

CS 61C: Great Ideas in Computer Architecture

Exceptions/Traps/Interrupts

Instructors:
Krste Asanovic, Randy H. Katz
<http://inst.eecs.Berkeley.edu/~cs61c/fa12>

Fall 2012 -- Lecture #35 1

Review

- Programmed I/O versus DMA
- Polling versus Interrupts
- Asynchronous interrupts versus synchronous traps
- Precise interrupt looks like execution stopped at exactly one instruction, every instruction before finished, no instruction after started.
 - Simplify software view of interrupted state

Fall 2012 -- Lecture #35 2

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

Harness Parallelism & Achieve High Performance

Fall 2012 -- Lecture #35 3

Interrupts

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.

Fall 2012 -- Lecture #35 4

Precise Interrupts

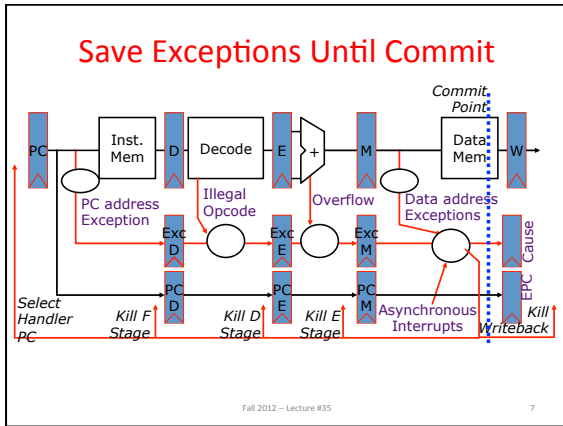
- Interrupt handler's view of machine state is that every instruction prior to the interrupted one has completed, and no instruction after the interrupt has executed.
 - Instruction taking interrupt might have written some special state but can be restarted.
- Implies that handler can return from interrupt by restoring user registers and jumping to EPC
 - Software doesn't need to understand the pipeline of the machine!
- Providing precise interrupts is tricky in a pipelined superscalar out-of-order processor!
 - But handling imprecise interrupts in software is even worse.

Fall 2012 -- Lecture #35 5

Exception Handling in 5-Stage Pipeline

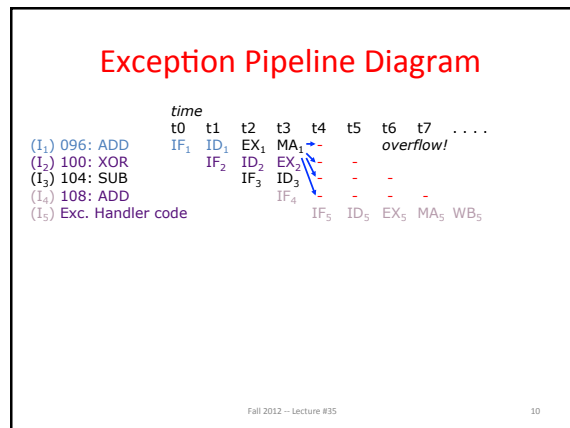
- How to handle multiple simultaneous exceptions in different pipeline stages?
- How and where to handle external asynchronous interrupts?

Fall 2012 -- Lecture #35 6



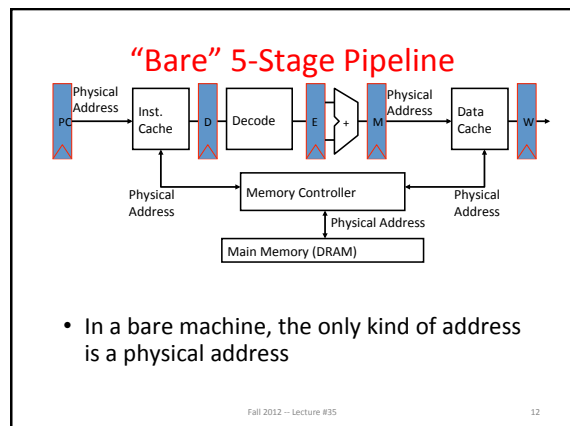
- ### Handling Exceptions in In-Order Pipeline
- Hold exception flags in pipeline until commit point (M stage)
 - Exceptions in earlier pipe stages override later exceptions *for a given instruction*
 - Inject external interrupts at commit point (override others)
 - If exception at commit: update Cause and EPC registers, kill all stages, inject handler PC into fetch stage
- Fall 2012 – Lecture #35 8

