

Virtual Memory

CS 256/456

Dept. of Computer Science, University of Rochester

2/21/2007

CSC 256/456 - Spring 2007

1

Virtual Memory

- **Virtual memory** - separation of user logical memory from physical memory (usually to save physical memory space).
 - Logical independent memory pieces may map to the same physical memory.
 - Allows physical memory sharing by several processes.
 - Copy-on-write: allows for more efficient process creation.
 - Some logical memory pieces may not map to any physical memory at all.
 - Allows a program to run with only part of its image in physical memory
- Paging makes virtual memory possible at fine-grain
- Demand paging
 - Make a physical instance of a page in memory only when needed.

2/21/2007

CSC 256/456 - Spring 2007

2

Backing Store

- With virtual memory, the whole address space of each process has a copy in the backing store (i.e., disk)
 - program code, data/stack
- Consider the whole program actually resides on the backing store, only part of it is cached in memory.
- With each page table entry a valid-invalid bit is associated (1 ⇒ in-memory, 0 ⇒ not-in-memory or invalid logical page)

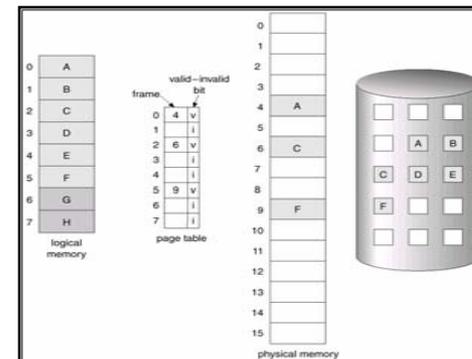
2/21/2007

CSC 256/456 - Spring 2007

3

Page Table with Virtual Memory

- With each page table entry a valid-invalid bit is associated (1 ⇒ in-memory, 0 ⇒ not-in-memory or invalid logical page)



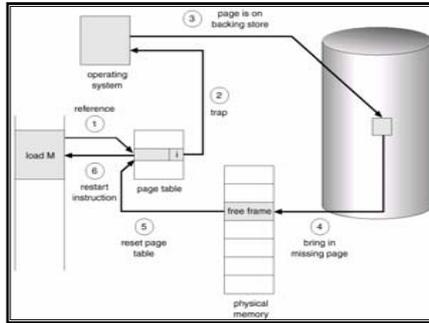
2/21/2007

CSC 256/456 - Spring 2007

4

Page Fault

- A reference to a page with the valid bit set to 0 will trap to OS ⇒ page fault
- Invalid logical page:
 - ⇒ abort.
- Just not in memory:
 - Get a free frame.
 - Swap page into the free frame.
 - Reset the page table entry, valid bit = 1.
 - Restart the program from the fault instruction.
- What if there is no free frame?



2/21/2007

CSC 256/456 - Spring 2007

5

Page Fault Overhead

- Page fault exception handling
- [swap page out]
- swap page in
- restart user program
- memory access

2/21/2007

CSC 256/456 - Spring 2007

6

Memory-Mapped Files

- Memory-mapped file I/O allows file I/O to be treated as routine memory access by *mapping* a disk block to a page in memory.
- At page fault:
 - A certain portion of the file is read from the file system into physical memory.
 - Subsequent reads/writes to/from the file are like ordinary memory accesses.
- Simplifies file access by treating file I/O through memory rather than **read()** **write()** system calls.

2/21/2007

CSC 256/456 - Spring 2007

7

Page Replacement

- Page replacement is necessary when no physical frames are available for demand paging
 - a victim page would be selected and replaced
- A dirty bit for each page
 - indicating if a page has been changed since last time loaded from the backing store
 - indicating whether swap-out is necessary for the victim page.
 - How is it maintained? Does it need to be in the page table entry?

2/21/2007

CSC 256/456 - Spring 2007

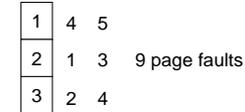
8

Page Replacement Algorithms

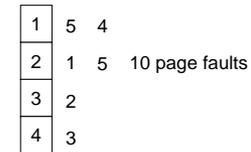
- Page replacement algorithm: the algorithm that picks the victim page.
 - low page-fault rate.
 - implementation cost/feasibility.
- Metric:
 - low page-fault rate.
 - implementation cost/feasibility.
- For the page-fault rate:
 - Evaluate an algorithm by running it on a particular string of memory references (reference string) and computing the number of page faults on that string.

First-In-First-Out (FIFO) Algorithm

- Reference string: 1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5
- 3 frames (3 pages can be in memory at a time)



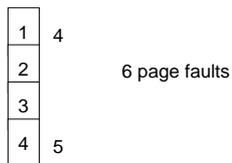
- 4 frames



- Anomaly for the FIFO Replacement
 - more frames not necessarily leading to less page faults

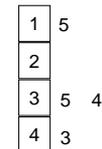
Optimal Algorithm

- Optimal algorithm:
 - Replace page that will not be used for longest period of time.
- 4 frames example
1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5



Least Recently Used (LRU) Algorithm

- Reference string: 1, 2, 3, 4, 1, 2, 5, 1, 2, 3, 4, 5



- Not always better than FIFO, but more frames always lead to less or equal page faults
 - imagine a virtual stack (infinite size) of pages
 - each page is moved to the top after being accessed
 - this virtual stack is independent of the number of frames
 - page fault number when there are N frames:
 - the number of accesses that do not hit the top N pages in the virtual stack.

Implementations

- FIFO implementation
- Time-of-use LRU implementation:
 - Every page entry has a time-of-use field; every time page is referenced through this entry, copy the clock into the field.
 - When a page needs to be changed, look at the time-of-use fields to determine which are to change.
- Stack LRU implementation - keep a stack of page numbers in a double link form:
 - Page referenced: move it to the top
 - Always replace at the bottom of the stack

Feasibility of the Implementations

- FIFO implementation.
- LRU implementations:
 - Time-of-use implementation
 - Stack implementation
- What needs to be done at each memory reference?
- What needs to be done at page loading or page replacement?

Disclaimer

- Parts of the lecture slides contain original work of Abraham Silberschatz, Peter B. Galvin, Greg Gagne, Andrew S. Tanenbaum, and Gary Nutt. The slides are intended for the sole purpose of instruction of operating systems at the University of Rochester. All copyrighted materials belong to their original owner(s).