

**Lecture 36  
 VM II**

2004-11-22



Lecturer PSOE Dan Garcia

[www.cs.berkeley.edu/~ddgarcia](http://www.cs.berkeley.edu/~ddgarcia)

**#4 Bears crush Stanford** →

In the 9<sup>th</sup>-longest rivalry in the US, we get the most dominant win (41-6) since 1930! JJ Arrington ran for 169yds, a school record for a single-season and is now the only RB in the US to have run for 100yds in every game this season.

We now must best Southern Miss on Dec 4...

[calbears.collegesports.com/sports/m-footbl/recaps/112004aac.html](http://calbears.collegesports.com/sports/m-footbl/recaps/112004aac.html)



CS61C L36 VM II (1)

Garcia, Fall 2004 © UCB

Review...

- 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, ...

• Virtual Memory

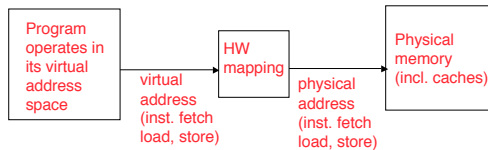
- Predates caches; each process thinks it has all the memory to itself; **protection!**



CS61C L36 VM II (2)

Garcia, Fall 2004 © UCB

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 L36 VM II (3)

Garcia, Fall 2004 © UCB

Analogy

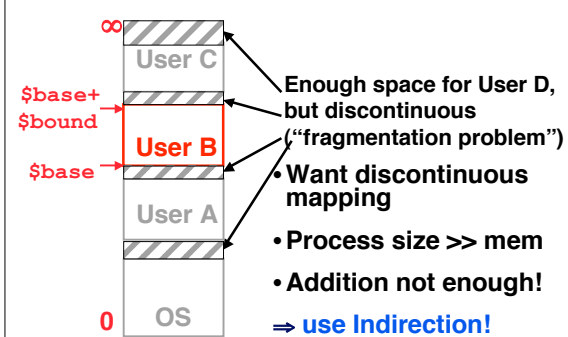
- Book title like **virtual address**
- Library of Congress call number like **physical address**
- Card catalogue like **page table**, mapping from book title to call #
- On card for book, in local library vs. in another branch like **valid bit** indicating in main memory vs. on disk
- On card, available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



CS61C L36 VM II (4)

Garcia, Fall 2004 © UCB

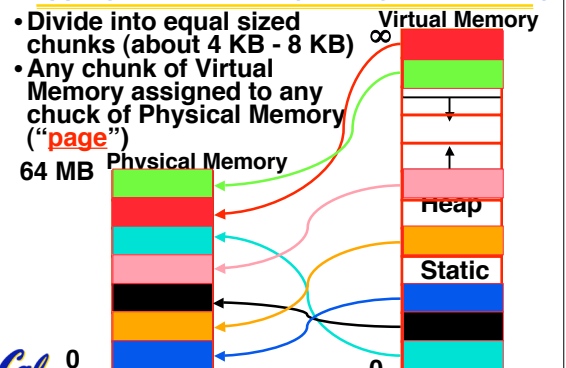
Simple Example: Base and Bound Reg



CS61C L36 VM II (5)

Garcia, Fall 2004 © UCB

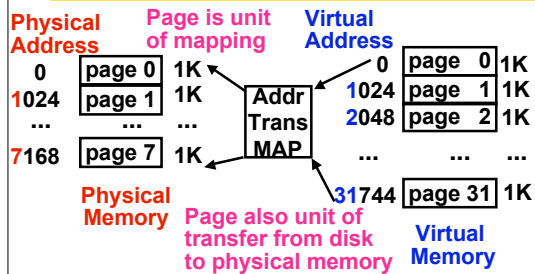
Mapping Virtual Memory to Physical Memory



CS61C L36 VM II (6)

Garcia, Fall 2004 © UCB

### Paging Organization (assume 1 KB pages)



CS61C L36 VM II (7)

Garcia, Fall 2004 © UCB

### Virtual Memory Mapping Function

- Cannot have simple function to predict arbitrary mapping
- Use table lookup of mappings
 

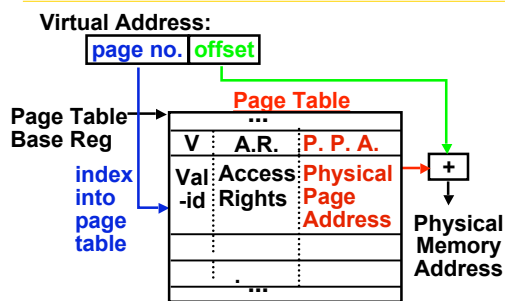
|             |        |
|-------------|--------|
| Page Number | Offset |
|-------------|--------|
- Use table lookup (“Page Table”) for mappings: Page number is index
- Virtual Memory Mapping Function
  - Physical Offset = Virtual Offset
  - Physical Page Number = PageTable[Virtual Page Number]
 (P.P.N. also called “Page Frame”)



CS61C L36 VM II (8)

Garcia, Fall 2004 © UCB

### Address Mapping: Page Table



Page Table located in physical memory

CS61C L36 VM II (9)

