

Chapter 4

Memory Management

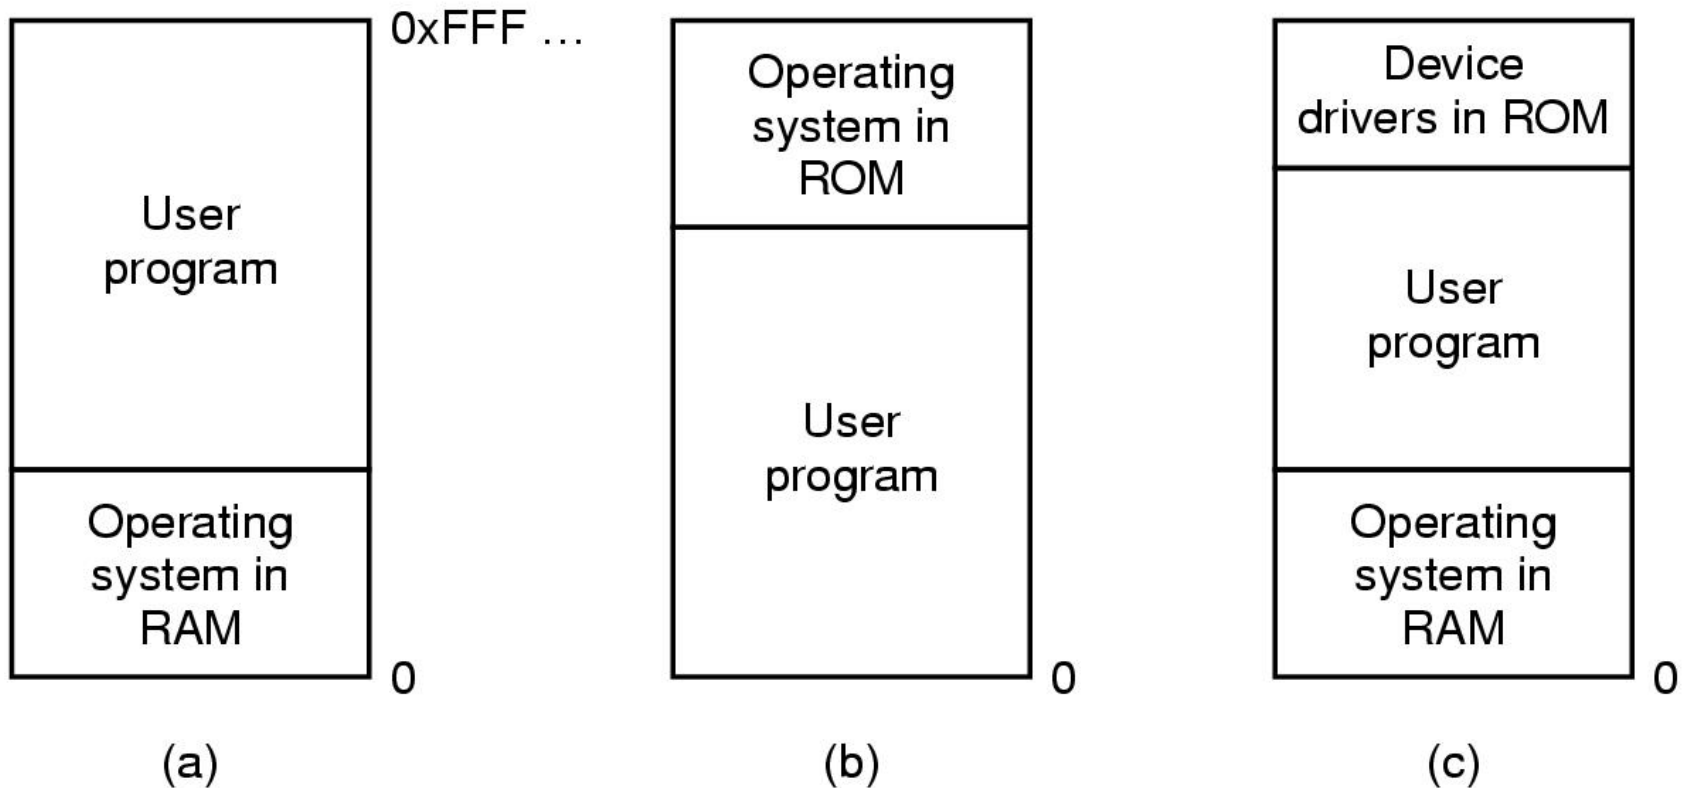
- 4.1 Basic memory management
- 4.2 Swapping
- 4.3 Virtual memory
- 4.4 Page replacement algorithms
- 4.5 Modeling page replacement algorithms
- 4.6 Design issues for paging systems
- 4.7 Implementation issues
- 4.8 Segmentation

Memory Management

- Ideally programmers want memory that is
 - large
 - fast
 - non volatile
- Memory hierarchy
 - small amount of fast, expensive memory – cache
 - some medium-speed, medium price main memory
 - gigabytes of slow, cheap disk storage
- Memory manager handles the memory hierarchy

Basic Memory Management

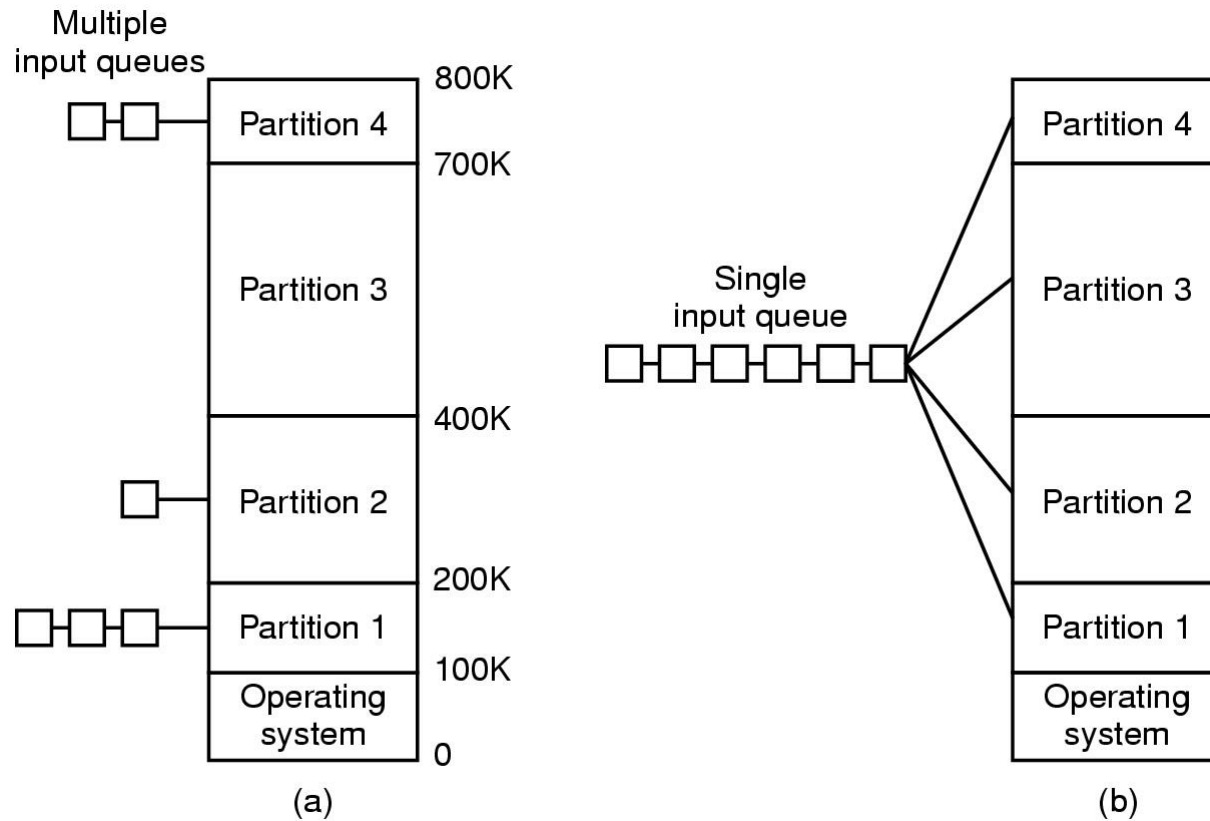
Monoprogramming (uniprogramming) without Swapping or Paging



