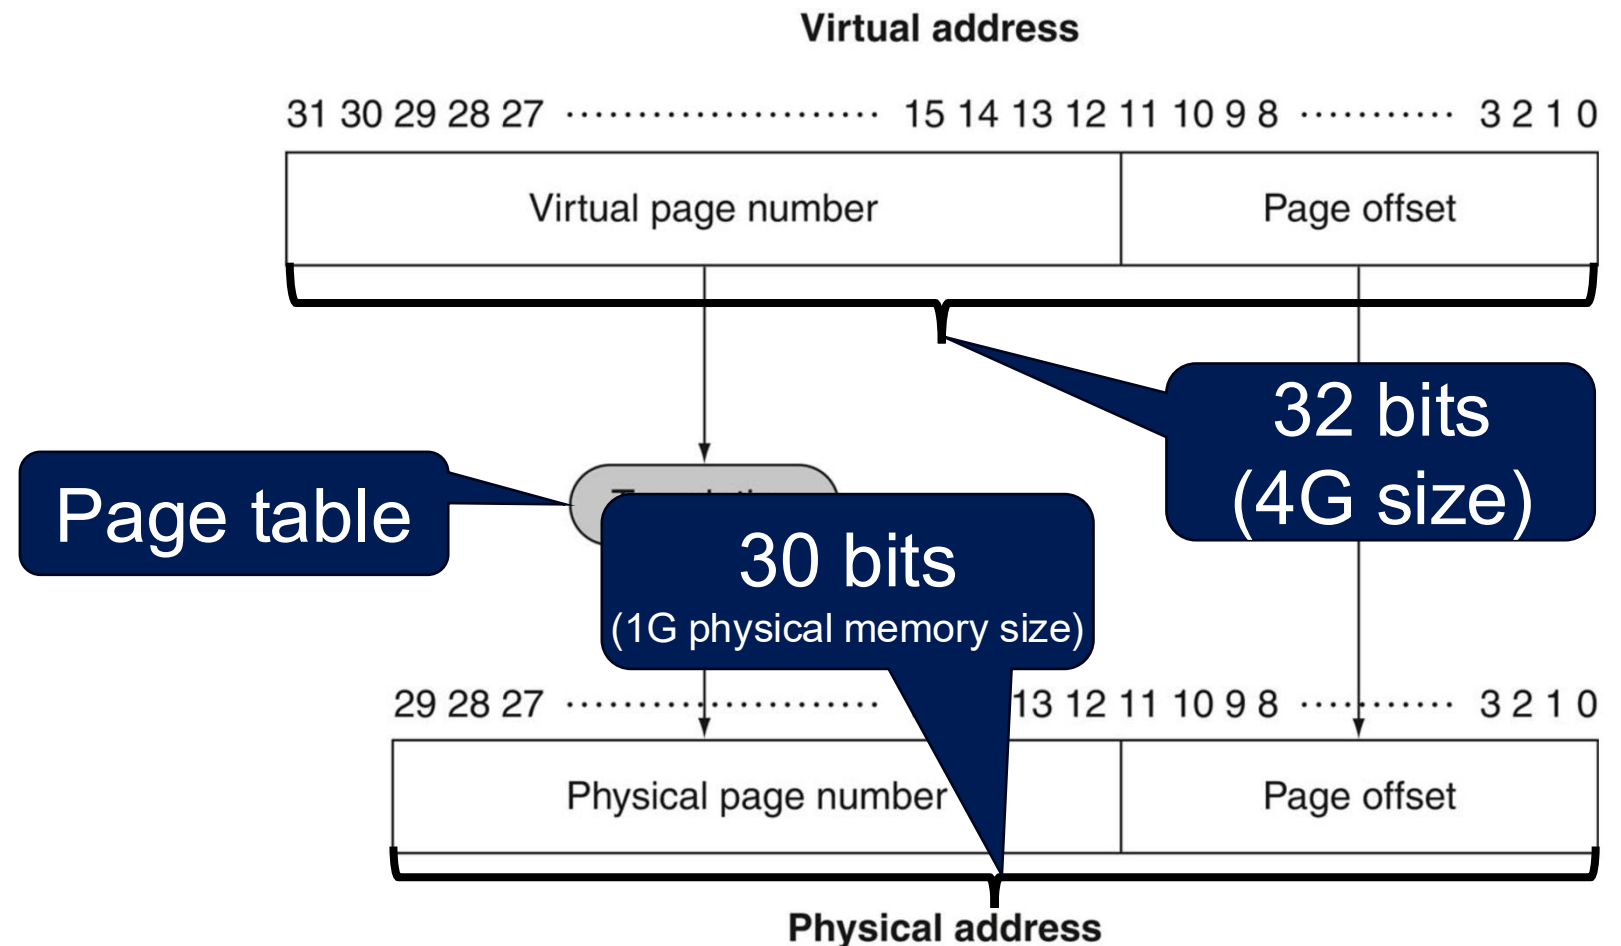
# Page Table: Address Translation

CPU converts virtual addresses into physical addresses via an OS-managed lookup table called **page table**



