

Chapter 8: Main Memory



Chapter 8: Memory Management

- Background
- Swapping
- Contiguous Memory Allocation
- Paging
- Structure of the Page Table
- Segmentation
- Example: The Intel Pentium



Objectives

- To provide a detailed description of various ways of organizing memory hardware
- To discuss various memory-management techniques, including paging and segmentation
- To provide a detailed description of the Intel Pentium, which supports both pure segmentation and segmentation with paging



Background

- Program must be brought (from disk) into memory and placed within a process for it to be run
- Main memory and registers are only storage CPU can access directly
- Register access in one CPU clock (or less)
- Main memory can take many cycles
- **Cache** sits between main memory and CPU registers
- Protection of memory required to ensure correct operation



Base and Limit Registers

■ A pair of **base** and **limit** registers define the logical address space

0
256000
300040
420940
880000
1024000

base
limit

Operating System Concepts – 8th Edition 8.5 Silberschatz, Galvin and Gagne ©2009

Binding of Instructions and Data to Memory

■ Address binding of instructions and data to memory addresses can happen at three different stages

- **Compile time:** If memory location known a priori, **absolute code** can be generated; must recompile code if starting location changes
- **Load time:** Must generate **relocatable code** if memory location is not known at compile time
- **Execution time:** Binding delayed until run time if the process can be moved during its execution from one memory segment to another. Need hardware support for address maps (e.g., base and limit registers)

Operating System Concepts – 8th Edition 8.6 Silberschatz, Galvin and Gagne ©2009

Multistep Processing of a User Program

source program
compiler or assembler
object module
linkage editor
load module
loader
in-memory binary memory image
executor
run-time image

Other object modules
system library
dynamically loaded system library
dynamic linking

compile time
load time
execution time (run time)

Operating System Concepts – 8th Edition 8.7 Silberschatz, Galvin and Gagne ©2009

Logical vs. Physical Address Space

■ The concept of a logical address space that is bound to a separate **physical address space** is central to proper memory management

- **Logical address** – generated by the CPU; also referred to as **virtual address**
- **Physical address** – address seen by the memory unit

■ Logical and physical addresses are the same in compile-time and load-time address-binding schemes; logical (virtual) and physical addresses differ in execution-time address-binding scheme

Operating System Concepts – 8th Edition 8.8 Silberschatz, Galvin and Gagne ©2009

Memory-Management Unit (MMU)

- Hardware device that maps virtual to physical address
- In MMU scheme, the value in the relocation register is added to every address generated by a user process at the time it is sent to memory
- The user program deals with *logical* addresses; it never sees the *real* physical addresses

Operating System Concepts – 8th Edition 8.9 Silberschatz, Galvin and Gagne ©2009

Dynamic relocation using a relocation register

```

graph LR
    CPU[CPU] -- "logical address  
346" --> MMU[MMU]
    subgraph MMU_Box [MMU]
        direction TB
        RR[relocation register  
14000]
        Add((+))
    end
    MMU_Box -- "physical address  
14346" --> Memory[memory]
  
```

Operating System Concepts – 8th Edition 8.10 Silberschatz, Galvin and Gagne ©2009

Dynamic Loading

- Routine is not loaded until it is called
- Better memory-space utilization; unused routine is never loaded
- Useful when large amounts of code are needed to handle infrequently occurring cases
- No special support from the operating system is required implemented through program design

Operating System Concepts – 8th Edition 8.11 Silberschatz, Galvin and Gagne ©2009


Dynamic Linking

- Linking postponed until execution time
- Small piece of code, *stub*, used to locate the appropriate memory-resident library routine
- Stub replaces itself with the address of the routine, and executes the routine
- Operating system needed to check if routine is in processes' memory address
- Dynamic linking is particularly useful for libraries
- System also known as **shared libraries**

Operating System Concepts – 8th Edition 8.12 Silberschatz, Galvin and Gagne ©2009

Swapping

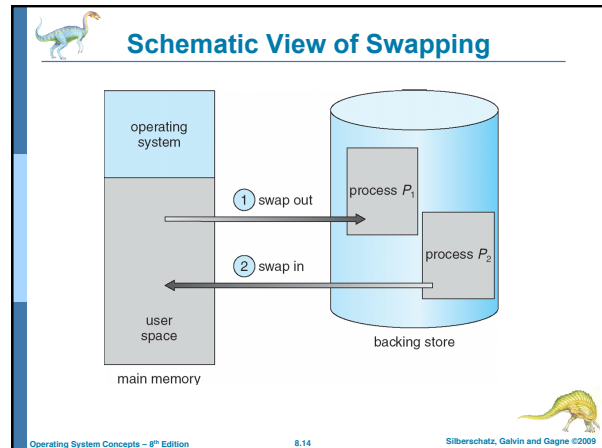
- A process can be swapped temporarily out of memory to a backing store, and then brought back into memory for continued execution
- **Backing store** – fast disk large enough to accommodate copies of all memory images for all users; must provide direct access to these memory images
- **Roll out, roll in** – swapping variant used for priority-based scheduling algorithms; lower-priority process is swapped out so higher-priority process can be loaded and executed
- Major part of swap time is transfer time; total transfer time is directly proportional to the amount of memory swapped
- Modified versions of swapping are found on many systems (i.e., UNIX, Linux, and Windows)
- System maintains a **ready queue** of ready-to-run processes which have memory images on disk