Three simple ways of organizing memory


- an operating system with one user process

Multiprogramming with Fixed Partitions



- Fixed memory partitions
 - separate input queues for each partition
 - single input queue

Modeling Multiprogramming




CPU utilization as a function of number of processes in memory

Relocation and Protection

- Cannot be sure where program will be loaded in memory
 - address locations of variables, code routines cannot be absolute
 - must keep a program out of other processes' partitions
- Use base and limit values
 - address locations added to base value to map to physical address
 - address locations larger than limit value is an error

Swapping (1)

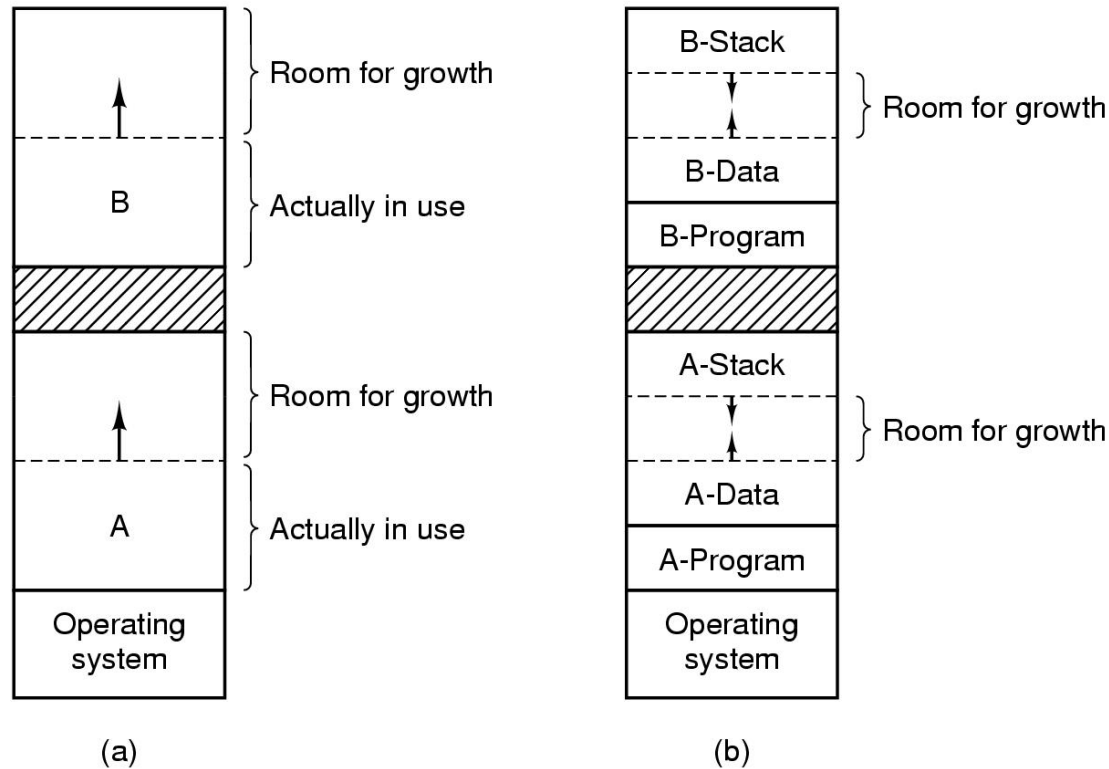


Memory allocation changes as

- processes come into memory
- leave memory

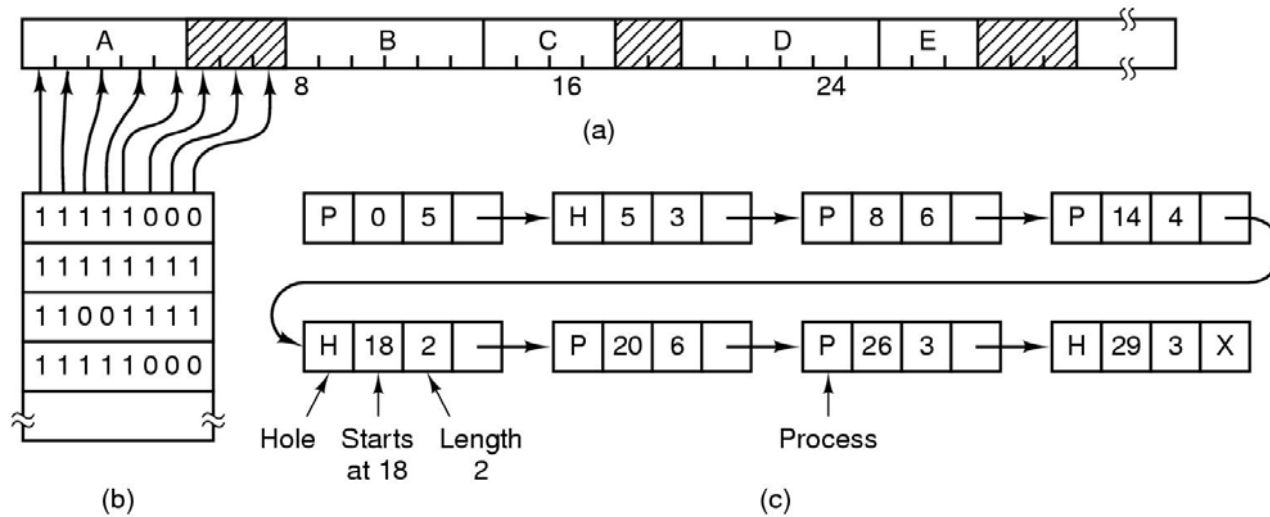
Shaded regions are unused memory

Swapping (2)