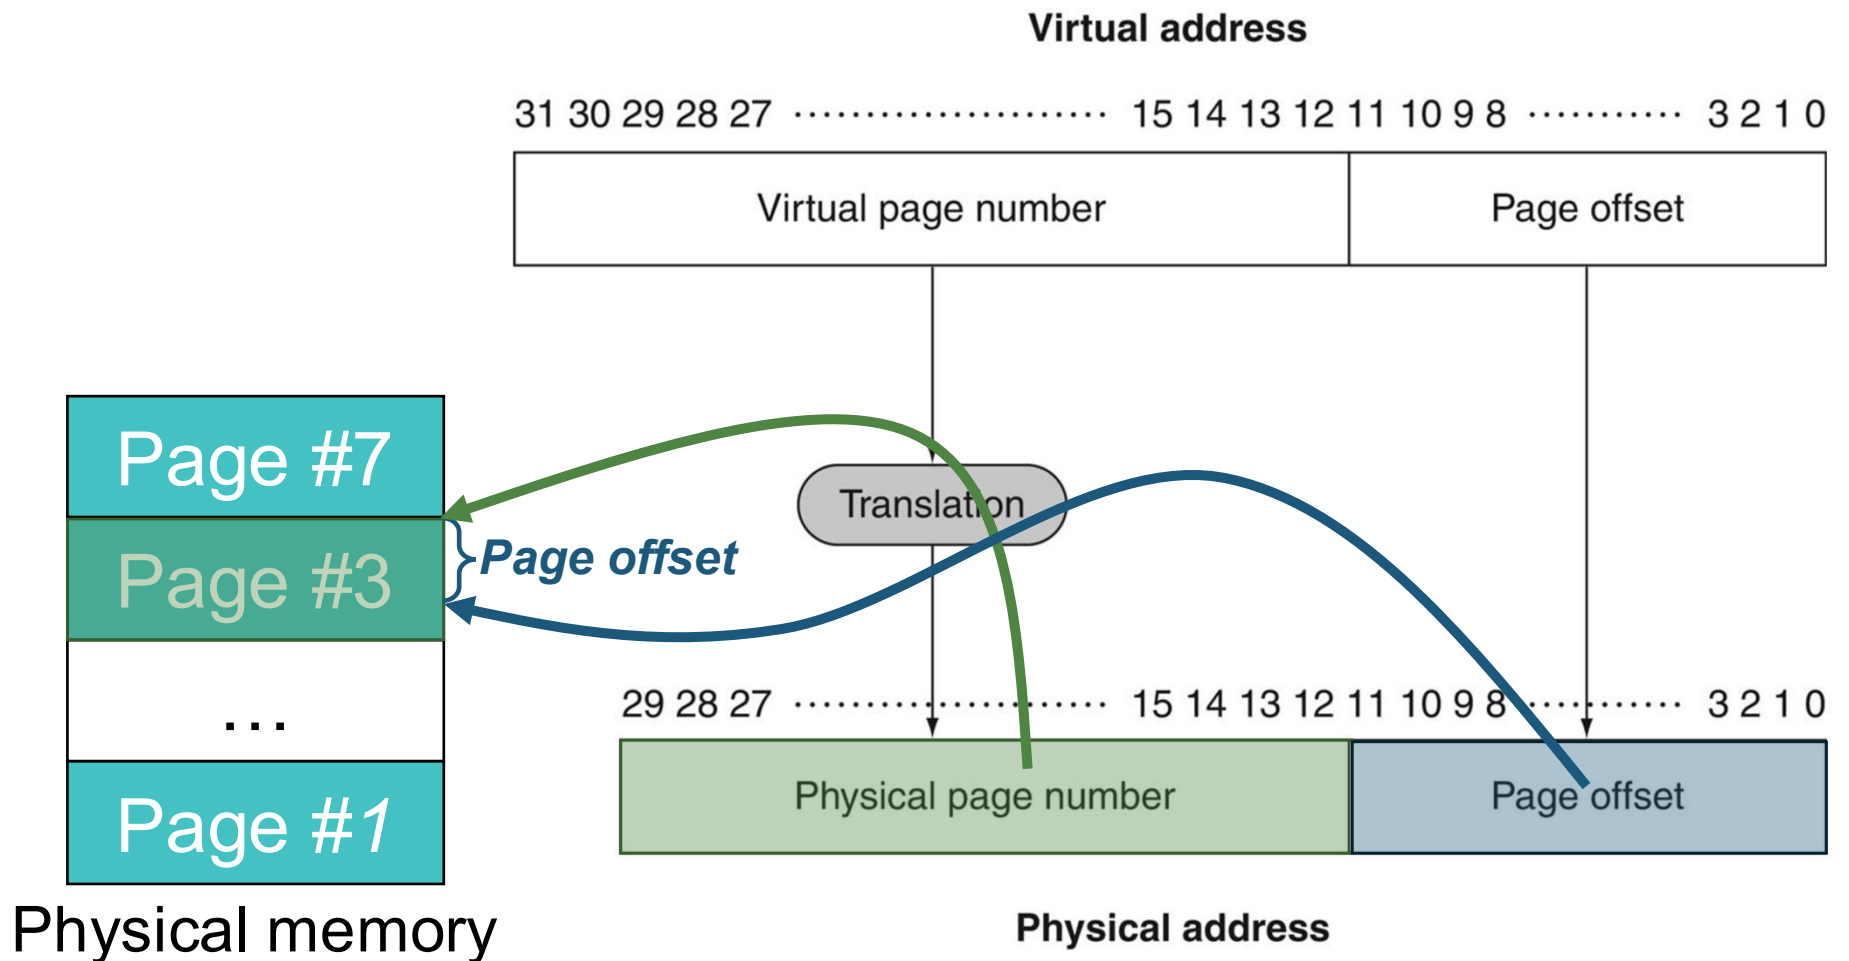
# Page Table: Address Translation

CPU converts virtual addresses into physical addresses via an OS-managed lookup table called **page table**



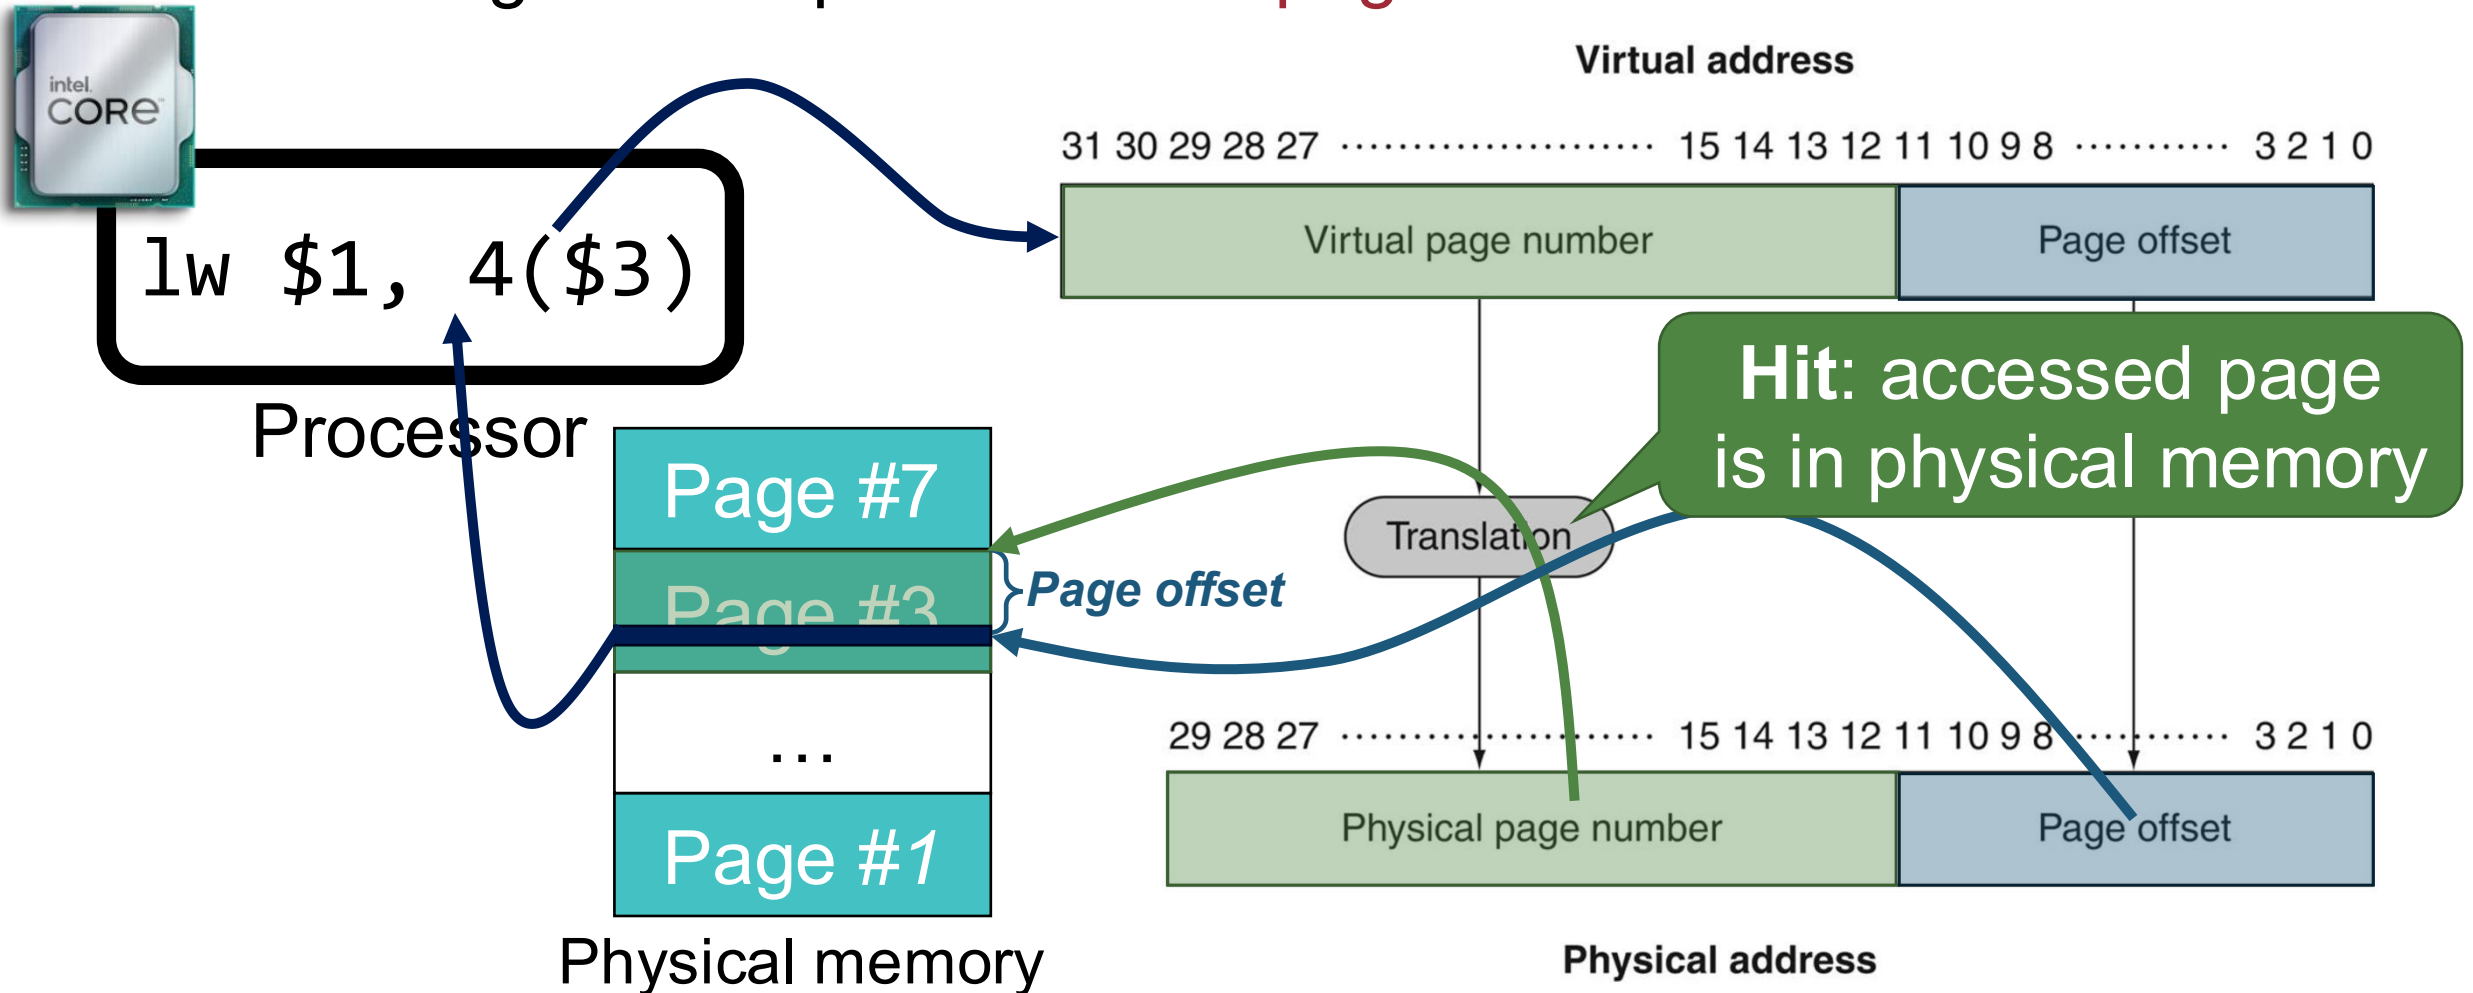
# Page Table: Address Translation

CPU converts virtual addresses into physical addresses via an OS-managed lookup table called **page table**



