


CS 61C: Great Ideas in Computer Architecture

Lecture 23: Virtual Memory Part 2

Instructor: Sagar Karandikar
sagark@eecs.berkeley.edu

<http://inst.eecs.berkeley.edu/~cs61c>



Review

- Programmed I/O
- Polling vs. Interrupts
- Booting a Computer
 - BIOS, Bootloader, OS Boot, Init
- Supervisor Mode, Syscalls
- Base and Bounds
 - Simple, but doesn't give us everything we want
- Intro to VM

You Are Here!

Software

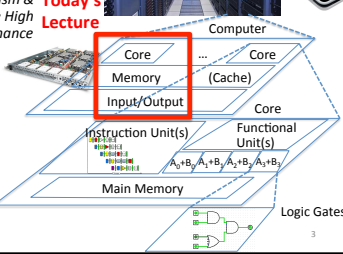
- Parallel Requests
Assigned to computer
e.g., Search "Katz"
- Parallel Threads
Assigned to core
e.g., Lookup, Ads
- Parallel Instructions
>1 instruction @ one time
e.g., 5 pipelined instructions
- Parallel Data
>1 data item @ one time
e.g., Add of 4 pairs of words
- Hardware descriptions
All gates @ one time
- Programming Languages

Hardware

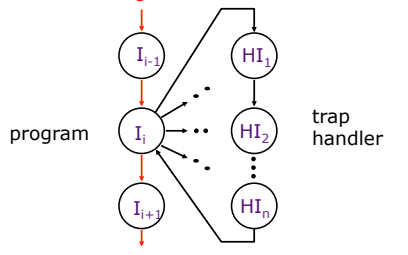
Warehouse Scale Computer

Smart Phone

Today's Lecture



Traps/Interrupts/Exceptions:
altering the normal flow of control



An external or internal event that needs to be processed by another (system) program. The event is usually unexpected or rare from program's point of view.

Terminology

In CS61C (you'll see other definitions in use elsewhere):

- Interrupt – caused by an event external to current running program (e.g. key press, mouse activity)
 - Asynchronous to current program, can handle interrupt on any convenient instruction
- Exception – caused by some event during execution of one instruction of current running program (e.g., page fault, illegal instruction)
 - Synchronous, must handle exception on instruction that causes exception
- Trap – action of servicing interrupt or exception by hardware jump to "trap handler" code

Precise Traps

- Trap handler's view of machine state is that every instruction prior to the trapped one has completed, and no instruction after the trap has executed.
- Implies that handler can return from an interrupt by restoring user registers and jumping to EPC
 - Interrupt handler software doesn't need to understand the pipeline of the machine, or what program was doing!
 - More complex to handle trap caused by an exception
- Providing precise traps is tricky in a pipelined superscalar out-of-order processor!
 - But handling imprecise interrupts in software is even worse.

Trap Handling in 5-Stage Pipeline

- How to handle multiple simultaneous exceptions in different pipeline stages?
- How and where to handle external asynchronous interrupts?

7

Save Exceptions Until Commit

8

Handling Traps in In-Order Pipeline

- Hold exception flags in pipeline until commit point (M stage)
- Exceptions in earlier pipe stages override later exceptions *for a given instruction*
- Inject external interrupts at commit point (override others)
- If exception/interrupt at commit: update Cause and EPC registers, kill all stages, inject handler PC into fetch stage

9

Trap Pipeline Diagram

	time	t0	t1	t2	t3	t4	t5	t6	t7	...
(I ₁) 096: ADD		IF ₁	ID ₁	EX ₁	MA ₁	overflow!				
(I ₂) 100: XOR			IF ₂	ID ₂	EX ₂					
(I ₃) 104: SUB				IF ₃	ID ₃					
(I ₄) 108: ADD					IF ₄					
(I ₅) Trap Handler code						IF ₅	ID ₅	EX ₅	MA ₅	WB ₅

10

Virtual Memory

11


"Bare" 5-Stage Pipeline

- In a bare machine, the only kind of address is a physical address

12

What do we need Virtual Memory for? Reason 1: Adding Disks to Hierarchy


- Need to devise a mechanism to “connect” memory and disk in the memory hierarchy



13

What do we need Virtual Memory for? Reason 2: Simplifying Memory for Apps

- Applications should see the straightforward memory layout we saw earlier ->
- User-space applications should think they own all of memory
- So we give them a **virtual view of memory**



14

What do we need Virtual Memory for? Reason 3: Protection Between Processes

- With a bare system, addresses issued with loads/stores are real **physical** addresses
- This means any program can issue any address, therefore can access any part of memory, even areas which it doesn't own
 - Ex: The OS data structures
- We should send all addresses through a mechanism that the OS controls, before they make it out to DRAM - a **translation mechanism**