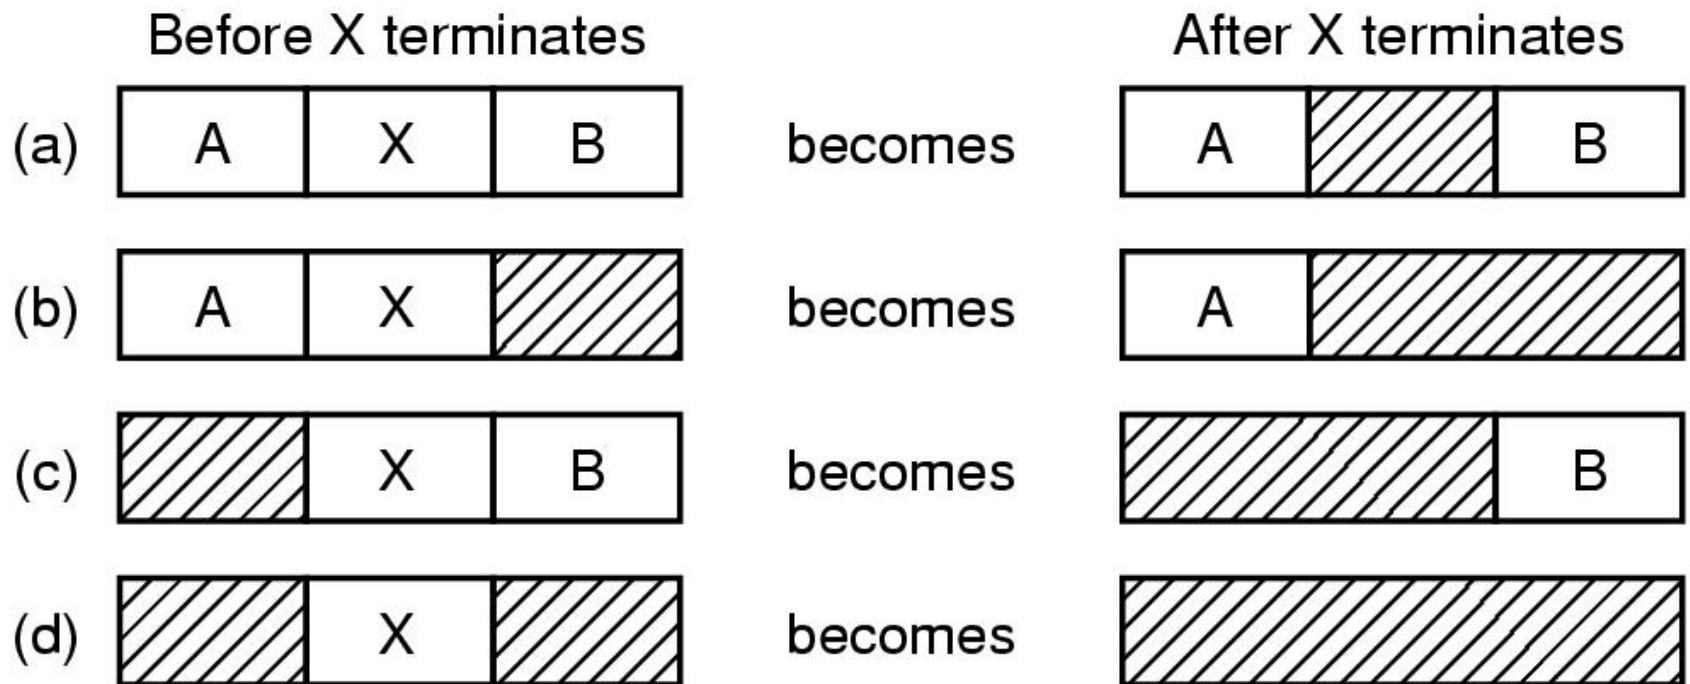
- Allocating space for growing data segment
- Allocating space for growing stack & data segment

Memory Management with Bit Maps



- Part of memory with 5 processes, 3 holes
 - tick marks show allocation units
 - shaded regions are free
- Corresponding bit map
- Same information as a list


Memory Management with Linked Lists



Four neighbor combinations for the terminating process X

Virtual Memory

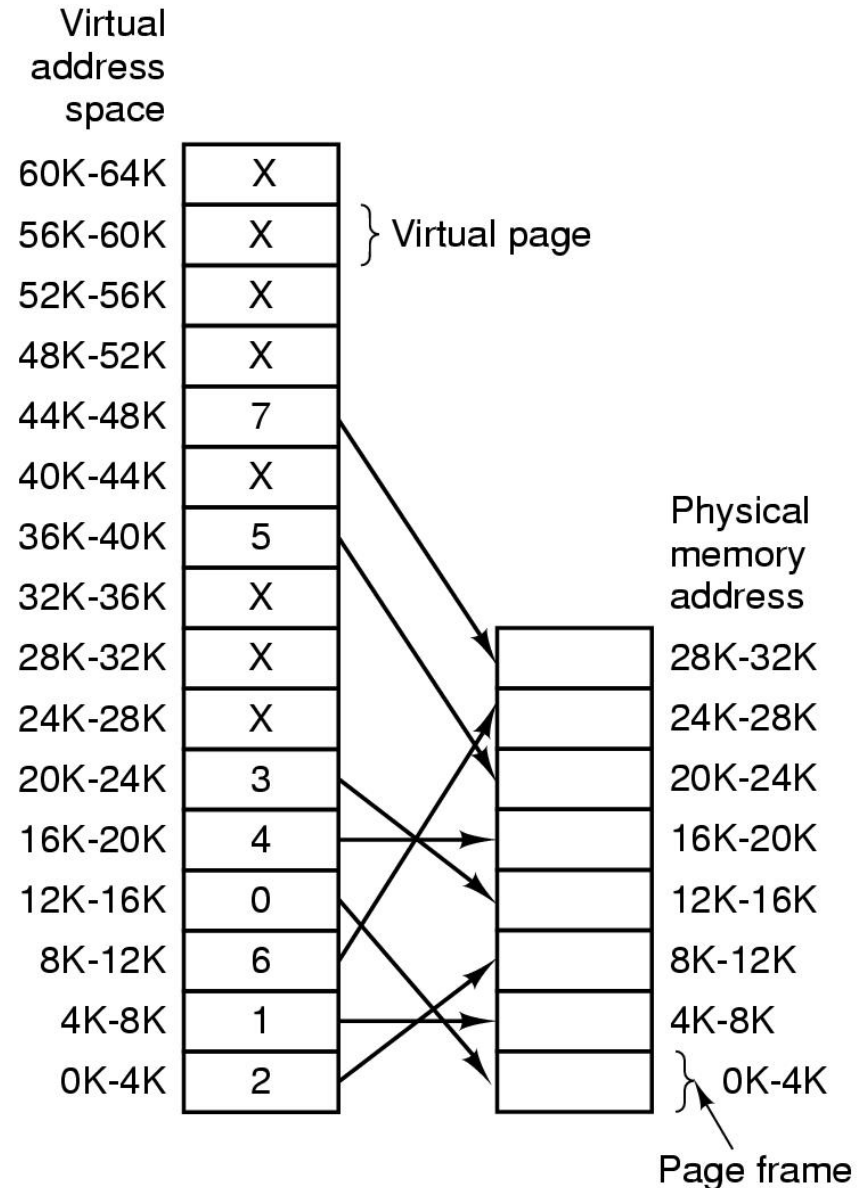
Paging (1)



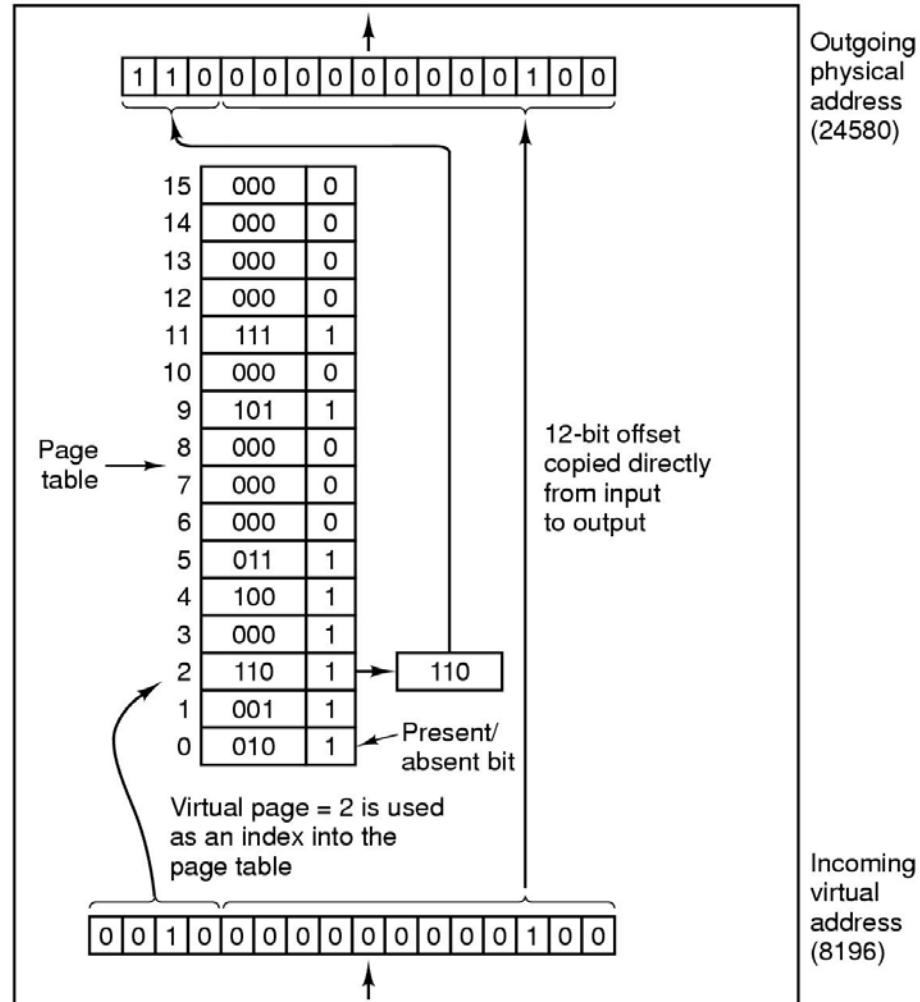
The position and function of the MMU

Paging (2)