# Address Translation Flow (Hit Case)

CPU converts virtual addresses into physical addresses via an OS-managed lookup table called **page table**



# Page Table: Address Translation

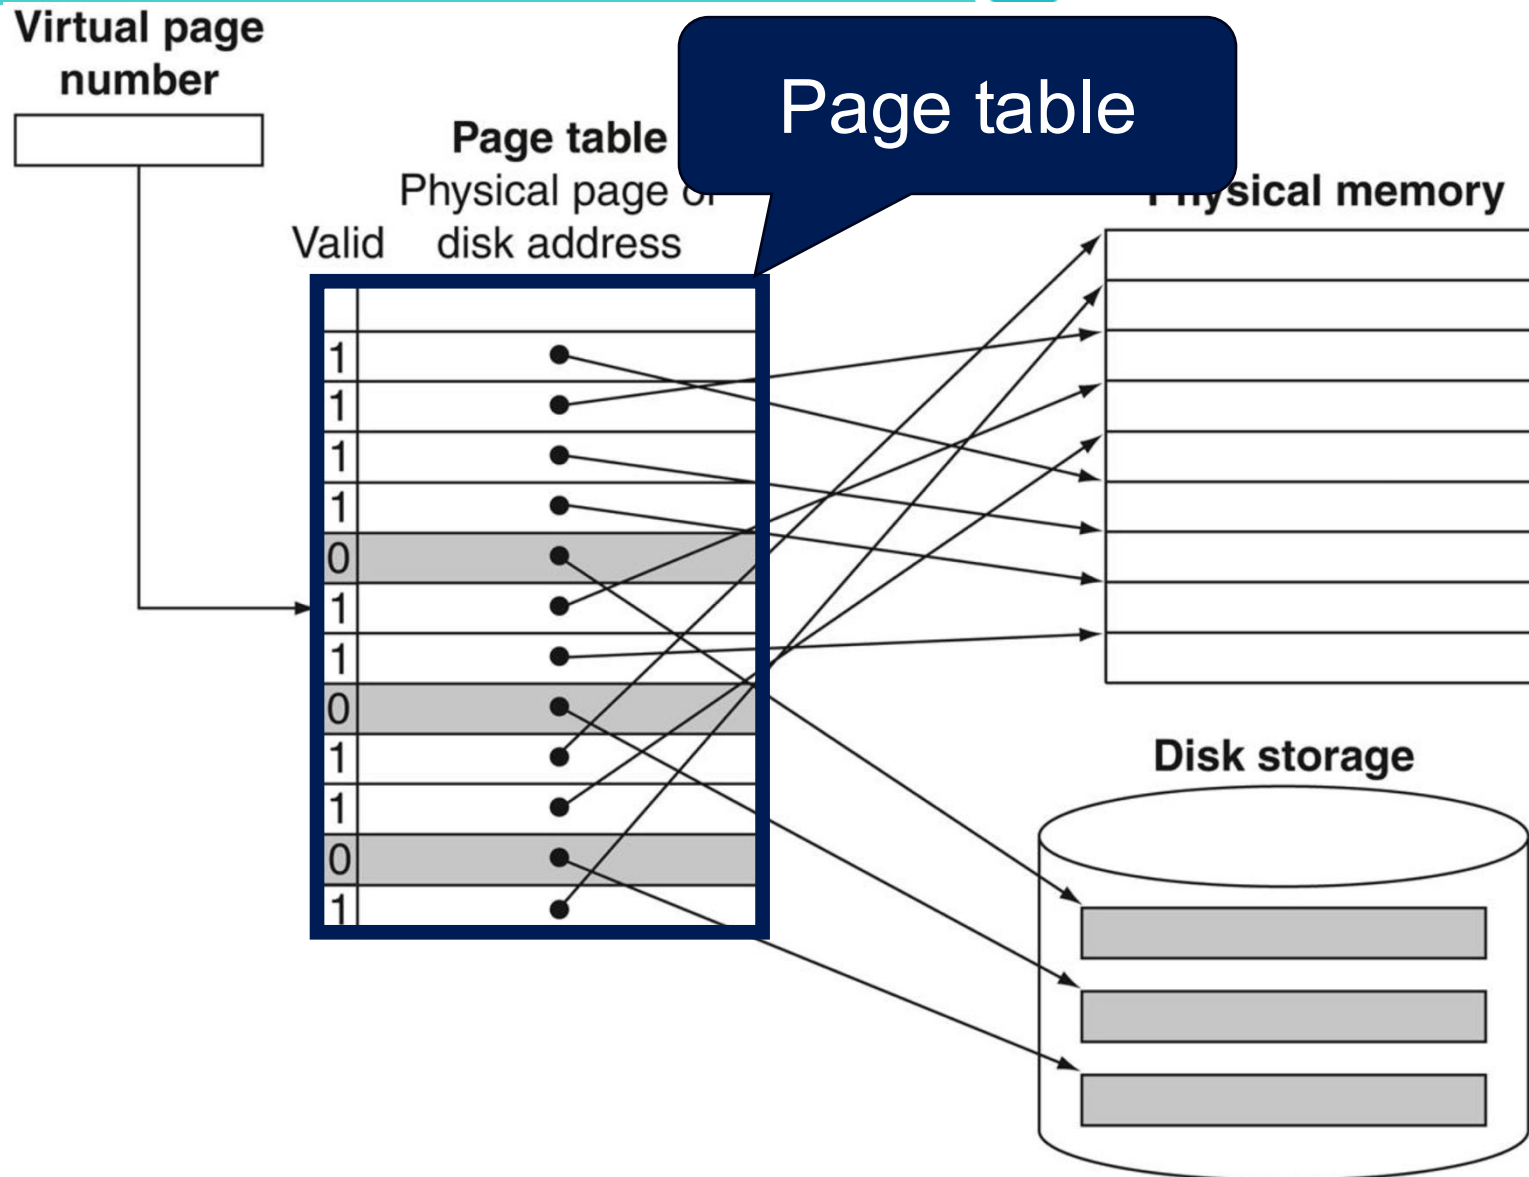
---



CPU converts virtual addresses into physical addresses via an OS-managed lookup table called **page table**

- **Mapping from a virtual to physical address**
  - **Virtual address** = virtual page number + page offset
  - **Physical address** = physical page number + page offset
- Each program has its own page table!

# Page Table: Details



# Page Table: Details

Virtual page number


#5

Page table


Physical page or  
disk address

	Valid	Physical page or disk address
#0	1	●
#1	1	●
#2	1	●
#3	1	●
#4	0	●
#5	1	●
#6	1	●
#7	0	●
.	1	●
.	1	●
.	0	●
#N	1	●

Physical memory



Disk storage



# Page Table: Details

Virtual page number

#5

Page table


Physical page or disk address

Valid disk address


Virtual page number	Valid	Physical page or disk address
#0	1	Physical page 1
#1	1	Physical page 2
#2	1	Physical page 3
#3	1	Physical page 4
#4	0	Invalid
#5	1	Physical page 5
#6	1	Physical page 6
#7	0	Invalid
.	1	Physical page 7
.	1	Physical page 8
.	0	Invalid
#N	1	Physical page N

Entry: Physical page number (or disk address)

