

Lecture 10


Virtual Memory

Peter Cheung
Imperial College London

URL: www.ee.imperial.ac.uk/pcheung/teaching/EE2_CAS/
E-mail: p.cheung@imperial.ac.uk

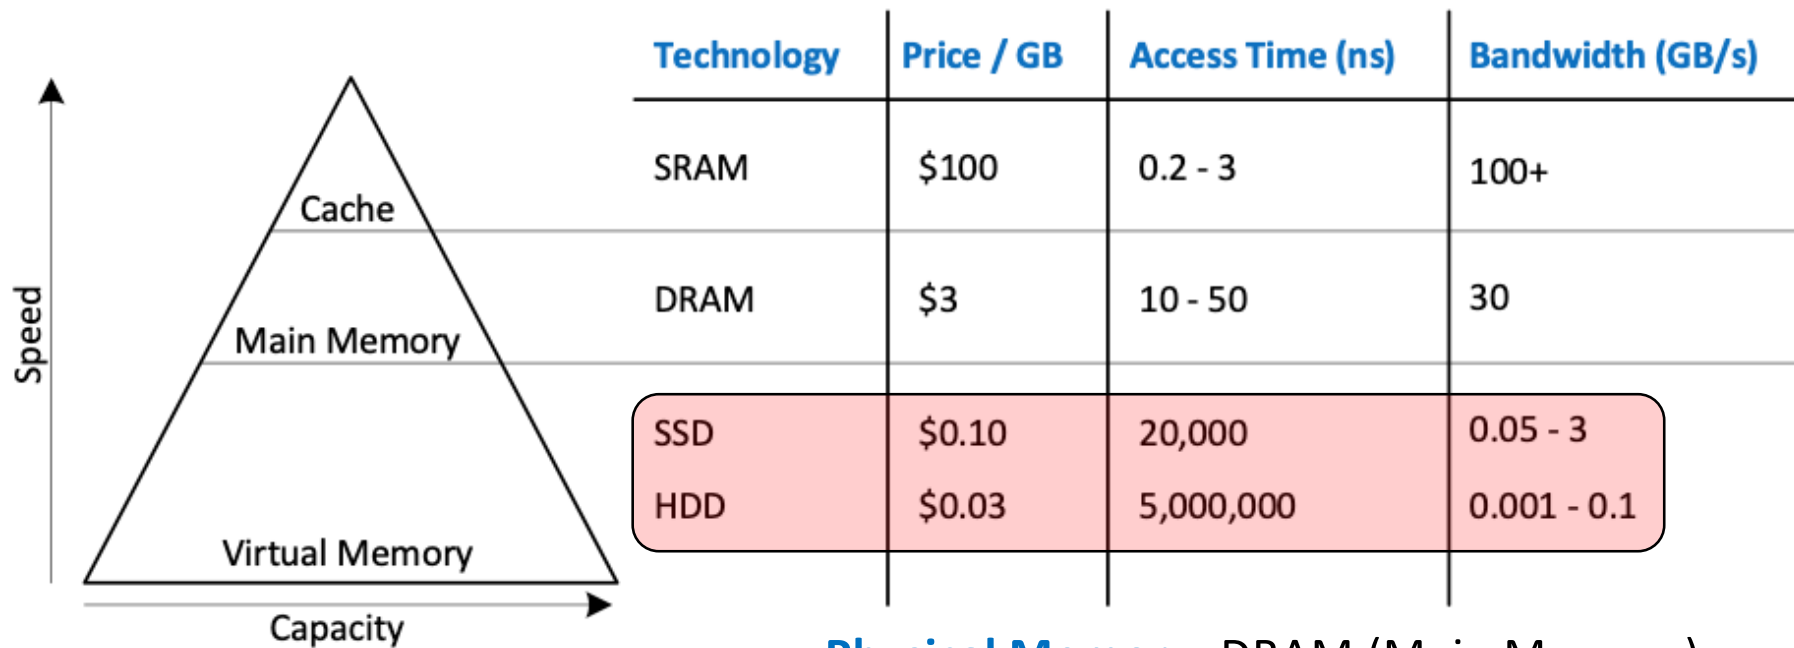

Virtual Memory – What is it?

