

CS 537 Lecture 12 Paging

Michael Swift

10/20/07

© 2007 Massachusetts Institute of Technology. All rights reserved.
Lecture 12: Paging

1

Notes

- Writing assignment 1 has been posted - due next Thursday
- I want to meet with groups this week. I'll have a sign-up sheet after class
- Today: Paging and TLBs
- Questions from last time:
 - What is virtual memory?
 - What does it do?
 - What is it good for?

10/20/07

© 2007 Massachusetts Institute of Technology. All rights reserved.
Lecture 12: Paging

2

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

10/20/07

© 2007 Massachusetts Institute of Technology. All rights reserved.
Lecture 12: Paging

3

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
 - need one PTE per page in virtual address space
 - 32 bit A3 with 4KB pages = 32K PTEs = 1,048,576 PTEs
 - 4 bytes/PTE = 4MB per page table
 - OSs typically have separate page tables per process
 - 25 processes = 100MB of page tables
 - solution: page the page tables (33)

10/20/07

© 2007 Massachusetts Institute of Technology. All rights reserved.
Lecture 12: Paging

4

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)

10/22/07

© 2005 Pearson Education, Inc. All rights reserved. This material is protected by copyright.

3

Page Frame Database

```

// Each physical page in the system has a struct page associated with it
// The struct page holds all information to be used when using the page. For up to
// several thousand pages in memory, to locate each struct page using
// a page.
struct page {
    int type; // type of page
    struct iovec *iovec; // IOVEC struct, see 21.10.
    struct file *file; // file descriptor, if any
    int flags; // flags
};

// Each physical page is also associated with a struct page_frame_database
// struct page_frame_database
struct page_frame_database {
    struct page *pages; // array of struct page pointers
    int n_pages; // number of pages
};

// Each process has a struct page_table associated with it
// struct page_table
struct page_table {
    struct page_frame_database *pfd; // pointer to struct page_frame_database
    int n_pages; // number of pages
};
    
```

10/22/07

© 2005 Pearson Education, Inc. All rights reserved. This material is protected by copyright.

4

Managing Page Tables

- Last lecture:
 - size of a page table for 32 bit AS with 4096 pages was 4MB!
 - far too much overhead
 - how can we reduce this?
 - observation: only need to map the portion of the address space that is actually being used (tiny fraction of address space)
 - only need page table entries for those portions
 - how can we do this?
 - make the page table structure dynamically extensible...
 - all problems in CS can be solved with a level of indirection
 - two-level page tables

10/22/07

© 2005 Pearson Education, Inc. All rights reserved. This material is protected by copyright.

5

Two-level page tables

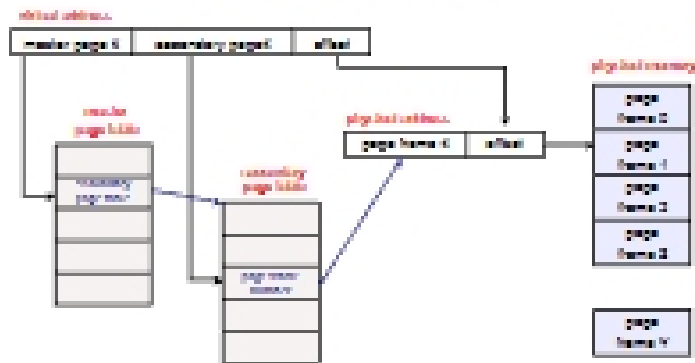
- With two-level PT's, virtual addresses have 3 parts:
 - master page number, secondary page number, offset
 - master PT maps master PN to secondary PT
 - secondary PT maps secondary PN to page frame number
 - offset + PFN = physical address
- Example:
 - 4096 pages, 4 bytes/PTe
 - how many bits in offset? need 12 bits for 4KB
 - want master PT in one page: 4096 bytes = 1024 PTe
 - hence, 1024 secondary page tables
 - so: master page number = 10 bits, offset = 12 bits
 - with a 32-bit address, that leaves 10 bits for secondary PN

10/22/07

© 2005 Pearson Education, Inc. All rights reserved. This material is protected by copyright.

6

Two level page tables




10/22/07

© 2007 Pearson Education, Inc. All rights reserved. This material is protected by copyright.

9

Inverted Page Table

- With all previous exam ples ("Forward Page Tables")
 - Size of page table is at least as large as amount of virtual memory allocated to process
 - Physical memory may be much less
 - Much of process space may be out on disk or not in use



- 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

10/22/07

© 2007 Pearson Education, Inc. All rights reserved. This material is protected by copyright.

10

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, addressing page tables requires translation
 - how do we break the recursion?
 - don't page the outer page table (called **wiring**)
- Question: can the kernel be paged?

10/22/07

© 2007 Pearson Education, Inc. All rights reserved. This material is protected by copyright.

11

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

10/22/07

© 2007 Pearson Education, Inc. All rights reserved. This material is protected by copyright.

12