Physical memory

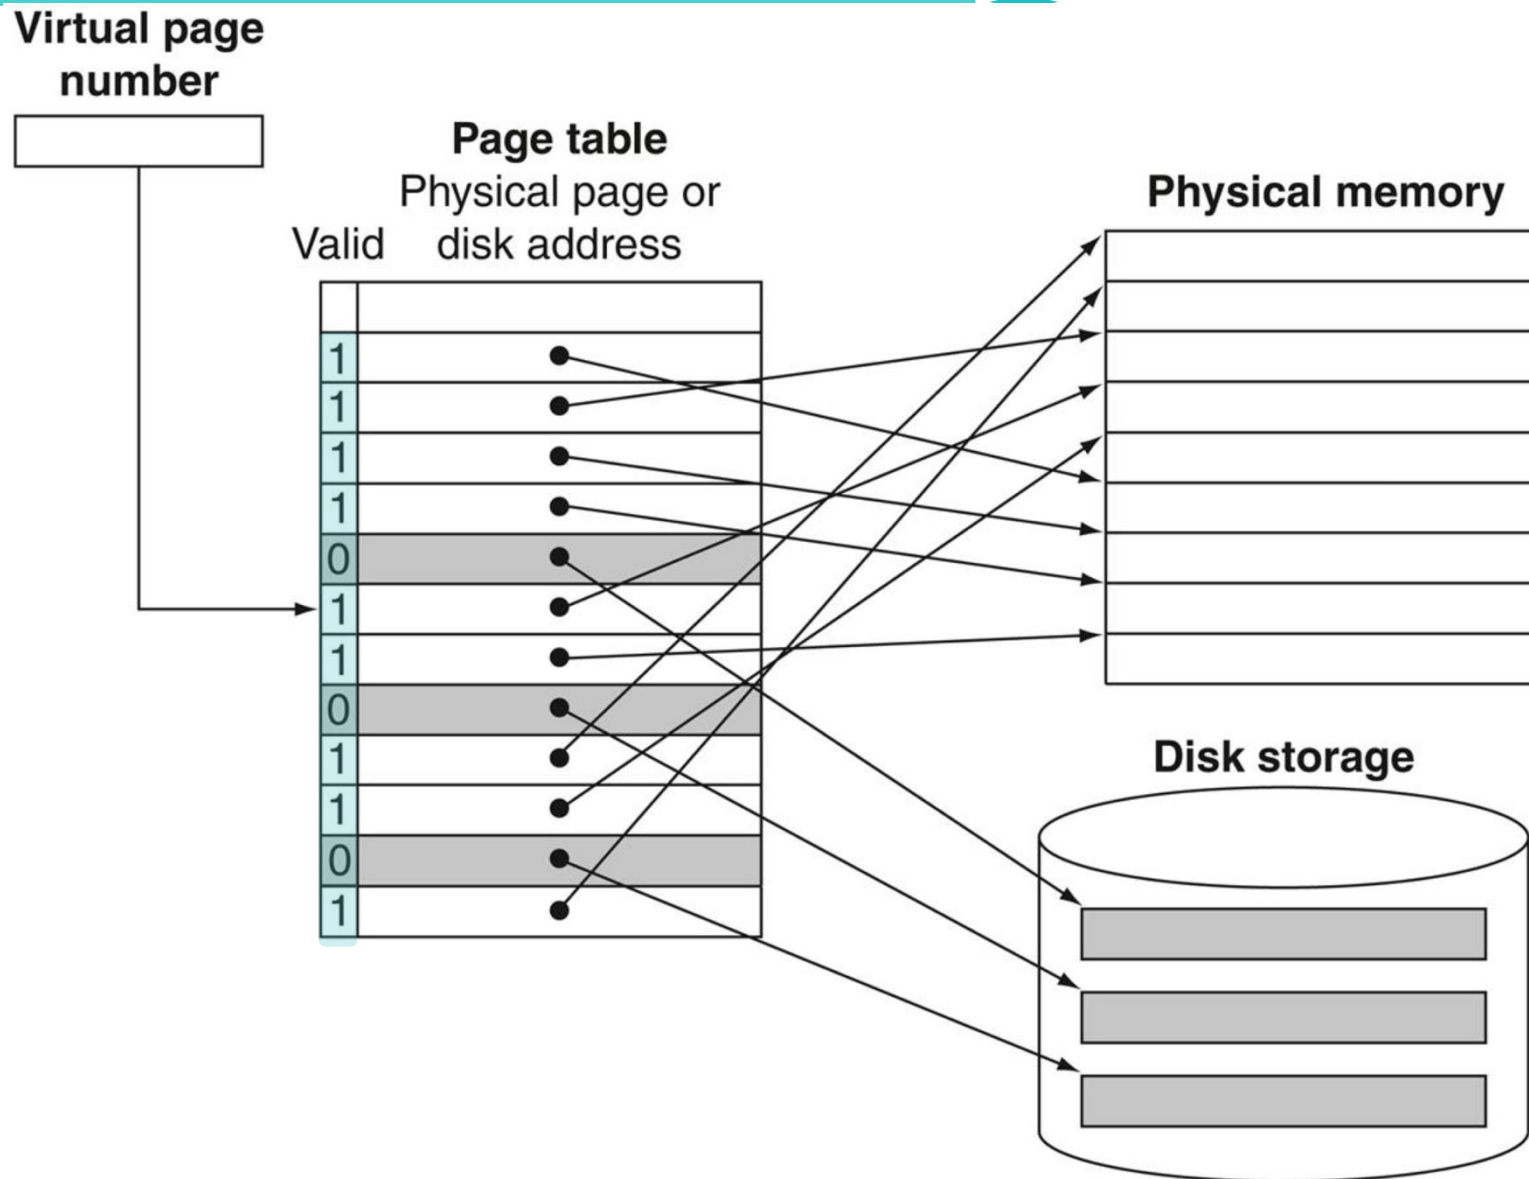


Disk storage



Index: virtual page number (as offset)  
(+ page table base address)