- Gives the illusion of memory bigger than physical size
- Main memory (DRAM) acts as cache for hard disk



Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

Memory Hierarchy - revisited




- **Physical Memory:** DRAM (Main Memory)
- **Virtual Memory:** Hard drive
 - Slow, Large, Cheap

Based on: "Digital Design and Computer Architecture (RISC-V Edition)" by Sarah Harris and David Harris (H&H),

Two common types of Disk memory

Hard Disk Drive



Solid State Drive



Relatively fast read, often slower write

Takes milliseconds to seek correct location on disk

Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

Virtual Memory Address

- **Virtual addresses**

- Programs use **virtual addresses**
- Entire virtual address space **stored on a hard drive**
- **Subset** of virtual address data in DRAM
- CPU **translates** virtual addresses into **physical addresses** (DRAM)
- Data not in DRAM **fetches** from hard drive

- **Memory Protection**

- Each program has own **virtual to physical mapping**
- Two programs can use same virtual address for different data
- Programs don't need to be aware others are running
- One program (or virus) can't corrupt memory used by another

Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

Virtual Memory Terminologies


Physical memory acts as cache for virtual memory

Cache	Virtual Memory
Block	Page
Block Size	Page Size
Block Offset	Page Offset
Miss	Page Fault
Tag	Virtual Page Number

- **Page size:** amount of memory transferred from hard disk to DRAM at once
- **Address translation:** determining physical address from virtual address
- **Page table:** lookup table used to translate virtual addresses to physical addresses

Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

Virtual to Physical Memory Mapping



© 2007 Elsevier, Inc. All rights reserved

Most accesses hit in physical memory

But programs have the **large capacity** of virtual memory

Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

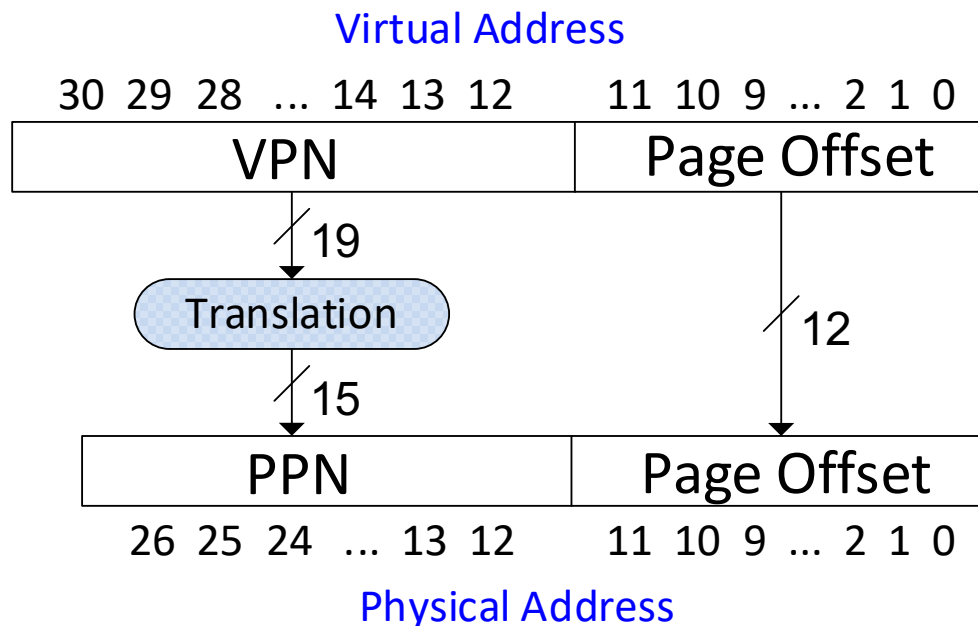
Address Translation Example

- **System:**

- Virtual memory size: 2 GB = 2^{31} bytes
- Physical memory size: 128 MB = 2^{27} bytes
- Page size: 4 KB = 2^{12} bytes

- **Organization:**

- Virtual address: **31** bits
- Physical address: **27** bits
- Page offset: **12** bits
- # Virtual pages = $2^{31}/2^{12} = 2^{19}$
(VPN = **19** bits)
- # Physical pages = $2^{27}/2^{12} = 2^{15}$
(PPN = **15** bits)



