MODERN OPERATING SYSTEMS

Third Edition

ANDREW S. TANENBAUM

Chapter 3
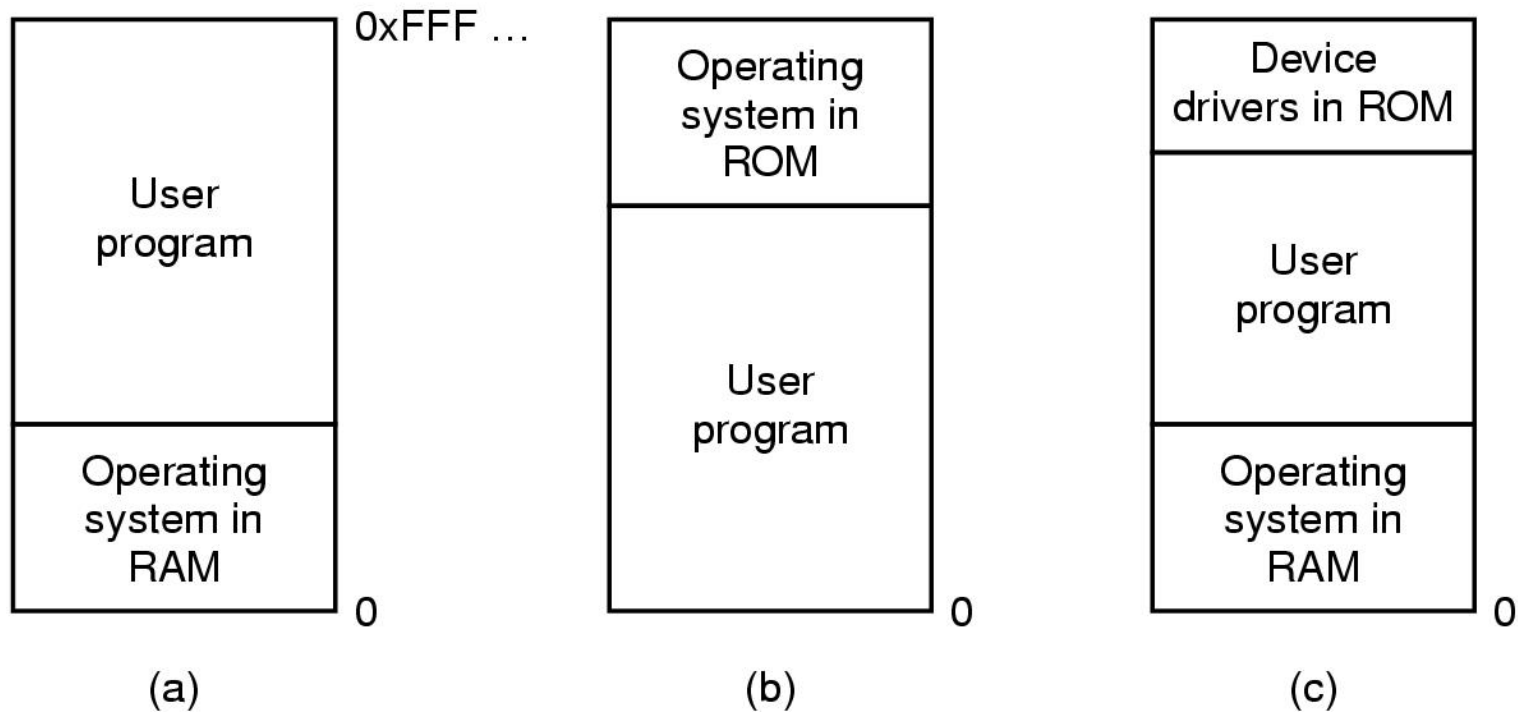# Memory Management

No Memory Abstraction

Figure 3-1. Three simple ways of organizing memory with an operating system and one user process.

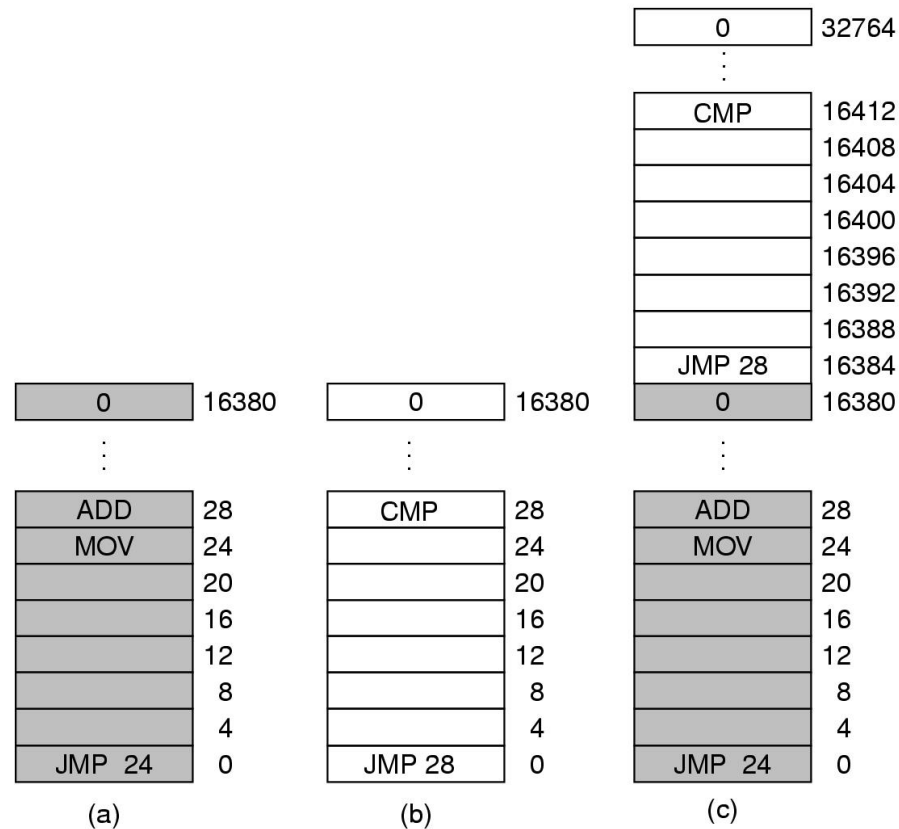# Multiple Programs Without Memory Abstraction

Figure 3-2. Illustration of the relocation problem.
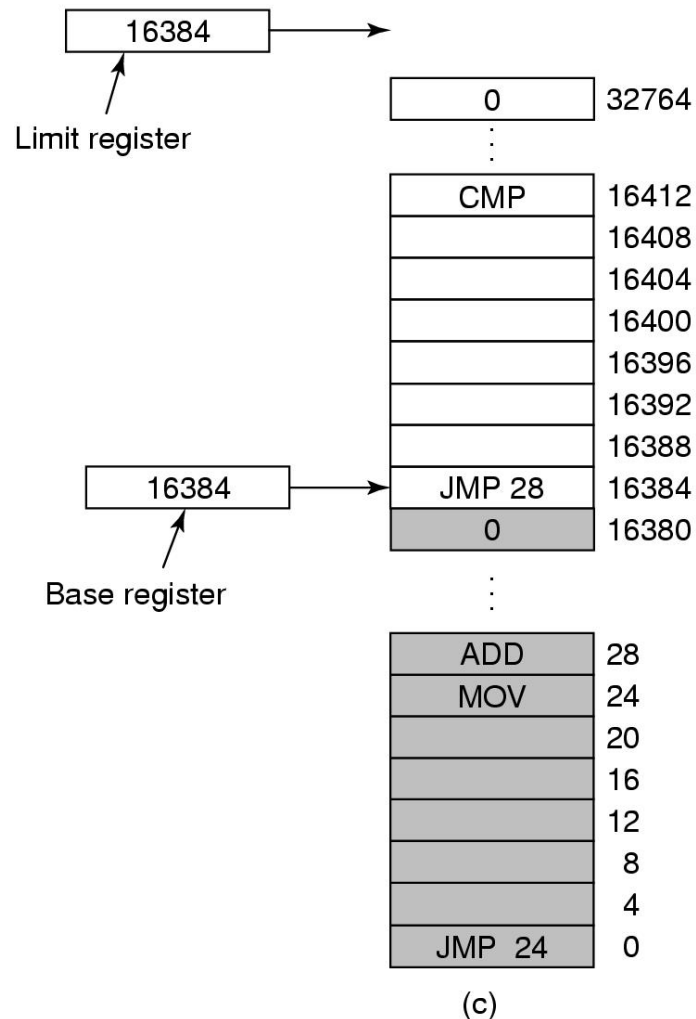
Base and Limit Registers

Figure 3-3. Base and limit registers can be used to give each process a separate address space.

Swapping (1)

Figure 3-4. Memory allocation changes as processes come into memory and leave it. The shaded regions are unused memory.

Swapping (2)

(a)   (b)

Figure 3-5. (a) Allocating space for growing data segment. (b) Allocating space for growing stack, growing data segment.

Memory Management with Bitmaps

Figure 3-6. (a) A part of memory with five processes and three holes. The tick marks show the memory allocation units. The shaded regions (0 in the bitmap) are free. (b) The corresponding bitmap. (c) The same information as a list.

Memory Management with Linked Lists
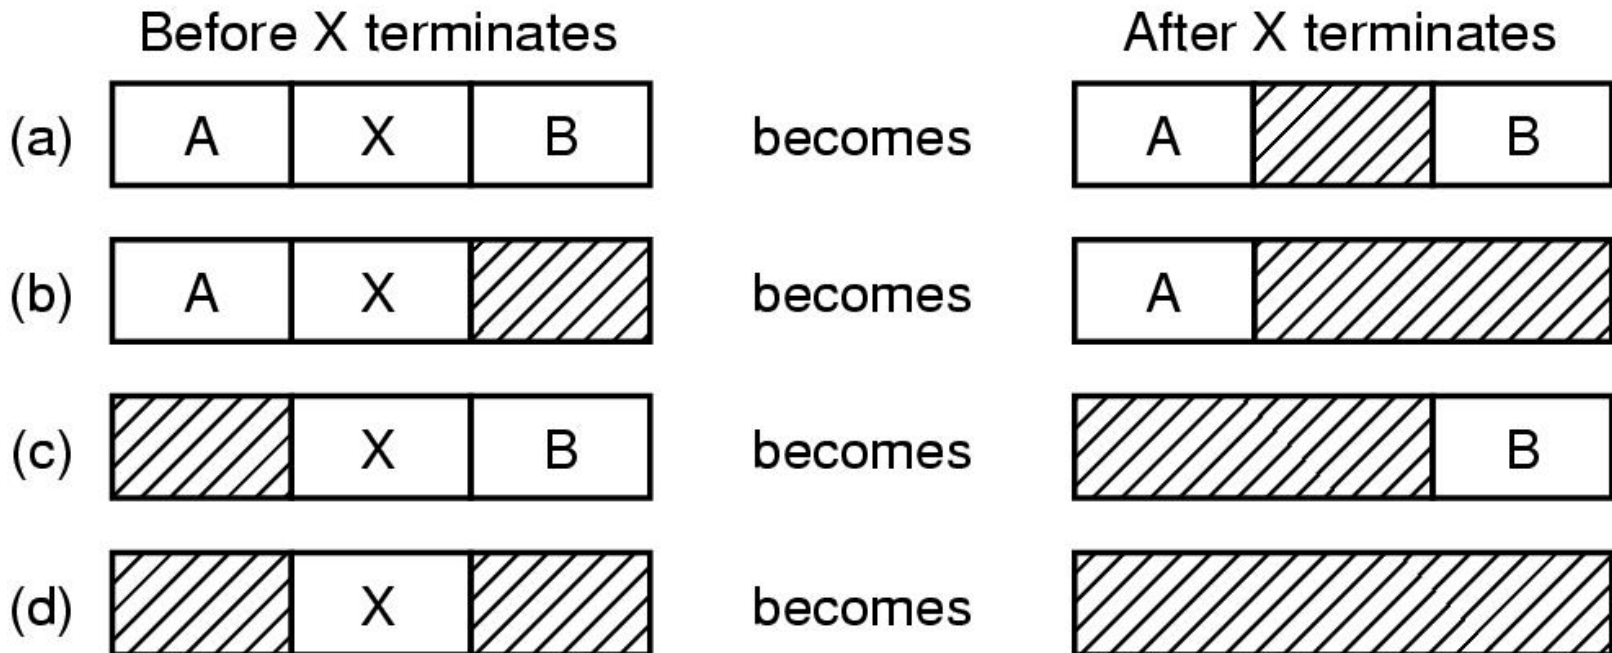
Figure 3-7. Four neighbor combinations
for the terminating process, X.