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

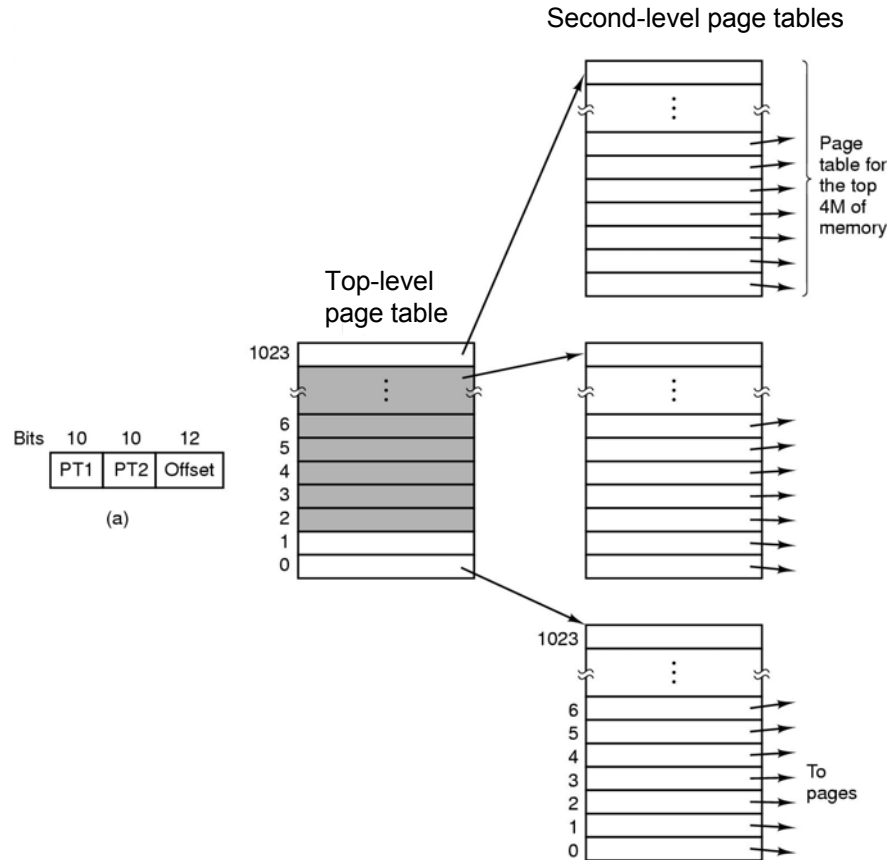


Page Tables (1)



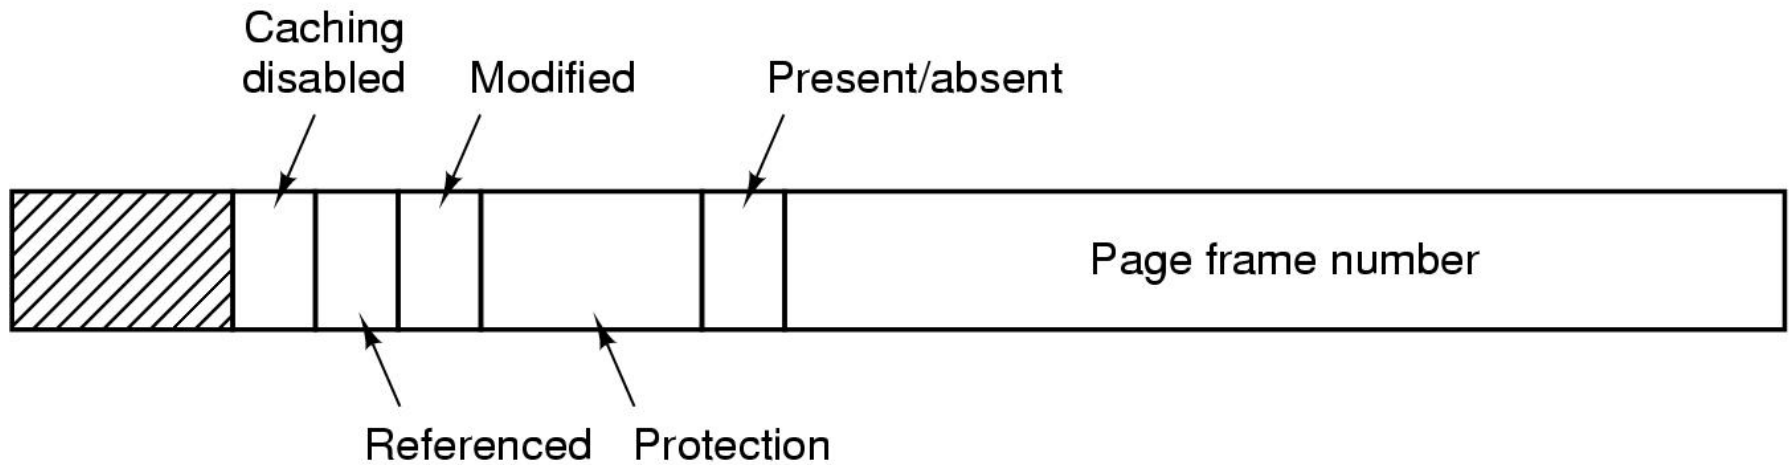
Internal operation of MMU with 16 4 KB pages

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

Traditional page table with an entry for each of the 2^{52} pages



256-MB physical memory has 2^{16} 4-KB page frames



Comparison of a traditional page table with an inverted page table

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

Second Chance Page Replacement Algorithm

Page loaded first




(a)



(b)

- Operation of a second chance
 - pages sorted in FIFO order
 - Page list if fault occurs at time 20, A has *R* bit set (numbers above pages are loading times)

The Clock Page Replacement Algorithm



When a page fault occurs, the page the hand is pointing to is inspected. The action taken depends on the R bit:

R = 0: Evict the page

R = 1: Clear R and advance hand

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			
	0	1	2	3
0	0	1	1	1
1	0	0	0	0
2	0	0	0	0
3	0	0	0	0

(a)

	Page			
	0	1	2	3
0	0	0	1	1
1	1	0	1	1
2	0	0	0	0
3	0	0	0	0

(b)

	Page			
	0	1	2	3
0	0	0	0	1
1	1	0	0	1
2	1	1	0	1
3	0	0	0	0

(c)

	Page			
	0	1	2	3
0	0	0	0	0
1	1	0	0	0
2	1	1	0	0
3	1	1	1	0

(d)

	Page			
	0	1	2	3
0	0	0	0	0
1	1	0	0	0
2	1	1	0	1
3	1	1	0	0

(e)

0	0	0	0
1	0	1	1
1	0	0	1
1	0	0	0

(f)

0	1	1	1
0	0	1	1
0	0	0	1
0	0	0	0

(g)

0	1	1	0
0	0	1	0
0	0	0	0
1	1	1	0

(h)

