

Virtual Memory: Systems

CSci 2021: Machine Architecture and Organization
Lecture #29-30, April 6-8th, 2015

Your instructor: Stephen McCamant

Based on slides originally by:

Randy Bryant, Dave O'Hallaron, Antonia Zhai

1

Outline

- Simple memory system example
- Case study: Core i7/Linux memory system
- Memory mapping

2



Review of Symbols

- Basic Parameters
 - $N = 2^n$: Number of addresses in virtual address space
 - $M = 2^m$: Number of addresses in physical address space
 - $P = 2^p$: Page size (bytes)
- Components of the virtual address (VA)
 - TLBI: TLB index
 - TLBT: TLB tag
 - VPO: Virtual page offset
 - VPN: Virtual page number
- Components of the physical address (PA)
 - PPO: Physical page offset (same as VPO)
 - PPN: Physical page number
 - CO: Byte offset within cache line
 - CI: Cache index
 - CT: Cache tag

3

Simple Memory System Example

- Addressing
 - 14-bit virtual addresses
 - 12-bit physical address
 - Page size = 64 bytes



4

Simple Memory System Page Table

Only show first 16 entries (out of 256)


VPN	PPN	Valid
00	28	1
01	-	0
02	33	1
03	02	1
04	-	0
05	16	1
06	-	0
07	-	0

VPN	PPN	Valid
08	13	1
09	17	1
0A	09	1
0B	-	0
0C	-	0
0D	2D	1
0E	11	1
0F	0D	1

5

Simple Memory System TLB

- 16 entries
- 4-way associative




Set	Tag	PPN	Valid	Tag	PPN	Valid	Tag	PPN	Valid	Tag	PPN	Valid
0	03	-	0	09	0D	1	00	-	0	07	02	1
1	03	2D	1	02	-	0	04	-	0	0A	-	0
2	02	-	0	08	-	0	06	-	0	03	-	0
3	07	-	0	03	0D	1	0A	34	1	02	-	0

6

Simple Memory System Cache


- 16 lines, 4-byte block size
- Physically addressed
- Direct mapped




Idx	Tag	Valid	B0	B1	B2	B3	Idx	Tag	Valid	B0	B1	B2	B3
0	19	1	99	11	23	11	8	24	1	3A	00	51	89
1	15	0	-	-	-	-	9	2D	0	-	-	-	-
2	1B	1	00	02	04	08	A	2D	1	93	15	DA	38
3	36	0	-	-	-	-	B	0B	0	-	-	-	-
4	32	1	43	6D	8F	09	C	12	0	-	-	-	-
5	0D	1	36	72	F0	1D	D	16	1	04	96	34	15
6	31	0	-	-	-	-	E	13	1	83	77	1B	D3
7	16	1	11	C2	DF	03	F	14	0	-	-	-	-

Address Translation Example #1

Virtual Address: 0x03D4




Physical Address




Address Translation Example #2

Virtual Address: 0x0B8F




Physical Address




Address Translation Example #3

Virtual Address: 0x0020



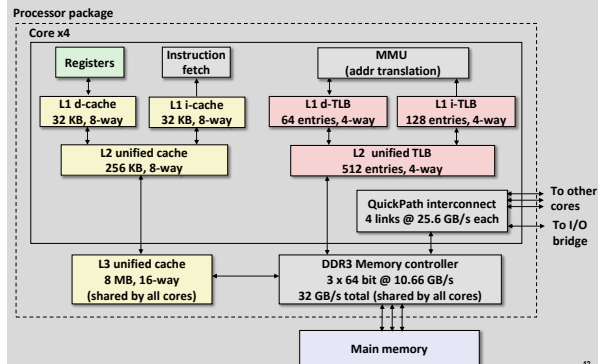
Physical Address



Today

- Simple memory system example
- Case study: Core i7/Linux memory system
- Memory mapping