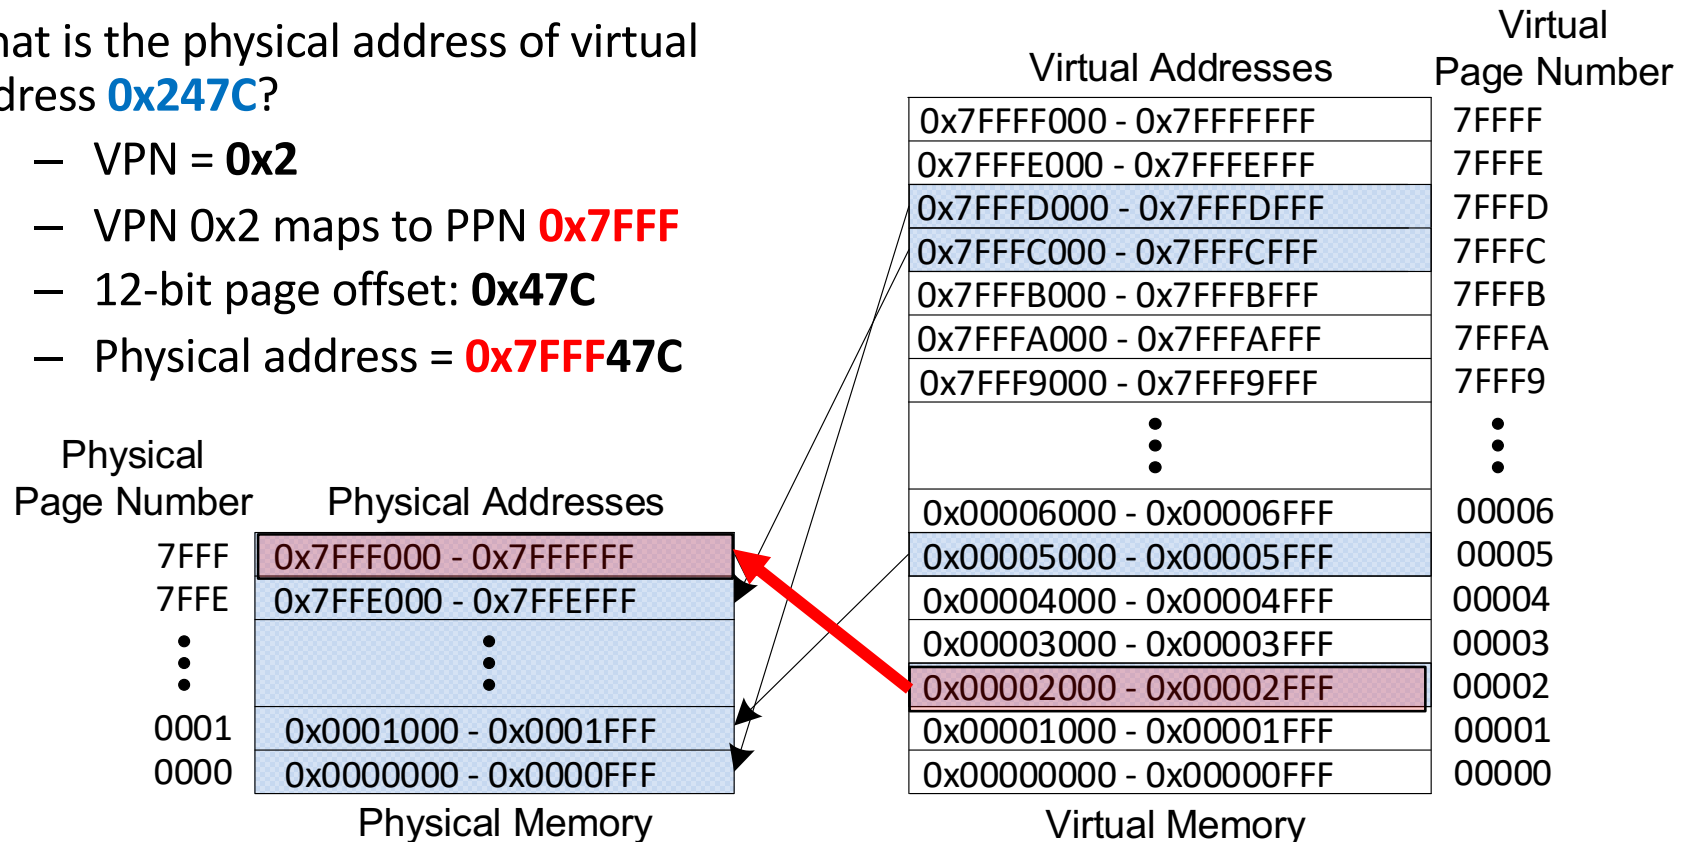
Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