0	1	0	0
0	0	0	0
1	1	0	1
1	1	0	0

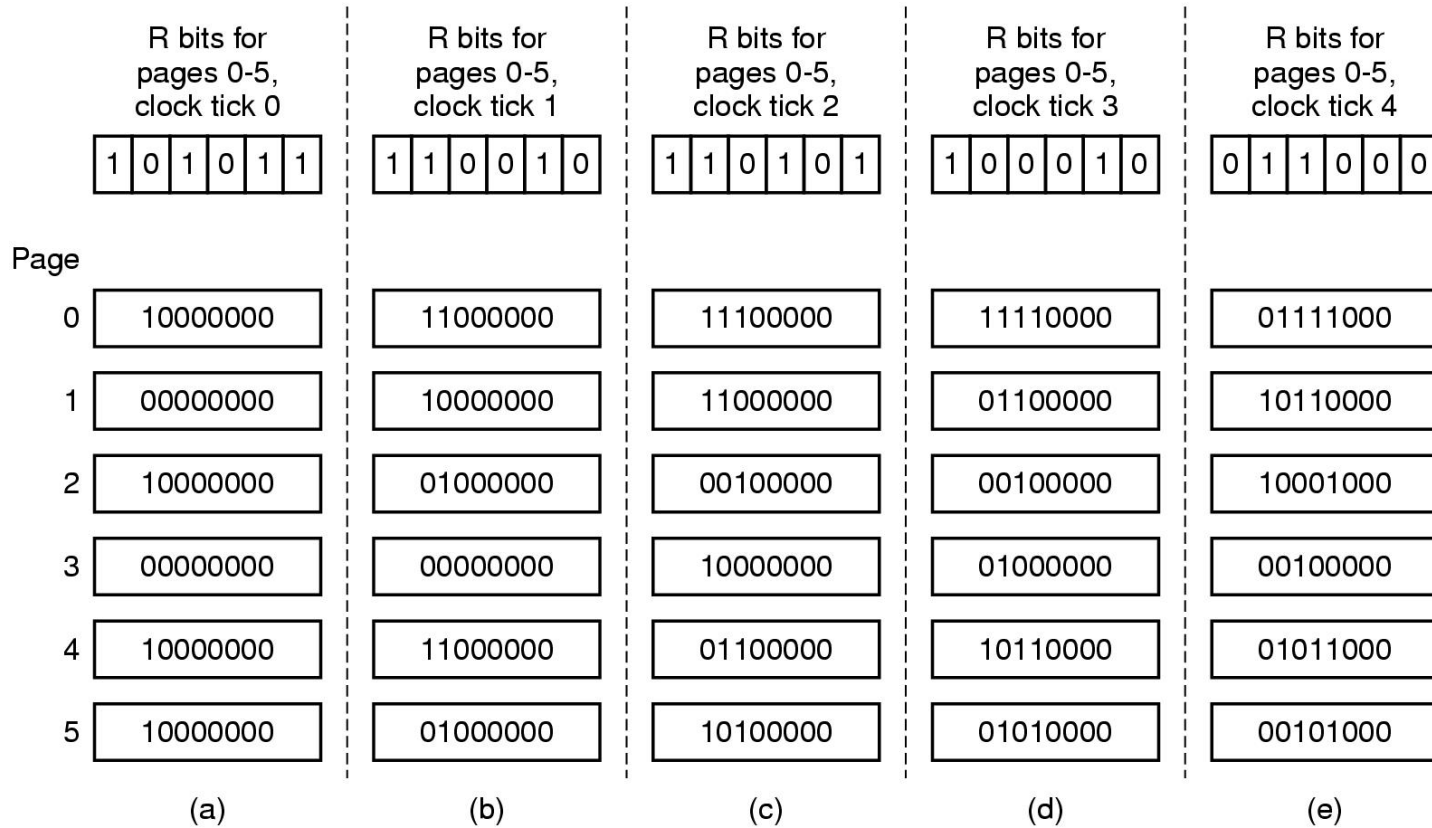
(i)

0	1	0	0
0	0	0	0
1	1	0	0
1	1	1	0

(j)


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

Simulating LRU in Software (2)



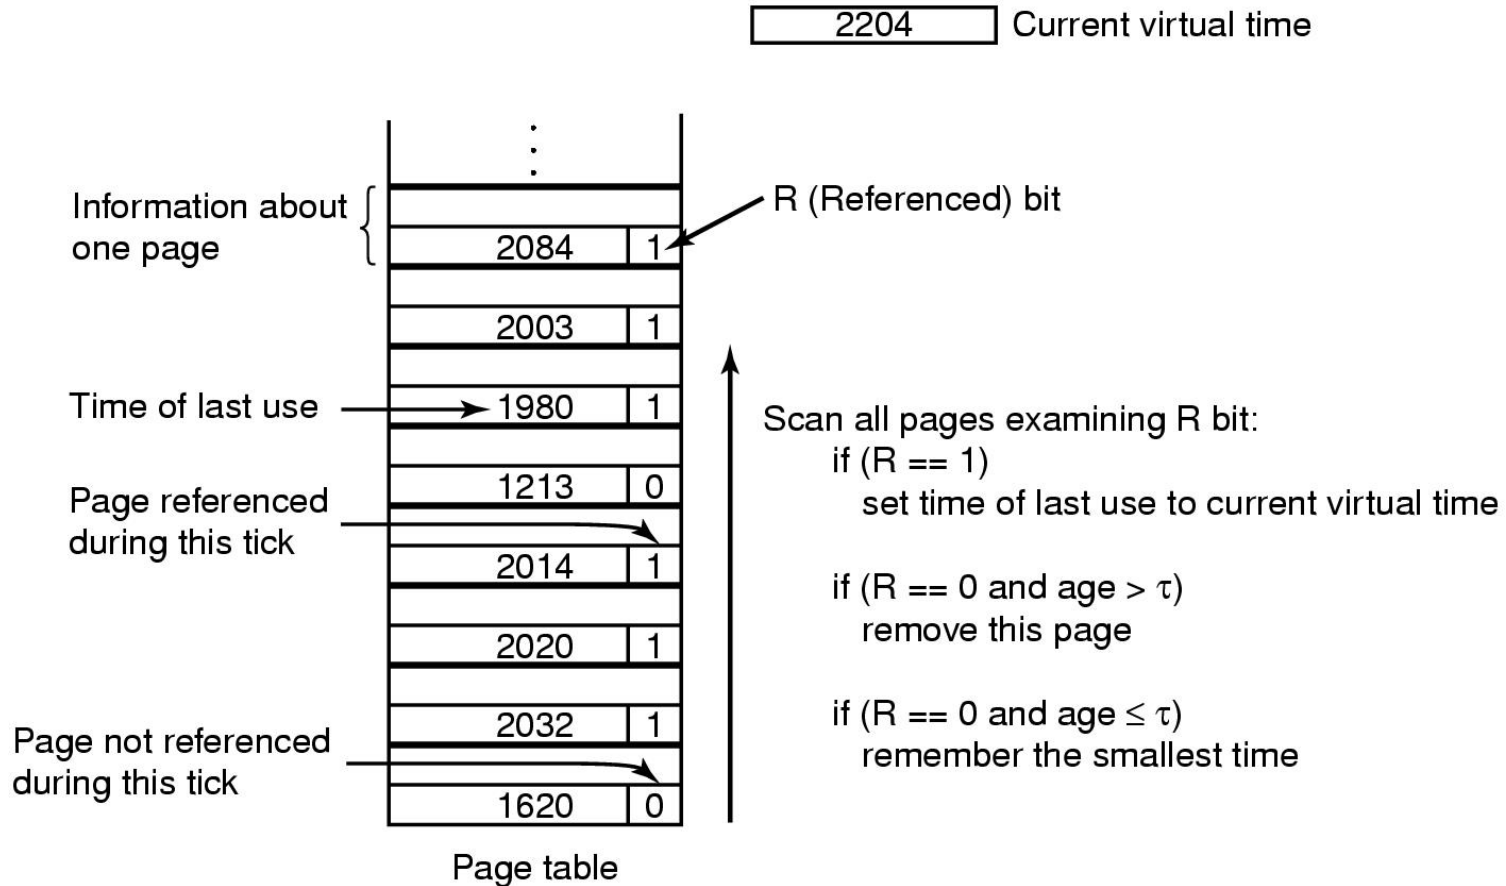
- The aging algorithm simulates LRU in software
- Note 6 pages for 5 clock ticks, (a) – (e)