Virtual Memory – Paging (1)

Figure 3-8. The position and function of the MMU – shown as being a part of the CPU chip (it commonly is nowadays). Logically it could be a separate chip, was in years gone by.

Paging (2)

Figure 3-9. Relation between virtual addresses and physical memory addresses given by page table.

Paging (3)
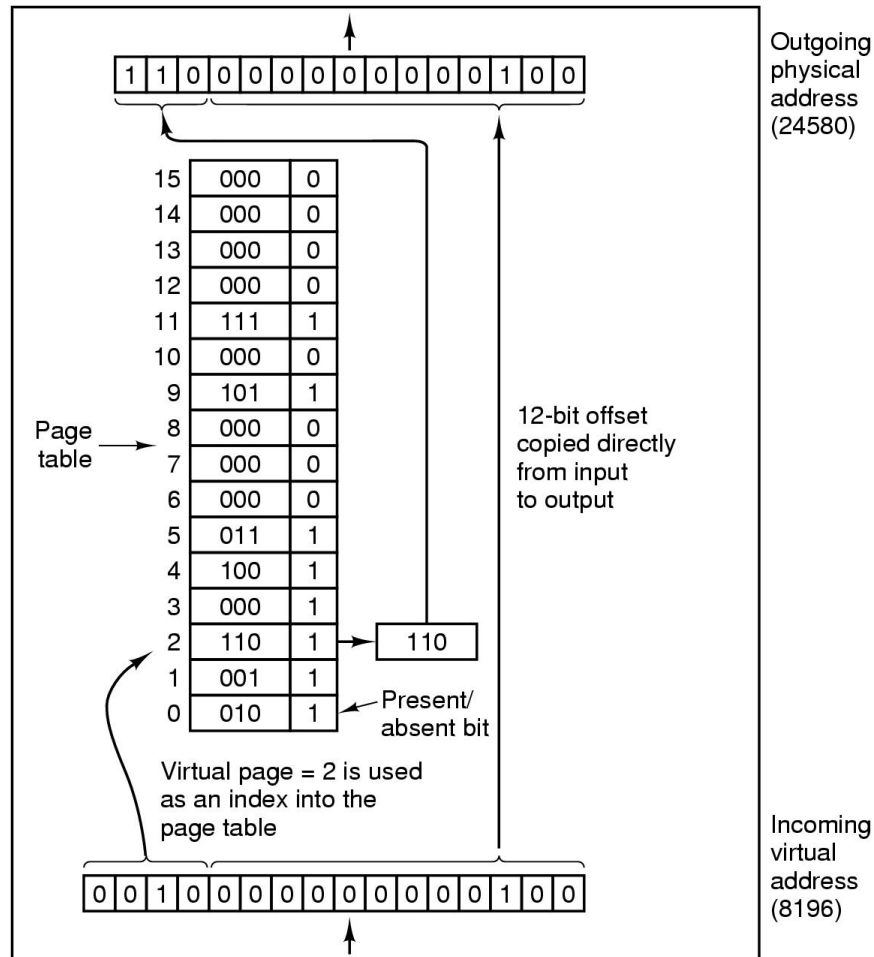
Figure 3-10. The internal operation of the MMU with 16 4-KB pages.

Structure of Page Table Entry
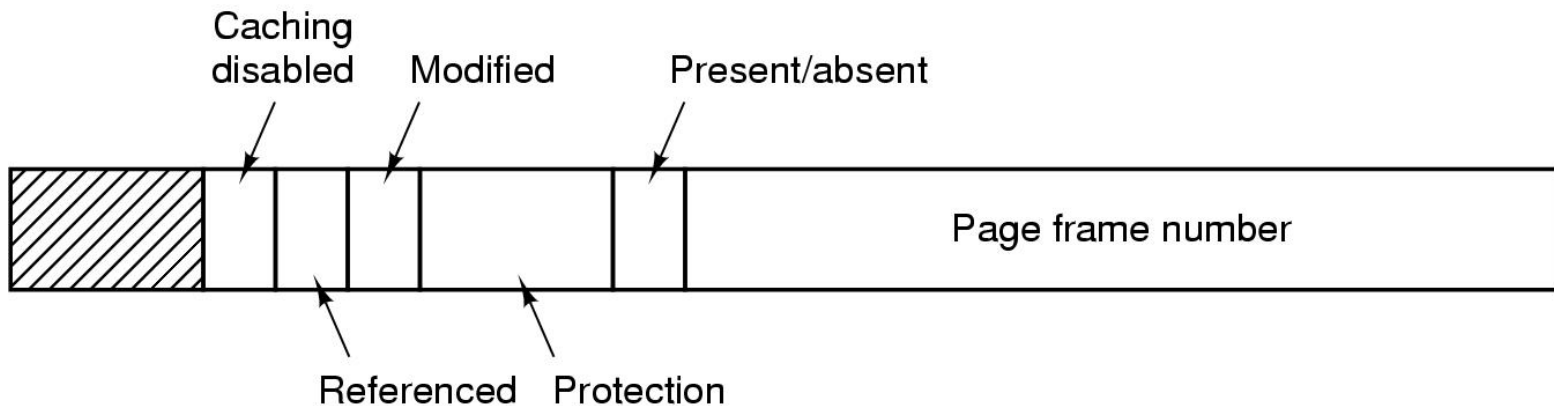
Figure 3-11. A typical page table entry.

Speeding Up Paging

Paging implementation issues:

- The mapping from virtual address to physical address must be fast.

- If the virtual address space is large, the page table will be large.

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

Figure 3-12. A TLB to speed up paging.

Multilevel Page Tables
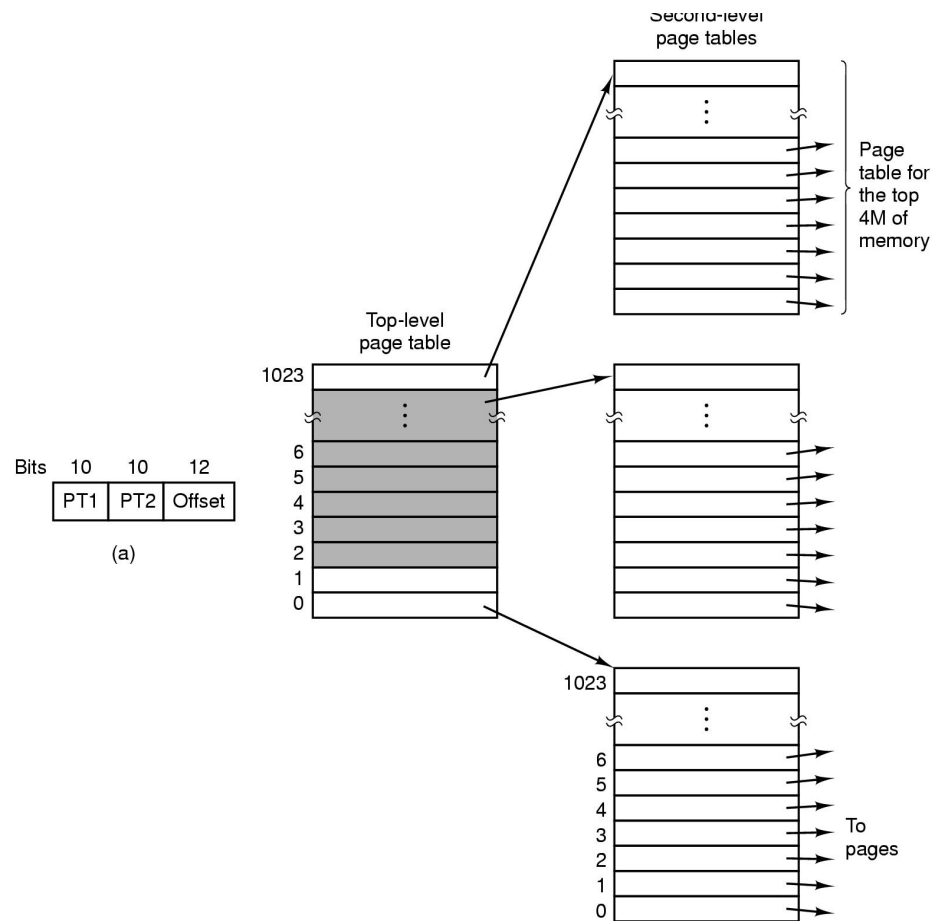
Figure 3-13. (a) A 32-bit address with two page table fields. (b) Two-level page tables.

Inverted Page Tables

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

$2^{52}-1$

0

Indexed by virtual page

1-GB physical memory has 2^{18} 4-KB page frames

$2^{18}-1$

0

Hash table

$2^{18}-1$

0

Indexed by hash on virtual page

Virtual page

Page frame

Figure 3-14. Comparison of a traditional page table with an inverted page table.

Page Replacement Algorithms

- Optimal page replacement algorithm
- Not recently used page replacement
- First-In, First-Out page replacement
- Second chance page replacement
- Clock page replacement
- Least recently used page replacement
- Working set page replacement
- WSClock page replacement

Second Chance Algorithm
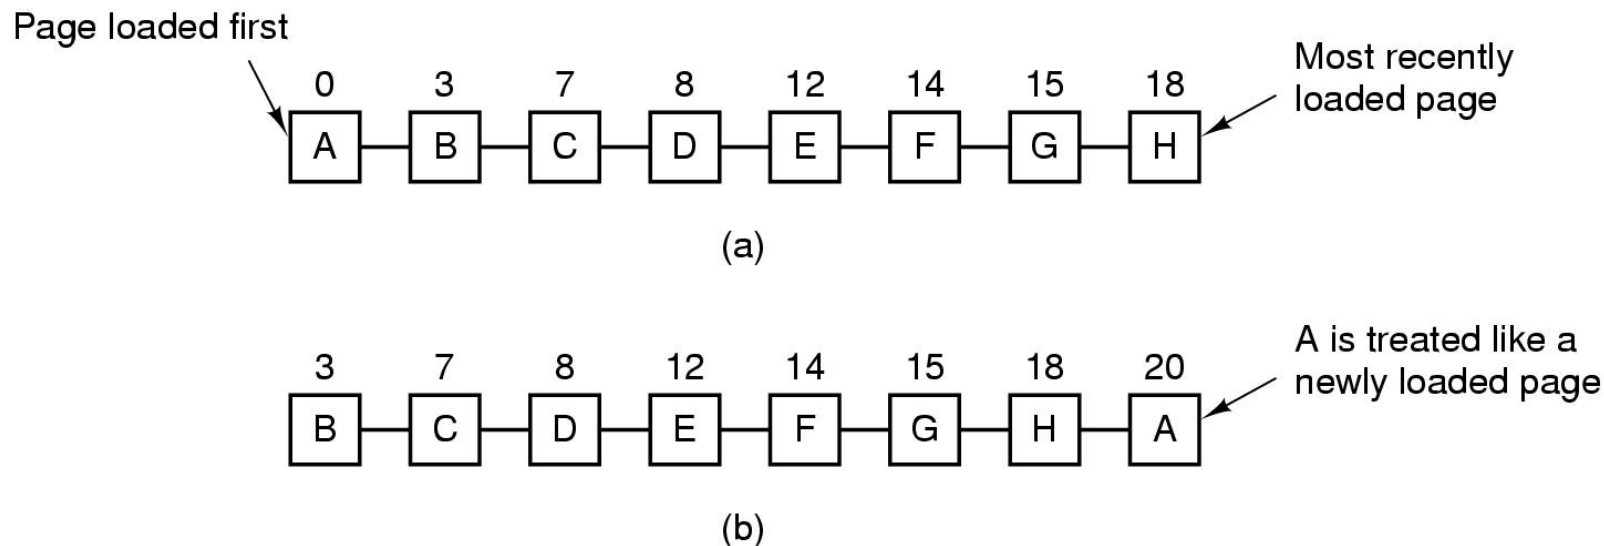
Figure 3-15. Operation of second chance.
(a) Pages sorted in FIFO order.
(b) Page list if a page fault occurs at time 20 and A has its R bit set. The numbers above the pages are their load times.

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

Figure 3-16. The clock page replacement algorithm.