Virtual Memory Example

- 19-bit virtual page numbers (VPN)
- 15-bit physical page numbers (PPN)

What is the physical address of virtual address **0x247C**?

- VPN = **0x2**
- VPN 0x2 maps to PPN **0x7FFF**
- 12-bit page offset: **0x47C**
- Physical address = **0x7FFF47C**

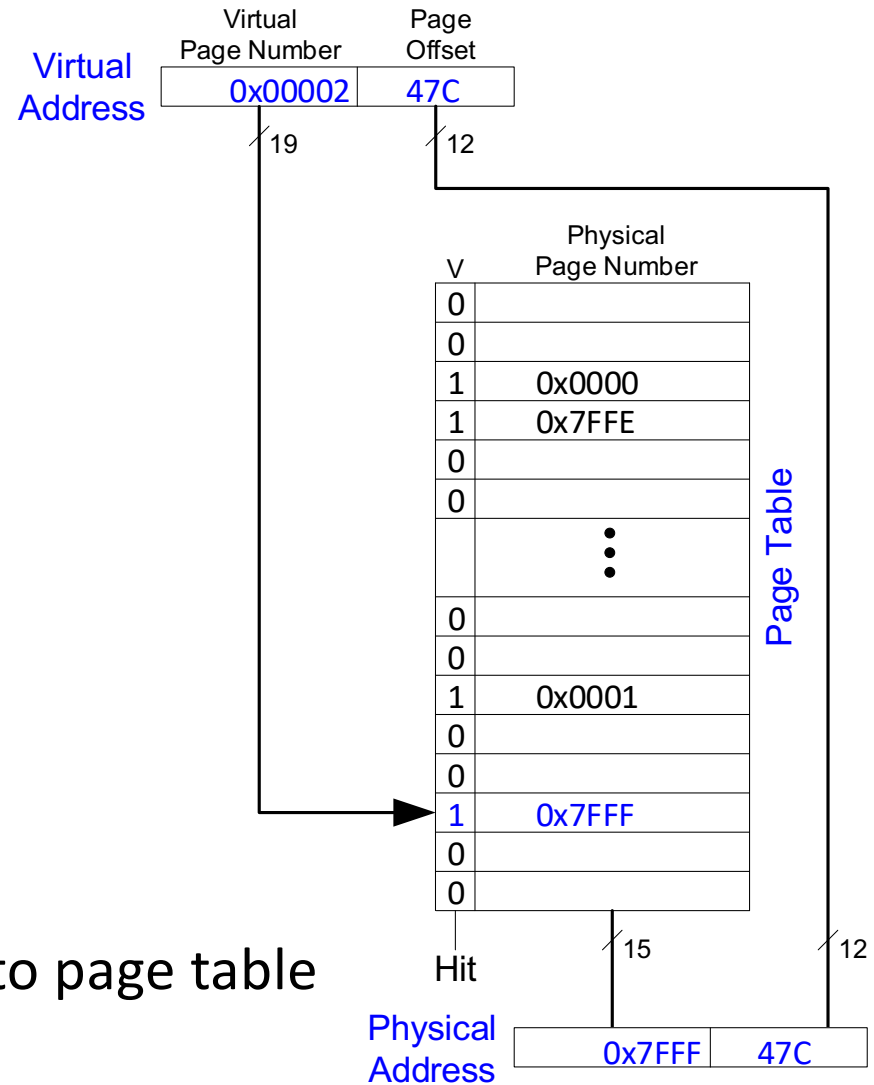


Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

Address Translation using Page Table

- **Page table**

- One entry for **each** virtual page
- Entry fields:
 - **Valid bit:** 1 if page in physical memory
 - **Physical page number:** where the page is located



