

Lecture 35
VM II



Lecturer PSOE Dan Garcia

www.cs.berkeley.edu/~ddgarcia

Happy 40th Moore's law! ⇒

Yesterday marked the 40th anniversary of the moment Gordon Moore offered a predication based on trends he'd seen about the semiconductor industry. It has served as an important yardstick!




Review: Caches

- Cache design choices:
 - size of cache: speed v. capacity
 - direct-mapped v. associative
 - for N-way set assoc: choice of N
 - block replacement policy
 - 2nd level cache?
 - Write through v. write back?
- Use performance model to pick between choices, depending on programs, technology, budget, ...



Another View of the Memory Hierarchy



Memory Hierarchy Requirements

- If Principle of Locality allows caches to offer (close to) speed of cache memory with size of DRAM memory, then recursively why not use at next level to give speed of DRAM memory, size of Disk memory?
- While we're at it, what other things do we need from our memory system?



Memory Hierarchy Requirements

- Share memory between multiple processes but still provide protection – don't let one program read/write memory from another
- Address space – give each program the illusion that it has its own private memory
 - Suppose code starts at address 0x40000000. But different processes have different code, both residing at the same address. So each program has a different view of memory.




Virtual Memory

- Called "Virtual Memory"
- Also allows OS to share memory, protect programs from each other
- Today, more important for protection vs. just another level of memory hierarchy
- Each process thinks it has all the memory to itself
- Historically, it predates caches



Virtual to Physical Addr. Translation



- Each program operates in its own virtual address space; ~only program running
- Each is protected from the other
- OS can decide where each goes in memory
- Hardware (HW) provides virtual ⇒ physical mapping



CS61C L35 VM I (7)

Garcia © UCB

Analogy

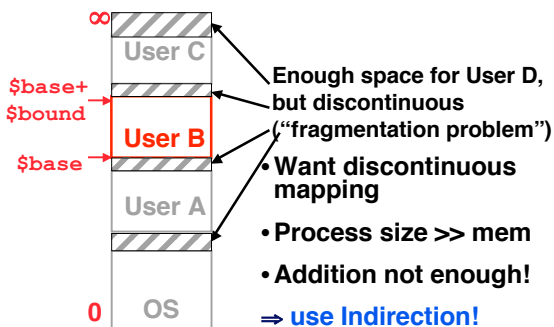
- Book title like **virtual address**
- Library of Congress call number like **physical address**
- Card catalogue like **page table**, mapping from book title to call #
- On card for book, in local library vs. in another branch like **valid bit** indicating in main memory vs. on disk
- On card, available for 2-hour in library use (vs. 2-week checkout) like **access rights**



CS61C L35 VM I (8)

Garcia © UCB


Simple Example: Base and Bound Reg



CS61C L35 VM I (9)

Garcia © UCB


Mapping Virtual Memory to Physical Memory



CS61C L35 VM I (10)

Garcia © UCB

Paging Organization (assume 1 KB pages)



CS61C L35 VM I (11)

Garcia © UCB

Virtual Memory Mapping Function

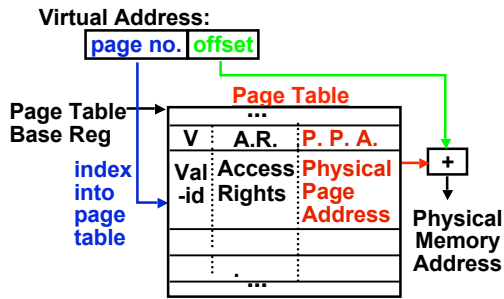
- Cannot have simple function to predict arbitrary mapping
- Use table lookup of mappings
- Use table lookup (“**Page Table**”) for mappings: Page number is index
- Virtual Memory Mapping Function
 - Physical Offset = Virtual Offset
 - Physical Page Number = PageTable[Virtual Page Number]
 - (P.P.N. also called “**Page Frame**”)



CS61C L35 VM I (12)

Garcia © UCB

Address Mapping: Page Table



Page Table located in physical memory



CS61C L35 VM | (13)

Garcia © UCB

Page Table

- A page table is an operating system structure which contains the mapping of virtual addresses to physical locations
 - There are several different ways, all up to the operating system, to keep this data around
- Each process running in the operating system has its own page table
 - “State” of process is PC, all registers, plus page table
 - OS changes page tables by changing contents of **Page Table Base Register**



CS61C L35 VM | (14)

Garcia © UCB

Requirements revisited

- Remember the motivation for VM:
- Sharing memory with protection
 - Different physical pages can be allocated to different processes (sharing)
 - A process can only touch pages in its own page table (protection)
- Separate address spaces
 - Since programs work only with virtual addresses, different programs can have different data/code at the same address!

What about the memory hierarchy?




CS61C L35 VM | (15)

Garcia © UCB

Page Table Entry (PTE) Format

- Contains either Physical Page Number or indication not in Main Memory
- OS maps to disk if Not Valid (V = 0)




- If valid, also check if have permission to use page: **Access Rights (A.R.)** may be Read Only, Read/Write, Executable



CS61C L35 VM | (16)

Garcia © UCB

Paging/Virtual Memory Multiple Processes



CS61C L35 VM | (17)

Garcia © UCB

Comparing the 2 levels of hierarchy

Cache Version	Virtual Memory vers.
Block or Line	Page
Miss	Page Fault
Block Size: 32-64B	Page Size: 4K-8KB
Placement: Direct Mapped, N-way Set Associative	Fully Associative
Replacement: LRU or Random	Least Recently Used (LRU)
Write Thru or Back	Write Back



CS61C L35 VM | (18)

Garcia © UCB

Notes on Page Table

- Solves Fragmentation problem: all chunks same size, so all holes can be used
- OS must reserve “**Swap Space**” on disk for each process
- To grow a process, ask Operating System
 - If unused pages, OS uses them first
 - If not, OS swaps some old pages to disk
 - (Least Recently Used to pick pages to swap)
- Each process has own Page Table
- Will add details, but Page Table is essence of Virtual Memory



Administrivia

- Friday office hours will be moved
- Project is out, due next Wednesday
 - Cache simulator (with a cool GUI)




Virtual Memory Problem #1

- Map every address \Rightarrow 1 indirection via Page Table in memory per virtual address \Rightarrow 1 virtual memory accesses = 2 physical memory accesses \Rightarrow SLOW!
- Observation: since locality in pages of data, there must be locality in **virtual address translations** of those pages
- Since small is fast, why not use a small cache of virtual to physical address translations to make translation fast?
- For historical reasons, cache is called a **Translation Lookaside Buffer**, or **TLB**



Translation Look-Aside Buffers (TLBs)

- TLBs usually small, typically 128 - 256 entries
- Like any other cache, the TLB can be direct mapped, set associative, or fully associative



On TLB miss, get page table entry from main memory



Peer Instruction

- Locality is important yet different for cache and virtual memory (VM): temporal locality for caches but spatial locality for VM
- Cache management is done by hardware (HW), page table management by the operating system (OS), but TLB management is either by HW or OS
- VM helps both with security and cost

	ABC
1:	FFF
2:	FFT
3:	FTF
4:	FTT
5:	TFF
6:	TFT
7:	FTF
8:	TTT



And in conclusion...

- Manage memory to disk? Treat as cache
 - Included protection as bonus, now critical
 - Use Page Table of mappings for each user vs. tag/data in cache
 - TLB is cache of Virtual \Rightarrow Physical addr trans
- Virtual Memory allows protected sharing of memory between processes
- Spatial Locality means Working Set of Pages is all that must be in memory for process to run fairly well