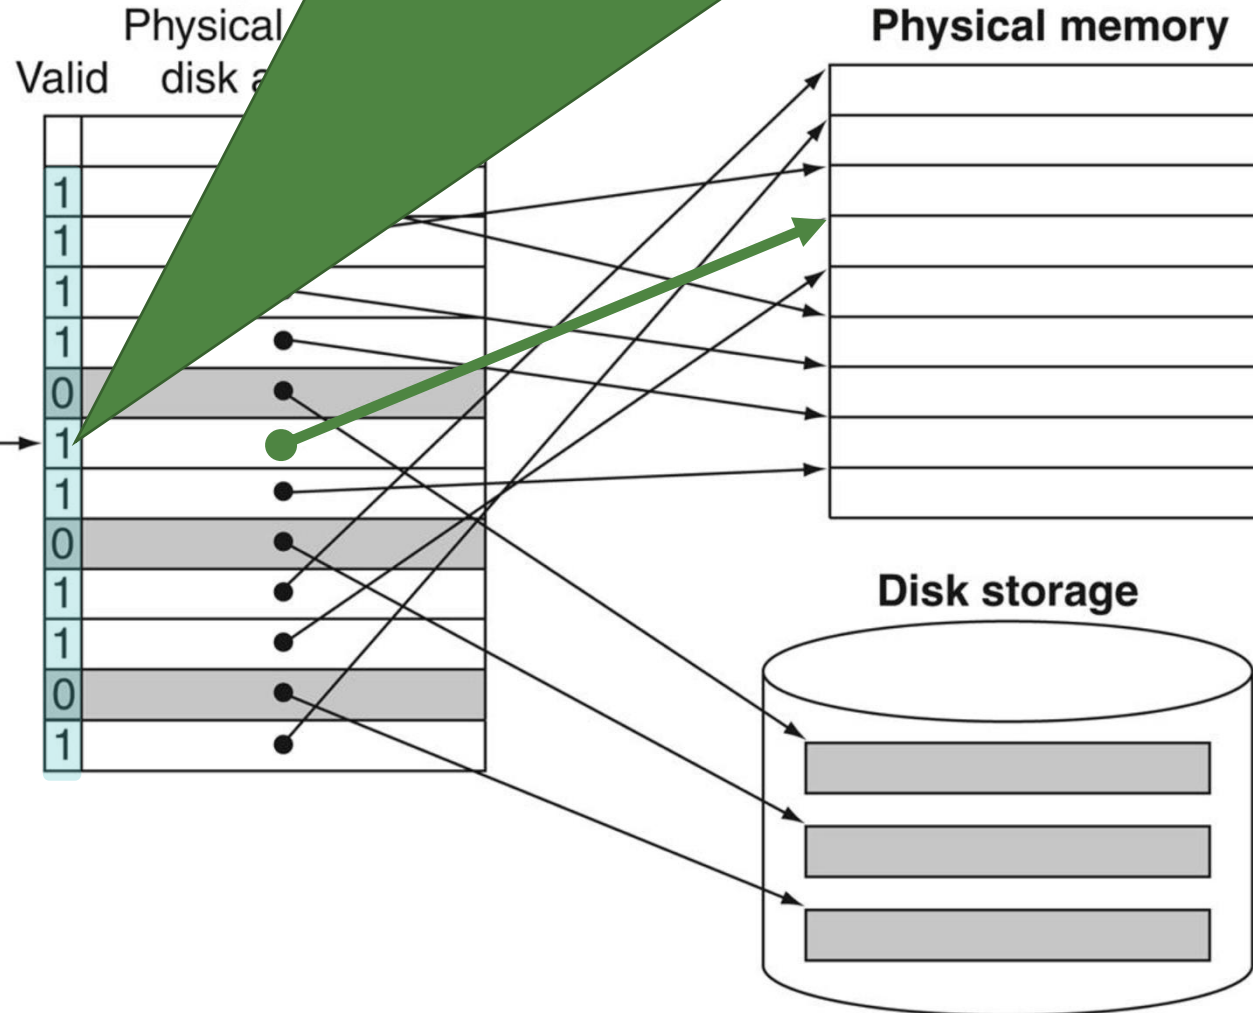
# Valid Bit in Page Table



# Valid Bit in Page Table

Virtual page number

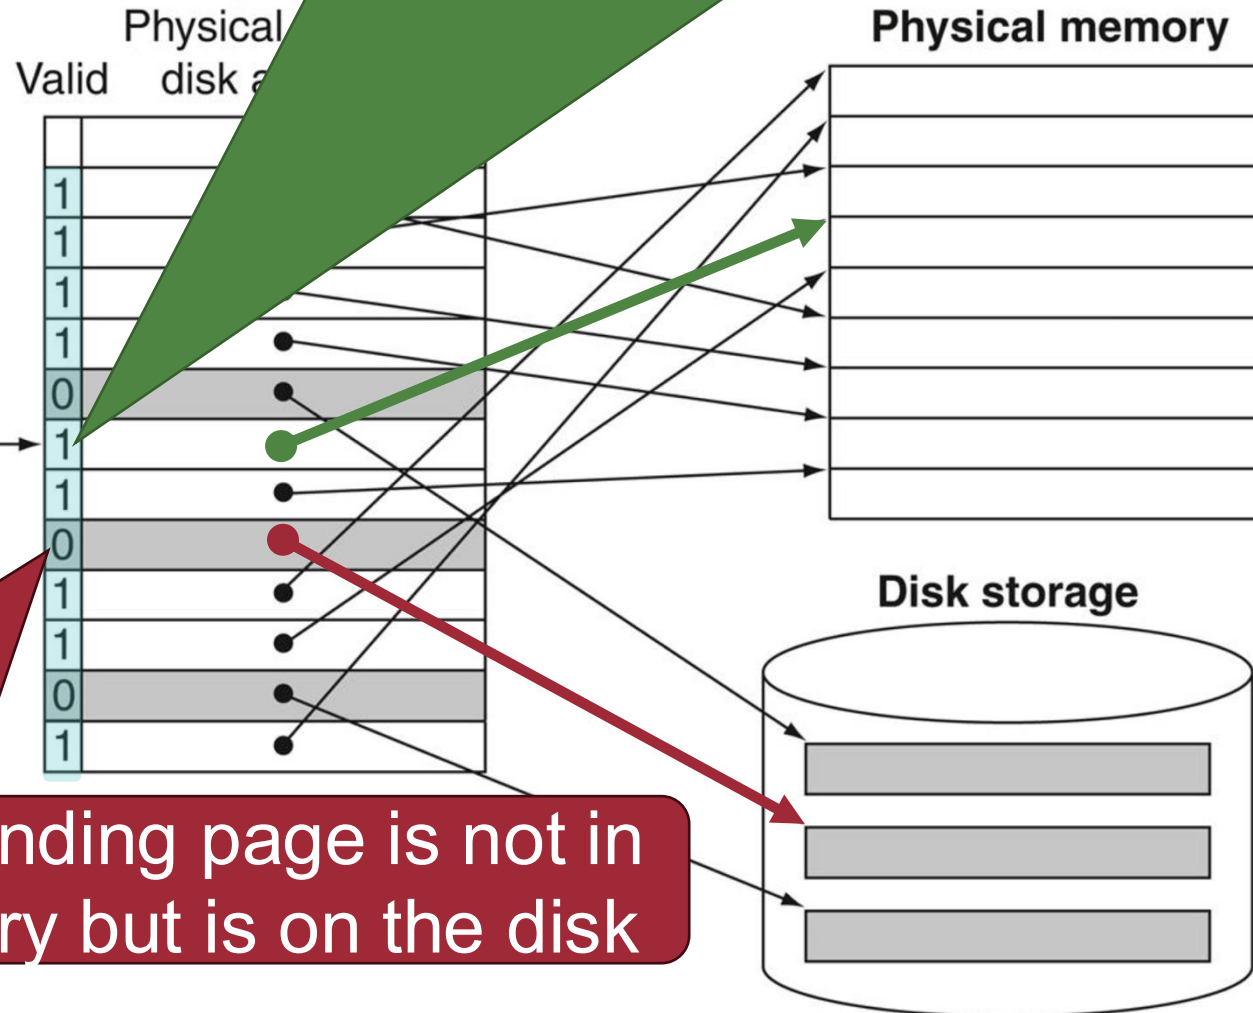
1: the corresponding page is in physical memory (i.e., mapped physical address exists in the entry)



# Valid Bit in Page Table

Virtual page number

**1:** the corresponding page is in physical memory (i.e., mapped physical address exists in the entry)



**0:** the corresponding page is not in physical memory but is on the disk

# Virtual Memory Hit

Virtual page number


#5

Page table

Physical page or disk address


	Valid	Physical page or disk address
#0	1	•
#1	1	•
#2	1	•
#3	1	•
#4	0	•
#5	1	•
#6	1	•
#7	0	•
.	1	•
.	1	•
.	0	•
#N	1	•