15

Address Spaces

- The set of addresses labeling all of memory that we can access
- Now, 2 kinds:
 - Virtual Address Space - the set of addresses that the user program knows about
 - Physical Address Space - the set of addresses that map to actual physical cells in memory
 - Hidden from user applications
- So, we need a way to map between these two address spaces

16

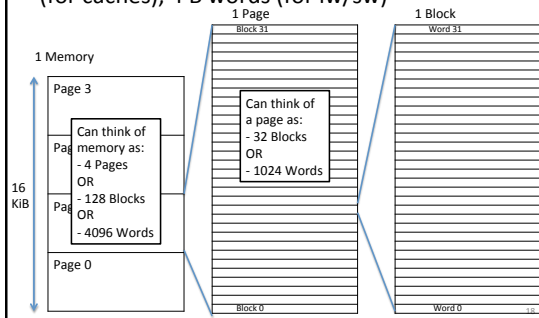
Blocks vs. Pages

- In caches, we dealt with individual blocks
 - Usually ~64B on modern systems
 - We could “divide” memory into a set of blocks
- In VM, we deal with individual pages
 - Usually ~4 KiB on modern systems
 - Now, we'll “divide” memory into a set of pages
- Common point of confusion: Bytes, Words, Blocks, Pages are all just different ways of looking at memory!

17

Bytes, Words, Blocks, Pages


16 KiB DRAM, 4 KiB Pages (for VM), 128 B blocks (for caches), 4 B words (for lw/sw)



18

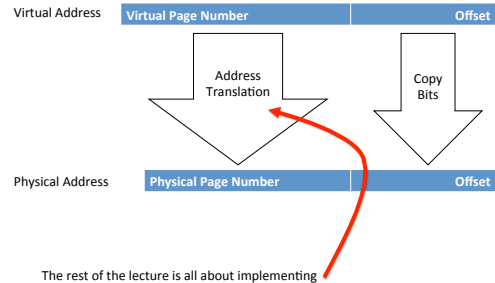
Address Translation

- So, what do we want to achieve at the hardware level?
 - Take a Virtual Address, that points to a spot in the Virtual Address Space of a particular program, and map it to a Physical Address, which points to a physical spot in DRAM of the whole machine



19


Address Translation



20

Paged Memory Systems

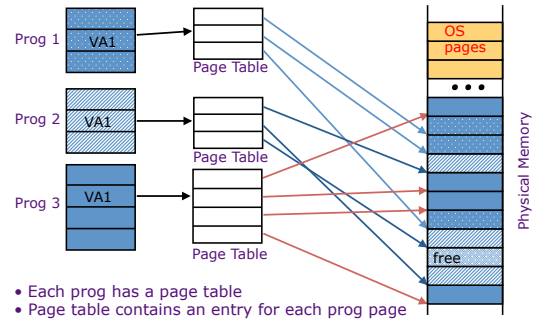
- Processor-generated address can be split into:
 - Virtual Page Number
 - Offset
- A page table contains the physical address of the base of each page



Page tables make it possible to store the pages of a program non-contiguously.

21

Private (Virtual) Address Space per Program



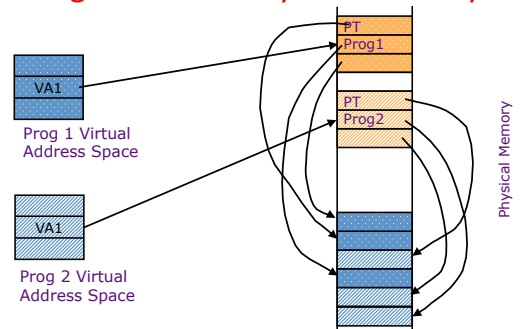
22

Where Should Page Tables Reside?

- Space required by the page tables (PT) is proportional to the address space, number of users, ...
 - ⇒ Too large to keep in registers inside CPU
- Idea: Keep page tables in the main memory
 - Needs one reference to retrieve the page base address and another to access the data word
 - ⇒ doubles the number of memory references! (but we can fix this using something we already know about...)