Intel Core i7 Memory System

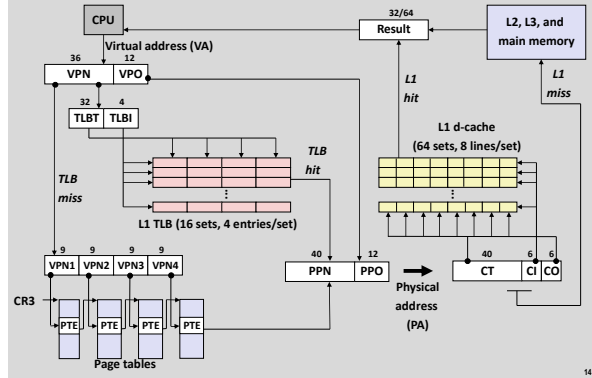


Review of Symbols

- **Basic Parameters**
 - $N = 2^n$: Number of addresses in virtual address space
 - $M = 2^m$: Number of addresses in physical address space
 - $P = 2^p$: Page size (bytes)
- **Components of the virtual address (VA)**
 - TLBI: TLB index
 - TLBT: TLB tag
 - VPO: Virtual page offset
 - VPN: Virtual page number
- **Components of the physical address (PA)**
 - PPO: Physical page offset (same as VPO)
 - PPN: Physical page number
 - CO: Byte offset within cache line
 - CI: Cache index
 - CT: Cache tag

13

End-to-end Core i7 Address Translation



14

Core i7 Level 1-3 Page Table Entries

63	62	52	51	12	11	9	8	7	6	5	4	3	2	1	0					
XD	Unused	Page table physical base address										Unused	G	PS	A	CD	WT	U/S	R/W	P=1
Available for OS (page table location on disk)																				
P=0																				

- Each entry references a 4K child page table**
- P: Child page table present in physical memory (1) or not (0).
- R/W: Read-only or read-write access permission for all reachable pages.
- U/S: User or supervisor (kernel) mode access permission for all reachable pages.
- WT: Write-through or write-back cache policy for the child page table.
- CD: Caching disabled or enabled for the child page table.
- A: Reference bit (set by MMU on reads and writes, cleared by software).
- PS: "Page size": if set, entry points to a large page (1GB or 2MB) not a page table.
- Page table physical base address:** 40 most significant bits of physical page table address (forces page tables to be 4KB aligned)

15

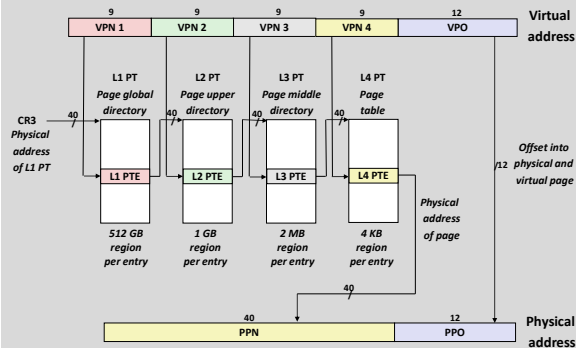
Core i7 Level 4 Page Table Entries

63	62	52	51	12	11	9	8	7	6	5	4	3	2	1	0					
XD	Unused	Page physical base address										Unused	G	D	A	CD	WT	U/S	R/W	P=1
Available for OS (page location on disk)																				
P=0																				

- Each entry references a 4K child page**
- P: Child page is present in memory (1) or not (0)
- R/W: Read-only or read-write access permission for child page
- U/S: User or supervisor (kernel) mode access
- WT: Write-through or write-back cache policy for this page
- CD: Caching disabled (1) or enabled (0)
- A: Reference bit (set by MMU on reads and writes, cleared by software)
- D: Dirty bit (set by MMU on writes, cleared by software)
- Page physical base address:** 40 most significant bits of physical page address (forces pages to be 4KB aligned)


16

Core i7 Page Table Translation



17

Cute Trick for Speeding Up L1 Access



- **Observation**
 - Bits that determine CI identical in virtual and physical address
 - Can index into cache while address translation taking place
 - Generally we hit in TLB, so PPN bits (CT bits) available next
 - "Virtually indexed, physically tagged"
 - Cache carefully sized to make this possible

