

CS 537 Lecture 8 Paging

Michael Swift

2/17/08

© 2008 Michael Swift. All rights reserved.
http://www.cs.cmu.edu/~mswift

1

Paging Advantages

- Easy to allocate physical memory
 - physical memory is allocated from free list of frames
 - to allocate a frame, just remove it from its free list
 - external fragmentation is not a problem!
 - complication for kernel contiguous physical memory allocation
 - many lists, each keeps track of free regions of particular size
 - regions' sizes are multiples of page sizes
 - "buddy algorithm"
- Easy to "page out" chunks of programs
 - all chunks are the same size (page size)
 - use valid bit to detect references to "paged-out" pages
 - also, page sizes are usually chosen to be convenient multiples of disk block sizes

2/17/08

© 2008 Michael Swift. All rights reserved.
http://www.cs.cmu.edu/~mswift

2

Paging Disadvantages

- Can still have internal fragmentation
 - process may not use memory in exact multiples of pages
- Memory reference overhead
 - 2 references per address lookup (page table, then memory)
 - solution: use a hardware cache to absorb page table lookups
 - translation lookaside buffer (TLB)
- Memory required to hold page tables can be large
 - needs one PTE per page in virtual address space
 - 32 bit addr with 4KB pages = 2^{22} PTEs = 1,048,576 PTEs
 - 4 bytes/PTE = 4MB per page table
 - OS's typically have separate page tables per process
 - 2G processes = 100MB of page tables
 - solution: page the page tables (??)

2/17/08

© 2008 Michael Swift. All rights reserved.
http://www.cs.cmu.edu/~mswift

3

Hardware and Kernel structures for paging

- Hardware:
 - Page table base register
 - TLB (will discuss soon)
- Software:
 - Page table
 - virtual \leftrightarrow physical or virtual \leftrightarrow disk mapping
 - Page frame database
 - One entry per physical page
 - information on page, owning process
 - Swap file / Section list (will discuss under page replacement)


2/17/08

© 2008 Michael Swift. All rights reserved.
http://www.cs.cmu.edu/~mswift

4

Inverted Page Table

- With all previous examples ("Forward Page Tables")
 - Size of page table is at least as large as amount of virtual memory allocated to processes
 - Physical memory may be much less
 - Much of process space may be active at any one time



- Answer: use a hash table
 - Called an "Inverted Page Table"
 - Size is independent of virtual address space
 - Directly related to amount of physical memory
 - Very attractive option for 64-bit address spaces
- Cons: Complexity of managing hash changes
 - Often in hardware

2/17/02

© 2002 Morgan Kaufmann Publishers, Inc. All rights reserved.

9

Addressing Page Tables

- Where are page tables stored?
 - and in which address space?
- Possibility #1: physical memory
 - easy to address, no translation required
 - but, page tables consume memory for lifetime of VAS
- Possibility #2: virtual memory (OS's VAS)
 - cold (unused) page table pages can be paged out to disk
 - but, addresses page tables requires translation
 - how does break the recursion?
 - don't page the outer page table (called **wiring**)
- Question: can the kernel be paged?

2/17/02

© 2002 Morgan Kaufmann Publishers, Inc. All rights reserved.

10

Generic PTE

- PTE maps virtual page to physical page
- Includes some page properties
 - Valid?, writable?, dirty?, cacheable?



Some acronyms used in this lecture:

- PTE = page table entry
- PDE = page directory entry
- VA = virtual address
- PA = physical address
- VPN = virtual page number
- (R,P)PN = (real, physical) page number

2/17/02

© 2002 Morgan Kaufmann Publishers, Inc. All rights reserved.

11

Real Page Tables

- Design requirements
 - Minimize memory use (PT are pure overhead)
 - Fast (logically accessed on every memory ref)
- Requirements lead to
 - Compact data structures
 - O(1) access (e.g. indexed lookup, hashtable)
- Examples: X86

2/17/02

© 2002 Morgan Kaufmann Publishers, Inc. All rights reserved.

12