Virtual Page Number (VPN) is index into page table

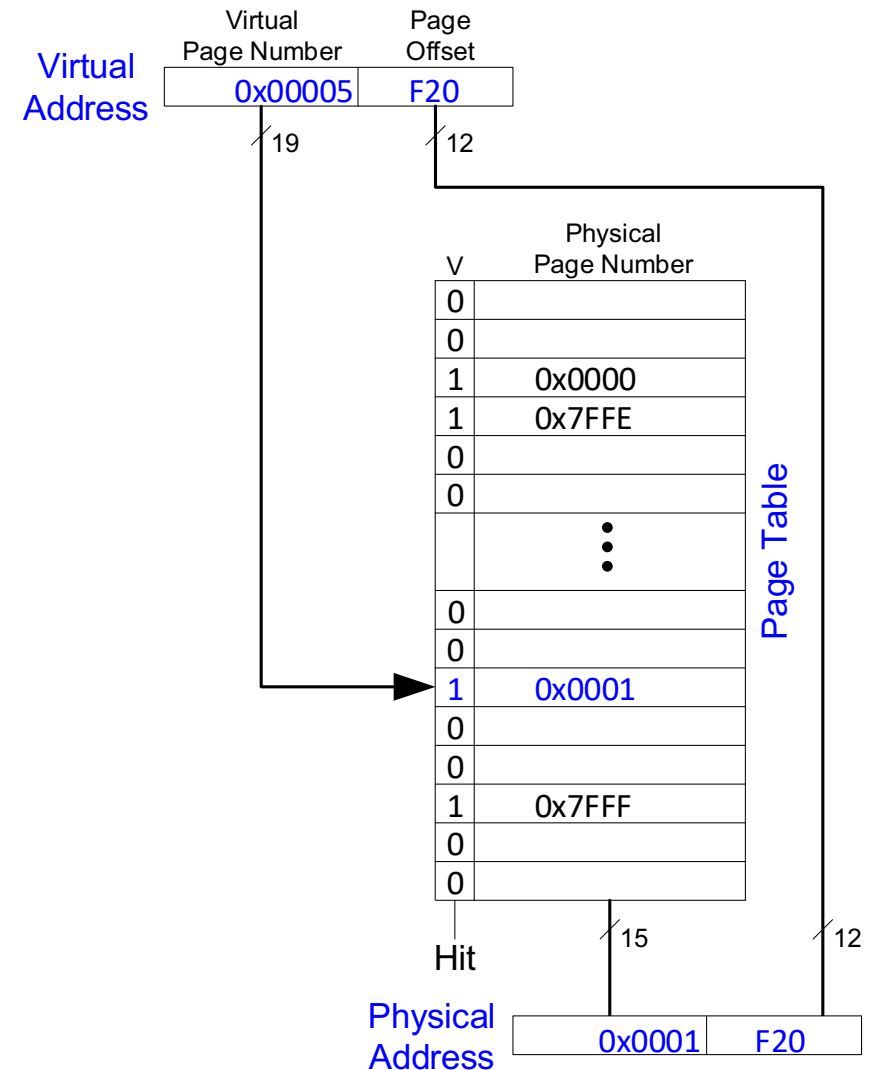
Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

Address Translation Example 1

Virtual address is **0x5F20**

What is the physical address?

- VPN = **5**
- Entry 5 in page table VPN 5
=> physical page **1**
- Physical address: **0x1F20**