LRU Page Replacement Algorithm

Figure 3-17. LRU using a matrix when pages are referenced in the order 0, 1, 2, 3, 2, 1, 0, 3, 2, 3.

Simulating LRU in Software

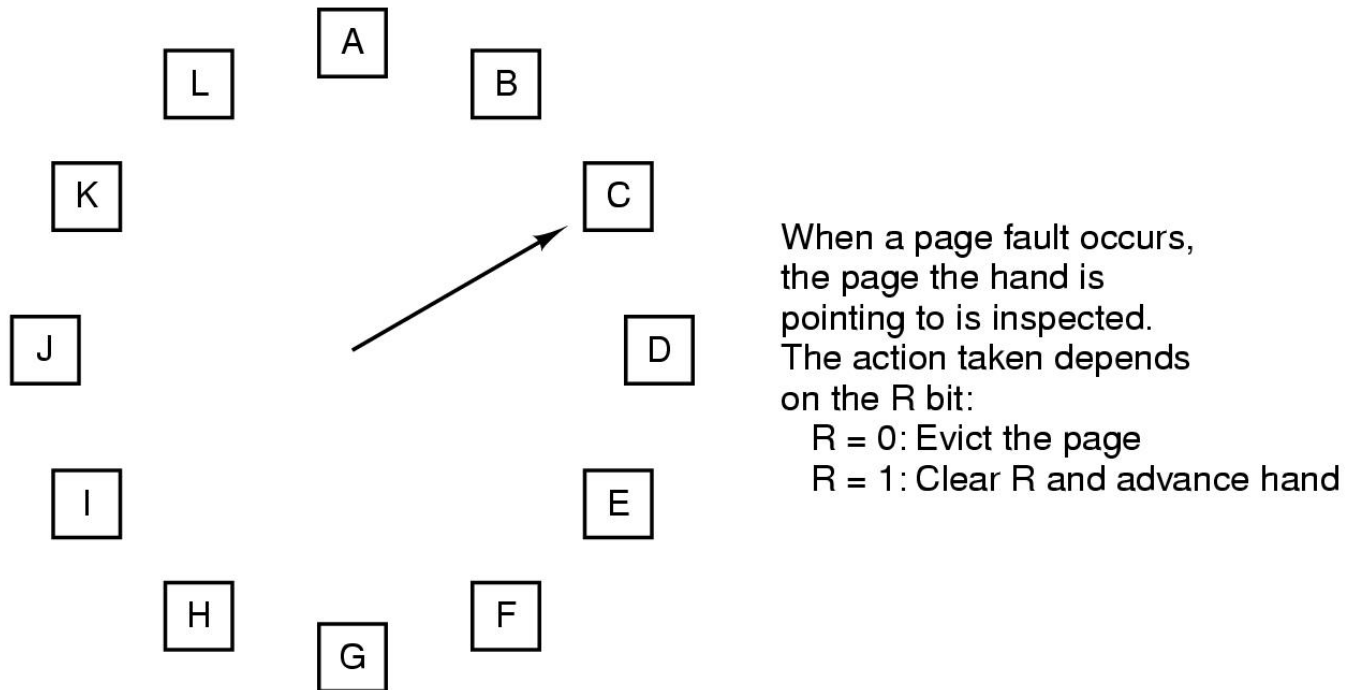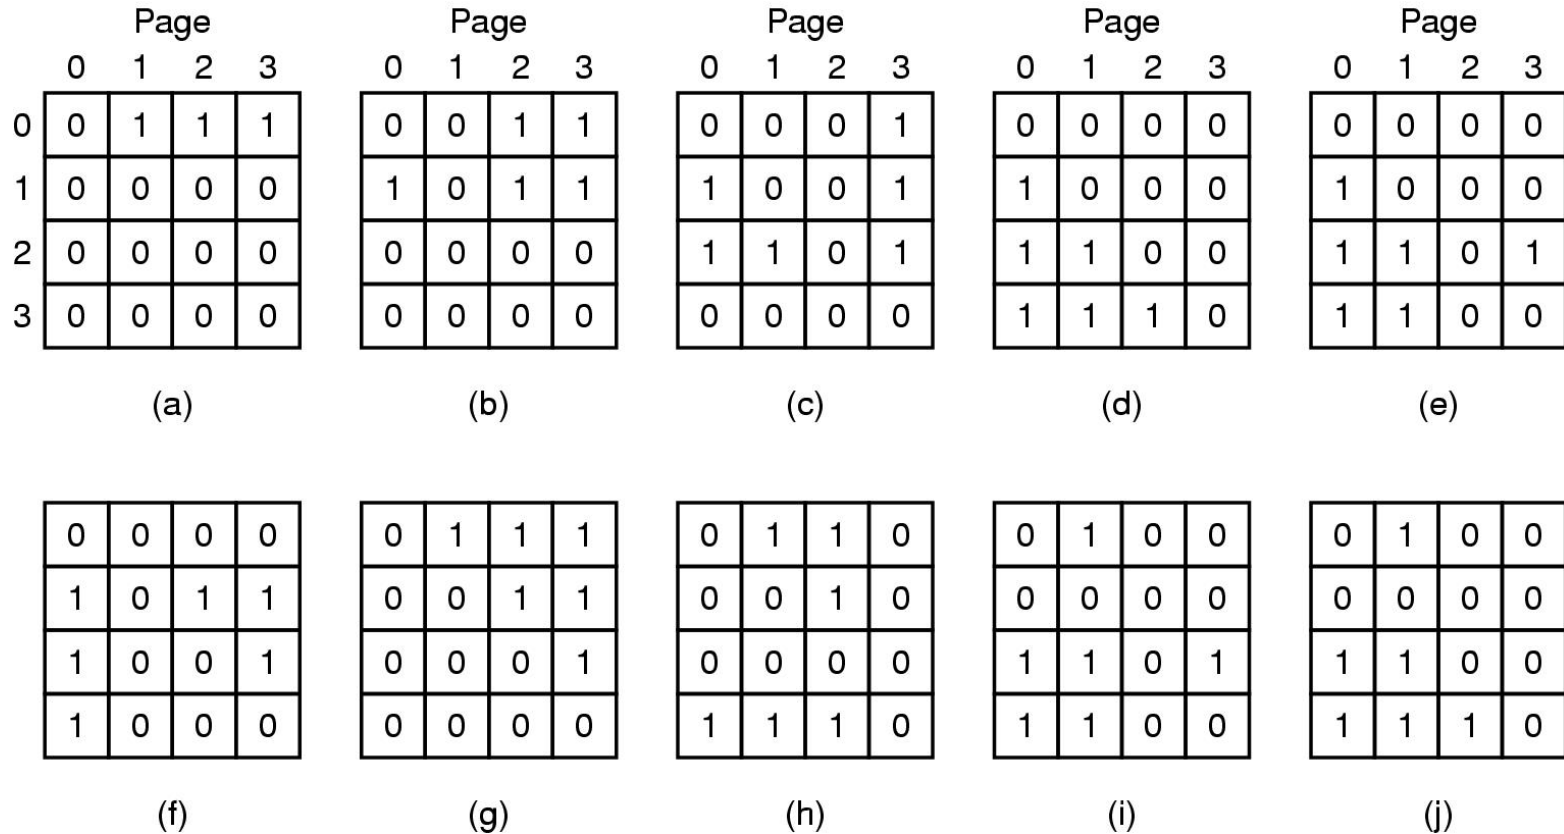| R bits for pages 0-5, clock tick 0	R bits for pages 0-5, clock tick 1	R bits for pages 0-5, clock tick 2	R bits for pages 0-5, clock tick 3	R bits for pages 0-5, clock tick 4
1 0 1 0 1 1	1 1 0 0 1 0	1 1 0 1 0 1	1 0 0 0 1 0	0 1 1 0 0 0

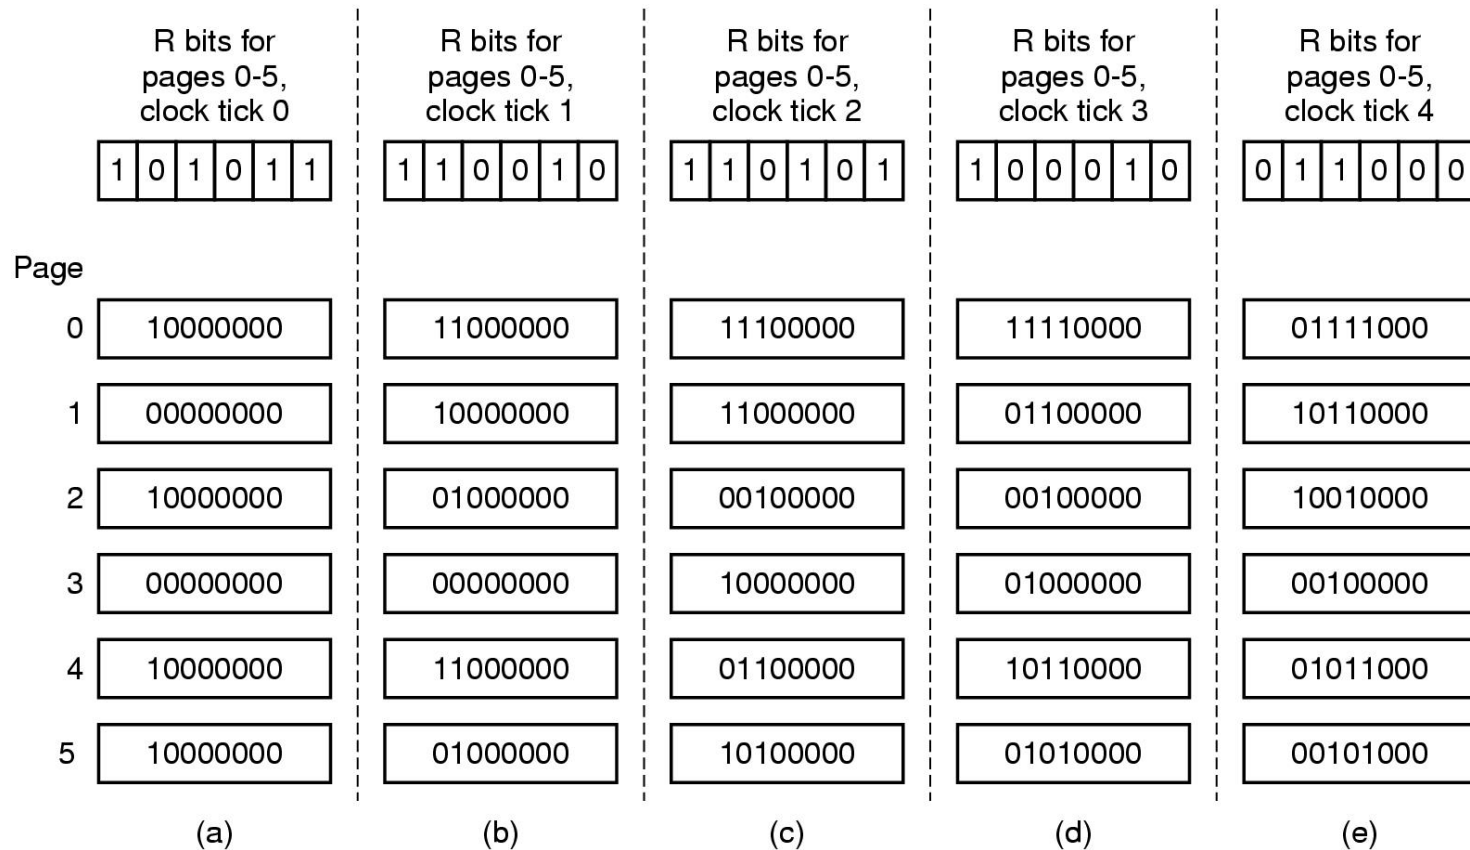| Page					
0	10000000	11000000	11100000	11110000	01111000
1	00000000	10000000	11000000	01100000	10110000
2	10000000	01000000	00100000	00100000	10010000
3	00000000	00000000	10000000	01000000	00100000
4	10000000	11000000	01100000	10110000	01011000
5	10000000	01000000	10100000	01010000	00101000
	(a)	(b)	(c)	(d)	(e)

Figure 3-18. The aging algorithm simulates LRU in software. Shown are six pages for five clock ticks. The five clock ticks are represented by (a) to (e).

