


# Virtual Memory Management


B.Ramamurthy

## Paging (2)


The relation between virtual addresses and physical memory addresses given by page table



## Demand Paging (contd.)




## Page Tables (1)




Internal operation of MMU with 16 4 KB pages

## Page Tables (2)



- ◆ 32 bit address with 2 page table fields
- ◆ Two-level page tables

## Page Tables (3)




Typical page table entry

## TLBs – Translation Lookaside Buffers

Valid	Virtual page	Modified	Protection	Page frame
1	140	1	RW	31
1	20	0	R X	38
1	130	1	RW	29
1	129	1	RW	62
1	19	0	R X	50
1	21	0	R X	45
1	860	1	RW	14
1	861	1	RW	75

A TLB to speed up paging

## Inverted Page Tables



Comparison of a traditional page table with an inverted page table

## Page Fault Handling (1)

- Hardware traps to kernel
- General registers saved
- OS determines which virtual page needed
- OS checks validity of address, seeks page frame
- If selected frame is dirty, write it to disk


## Page Fault Handling (2)

- OS brings schedules new page in from disk
- Page tables updated
- Faulting instruction backed up to when it began
- Faulting process scheduled
- Registers restored
- Program continues

## Locking Pages in Memory

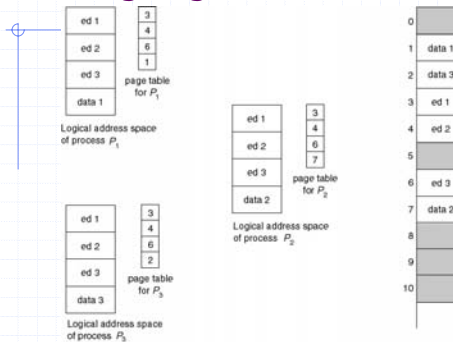
- ◆ Virtual memory and I/O occasionally interact
- ◆ Proc issues call for read from device into buffer
  - while waiting for I/O, another processes starts up
  - has a page fault
  - buffer for the first proc may be chosen to be paged out
- ◆ Need to specify some pages locked
  - exempted from being target pages

## Backing Store



(a) Paging to static swap area  
 (b) Backing up pages dynamically

## Sharing Pages: a text editor



## Implementation Issues

Operating System Involvement with Paging  
 Four times when OS involved with paging

1. Process creation
  - determine program size
  - create page table
2. Process execution
  - MMU reset for new process
  - TLB flushed
3. Page fault time
  - determine virtual address causing fault
  - swap target page out, needed page in
4. Process termination time
  - release page table, pages

## Page Replacement Algorithms

- ◆ Page fault forces choice
  - which page must be removed
  - make room for incoming page
- ◆ Modified page must first be saved
  - unmodified just overwritten
- ◆ Better not to choose an often used page
  - will probably need to be brought back in soon

## Optimal Page Replacement Algorithm

- ◆ Replace page needed at the farthest point in future
  - Optimal but unrealizable
- ◆ Estimate by ...
  - logging page use on previous runs of process
  - although this is impractical

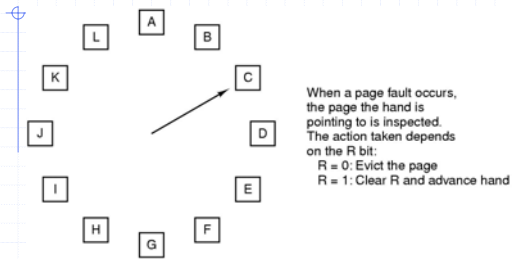
## Not Recently Used Page Replacement Algorithm

- ◆ Each page has Reference bit, Modified bit
  - bits are set when page is referenced, modified
- ◆ Pages are classified
  1. not referenced, not modified
  2. not referenced, modified
  3. referenced, not modified
  4. referenced, modified
- ◆ NRU removes page at random
  - from lowest numbered non empty class