The Working Set Page Replacement Algorithm (1)



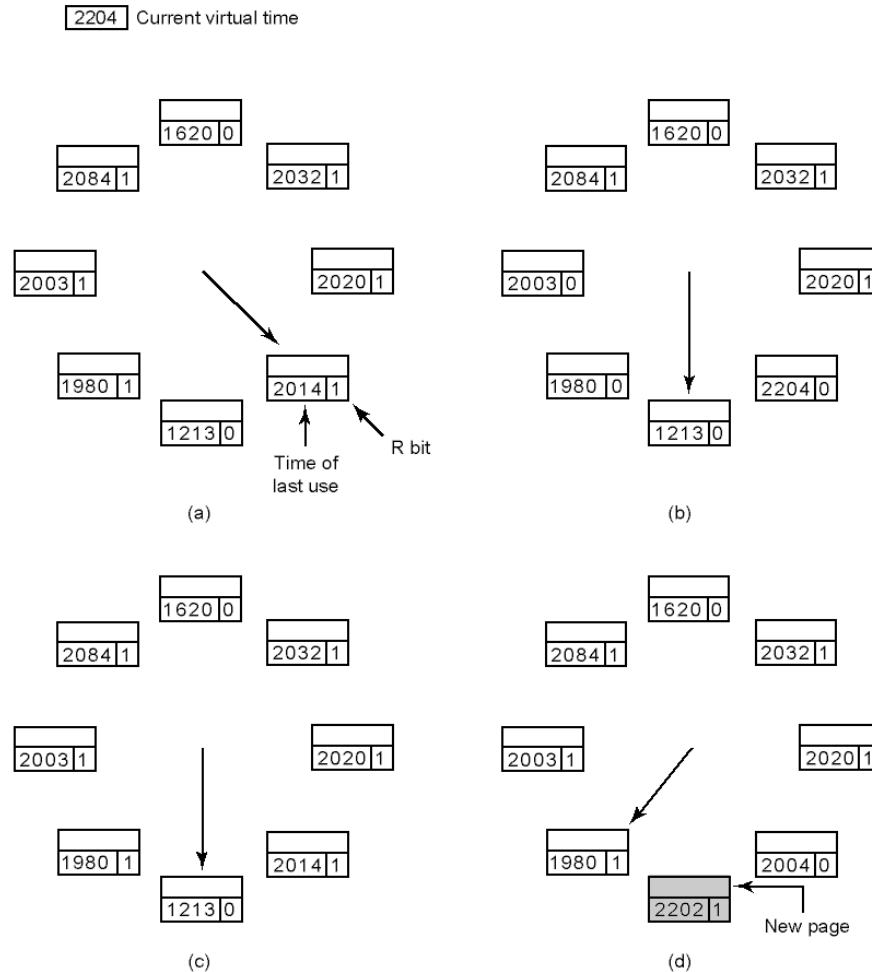
- The working set is the set of pages used by the k most recent memory references
- $w(k,t)$ is the size of the working set at time, t

The Working Set Page Replacement Algorithm (2)



The working set algorithm

The WSClock Page Replacement Algorithm



Operation of the WSClock algorithm

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

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		

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	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


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	4	5	5
						0	2	2	1	1	1	1	1	1	1	1	6	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	2
									0	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


Design Issues for Paging Systems

Local versus Global Allocation Policies (1)



- Original configuration
- Local page replacement
- Global page replacement

Local versus Global Allocation Policies (2)



Page fault rate as a function of the number of page frames assigned

Load Control

- Despite good designs, system may still thrash
- When PFF algorithm indicates
 - some processes need more memory
 - but no processes need less
- Solution :
 - Reduce number of processes competing for memory
 - swap one or more to disk, divide up pages they held
 - reconsider degree of multiprogramming

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

Page Size (2)

- Overhead due to page table and internal fragmentation

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

The diagram shows the equation $\text{overhead} = \frac{s \cdot e}{p} + \frac{p}{2}$. The first term, $\frac{s \cdot e}{p}$, is enclosed in an oval with an arrow pointing to a box labeled "page table space". The second term, $\frac{p}{2}$, is also enclosed in an oval with an arrow pointing to a box labeled "internal fragmentation".

- Where

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

Optimized when

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

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 Fault Handling (1)

1. Hardware traps to kernel
2. General registers saved
3. OS determines which virtual page needed
4. OS checks validity of address, seeks page frame
5. If selected frame is dirty, write it to disk

Page Fault Handling (2)