Working Set Page Replacement (1)

Figure 3-19. The working set is the set of pages used by the k most recent memory references. The function w(k, t) is the size of the working set at time t.
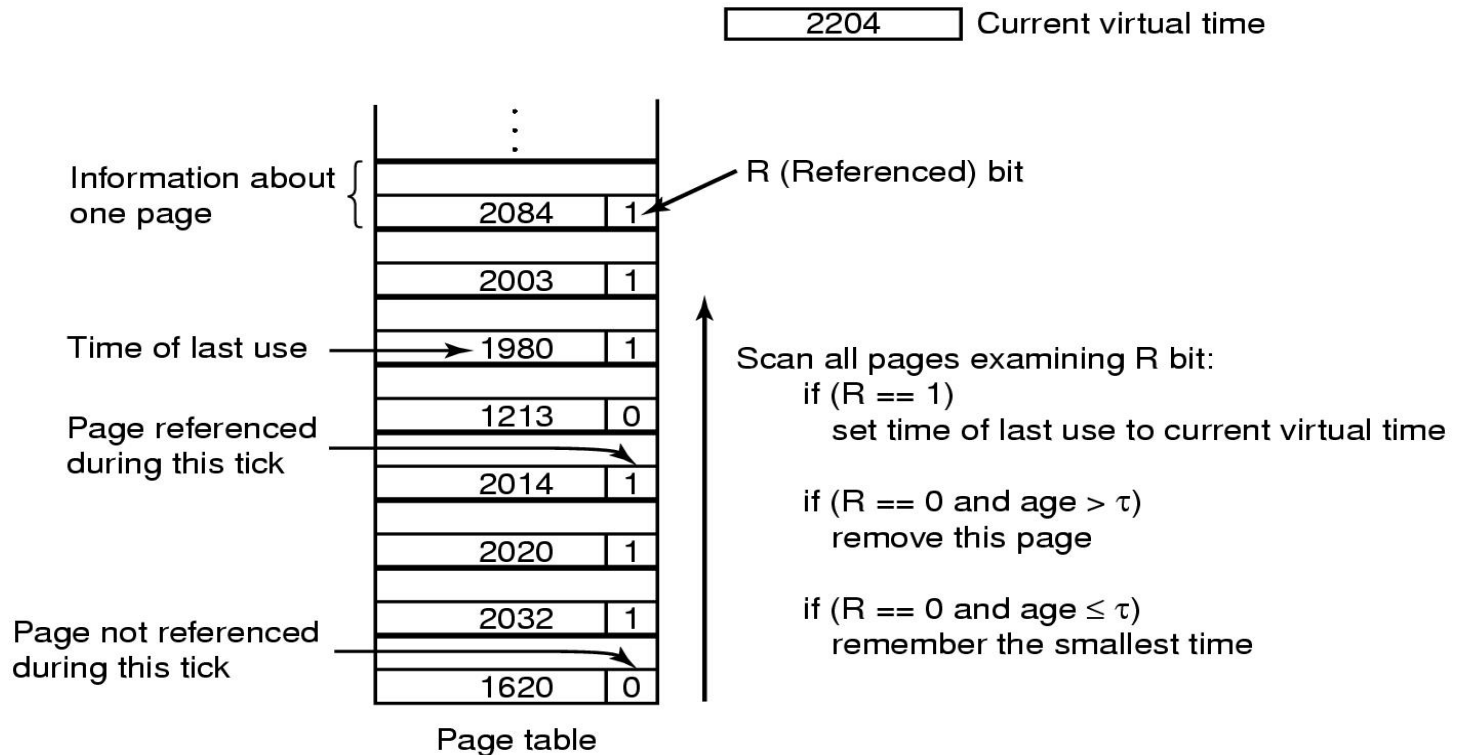
Working Set Page Replacement (2)

Figure 3-20. The working set algorithm.

The WSClock Page Replacement Algorithm (1)

When the hand comes all the way around to its starting point there are two cases to consider:

- At least one write has been scheduled.
- No writes have been scheduled.
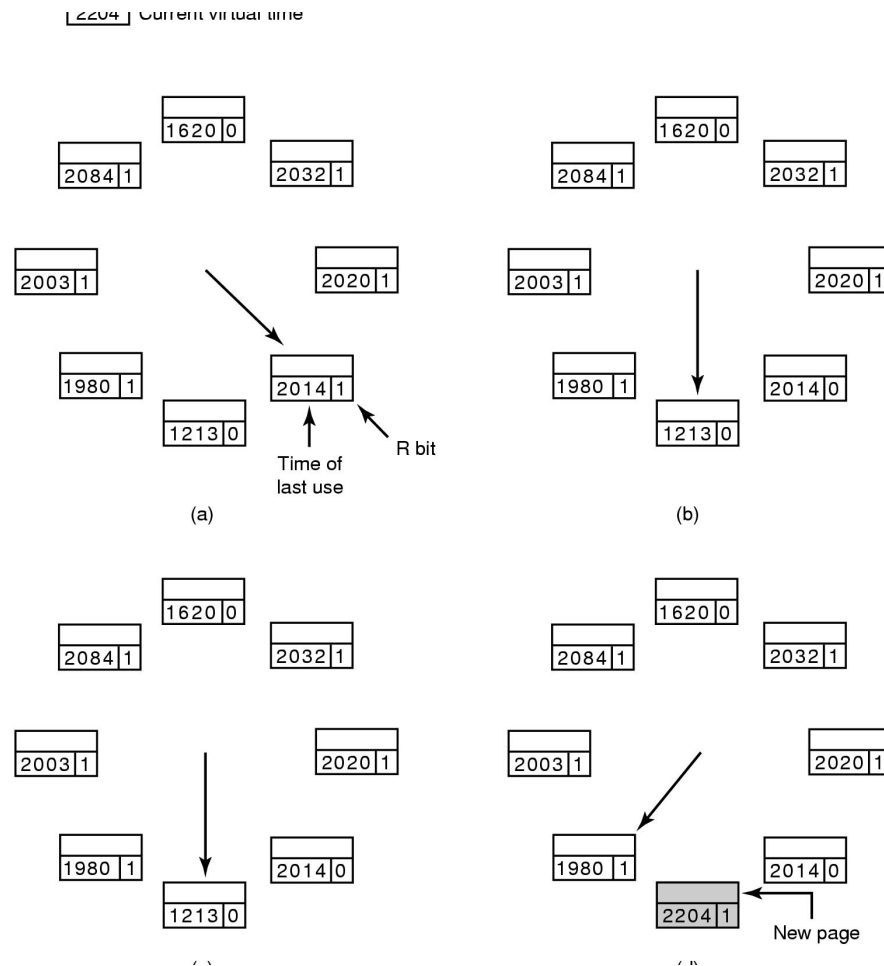
The WSClock Page Replacement Algorithm (2)

2204 Current virtual time

1620 0
2084 1   2032 1
2003 1   2020 1
1980 1   2014 1
        1213 0   R bit
        Time of
        last use
        (a)

1620 0
2084 1   2032 1
2003 1   2020 1
1980 1   2014 0
        1213 0
        (b)

1620 0
2084 1   2032 1
2003 1   2020 1
1980 1   2014 0
        1213 0

1620 0
2084 1   2032 1
2003 1   2020 1
1980 1   2014 0
        2204 1
        New page

Figure 3-21. Operation of the WSClock algorithm. (a) and (b) give an example of what happens when R = 1.

2204 Current virtual time

1620 0
2084 1
2032 1
2003 1
2020 1
1980 1
2014 1
1213 0
R bit
Time of last use

(a)

1620 0
2084 1
2032 1
2003 1
2020 1
1980 1
2014 0
1213 0

(b)

1620 0
2084 1
2032 1
2003 1
2020 1
1980 1
2014 0
1213 0

1620 0
2084 1
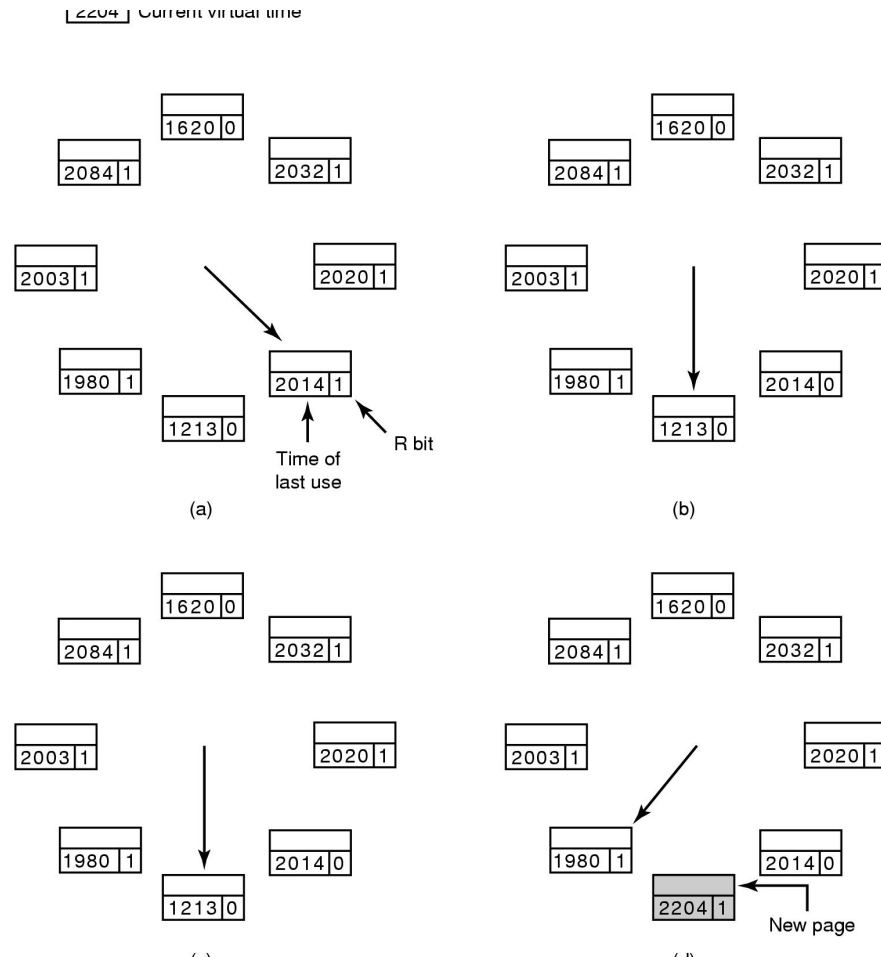2032 1
2003 1
2020 1
1980 1
2014 0
2204 1
New page

Figure 3-21. Operation of the WSClock algorithm.
(c) and (d) give an example of R = 0.

Summary of Page Replacement Algorithms

Algorithm	Comment
Optimal	Not implementable, but useful as a benchmark
NRU (Not Recently Used)	Very crude approximation of LRU
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

Figure 3-22. Page replacement algorithms discussed in the text.

Local versus Global Allocation Policies (1)

	Age
A0	10
A1	7
A2	5
A3	4
A4	6
A5	3
B0	9
B1	4
B2	6
B3	2
B4	5
B5	6
B6	12
C1	3
C2	5
C3	6

(a)

A0
A1
A2
A3
A4
(A6)
B0
B1
B2
B3
B4
B5
B6
C1
C2
C3

(b)

A0
A1
A2
A3
A4
A5
B0
B1
B2
(A6)
B4
B5
B6
C1
C2
C3

(c)

(a) Original configuration. (b) Local page replacement. (c) Global page replacement.

Local versus Global Allocation Policies (2)

Figure 3-24. Page fault rate as a function
of the number of page frames assigned.

# Separate Instruction and Data Spaces

Figure 3-25. (a) One address space.
(b) Separate I and D spaces.

Shared Pages

Figure 3-26. Two processes sharing the same program sharing its page table.

Shared Libraries

Figure 3-27. A shared library being used by two processes.

Page Fault Handling (1)

- The hardware traps to the kernel, saving the program counter on the stack.

- An assembly code routine is started to save the general registers and other volatile information.

- The operating system discovers that a page fault has occurred, and tries to discover which virtual page is needed.

- Once the virtual address that caused the fault is known, the system checks to see if this address is valid and the protection consistent with the access

Page Fault Handling (2)

- If the page frame selected is dirty, the page is scheduled for transfer to the disk, and a context switch takes place.

- When page frame is clean, operating system looks up the disk address where the needed page is, schedules a disk operation to bring it in.

- When disk interrupt indicates page has arrived, page tables updated to reflect position, frame marked as being in normal state.

