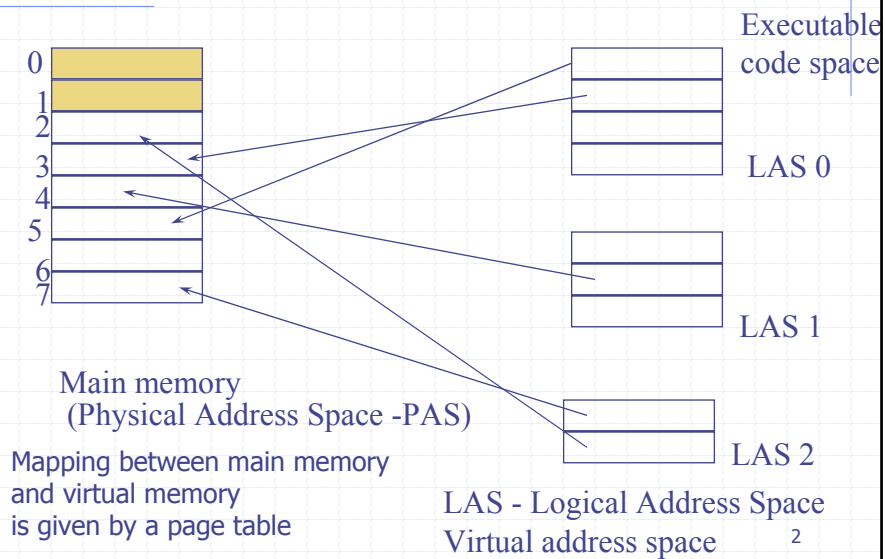


Virtual Memory Management

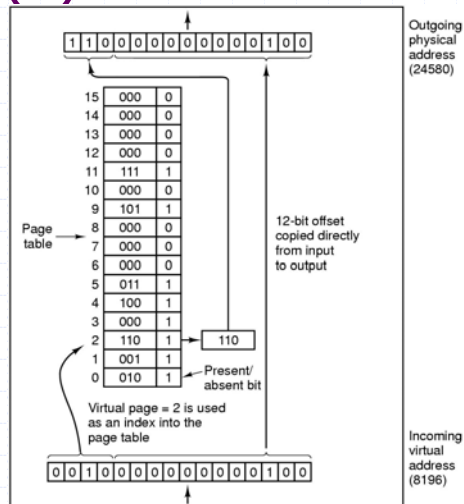
B.Ramamurthy

1

Demand Paging




Page Tables (1)




Internal operation of MMU with 16 4 KB pages ₃

Page Tables (2)



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

Page Tables (3)