Physical memory

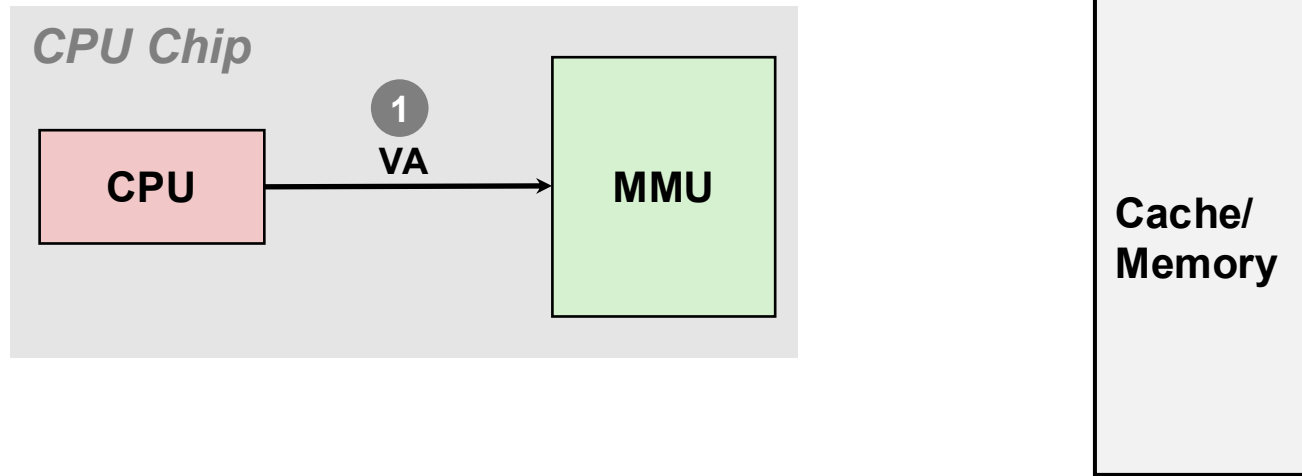


Do not need to access to the disk

Disk storage



# Summary+Details: Virtual Memory Hit



1) Processor sends virtual address to MMU

# FYI: Memory Management Unit (MMU)

---

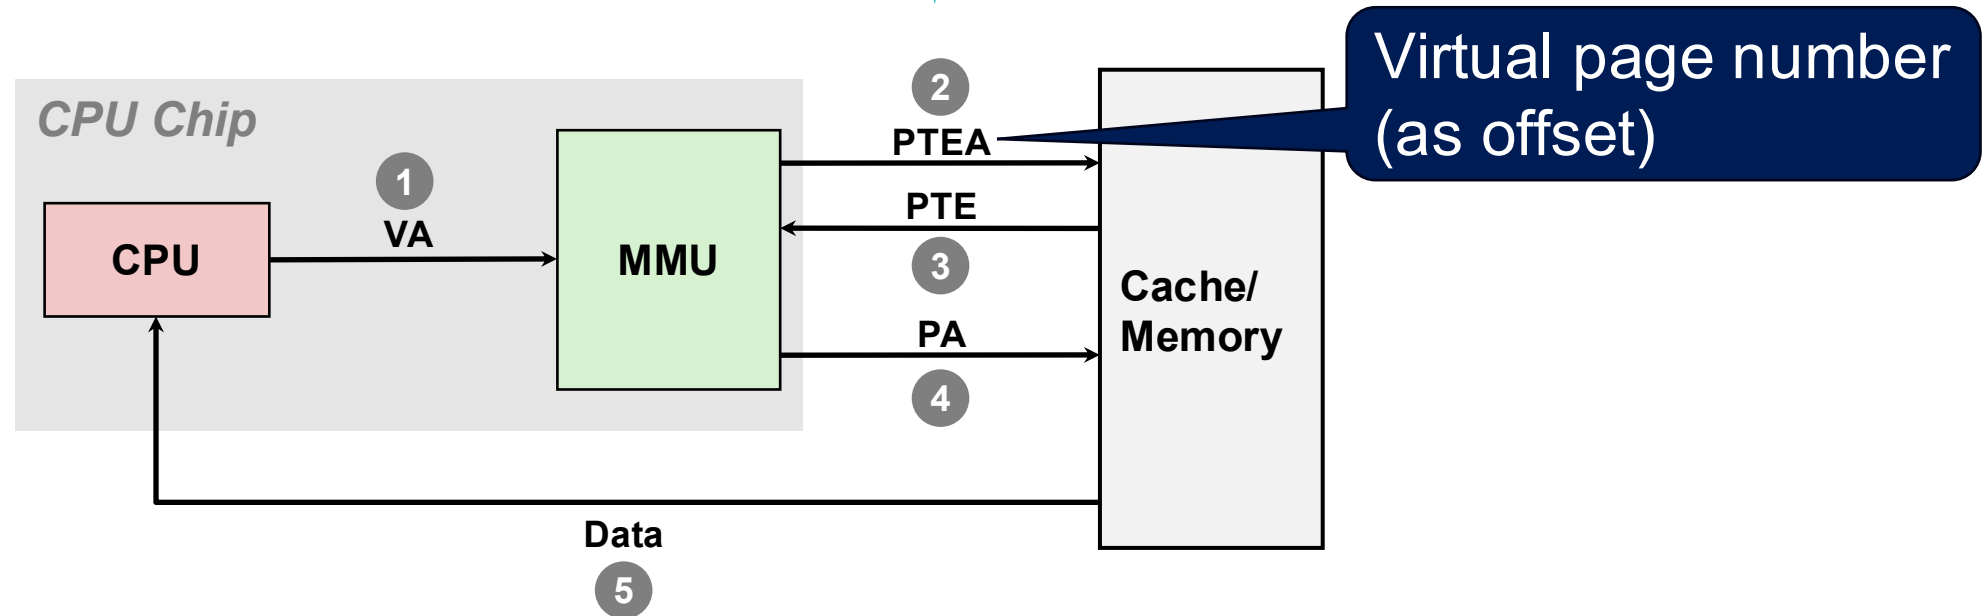
61



A computer hardware unit that translates the virtual address being referenced into physical addresses

# Summary+Details: Virtual Memory Hit

62



- 1) Processor sends virtual address to MMU
- 2-3) MMU fetches page table entry (PTE) from page table in memory
- 4) MMU sends physical address to cache/memory
- 5) Cache/memory sends data word to processor

# Virtual Memory Miss (Page Fault)

Virtual page number


#7

Page table


Physical page or  
disk address

	Valid	Physical page or disk address
#0	1	●
#1	1	●
#2	1	●
#3	1	●
#4	0	●
#5	1	●
#6	1	●
#7	0	●
·	1	●
·	1	●
·	0	●
#N	1	●

Physical memory



Disk storage



**Miss (valid bit 0):** access on memory space not in physical memory (but in disk)

# Virtual Memory Miss (Page Fault)

Virtual page number


#7

Page table


Physical page or  
Valid disk address

Virtual page number	Valid	Physical page or disk address
#0	1	•
#1	1	•
#2	1	•
#3	1	•
#4	0	•
#5	1	•
#6	1	•
#7	0	•
·	1	•
·	1	•
·	0	•
#N	1	•

Physical memory



Disk storage



*Paging  
(swapping)*

*(1) Fetch "page" to the  
physical memory  
(write back depending  
on the situation)*

# Virtual Memory Miss (Page Fault)

Virtual page number

#7


Page table

Physical page or  
Valid disk address


Virtual page number	Valid	Physical page or disk address
#0	1	•
#1	1	•
#2	0	•
#3	1	•
#4	0	•
#5	1	•
#6	1	•
#7	1	•
·	1	•
·	1	•
·	0	•
#N	1	•

*(2) Update the address translation table*

Physical memory

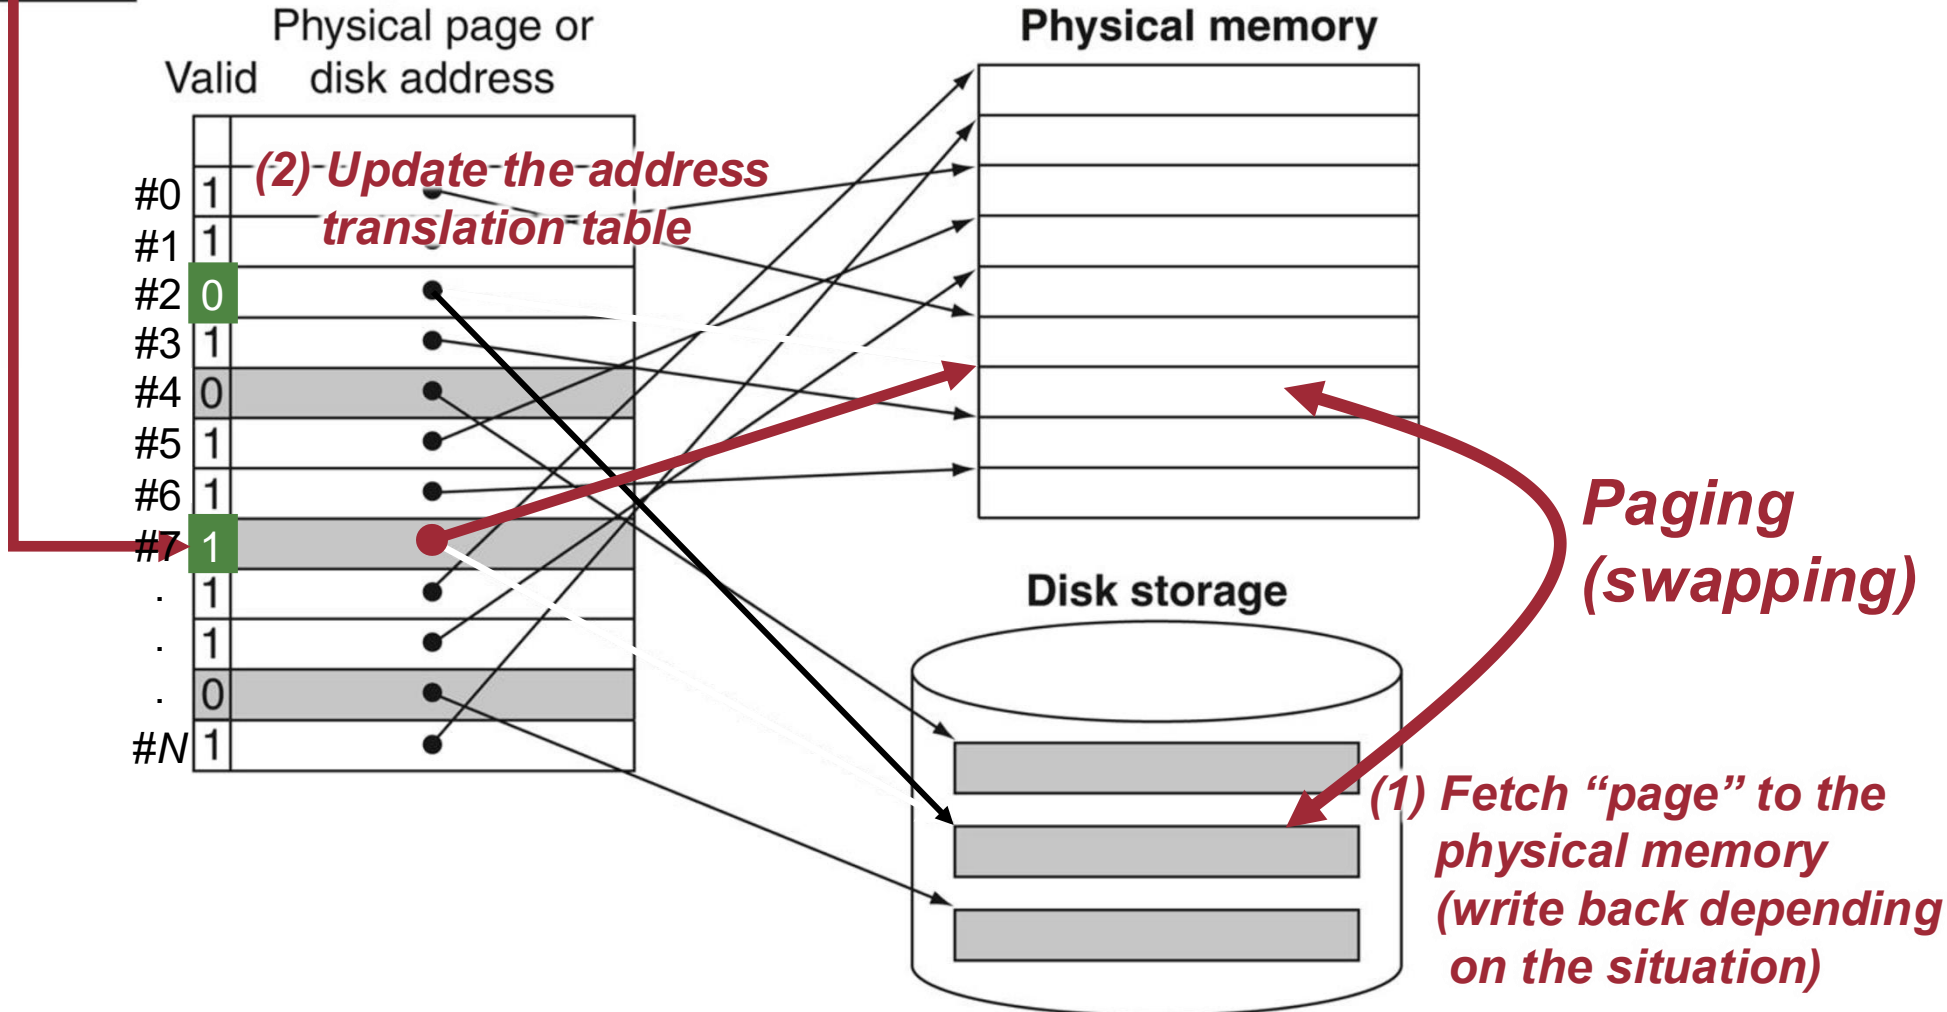


Disk storage




*Paging (swapping)*

*(1) Fetch "page" to the physical memory (write back depending on the situation)*