Page Fault Handling (3)

- Faulting instruction backed up to state it had when it began and program counter reset to point to that instruction.

- Faulting process scheduled, operating system returns to the (assembly language) routine that called it.

- This routine reloads registers and other state information and returns to user space to continue execution, as if no fault had occurred.

Instruction Backup

MOVE.L #6(A1), 2(A0)

	16 Bits	
1000	MOVE	} Opcode
1002	6	} First operand
1004	2	} Second operand

Figure 3-28. An instruction causing a page fault.
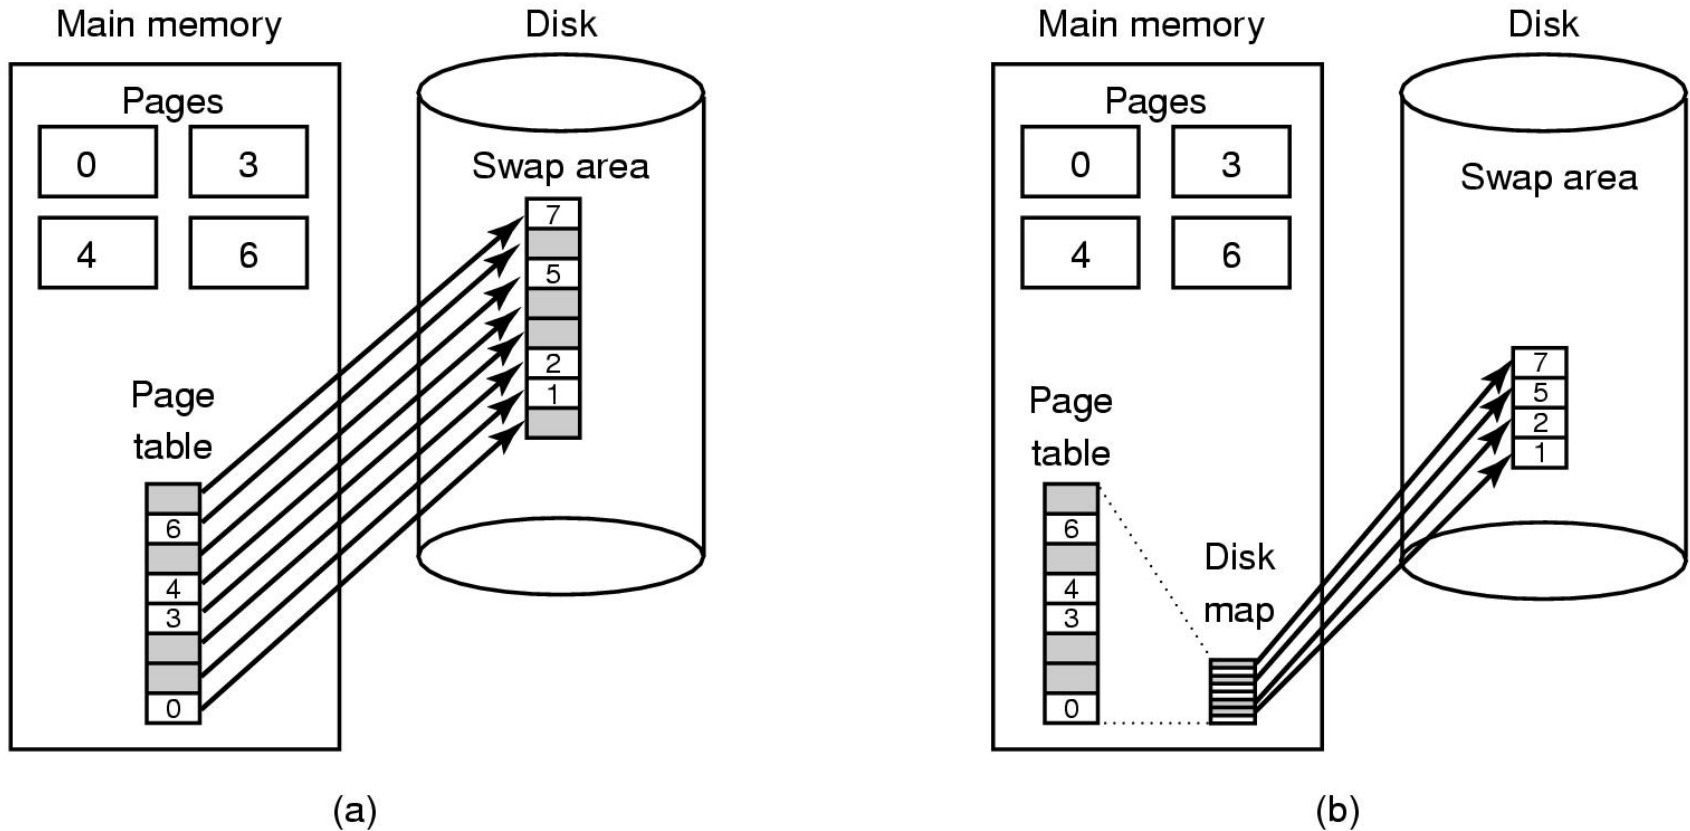
Backing Store (1)

Figure 3-29. (a) Paging to a static swap area.

Backing Store (2)

Figure 3-29. (b) Backing up pages dynamically.

Separation of Policy and Mechanism (1)

Memory management system is divided into three parts:

- A low-level MMU handler.
- A page fault handler that is part of the kernel.
- An external pager running in user space.
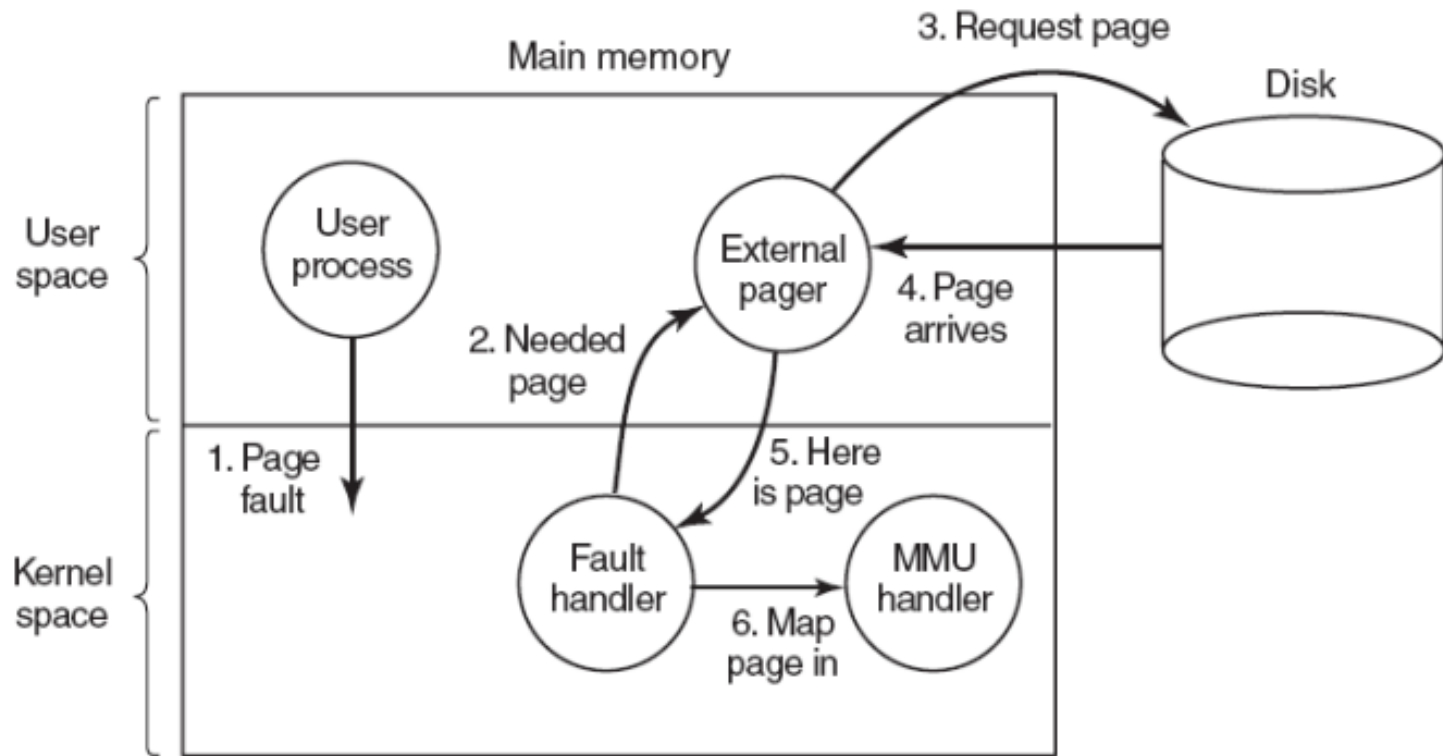
Separation of Policy and Mechanism (2)

Figure 3-30. Page fault handling with an external pager.

Segmentation (1)

A compiler has many tables that are built up as compilation proceeds, possibly including:

- The source text being saved for the printed listing (on batch systems).
- The symbol table – the names and attributes of variables.
- The table containing integer, floating-point constants used.
- The parse tree, the syntactic analysis of the program.
- The stack used for procedure calls within the compiler.
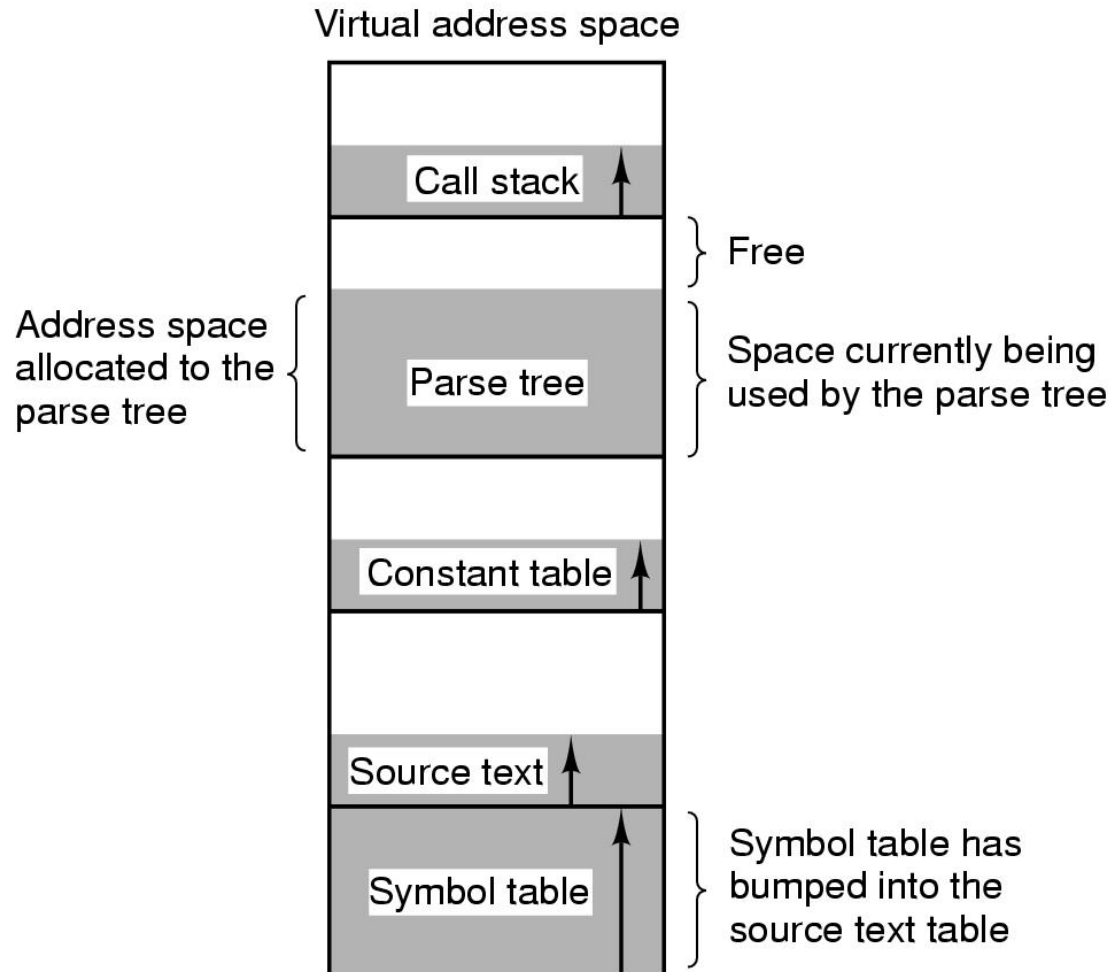
Segmentation (2)