23

Page Tables in Physical Memory



24

Administrivia

- Project 3-2 due tonight
- Project 4 out now
 - Competition for glory + EC, awards during last lecture
- HW6 Out – VM, I/O, Parity, ECC
- Guerrilla Section on VM, I/O, ECC, on Thursday from 5-7pm, Woz
- Sign up for tutoring on Piazza

25

Administrivia

- Upcoming Lecture Schedule
 - 8/03: VM (today)
 - 8/04: I/O: DMA, Disks, Networking
 - 8/05: Dependability: Parity, ECC, RAID
 - Last day of new material
 - 8/06: Final Exam Review, Day 1 (Formerly GPUs)
 - 8/10: Final Exam Review, Day 2 (Formerly tools)
 - 8/11: Summary, What's Next? (+ HKN reviews)
 - Project 4 Competition Winners Announced
 - 8/12: No Lecture, I'll have OH in this room

26

Administrivia

- Final Exam is next Thursday (8/13)
 - **9am**-12pm, 10 Evans
 - More info soon

27

CS61C In the News: Moore's Law's 50th Anniversary this year!




Fig. 2 Number of components per integrated function (at minimum cost per component extrapolated vs time).

Gordon Moore's paper appeared in 19 April 1965 issue of *Electronics*.

"With unit cost falling as the number of components per circuit rises, by 1975 economics may dictate squeezing as many as 65,000 components on a single silicon chip."

Today: 18-core Intel Xeon Haswell-E5 5.56 **billion** transistors on a chip

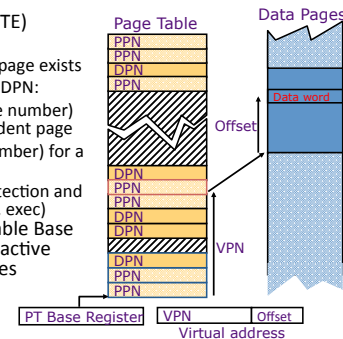
28

Break

29

Linear Page Table

- Page Table Entry (PTE) contains:
 - 1 bit to indicate if page exists
 - And either PPN or DPN:
 - PPN (physical page number) for a memory-resident page
 - DPN (disk page number) for a page on the disk
 - Status bits for protection and usage (read, write, exec)
- OS sets the Page Table Base Register whenever active user process changes



30

What if the page isn't in DRAM?

- We get a "page fault":
 - Initiate transfer of the page we're requesting from disk to DRAM, should place it into an unused page
 - If no unused page is left, a *page currently in DRAM is selected to be replaced* (based on usage)
 - The replaced page is written back to disk, page table entry that maps that VPN->PPN is marked as invalid
 - Page table entry of the page we're requesting is updated with a (now) valid PPN

31

Size of Linear Page Table

With 32-bit addresses, 4-KB pages & 4-byte PTEs:
 ⇒ 2^{20} PTEs, i.e, 4 MB page table per user
 ⇒ 4 GB of swap needed to back up full virtual address space

Larger pages?

- Internal fragmentation (Not all memory in page is used)
- Larger page fault penalty (more time to read from disk)

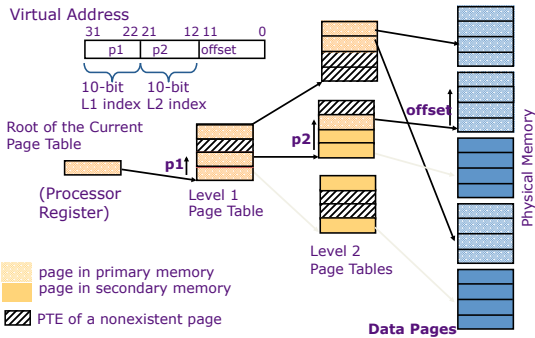
What about 64-bit virtual address space???

- Even 1MB pages would require 2^{44} 8-byte PTEs (35 TB!)

What is the "saving grace" ?

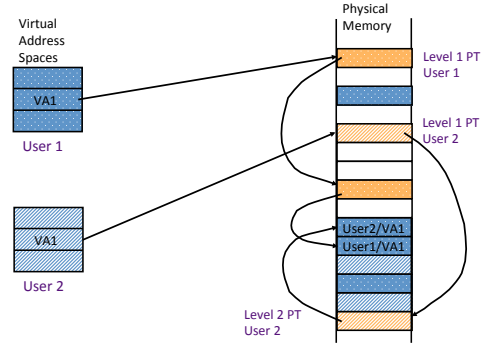
32

Hierarchical Page Table




33

Two-Level Page Tables in Physical Memory



34

Address Translation & Protection



- Every instruction and data access needs address translation and protection checks

A good VM design needs to be fast (~ one cycle) and space efficient


35

Translation Lookaside Buffers (TLB)

Address translation is very expensive!
 In a two-level page table, each reference becomes several memory accesses

Solution: Cache translations in TLB

- TLB hit ⇒ Single-Cycle Translation
- TLB miss ⇒ Page-Table Walk to refill



36

TLB Designs


- Typically 32-128 entries, usually fully associative
 - Each entry maps a large page, hence less spatial locality across pages → more likely that two entries conflict
 - Sometimes larger TLBs (256-512 entries) are 4-8 way set-associative
 - Larger systems sometimes have multi-level (L1 and L2) TLBs
- Random or FIFO replacement policy
- No process information in TLB?
- TLB Reach: Size of largest virtual address space that can be simultaneously mapped by TLB

Example: 64 TLB entries, 4KB pages, one page per entry

TLB Reach = _____?