Typical page table entry

5

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


6

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

7

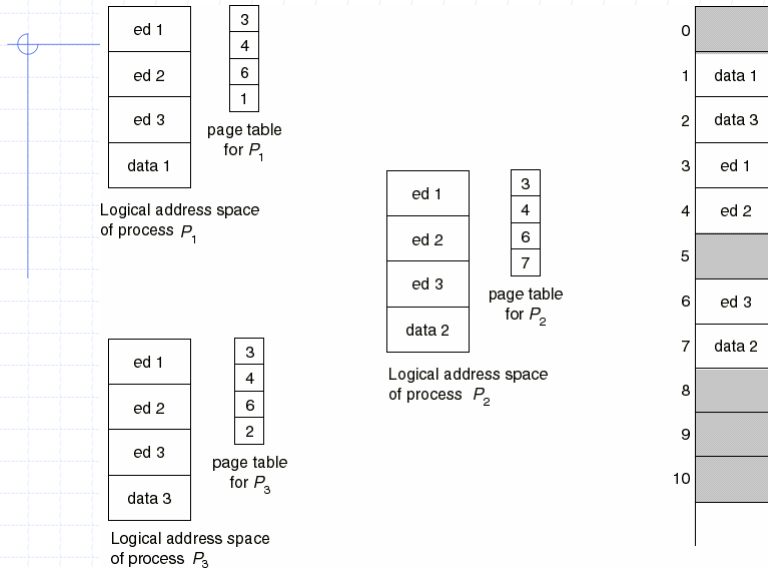
Backing Store



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

8

Sharing Pages: a text editor



9

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

10

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

11

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

12

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

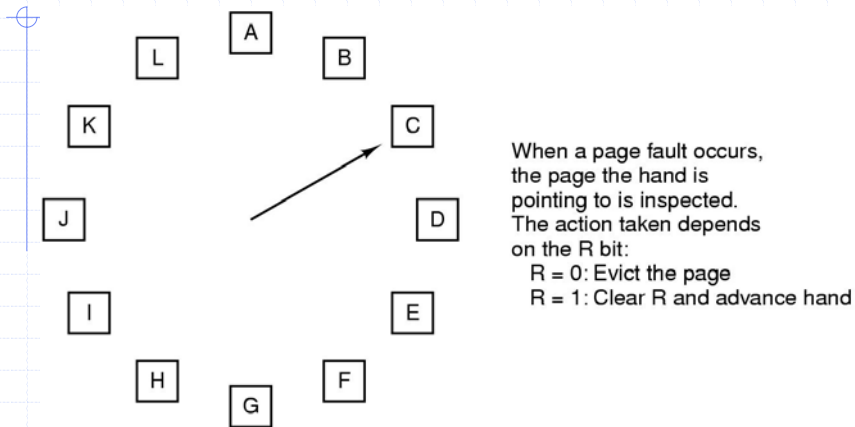
13

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

14

The Clock Page Replacement Algorithm



15

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

16

Simulating LRU in Software (1)

Figure 1 displays eight 4x4 grids, labeled (a) through (h), representing the evolution of a 2D Ising spin system. Each grid has a 'Page' header with indices 0, 1, 2, 3. The grids show the state of the system at different stages of the simulation, with spins represented by 0 (up) and 1 (down).

- (a) Initial state: All spins are up (0).
- (b) Initial state: All spins are up (0).
- (c) State after iteration 1: Spins at (1,1), (1,2), (2,1), and (2,2) are down (1).
- (d) State after iteration 2: Spins at (1,1), (1,2), (2,1), and (2,2) are down (1).
- (e) State after iteration 3: Spins at (1,1), (1,2), (2,1), and (2,2) are down (1).
- (f) State after iteration 4: Spins at (1,1), (1,2), (2,1), and (2,2) are down (1).
- (g) State after iteration 5: Spins at (1,1), (1,2), (2,1), and (2,2) are down (1).
- (h) State after iteration 6: Spins at (1,1), (1,2), (2,1), and (2,2) are down (1).
- (i) State after iteration 7: Spins at (1,1), (1,2), (2,1), and (2,2) are down (1).
- (j) State after iteration 8: Spins at (1,1), (1,2), (2,1), and (2,2) are down (1).

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

17

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	1000000	1100000	1110000	1111000	01111000
1	0000000	1000000	1100000	0110000	10110000
2	1000000	0100000	0010000	0010000	10001000
3	0000000	0000000	1000000	0100000	00100000
4	1000000	1100000	0110000	1011000	01011000
5	1000000	0100000	1010000	0101000	00101000
	(a)	(b)	(c)	(d)	(e)

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

18

Working-Set Model



- ◆ $\Delta \equiv$ working-set window \equiv a fixed number of page references
Example: 10,000 instruction
- ◆ WSS_i (working set of Process P_i) =
total number of pages referenced in the most recent Δ (varies in time)
 - if Δ too small will not encompass entire locality.
 - if Δ too large will encompass several localities.
 - if $\Delta = \infty \Rightarrow$ will encompass entire program.
- ◆ $D = \sum WSS_i \equiv$ total demand frames
- ◆ if $D > m \Rightarrow$ Thrashing
- ◆ Policy if $D > m$, then suspend one of the processes.

19

Working-set model

page reference table

... 2 6 1 5 7 7 7 7 5 1 6 2 3 4 1 2 3 4 4 4 3 4 3 4 4 4 4 1 3 2 3 4 4 4 3 4 4 4 ...



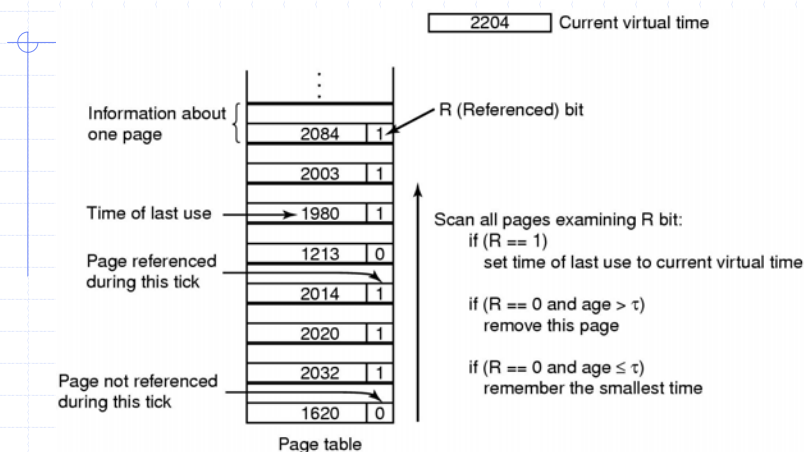
20

Keeping Track of the Working Set

- ◆ Approximate with interval timer + a reference bit
- ◆ Example: $\Delta = 10,000$
 - Timer interrupts after every 5000 time units.
 - Keep in memory 2 bits for each page.
 - Whenever a timer interrupts copy and sets the values of all reference bits to 0.
 - If one of the bits in memory = 1 \Rightarrow page in working set.
- ◆ Why is this not completely accurate?
- ◆ Improvement = 10 bits and interrupt every 1000 time units.