Figure 3-31. In a one-dimensional address space with growing tables, one table may bump into another.
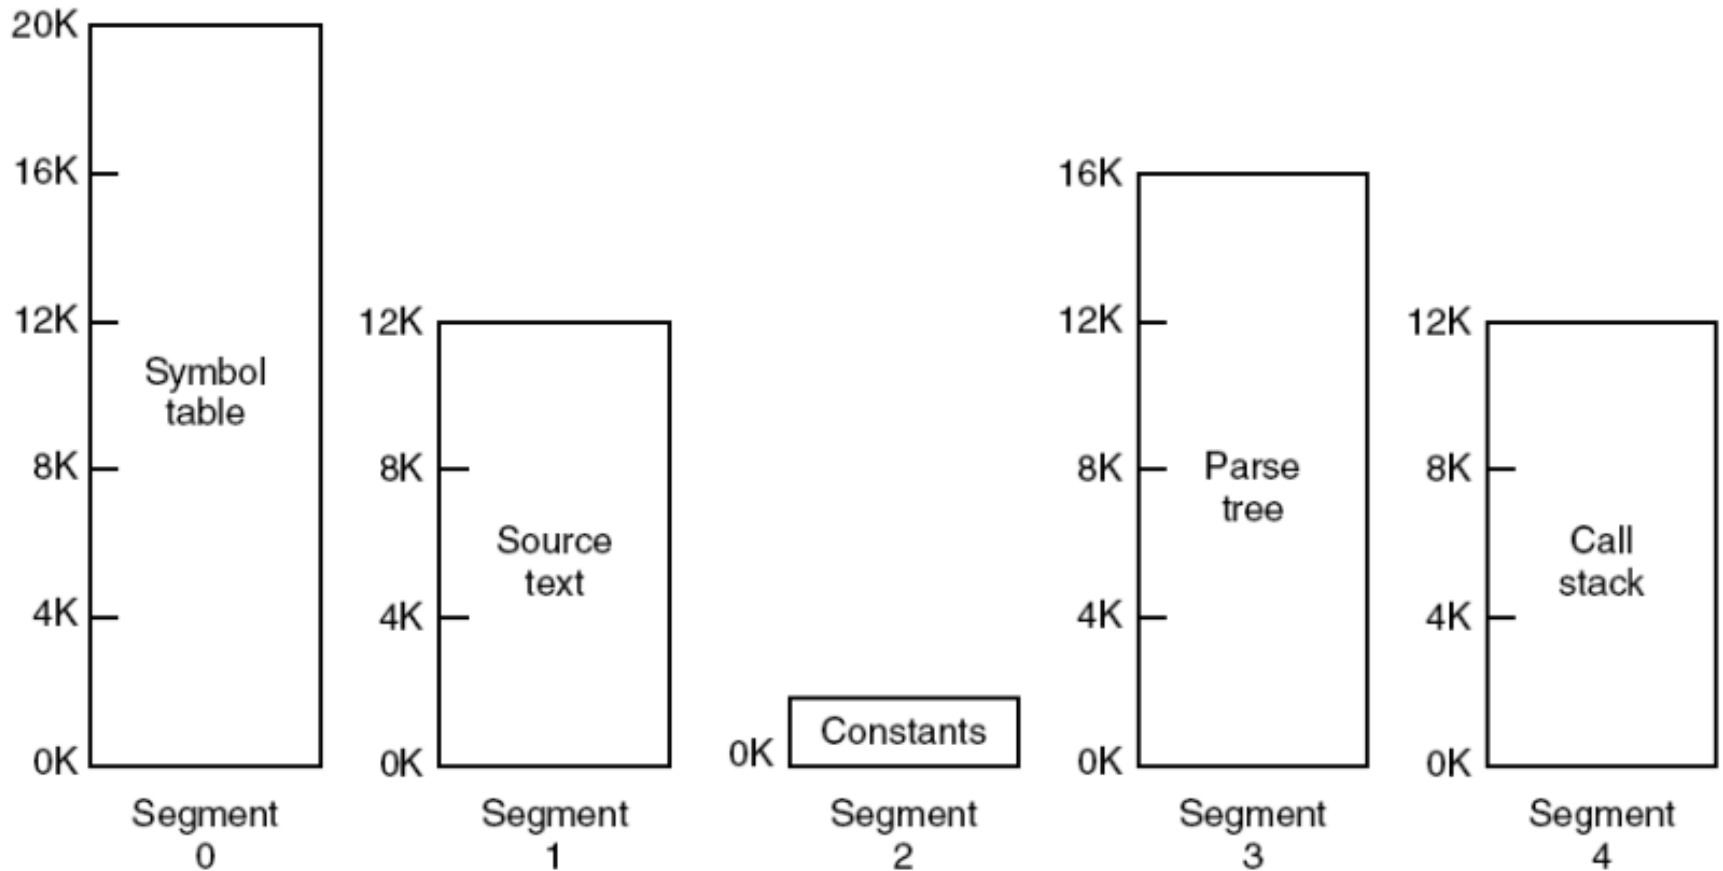
Segmentation (3)

Figure 3-32. A segmented memory allows each table to grow or shrink independently of the other tables.

Implementation of Pure Segmentation

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

Figure 3-33. Comparison of paging and segmentation.
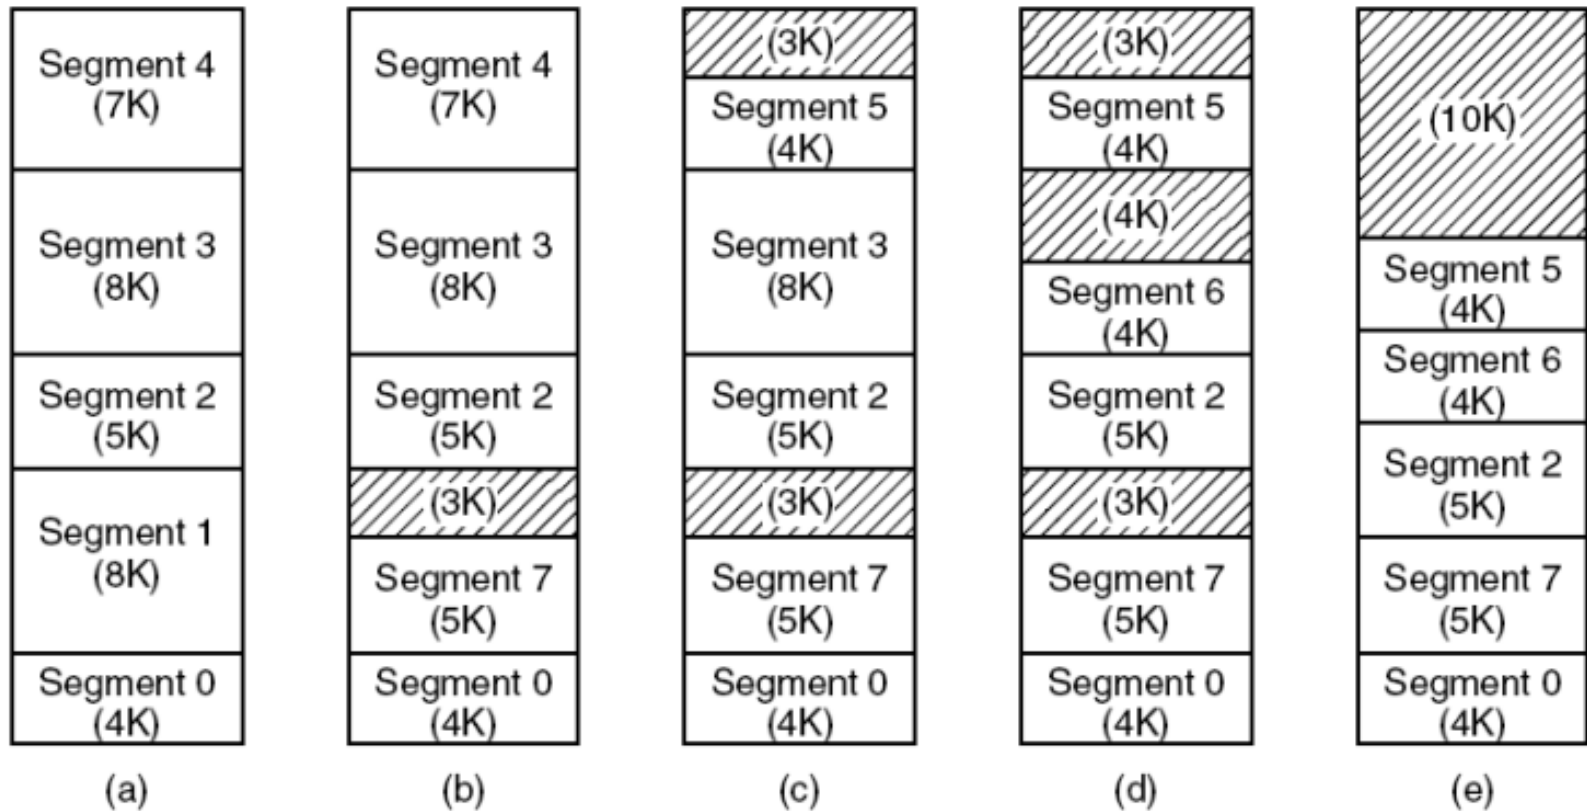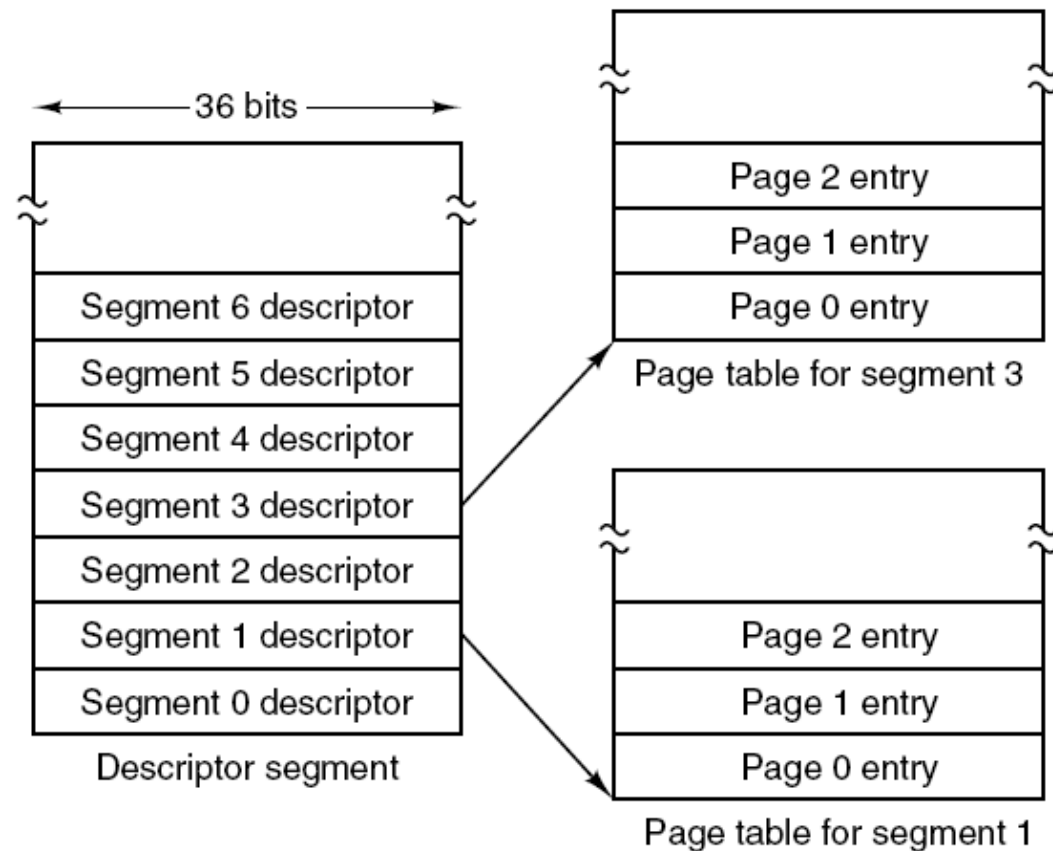
Segmentation with Paging: MULTICS (1)