# Virtual Memory Miss (Page Fault)

- If a page is not in physical memory but disk (if valid == 0)
  - Page table entry (especially, valid bit) indicates that the page not in memory
  - **Page fault exception** is occurred! **OS trap handler** invoked to move data from disk into memory
    - OS has full control over placement!



# Page Fault Penalty



- On page fault, the page must be fetched from disk
  - Takes millions of clock cycles
- Try to minimize page fault rate and fault penalty
  - Prefer **write back** (write through is impractical)
    - **Dirty** bit in each page table entry

Reference	Dirty	Valid	Physical page number
-----------	-------	-------	----------------------

Page table entry

# Recap: Write Policies

---



- Write **hits**
  - Cache and memory would be inconsistent! What can we do?
  - (1) **Write through**: On each write hit, the information is written to both in the cache and in the memory
    - **Pros**: faster processing in case of 'miss'
    - **Cons**: make writes take longer (whenever data cache is updated, there should a memory write)
  - (2) **Write back**: On each write hit, update only the block in cache. The modified cache block is written to memory **only when it is replaced**
    - Keep track of whether each block is *dirty (additional bit)*.
    - **Pros**: faster processing in case of 'hit' (No repeated writes to memory)
    - **Cons**: slower in case of 'miss'

# Page Fault Penalty

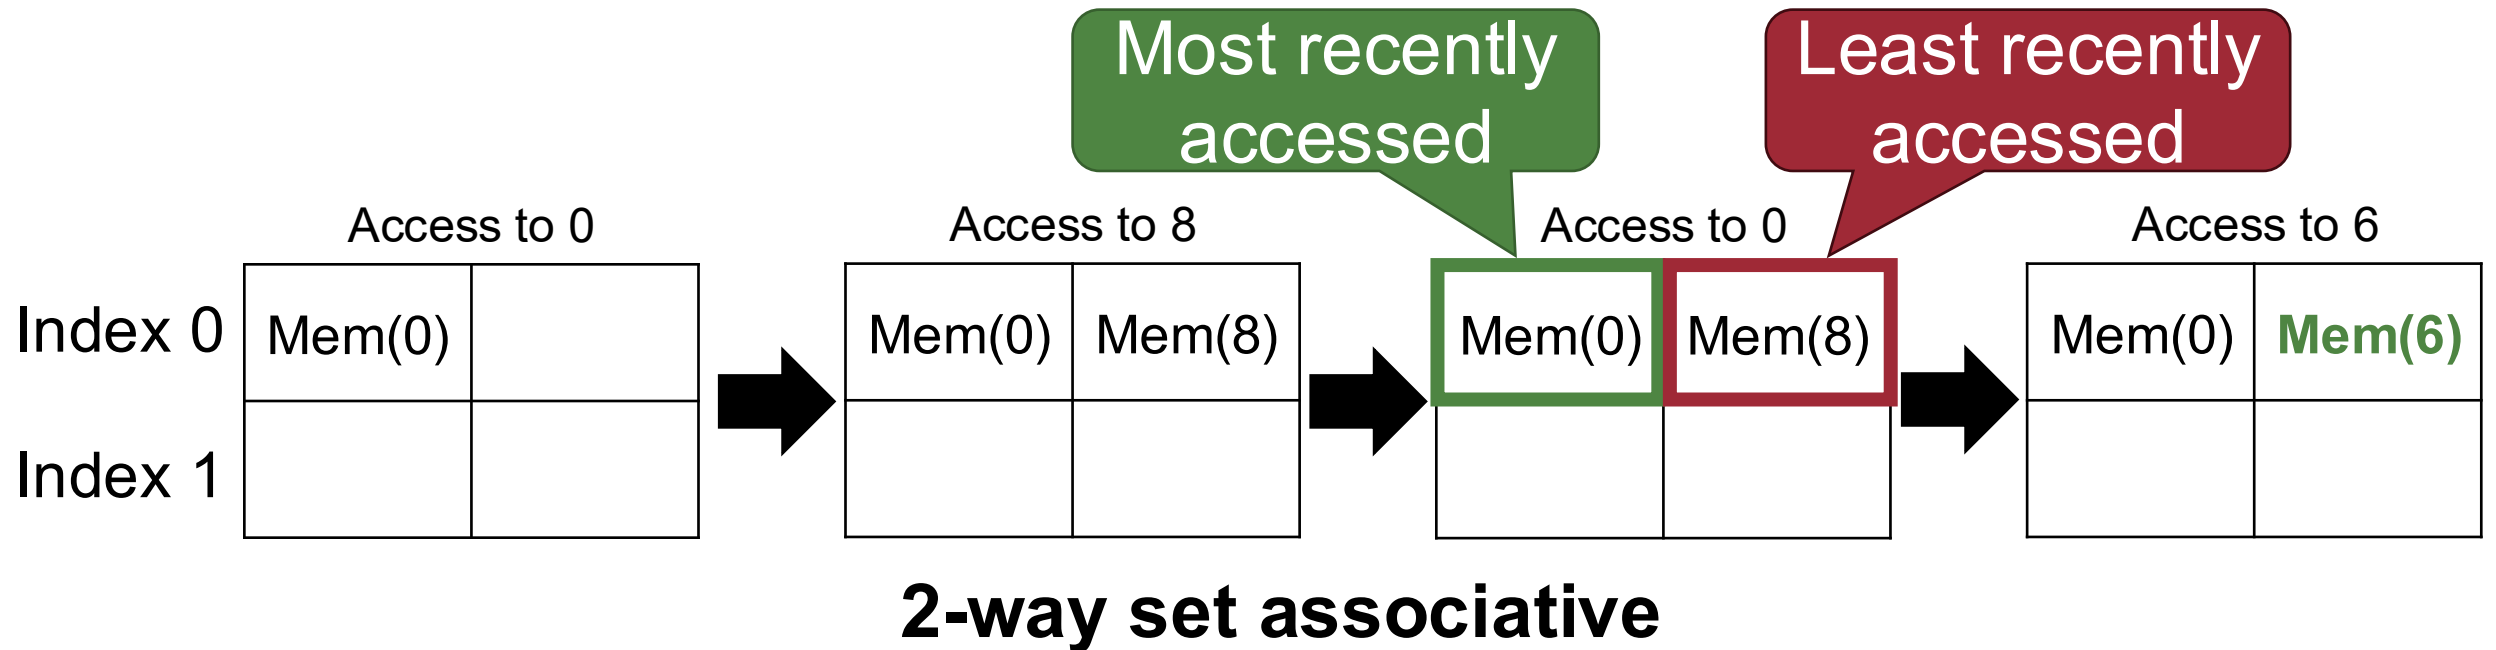
- On page fault, the page is read from disk
  - Takes millions of clock cycles
- Set (1) if the corresponding page is written
- Cleared (0) when the page is replaced
- Try to minimize page fault rate and fault penalty
  - Prefer **write back** (write through is impractical)
    - **Dirty** bit in each page table entry
  - Prefer **LRU** replacement
    - Reference bit in each page table set to 1 on access to page
    - Periodically cleared to 0 by OS
    - A page with **reference** bit = 0 has NOT been used recently

Reference	Dirty	Valid	Physical page number
-----------	-------	-------	----------------------