- ### Speculating on Exceptions to avoid Control Hazard
- Prediction mechanism
 - Exceptions are rare, so simply predicting no exceptions is very accurate!
 - Check prediction mechanism
 - Exceptions detected at end of instruction execution pipeline, special hardware for various exception types
 - Recovery mechanism
 - Only write architectural state at commit point, so can throw away partially executed instructions after exception
 - Launch exception handler after flushing pipeline
 - Bypassing allows use of uncommitted instruction results by following instructions
- Fall 2012 – Lecture #35 9



Virtual Memory

Fall 2012 – Lecture #35 11



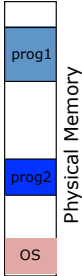
Dynamic Address Translation

Motivation
 In early machines, I/O operations were slow and each word transferred involved the CPU
 Higher throughput if CPU and I/O of 2 or more programs were overlapped.
 How? ⇒ multiprogramming with DMA I/O devices, interrupts

Location-independent programs
 Programming and storage management ease
 ⇒ need for a *base register*

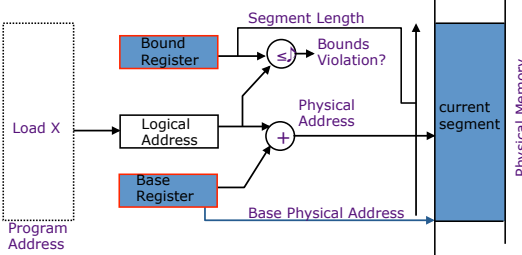
Protection
 Independent programs should not affect each other inadvertently
 ⇒ need for a *bound register*

Multiprogramming drives requirement for resident *supervisor* software to manage context switches between multiple programs



Fall 2012 – Lecture #35 13

Simple Base and Bound Translation

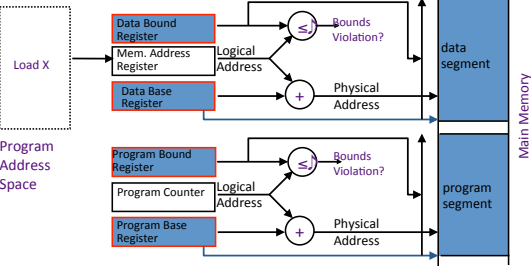


Base and bounds registers are visible/accessible only when processor is running in *kernel mode*

Fall 2012 – Lecture #35 14

Separate Areas for Program and Data

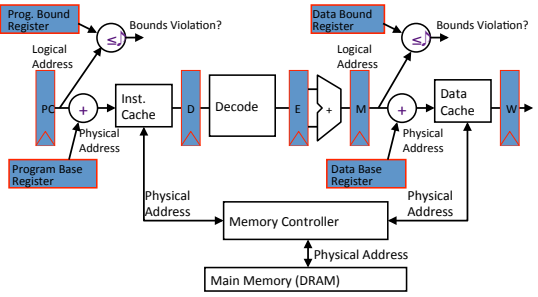
(Scheme used on all Cray vector supercomputers prior to X1, 2002)



What is an advantage of this separation?

Fall 2012 – Lecture #35 15

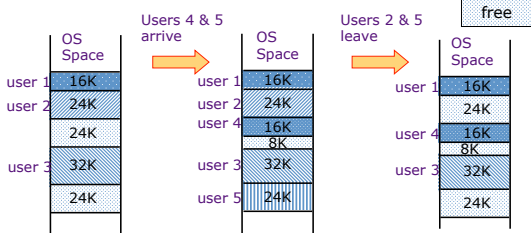
Base and Bound Machine



[Can fold addition of base register into (register+immediate) address calculation using a carry-save adder (sums three numbers with only a few gate delays more than adding two numbers)]

Fall 2012 – Lecture #35 16

Memory Fragmentation

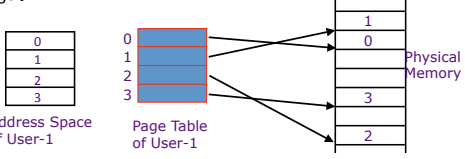


As users come and go, the storage is "fragmented". Therefore, at some stage programs have to be moved around to compact the storage.