Figure 3-34. (a)-(d) Development of checkerboarding. (e) Removal of the checkerboarding by compaction.
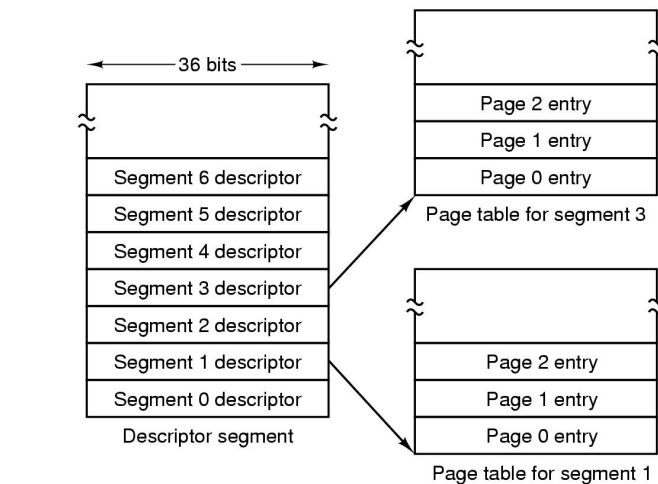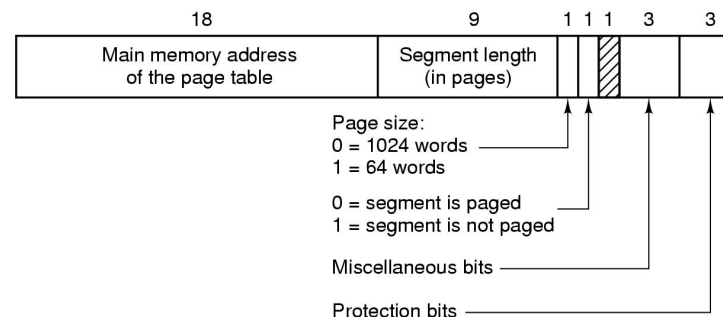
Figure 3-35. The MULTICS virtual memory. (a) The descriptor segment points to the page tables.

Segmentation with Paging: MULTICS (5)

(a)

```
          18                              9        1 1 1   3      3
┌──────────────────────────┬──────────────────┬─┬─┬─┬─────┬─────┐
│   Main memory address    │  Segment length  │ │▨│ │     │     │
│    of the page table     │    (in pages)    │ │ │ │     │     │
└──────────────────────────┴──────────────────┴─┴─┴─┴─────┴─────┘
```

Page size:
0 = 1024 words
1 = 64 words

0 = segment is paged
1 = segment is not paged

Miscellaneous bits

Protection bits

(b)

Figure 3-35. The MULTICS virtual memory. (b) A segment descriptor. The numbers are the field lengths.

Segmentation with Paging: MULTICS (6)

When a memory reference occurs, the following algorithm is carried out:

- The segment number used to find segment descriptor.
- Check is made to see if the segment's page table is in memory.
  - If not, segment fault occurs.
  - If there is a protection violation, a fault (trap) occurs.

Segmentation with Paging: MULTICS (7)

- Page table entry for the requested virtual page examined.
  - If the page itself is not in memory, a page fault is triggered.
  - If it is in memory, the main memory address of the start of the page is extracted from the page table entry

- The offset is added to the page origin to give the main memory address where the word is located.

- The read or store finally takes place.

Segmentation with Paging: MULTICS (8)

Address within
the segment

Segment number		Page number	Offset within the page
18		6	10

Figure 3-36. A 34-bit MULTICS virtual address.

Segmentation with Paging: MULTICS (9)

MULTICS virtual address

| Segment number | | Page number | Offset |
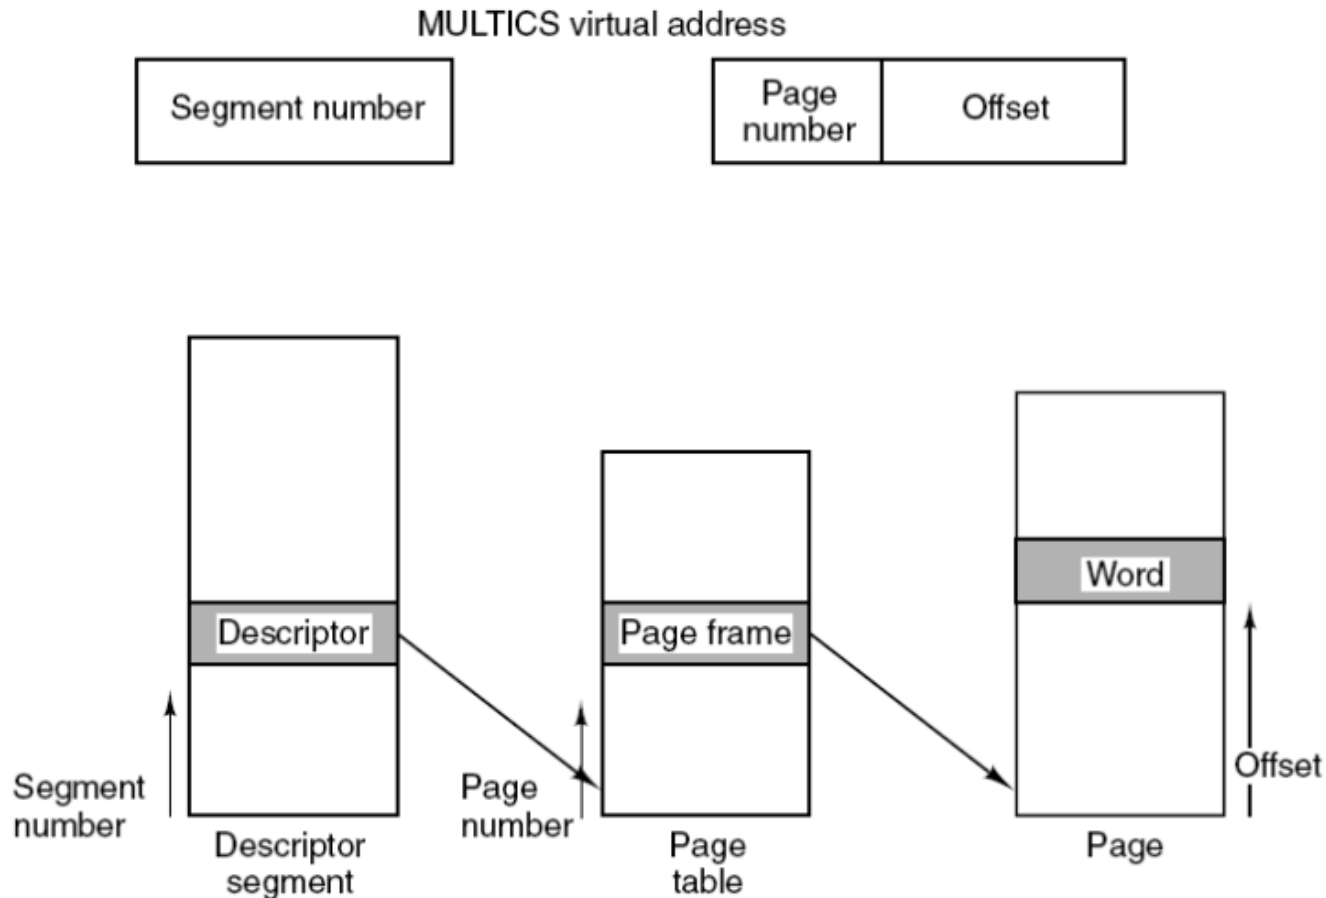
Figure 3-37. Conversion of a two-part MULTICS address into a main memory address.

Comparison field		Page frame	Protection	Age	Is this entry used?
Segment number	Virtual page				
4	1	7	Read/write	13	1
6	0	2	Read only	10	1
12	3	1	Read/write	2	1
					0
2	1	0	Execute only	7	1
2	2	12	Execute only	9	1

Figure 3-38. A simplified version of the MULTICS TLB. The existence of two page sizes makes the actual TLB more complicated.

Segmentation with Paging: The Pentium (1)
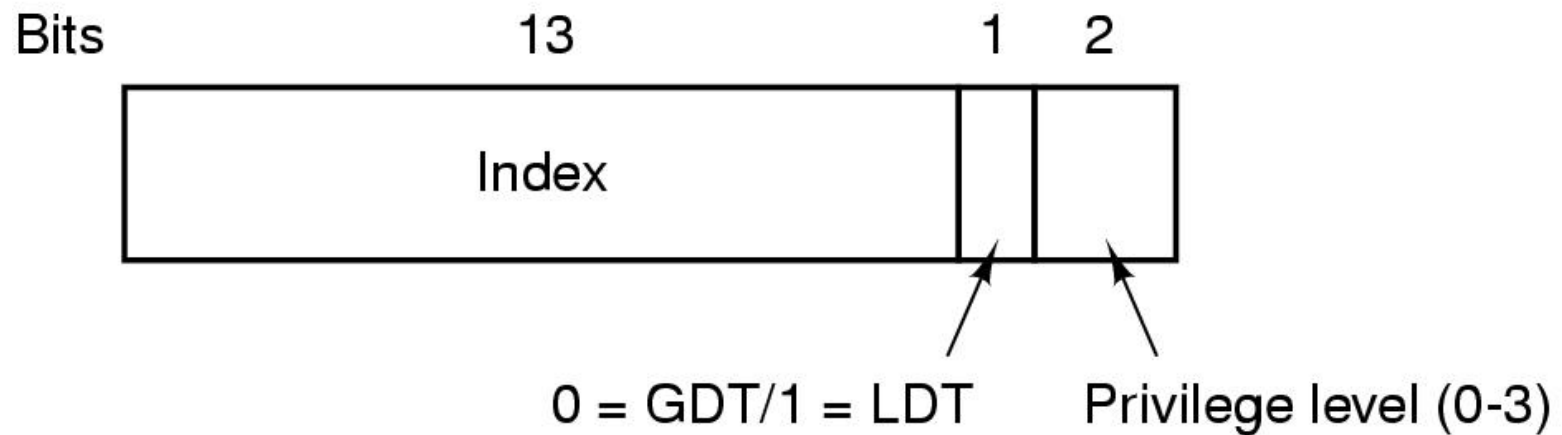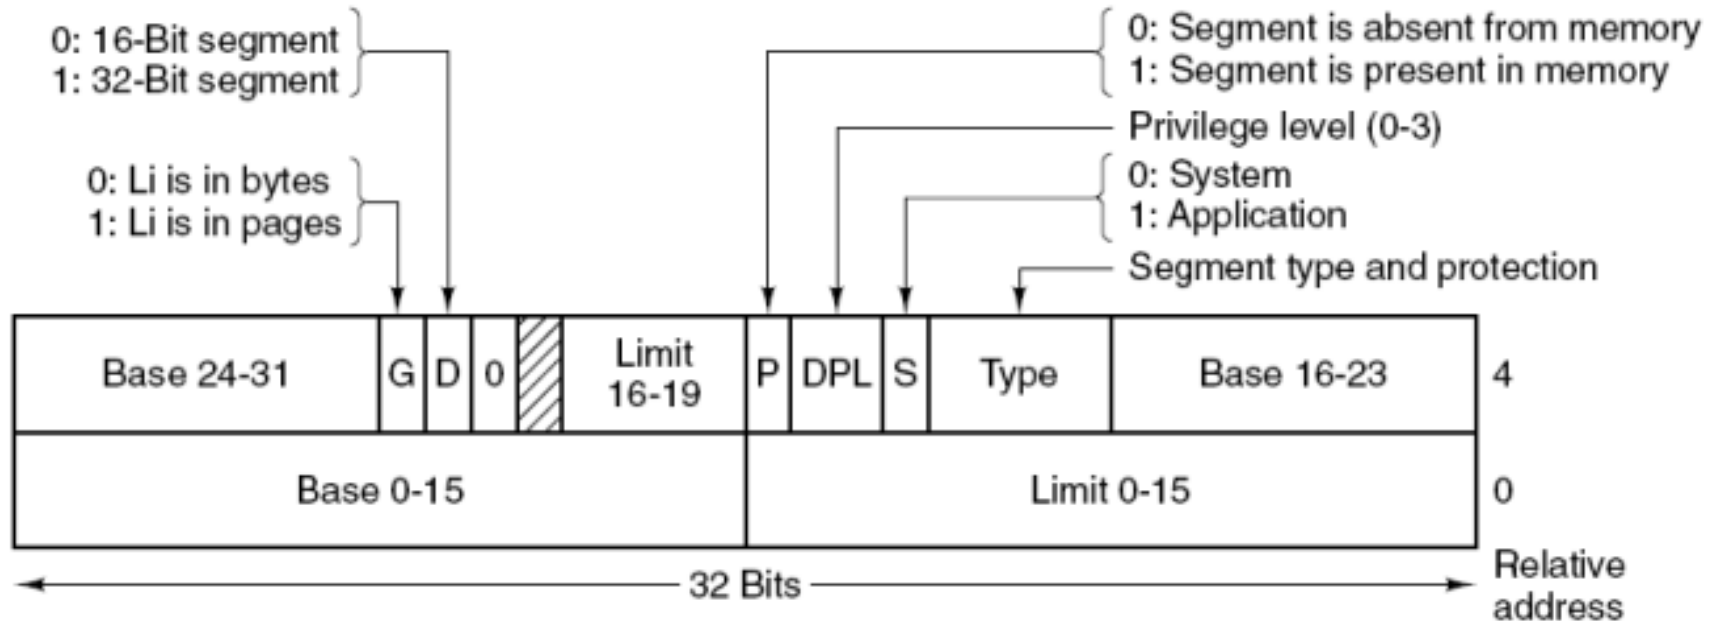

Figure 3-39. A Pentium selector.