## FIFO Page Replacement Algorithm

- ◆ Maintain a linked list of all pages
  - in order they came into memory
- ◆ Page at beginning of list replaced
- ◆ Disadvantage
  - page in memory the longest may be often used

## The Clock Page Replacement Algorithm



## Least Recently Used (LRU)

- ◆ Assume pages used recently will be used again soon
  - throw out page that has been unused for longest time
- ◆ Must keep a linked list of pages
  - most recently used at front, least at rear
  - update this list every memory reference !!
- ◆ Alternatively keep counter in each page table entry
  - choose page with lowest value counter
  - periodically zero the counter

## Simulating LRU in Software (1)

	Page				Page				Page				Page				Page			
	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3
0	0	1	1	1	0	0	1	1	0	0	0	1	0	0	0	0	0	0	0	0
1	0	0	0	0	1	0	1	1	1	0	0	1	1	0	0	0	1	0	0	0
2	0	0	0	0	0	0	0	0	1	1	0	1	1	1	0	0	1	1	0	1
3	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	0	1	1	1	0
	(a)	(b)	(c)	(d)	(e)															
	0	0	0	0	0	1	1	1	0	1	1	0	0	1	0	0	0	1	0	0
	1	0	1	1	0	0	1	1	0	0	1	0	0	0	0	0	0	0	0	0
	1	0	0	1	0	0	0	1	0	0	0	0	1	1	0	1	1	1	0	0
	1	0	0	0	0	0	0	0	1	1	1	0	1	1	1	0	1	1	1	0
	(f)	(g)	(h)	(i)	(j)															


LRU using a matrix – pages referenced in order 0,1,2,3,2,1,0,3,2,3

## Simulating LRU in Software (2)

	R bits for pages 0-5, clock tick 0	R bits for pages 0-5, clock tick 1	R bits for pages 0-5, clock tick 2	R bits for pages 0-5, clock tick 3	R bits for pages 0-5, clock tick 4	
	1 0 1 0 1 1	1 1 0 0 1 0	1 1 0 1 0 1	1 0 0 0 1 0	0 1 1 0 0 0	
Page	0	1	2	3	4	5
0	10000000	11000000	11100000	11110000	01111000	
1	00000000	10000000	11000000	01100000	10110000	
2	10000000	01000000	00100000	00100000	10001000	
3	00000000	00000000	10000000	01000000	00100000	
4	10000000	11000000	01100000	10110000	01011000	
5	10000000	01000000	10100000	01010000	00101000	
	(a)	(b)	(c)	(d)	(e)	

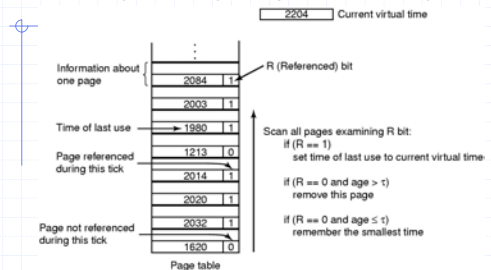
- ◆ The aging algorithm simulates LRU in software
- ◆ Note 6 pages for 5 clock ticks, (a) – (e)

## The Working Set Page Replacement Algorithm (1)



- ◆ The working set is the set of pages used by the  $k$  most recent memory references
- ◆  $w(k,t)$  is the size of the working set at time,  $t$

## The Working Set Page Replacement Algorithm (2)



The working set algorithm



## Page Size (2)

- ◆ Overhead due to page table and internal fragmentation

$$\text{overhead} = \frac{s \cdot e}{p} + \frac{p}{2}$$

Annotations:  $\frac{s \cdot e}{p}$  is labeled "page table space" and  $\frac{p}{2}$  is labeled "internal fragmentation".


- ◆ Where

- $s$  = average process size in bytes
- $p$  = page size in bytes
- $e$  = page entry

Optimized when

$$p = \sqrt{2se}$$

## Separate Instruction and Data Spaces



- ◆ One address space
- ◆ Separate I and D spaces