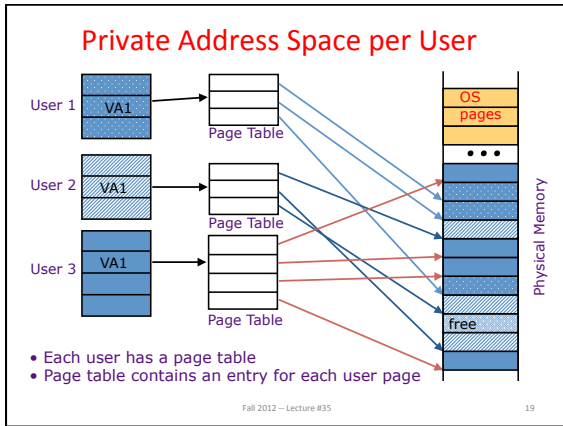
Fall 2012 – Lecture #35 17

Paged Memory Systems

- Processor-generated address can be split into:
 - 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.



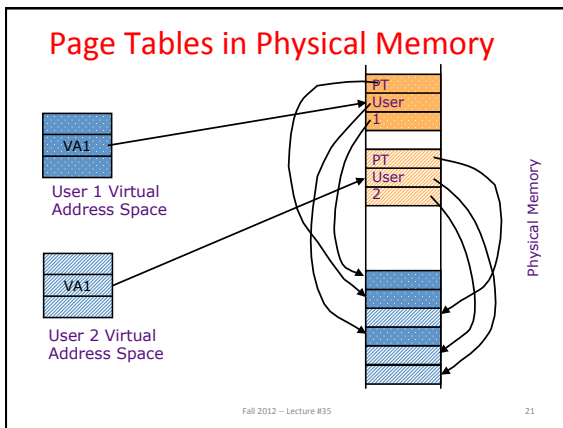
Fall 2012 – Lecture #35 18



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
- Idea: Keep PTs 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!

Fall 2012 – Lecture #35 20



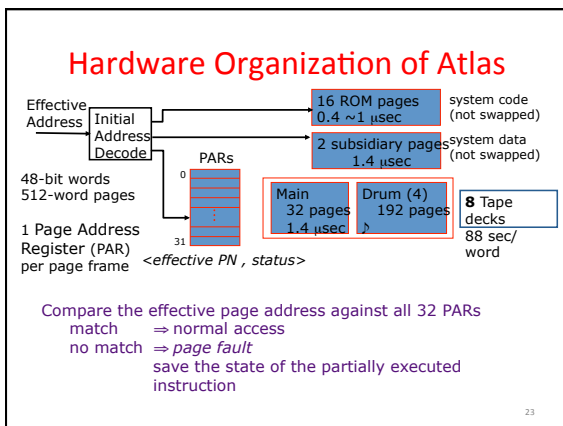
Demand Paging in Atlas (1962)

“A page from secondary storage is brought into the primary storage whenever it is (implicitly) demanded by the processor.”
Tom Kilburn

Primary memory as a *cache* for secondary memory

User sees $32 \times 6 \times 512$ words of storage

22



Atlas Demand Paging Scheme

- On a page fault:
 - Input transfer into a free page is initiated
 - The Page Address Register (PAR) is updated
 - If no free page is left, a page is selected to be replaced (based on usage)
 - The replaced page is written on the drum
 - to minimize drum latency effect, the first empty page on the drum was selected
 - The page table is updated to point to the new location of the page on the drum

24

Administrivia

- Regrade request deadline Monday Nov 26
 - For everything up to Project 4

CS61C In the News:

“Texas Instruments Cuts 1,700 Jobs and Winds Down Tablet Chips”, NY Times 11/14/2012

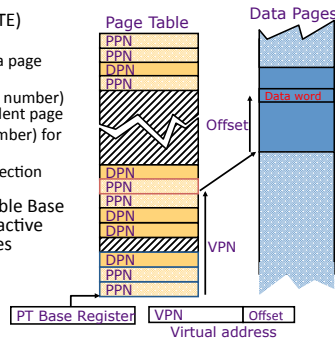
“Texas Instruments is eliminating 1,700 jobs, as it winds down its mobile processor business to focus on chips for more profitable markets like cars and home appliances. Texas Instruments said in September it would halt costly investments in the increasingly competitive smartphone and tablet chip business, leading Wall Street to speculate that part of the company’s processor unit, called OMAP, could be sold. The layoffs are equivalent to nearly 5 percent of the Austin, Texas-based company’s global workforce.