Figure 3-40. Pentium code segment descriptor.
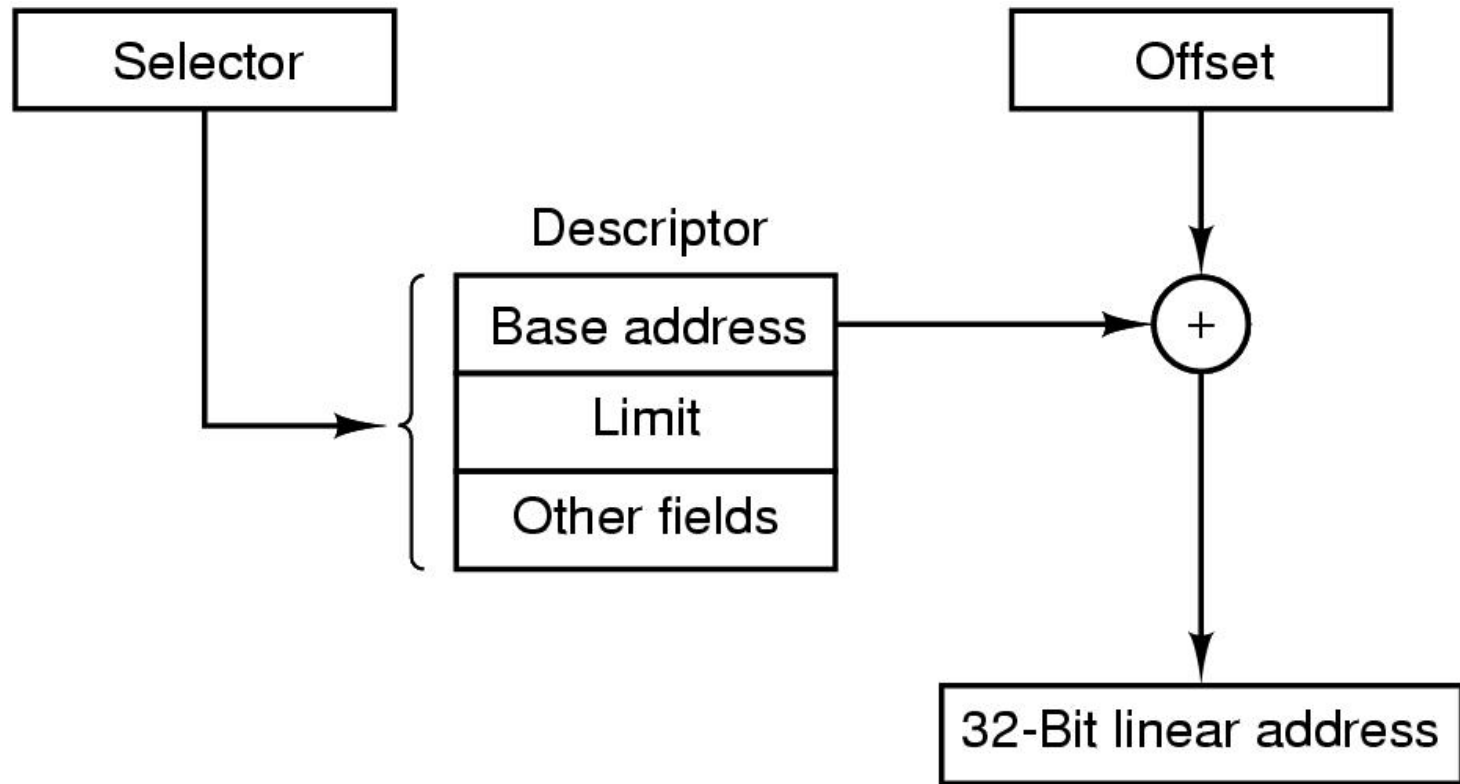Data segments differ slightly.

Figure 3-41. Conversion of a (selector, offset)
pair to a linear address.

Segmentation with Paging: The Pentium (4)

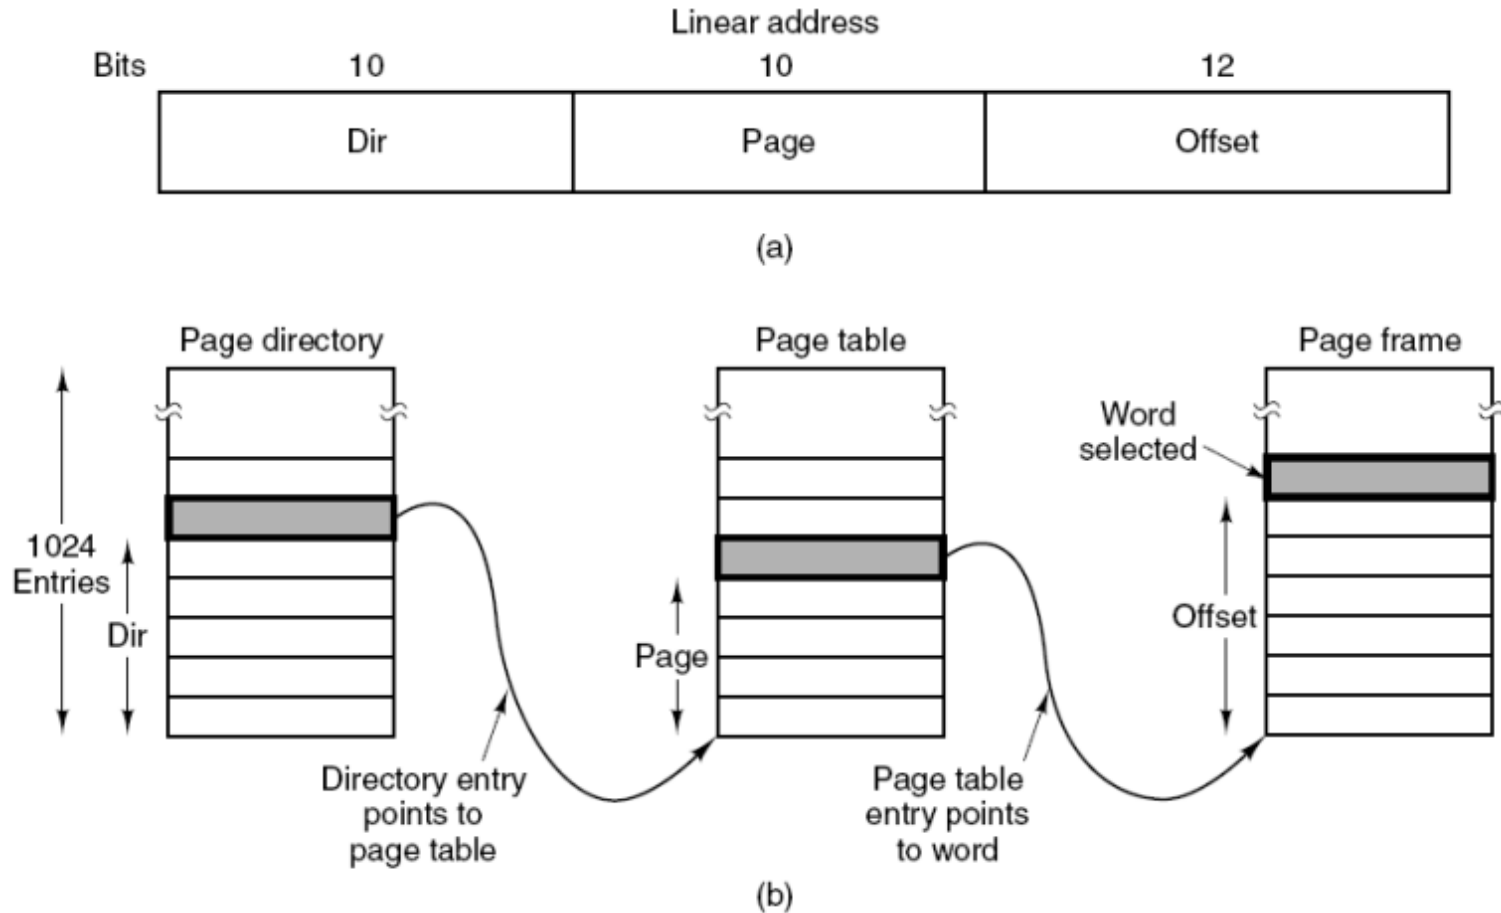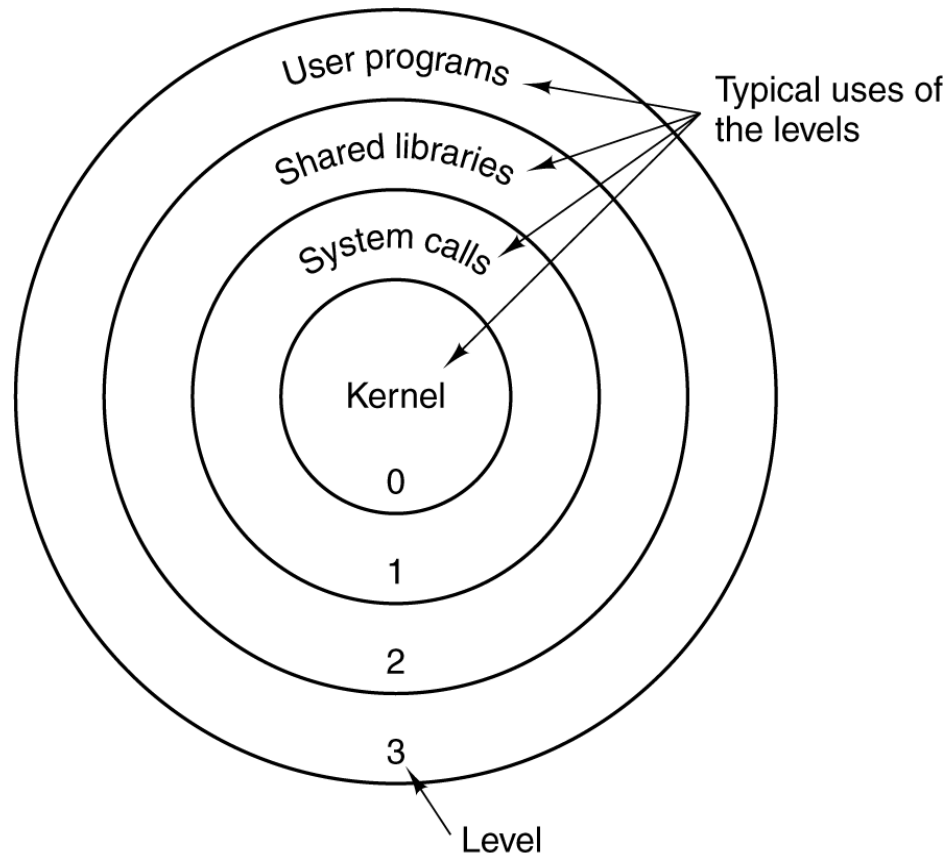Figure 3-42. Mapping of a linear address onto a physical address.

Figure 3-43. Protection on the Pentium.