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

Instruction Backup

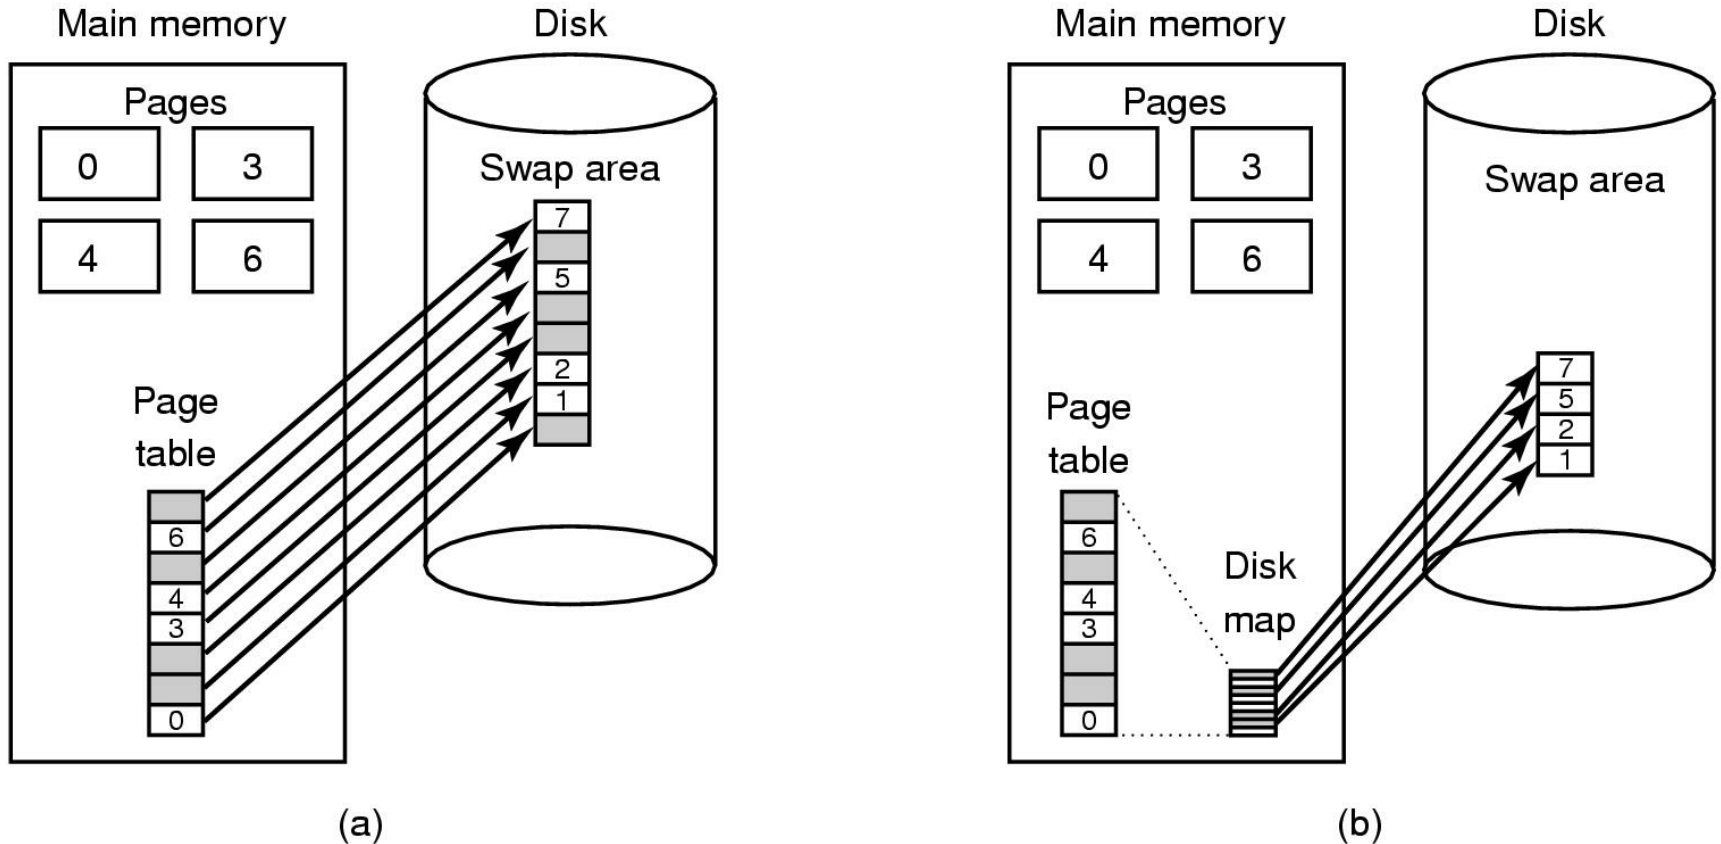


An instruction causing a page fault

Locking Pages in Memory

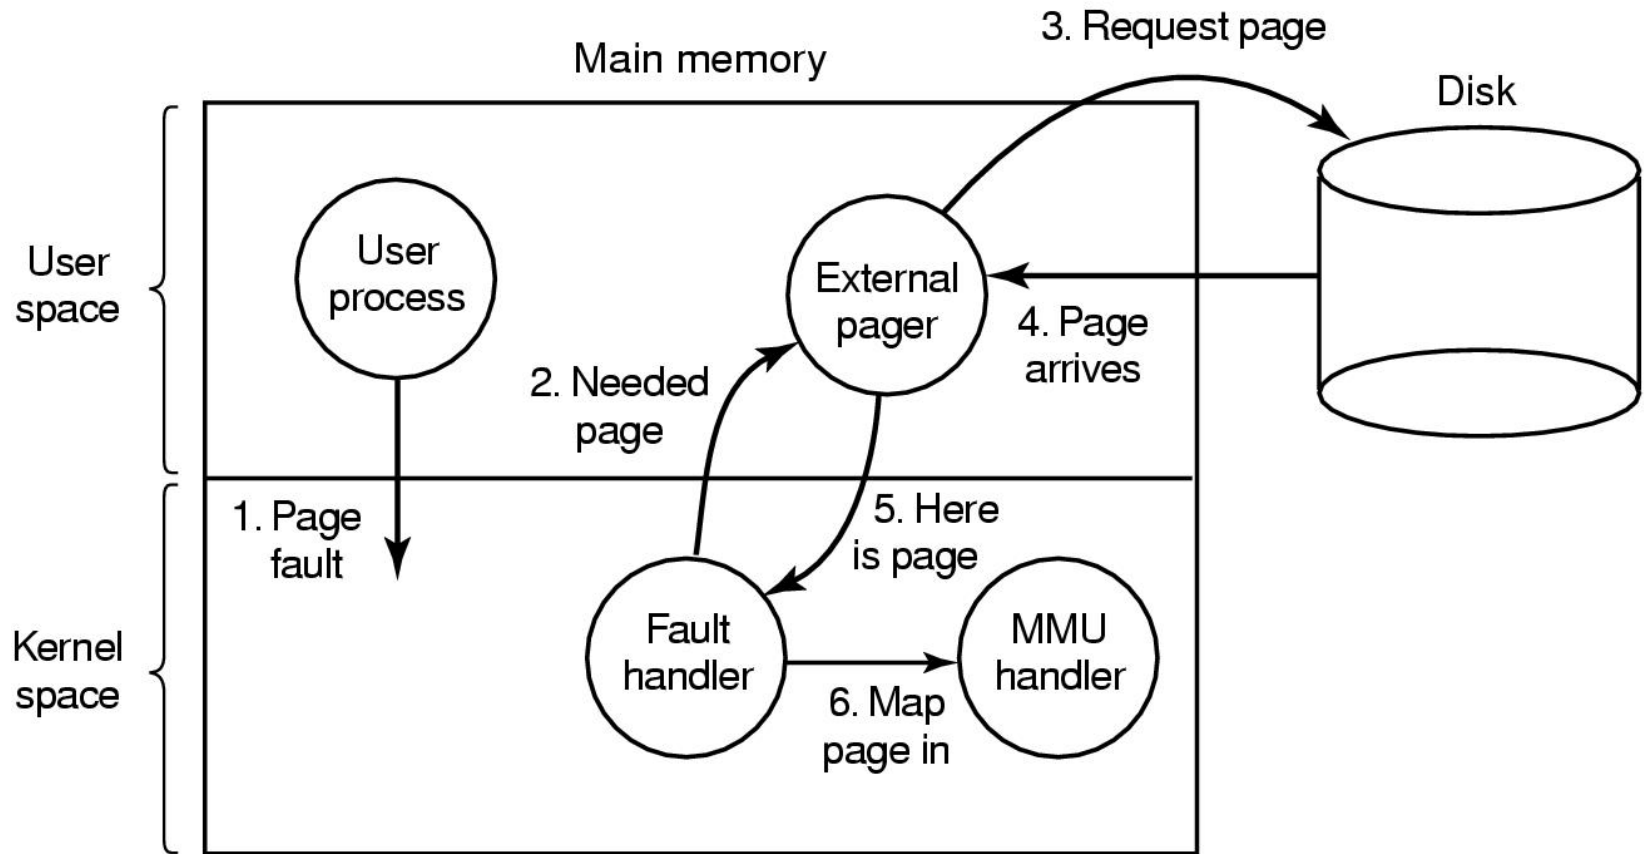
- Virtual memory and I/O occasionally interact
- Proc issues call for read from device into buffer
 - while waiting for I/O, another processes starts up
 - has a page fault
 - buffer for the first proc may be chosen to be paged out
- Need to specify some pages locked
 - exempted from being target pages