TI has been under pressure in mobile processors, where it has lost ground to rival Qualcomm Inc. Leading smartphone makers Apple Inc and Samsung Electronics Co Ltd have been developing their own chips instead of buying them from suppliers like TI. Instead of competing in phones and tablets, TI wants to sell its OMAP processors in markets that require less investment, like industrial clients like carmakers. TI is expected to continue selling existing tablet and phone processors for products like Amazon.Com Inc’s Kindle tablets for as long as demand remains, but stop developing new chips.”

[Rumors of Amazon being interested in buying this OMAP unit from TI]

Linear Page Table

- Page Table Entry (PTE) contains:
 - A bit to indicate if a page exists
 - 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
- OS sets the Page Table Base Register whenever active user process changes



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?

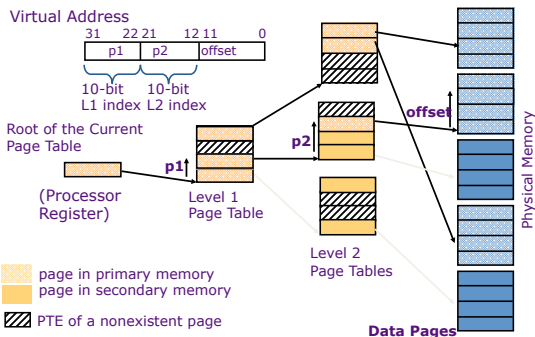
- Internal fragmentation (Not all memory in page is used)
- Larger page fault penalty (more time to read from disk)

What about 64-bit virtual address space???


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

What is the “saving grace”?

Hierarchical Page Table



Two-Level Page Tables in Physical Memory



Address Translation & Protection

Virtual Address: Virtual Page No. (VPN) | offset

Kernel/User Mode, Read/Write → Protection Check

Address Translation

Exception?

Physical Address: Physical Page No. (PPN) | offset

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

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

Fall 2012 -- Lecture #35 31

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

virtual address: VPN | offset

hit? physical address: PPN | offset

(VPN = virtual page number)
(PPN = physical page number)

Fall 2012 -- Lecture #35 32

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 = _____?

Fall 2012 -- Lecture #35 33

Handling a TLB Miss

Software (MIPS, Alpha)
TLB miss causes an exception and the operating system walks the page tables and reloads TLB. A privileged "untranslated" addressing mode used for walk

Hardware (SPARC v8, x86, PowerPC, RISC-V)
A memory management unit (MMU) walks the page tables and reloads the TLB

If a missing (data or PT) page is encountered during the TLB reloading, MMU gives up and signals a Page-Fault exception for the original instruction

Fall 2012 -- Lecture #35 35

Hierarchical Page Table Walk: SPARC v8

Virtual Address: Index 1 | Index 2 | Index 3 | Offset

Context Table Register → Context Table → root ptr → L1 Table → PTP → L2 Table → PTP → L3 Table → PTE → Physical Address: PPN | Offset

MMU does this table walk in hardware on a TLB miss

Fall 2012 -- Lecture #35 36

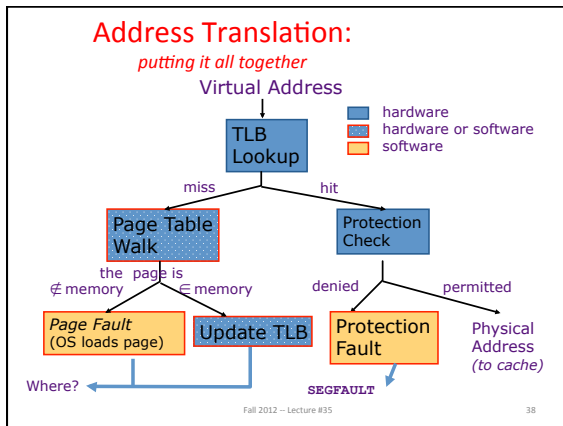
Page-Based Virtual-Memory Machine (Hardware Page-Table Walk)

Virtual Address → Inst. TLB → Inst. Cache → Decode → E → M → Data TLB → Data Cache → W

Page-Table Base Register, Hardware Page Table Walker, Memory Controller, Main Memory (DRAM)

Miss?, Page Fault? Protection violation?

Fall 2012 -- Lecture #35 37



Acknowledgements

- These slides contain material developed and copyright by:
 - Arvind (MIT)
 - Krste Asanovic (MIT/UCB)
 - Joel Emer (Intel/MIT)
 - James Hoe (CMU)
 - John Kubiatowicz (UCB)
 - David Patterson (UCB)
- MIT material derived from course 6.823
- UCB material derived from course CS252