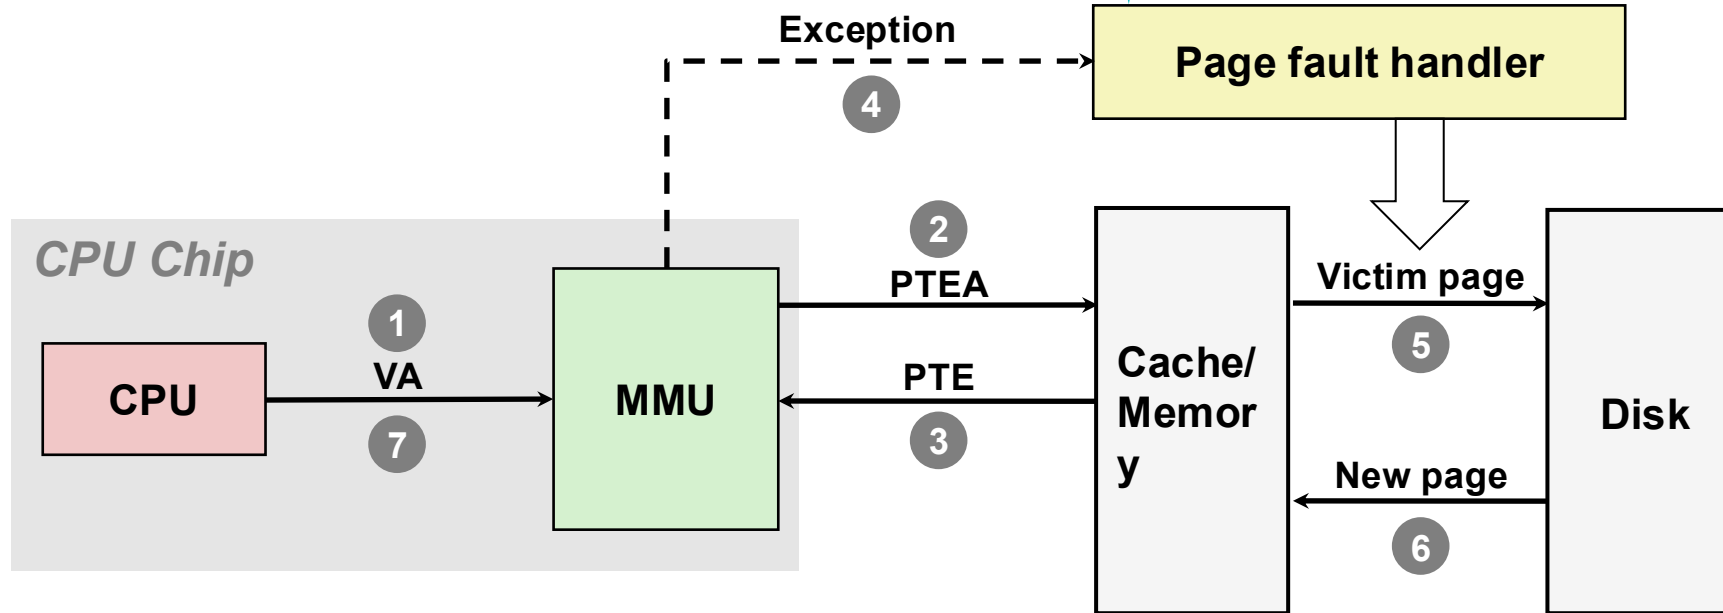
Page table entry

# Recap: LRU

- **Least Recently Used (LRU):** replace the one NOT used (accessed) for the longest time
  - Temporal locality of access is considered
  - Need a reference history information




# Summary+Details: Page Fault



- 1) Processor sends virtual address to MMU
- 2-3) MMU fetches PTE from page table in memory
- 4) Valid bit is zero, so MMU triggers page fault exception
- 5) Handler identifies victim (and, if dirty, pages it out to disk)
- 6) Handler pages in new page and updates PTE in memory
- 7) Handler returns to original process, restarting faulting instruction

# Place of the Page Table?

Virtual page number




Valid	Physical page or disk address
1	•
1	•
1	•
1	•
0	•
1	•
1	•
0	•
1	•
1	•
0	•
1	•

Physical memory



Disk storage



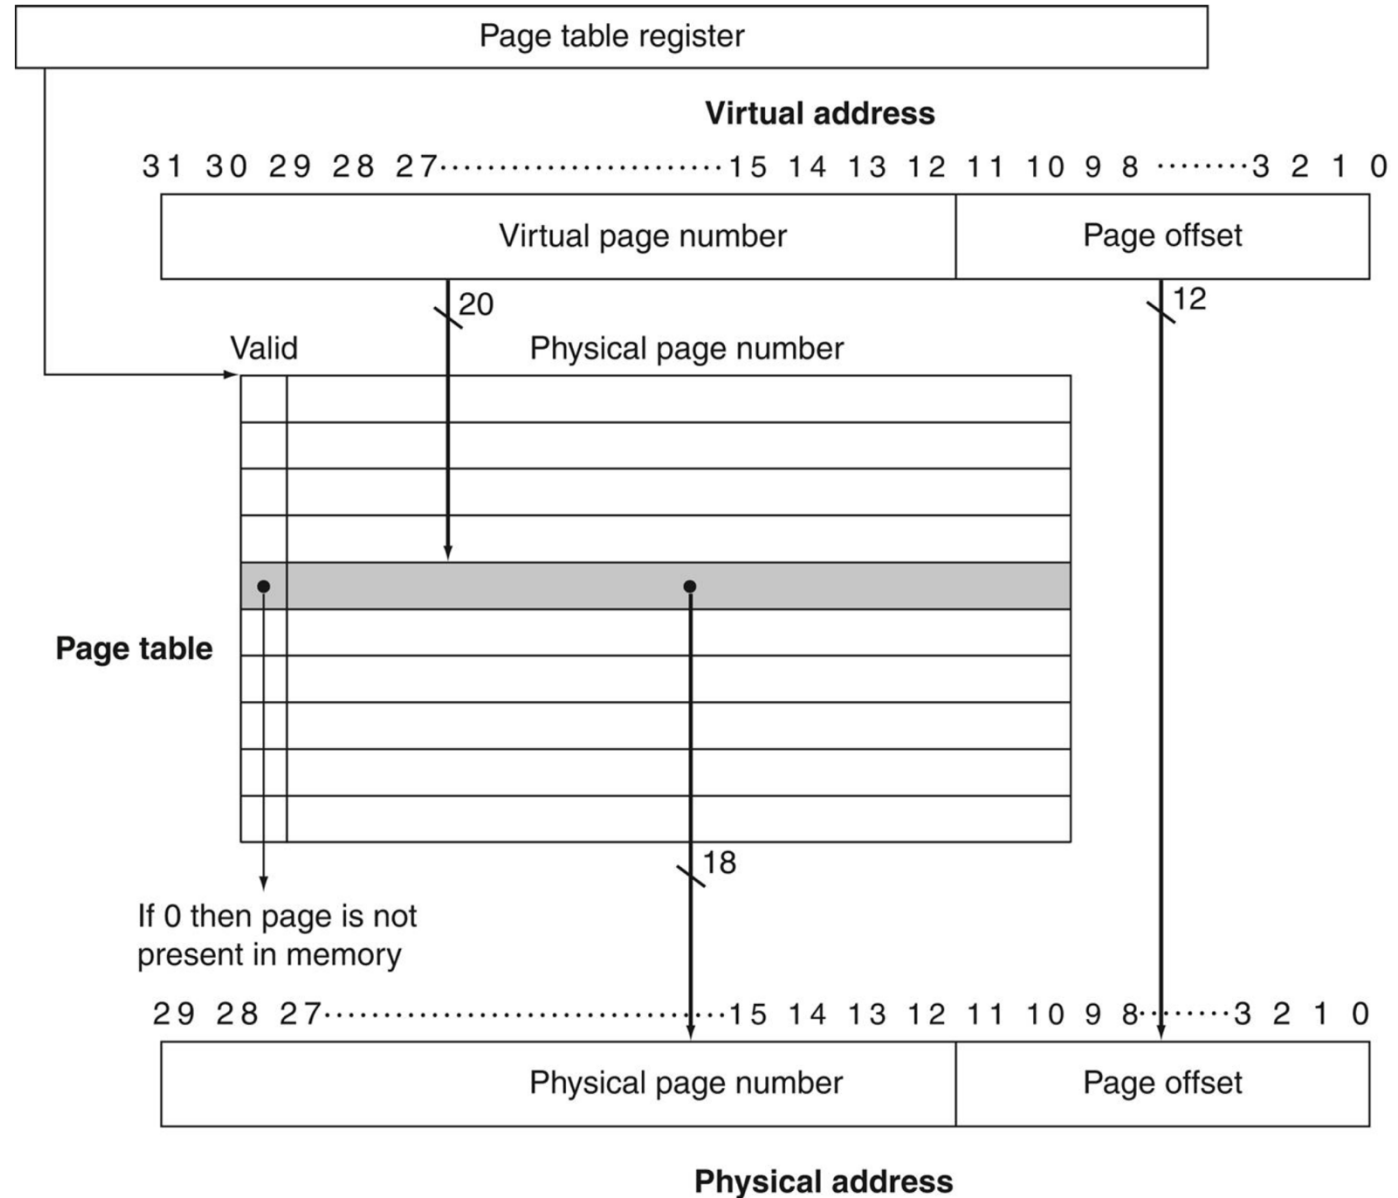
*Where does the page table exist?*

*Physical memory!*

*Then, how can we know the location of the page table itself?*

*Use a page table (base) register*



# Page Table Register



# Page Table Register

CPU register pointing to page table in physical memory (Absolute address)

**Index:** value of the page table register + virtual page number



**Question?**