Backing Store



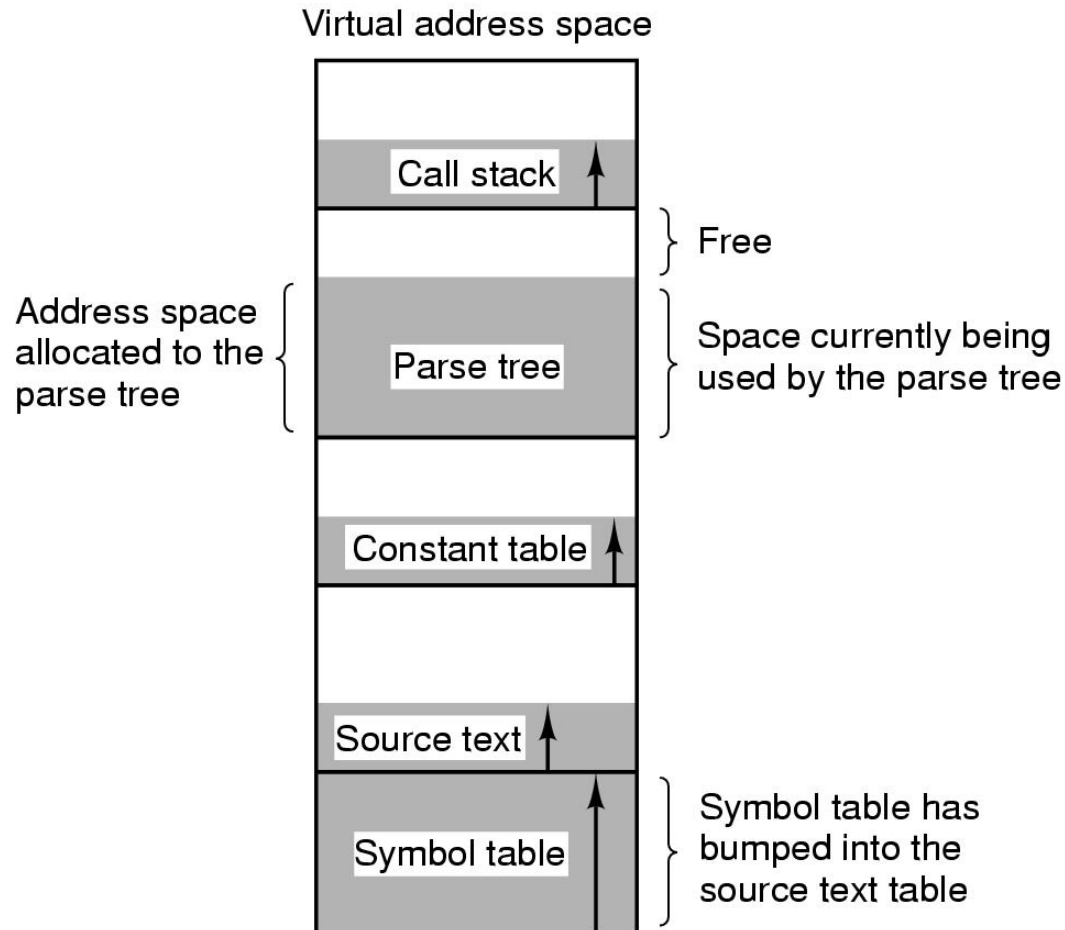
- (a) Paging to static swap area
- (b) Backing up pages dynamically

Separation of Policy and Mechanism



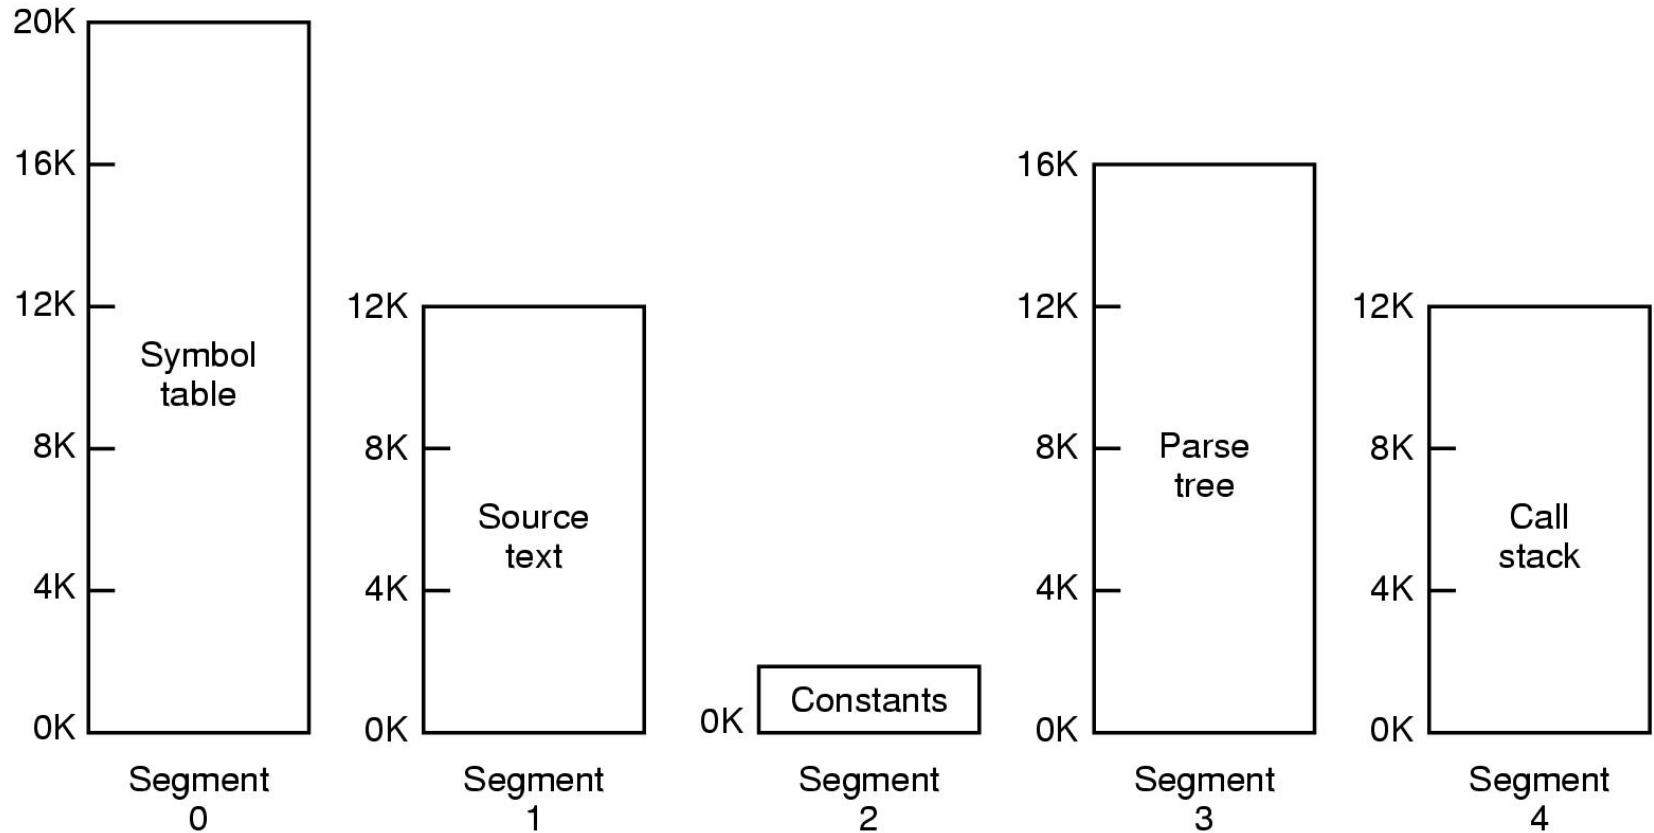
Page fault handling with an external pager

Segmentation (1)



- One-dimensional address space with growing tables
- One table may bump into another

Segmentation (2)



Allows each table to grow or shrink, independently

Segmentation (3)

Consideration	Paging	Segmentation
Need the programmer be aware that this technique is being used?	No	Yes
How many linear address spaces are there?	1	Many
Can the total address space exceed the size of physical memory?	Yes	Yes
Can procedures and data be distinguished and separately protected?	No	Yes
Can tables whose size fluctuates be accommodated easily?	No	Yes
Is sharing of procedures between users facilitated?	No	Yes
Why was this technique invented?	To get a large linear address space without having to buy more physical memory	To allow programs and data to be broken up into logically independent address spaces and to aid sharing and protection

Comparison of paging and segmentation

Summary

- We studied Memory management methods paging and segmentation.
- We also studied the effect of multiprogramming on memory needs.
- We will reinforce the concepts studied by implementing multiprogramming with simple paging in project 2 and demand paging in project 3.