37

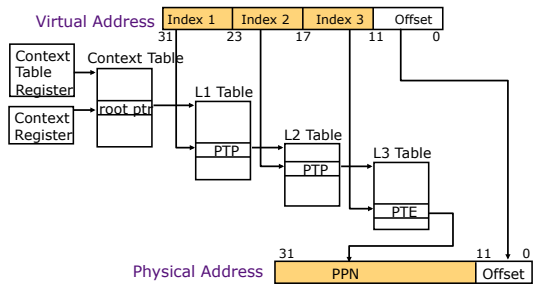
VM-related events in pipeline



- Handling a TLB miss needs a hardware or software mechanism to refill TLB
 - usually done in hardware now
- Handling a page fault (e.g., page is on disk) needs a *precise* trap so software handler can easily resume after retrieving page
- Handling protection violation may abort process

38

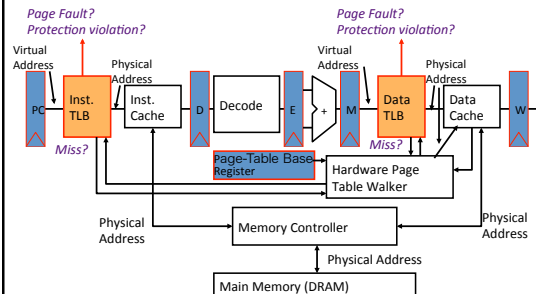
Hierarchical Page Table Walk: SPARC v8



MMU does this table walk in hardware on a TLB miss

39

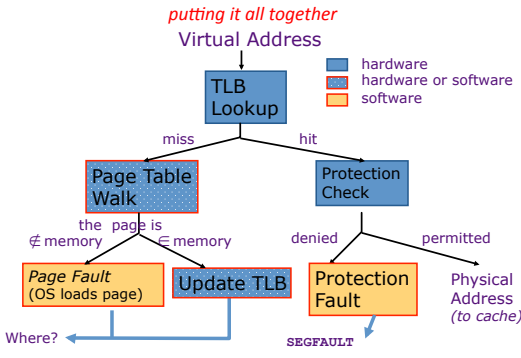
Page-Based Virtual-Memory Machine (Hardware Page-Table Walk)



- Assumes page tables held in untranslated physical memory

40

Address Translation: putting it all together



41

Modern Virtual Memory Systems


Illusion of a large, private, uniform store

Protection & Privacy
 several users, each with their private address space and one or more shared address spaces
 page table = name space

Demand Paging
 Provides the ability to run programs larger than the primary memory

Hides differences in machine configurations

The price is address translation on each memory reference



42

Clicker Question

Let's try to extrapolate from caches... Which one is false?

- A. # offset bits in V.A. = $\log_2(\text{page size})$
- B. # offset bits in P.A. = $\log_2(\text{page size})$
- C. # VPN bits in V.A. = $\log_2(\# \text{ of physical pages})$
- D. # PPN bits in P.A. = $\log_2(\# \text{ of physical pages})$
- E. A single-level page table contains a PTE for every possible VPN in the system

43

Conclusion: VM features track historical uses

- **Bare machine, only physical addresses**
 - One program owned entire machine
- **Batch-style multiprogramming**
 - Several programs sharing CPU while waiting for I/O
 - Base & bound: translation and protection between programs (not virtual memory)
 - Problem with external fragmentation (holes in memory), needed occasional memory defragmentation as new jobs arrived
- **Time sharing**
 - More interactive programs, waiting for user. Also, more jobs/second.
 - Motivated move to fixed-size page translation and protection, no external fragmentation (but now internal fragmentation, wasted bytes in page)
 - Motivated adoption of virtual memory to allow more jobs to share limited physical memory resources while holding working set in memory
- **Virtual Machine Monitors**
 - Run multiple operating systems on one machine
 - Idea from 1970s IBM mainframes, now common on laptops
 - e.g., run Windows on top of Mac OS X
 - Hardware support for two levels of translation/protection
 - Guest OS virtual -> Guest OS physical -> Host machine physical
 - Also basis of Cloud Computing
 - Virtual machine instances on EC2 for Lab 13

44