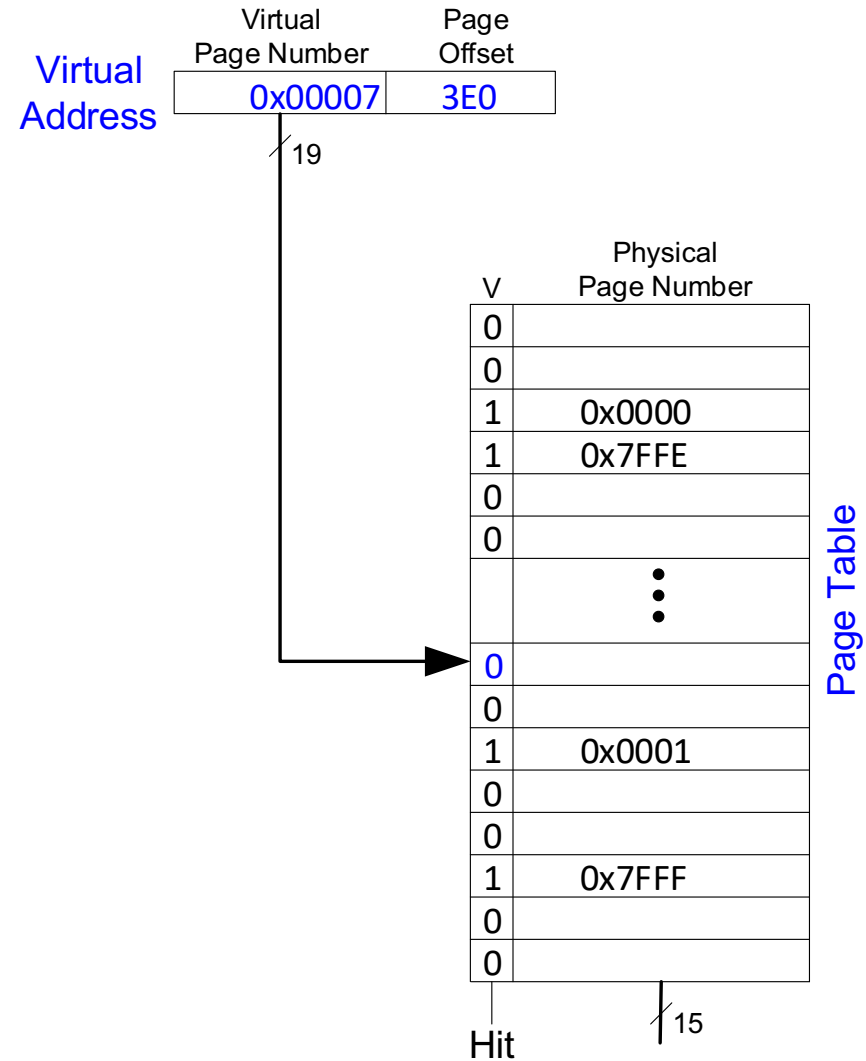
Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

Address Translation Example 2

Virtual address is **0x73E4**

What is the physical address?

- VPN = **7**
- Entry 7 is invalid
- Virtual page must be *paged* into physical memory from disk



Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

Page Table Challenges

- Page table is **large**
 - usually located in physical memory
- Load/store requires **2 main memory accesses**:
 - one for translation (page table read)
 - one to access data (after translation)
- Cuts memory performance in half
 - *Unless we get clever...*

- **Cache most recent translations**
- Reduces number of memory accesses for *most* loads/stores from 2 to **1**


Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

Translation Lookaside Buffer (TLB)

- **Page table accesses:** high temporal locality
 - Large page size, so consecutive loads/stores likely to access same page
- **TLB**
 - **Small:** accessed in < 1 cycle
 - Typically **16 - 512 entries**
 - **Fully associative**
 - **> 99%** hit rates typical
 - **Reduces number of memory accesses** for most loads/stores from 2 to 1

Based on: "*Digital Design and Computer Architecture (RISC-V Edition)*"
by Sarah Harris and David Harris (H&H),

A 2-entry TLB



Based on: "Digital Design and Computer Architecture (RISC-V Edition)"
by Sarah Harris and David Harris (H&H),

Virtual memory Summary

- **Multiple processes** (programs) run at once
- Each process has its **own page table**
- Each process can use **entire virtual address space**
- A process can only access a **subset of physical pages**: those mapped in its own page table
- Virtual memory increases **capacity**
- A subset of virtual pages in physical memory
- **Page table** maps virtual pages to physical pages – address translation
- A **TLB** speeds up address translation
- Different page tables for different programs provides **memory protection**

Based on: "*Digital Design and Computer Architecture (RISC-V Edition)*"
by Sarah Harris and David Harris (H&H),