

CS162
Operating Systems and
Systems Programming
Lecture 13


Address Translation (con't)
Caches and TLBs

March 2, 2010

Ion Stoica

<http://inst.eecs.berkeley.edu/~cs162>

Review: Single-Level Translation

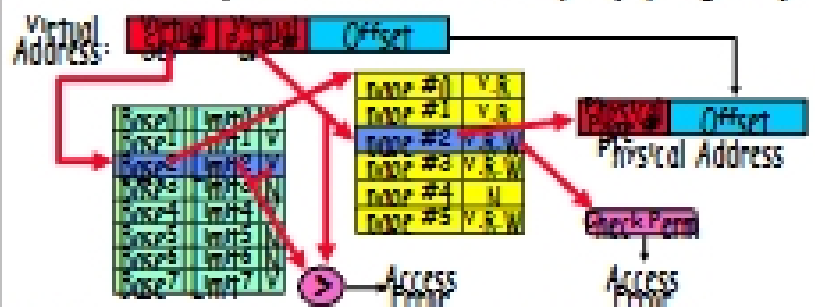


- Advantages
 - Low translation overhead
 - Simplicity
- Disadvantages
 - Large page tables
 - > E.g., 32b address space, 4KB pages → up to 2^{21} = 2M page entries for each process
 - Expensive to share memory
 - > E.g., 4KB pages, want to share 100MB → need to update 25,000 entries in page table

3/2/10 cs162 @ UC Berkeley 13.2 Lec 13.2

Review: Multi-level Translation

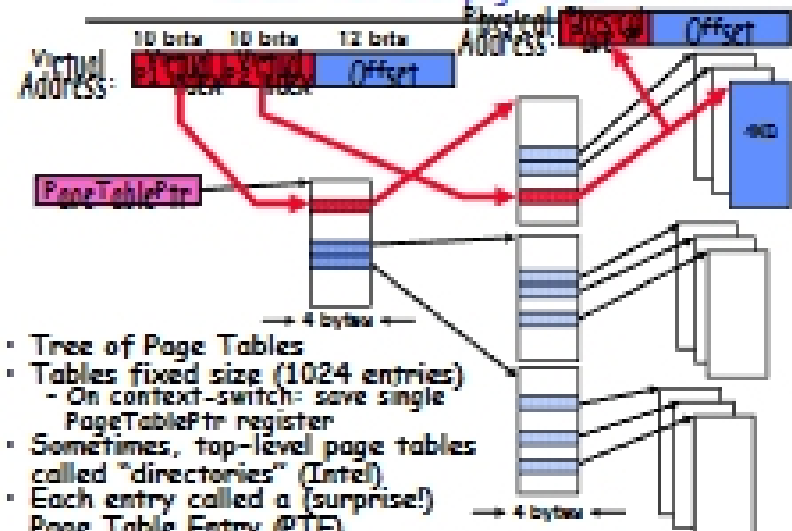
- What about a tree of tables?
 - Lowest level page table → memory still allocated with bitmap
 - Higher levels often segmented
- Could have any number of levels. Example (top segment):



- What must be saved/restored on context switch?
 - Contents of top-level segment registers (for this example)
 - Pointer to top-level table (page table)

3/2/10 cs162 @ UC Berkeley 13.3 Lec 13.3

Review: Two-level page table



- Tree of Page Tables
- Tables fixed size (1024 entries)
 - On context-switch: save single PageTablePtr register
- Sometimes, top-level page tables called "directories" (Intel).
- Each entry called a (surprise!) Page Table Entry (PTE)

3/2/10 cs162 @ UC Berkeley 13.4 Lec 13.4

Dual-Mode Operation

- Can Application modify its own translation tables?
 - If it could, could get access to all of physical memory
 - Has to be restricted somehow
- To Assist with Protection, **Hardware** provides at least two modes (Dual-Mode Operation):
 - "Kernel" mode (or "supervisor" or "protected")
 - "User" mode (Normal program mode)
 - Mode set with bits in special control register only accessible in kernel-mode
- Intel processor actually has four "rings" of protection:
 - PL (Privilege Level) from 0 - 3
 - > PL0 has full access. PL3 has least
 - Privilege Level set in code segment descriptor (CS)
 - Typical OS kernels on Intel processors only use PL0 ("user") and PL3 ("kernel")

3/2/10

CS 442 @ UCSD Spring 2010

Lec 13.9

For Protection, Lock User-Programs in Asylum

- Idea: Lock user programs in padded cell with no exit or sharp objects
 - Cannot change mode to kernel mode
 - User cannot modify page table mapping
 - Limited access to memory: cannot adversely affect other processes
 - > Side-effect: Limited access to memory-mapped I/O operations (I/O that occurs by reading/writing memory locations)
 - Limited access to interrupt controller
- A couple of issues
 - How to share GPU between kernel and user programs?
 - > Kinda like both the inmates and the warden in asylum are the same person. How do you manage this?
 - How do programs interact?
 - How does one switch between kernel and user modes?
 - > OS → user (kernel → user mode): getting into cell
 - > User → OS (user → kernel mode): getting out of cell



3/2/10

CS 442 @ UCSD Spring 2010

Lec 13.10

How to get from Kernel→User

- What does the kernel do to create a new user process?
 - Allocate and initialize address-space control block
 - Read program off disk and store in memory
 - Allocate and initialize translation table
 - > Point at code in memory so program can execute
 - > Possibly point at statically initialized data
 - Run Program:
 - > Set machine registers
 - > Set hardware pointer to translation table
 - > Set processor status word for user mode
 - > Jump to start of program
- How does kernel switch between processes?
 - Same saving/restoring of registers as before
 - Save/restore PSL (hardware pointer to translation table)


3/2/10

CS 442 @ UCSD Spring 2010

Lec 13.11

User→Kernel (System Call)

- Can't let inmate (user) get out of padded cell on own
 - Would defeat purpose of protection!
 - So, how does the user program get back into kernel?



- **System call:** Voluntary procedure call into kernel
 - Hardware for controlled User→Kernel transition
 - Can any kernel routine be called?
 - > No! Only specific ones.
 - System call ID encoded into system call instruction
 - > Index forces well-defined interface with kernel

3/2/10

CS 442 @ UCSD Spring 2010

Lec 13.12