21

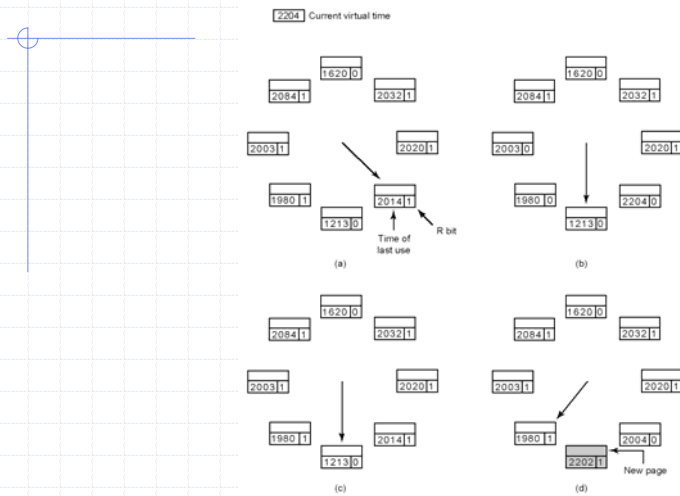
The Working Set Page Replacement Algorithm (2)



The working set algorithm

22

The WSClock Page Replacement Algorithm



Operation of the WSClock algorithm

23

Review of Page Replacement Algorithms

Algorithm	Comment
Optimal	Not implementable, but useful as a benchmark
NRU (Not Recently Used)	Very crude
FIFO (First-In, First-Out)	Might throw out important pages
Second chance	Big improvement over FIFO
Clock	Realistic
LRU (Least Recently Used)	Excellent, but difficult to implement exactly
NFU (Not Frequently Used)	Fairly crude approximation to LRU
Aging	Efficient algorithm that approximates LRU well
Working set	Somewhat expensive to implement
WSClock	Good efficient algorithm

24

Modeling Page Replacement Algorithms

Belady's Anomaly

All pages frames initially empty

	0	1	2	3	0	1	4	0	1	2	3	4
Youngest page	0	1	2	3	0	1	4	4	4	2	3	3
		0	1	2	3	0	1	1	1	4	2	2
Oldest page			0	1	2	3	0	0	0	1	4	4
	P	P	P	P	P	P	P			P	P	

9 Page faults

(a)

	0	1	2	3	0	1	4	0	1	2	3	4
Youngest page	0	1	2	3	3	3	4	0	1	2	3	4
		0	1	2	2	2	3	4	0	1	2	3
Oldest page			0	1	1	1	2	3	4	0	1	2
			0	0	0	0	1	2	3	4	0	1
	P	P	P	P			P	P	P	P	P	P

10 Page faults

(b)

- ◆ FIFO with 3 page frames
- ◆ FIFO with 4 page frames
- ◆ *P*s show which page references show page faults

25

Stack Algorithms

Reference string	0	2	1	3	5	4	6	3	7	4	7	3	3	5	5	3	1	1	1	7	1	3	4	1	
	0	2	1	3	5	4	6	3	7	4	7	3	3	5	5	3	1	1	1	7	1	3	4	1	
		0	2	1	3	5	4	6	3	7	4	7	7	3	3	5	3	3	3	1	7	1	3	4	
			0	2	1	3	5	4	6	3	3	4	4	7	7	7	5	5	5	3	3	7	1	3	
				0	2	1	3	5	4	6	6	6	6	4	4	4	7	7	7	5	5	5	7	7	
					0	2	1	1	5	5	5	5	5	6	6	6	4	4	4	4	4	4	5	5	
						0	2	2	1	1	1	1	1	1	1	1	6	6	6	6	6	6	6	6	
							0	0	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	
									0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Page faults	P	P	P	P	P	P	P		P					P		P						P			
Distance string	∞	∞	∞	∞	∞	∞	∞	4	∞	4	2	3	1	5	1	2	6	1	1	4		7	4	6	5

State of memory array, M , after each item in reference string is processed

26

Page Size (1)

Small page size

◆ Advantages

- less internal fragmentation
- better fit for various data structures, code sections
- less unused program in memory

◆ Disadvantages

- programs need many pages, larger page tables

27

Page Size (2)

◆ Overhead due to page table and internal fragmentation

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

Diagram illustrating the overhead formula:

- The term $\frac{s \cdot e}{p}$ is circled and labeled "page table space".
- The term $\frac{p}{2}$ is circled and labeled "internal fragmentation".

◆ Where

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

Optimized when

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

28


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

29

Inverted Page Tables



Comparison of a traditional page table with an inverted page table

30