18


Discussion Point: TLBs and Processes

- Our system will have multiple processes running at once, each with their own address space. How does this affect the way the TLB has to work? (There are several choices)

 1. Clear the TLB when you switch processes: simple, but hurts performance
 2. Include a process/address space identifier as part of TLB entry tags (“tagged TLB”): needs larger TLB
 3. Global page flag: non-global entries are cleared on process switch
 - Compromise makes kernel faster
 - Used in current x86 processors

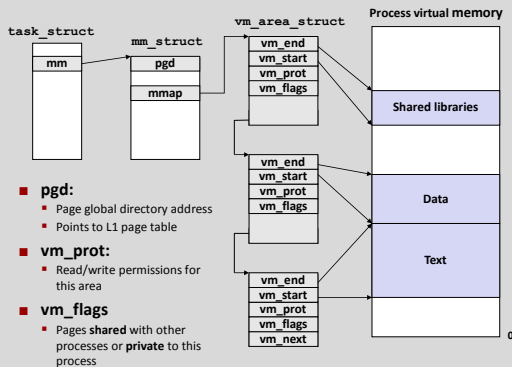
19

Virtual Memory of a Linux Process



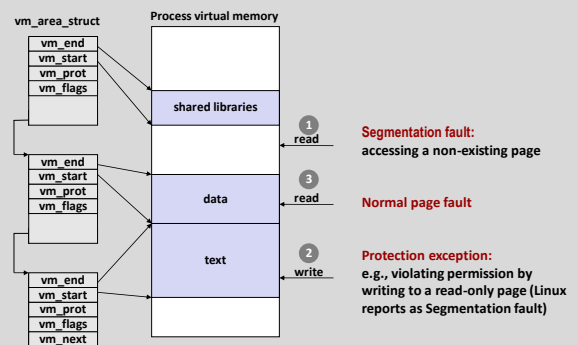
20

Linux Organizes VM as Collection of “Areas”



21

Linux Page Fault Handling



22

Today

- Simple memory system example
- Case study: Core i7/Linux memory system
- Memory mapping

23

Memory Mapping

- VM areas initialized by associating them with disk objects.
 - Process is known as **memory mapping**.
- Area can be backed by (i.e., get its initial values from) :
 - **Regular file** on disk (e.g., an executable object file)
 - Initial page bytes come from a section of a file
 - **Anonymous file** (e.g., nothing)
 - First fault will allocate a physical page full of 0's (**demand-zero page**)
 - Once the page is written to (**dirtied**), it is like any other page
- Dirty pages are copied back and forth between memory and a special **swap file**.


24

Demand paging

- **Key point:** no virtual pages are copied into physical memory until they are referenced!
 - Known as *demand paging*
- Improves time and space efficiency

25


Sharing Revisited: Shared Objects



- Process 1 maps the shared object.

26


Sharing Revisited: Shared Objects



- Process 2 maps the shared object.
- Notice how the virtual addresses can be different.

27

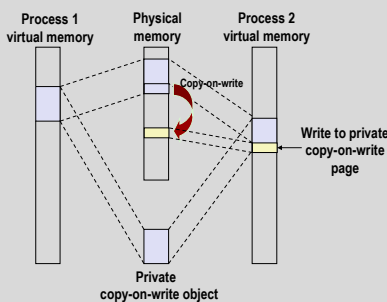
Sharing Revisited: Private Copy-on-write (COW) Objects



- Two processes mapping a *private copy-on-write (COW)* object.
- Area flagged as private copy-on-write
- PTEs in private areas are flagged as read-only

28

Sharing Revisited: Private Copy-on-write (COW) Objects



- Instruction writing to private page triggers protection fault.
- Handler creates new R/W page.
- Instruction restarts upon handler return.
- Copying deferred as long as possible!

29