Garcia, Fall 2004 © 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 L36 VM II (10)

Garcia, Fall 2004 © 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 L36 VM II (11)

Garcia, Fall 2004 © UCB

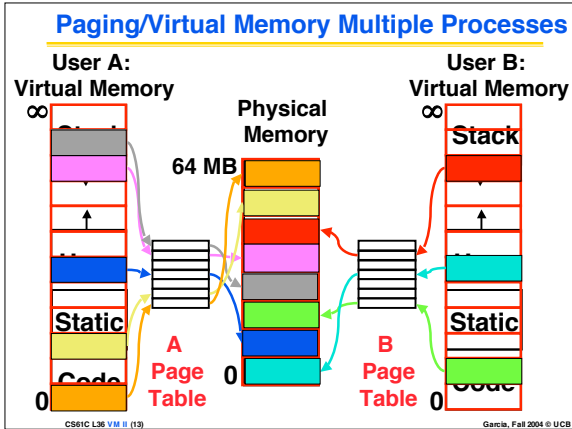
### Page Table Entry (PTE) Format

- Contains either Physical Page Number or indication not in Main Memory
  - OS maps to disk if Not Valid (V = 0)
- | V        | A.R.           | P. P. N.             |
|----------|----------------|----------------------|
| Val: -id | Access: Rights | Physical Page Number |
| V        | A.R.           | P. P. N.             |
- ← P.T.E.
- If valid, also check if have permission to use page: Access Rights (A.R.) may be Read Only, Read/Write, Executable



CS61C L36 VM II (12)

Garcia, Fall 2004 © UCB



### Comparing the 2 levels of hierarchy

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

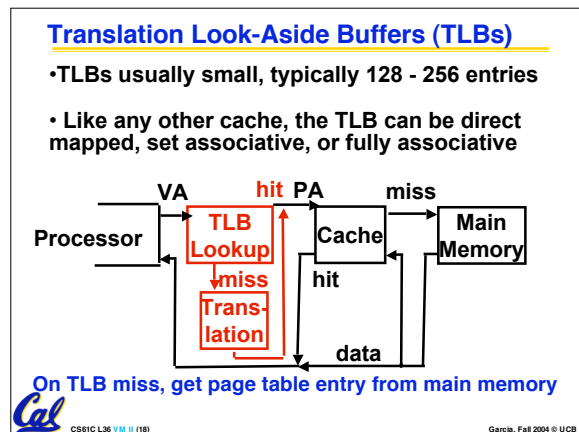
CS61C L36 VM II (14) Garcia, Fall 2004 © UCB

- ### 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
- CS61C L36 VM II (15) Garcia, Fall 2004 © UCB

### Administrivia?

CS61C L36 VM II (16) Garcia, Fall 2004 © UCB

- ### 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**
- CS61C L36 VM II (17) Garcia, Fall 2004 © UCB



### Typical TLB Format

| Virtual Address | Physical Address | Dirty | Ref | Valid | Access Rights |
|-----------------|------------------|-------|-----|-------|---------------|
|                 |                  |       |     |       |               |

- TLB just a cache on the page table mappings
- TLB access time comparable to cache (much less than main memory access time)
- **Dirty**: since use write back, need to know whether or not to write page to disk when replaced
- **Ref**: Used to help calculate LRU on replacement
  - Cleared by OS periodically, then checked to see if page was **referenced**



CS61C L36 VM II (19)

Garcia, Fall 2004 © UCB

### What if not in TLB?

- **Option 1**: Hardware checks page table and loads new Page Table Entry into TLB
- **Option 2**: Hardware traps to OS, up to OS to decide what to do
  - MIPS follows Option 2: Hardware knows nothing about page table



CS61C L36 VM II (20)

Garcia, Fall 2004 © UCB

### What if the data is on disk?

- We load the page off the disk into a free block of memory, using a DMA (Direct Memory Access – very fast!) transfer
  - Meantime we switch to some other process waiting to be run
- When the DMA is complete, we get an interrupt and update the process's page table
  - So when we switch back to the task, the desired data will be in memory



CS61C L36 VM II (21)

Garcia, Fall 2004 © UCB

### What if we don't have enough memory?

- We chose some other page belonging to a program and transfer it onto the disk if it is dirty
  - If clean (disk copy is up-to-date), just overwrite that data in memory
  - We chose the page to evict based on replacement policy (e.g., LRU)
- And update that program's page table to reflect the fact that its memory moved somewhere else
- If continuously swap between disk and memory, called **Thrashing**



CS61C L36 VM II (22)

Garcia, Fall 2004 © UCB

### 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

|    |    |    |
|----|----|----|
| A  | B  | C  |
| 1: | FF | FF |
| 2: | FF | TT |
| 3: | FT | FF |
| 4: | FT | TT |
| 5: | TF | FF |
| 6: | TF | TT |
| 7: | TF | FF |
| 8: | TT | TT |



CS61C L36 VM II (23)

Garcia, Fall 2004 © UCB

### 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



CS61C L36 VM II (25)

Garcia, Fall 2004 © UCB