Operating System Concepts – 8th Edition

8.13

Silberschatz, Galvin and Gagne ©2009




Operating System Concepts – 8th Edition

8.14

Silberschatz, Galvin and Gagne ©2009

Contiguous Allocation

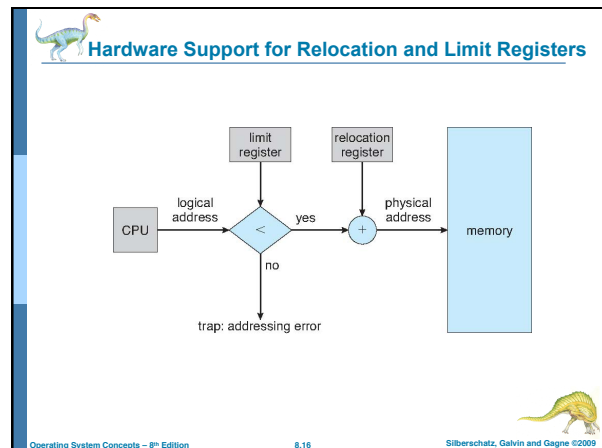
- Main memory usually into two partitions:
 - Resident operating system, usually held in low memory with interrupt vector
 - User processes then held in high memory
- Relocation registers used to protect user processes from each other, and from changing operating-system code and data
 - Base register contains value of smallest physical address
 - Limit register contains range of logical addresses – each logical address must be less than the limit register
 - MMU maps logical address *dynamically*



Operating System Concepts – 8th Edition

8.15


Silberschatz, Galvin and Gagne ©2009



Operating System Concepts – 8th Edition

8.16

Silberschatz, Galvin and Gagne ©2009



Contiguous Allocation (Cont)

- Multiple-partition allocation
 - Hole – block of available memory; holes of various size are scattered throughout memory
 - When a process arrives, it is allocated memory from a hole large enough to accommodate it
 - Operating system maintains information about:
 - a) allocated partitions b) free partitions (hole)

| |
|-----------|
| OS |
| process 5 |
| process 8 |
| process 2 |

→


| |
|-----------|
| OS |
| process 5 |
| |
| process 2 |

→


| |
|-----------|
| OS |
| process 5 |
| process 9 |
| |
| process 2 |

→

| |
|------------|
| OS |
| process 5 |
| process 9 |
| process 10 |
| process 2 |



Operating System Concepts – 8th Edition 8.17 Silberschatz, Galvin and Gagne ©2009




Dynamic Storage-Allocation Problem


How to satisfy a request of size n from a list of free holes

- **First-fit:** Allocate the *first* hole that is big enough
- **Best-fit:** Allocate the *smallest* hole that is big enough; must search entire list, unless ordered by size
 - Produces the smallest leftover hole
- **Worst-fit:** Allocate the *largest* hole; must also search entire list
 - Produces the largest leftover hole

First-fit and best-fit better than worst-fit in terms of speed and storage utilization




Operating System Concepts – 8th Edition 8.18 Silberschatz, Galvin and Gagne ©2009




Fragmentation

- **External Fragmentation** – total memory space exists to satisfy a request, but it is not contiguous
- **Internal Fragmentation** – allocated memory may be slightly larger than requested memory; this size difference is memory internal to a partition, but not being used
- Reduce external fragmentation by **compaction**
 - Shuffle memory contents to place all free memory together in one large block
 - Compaction is possible *only* if relocation is dynamic, and is done at execution time
 - I/O problem
 - Latch job in memory while it is involved in I/O
 - Do I/O only into OS buffers




Operating System Concepts – 8th Edition 8.19 Silberschatz, Galvin and Gagne ©2009



Paging

- Logical address space of a process can be noncontiguous; process is allocated physical memory whenever the latter is available
- Divide physical memory into fixed-sized blocks called **frames** (size is power of 2, between 512 bytes and 8,192 bytes)
- Divide logical memory into blocks of same size called **pages**
- Keep track of all free frames
- To run a program of size n pages, need to find n free frames and load program
- Set up a page table to translate logical to physical addresses
- Internal fragmentation



Operating System Concepts – 8th Edition 8.20 Silberschatz, Galvin and Gagne ©2009

Address Translation Scheme

- Address generated by CPU is divided into:
 - Page number (p)** – used as an index into a **page table** which contains base address of each page in physical memory
 - Page offset (d)** – combined with base address to define the physical memory address that is sent to the memory unit

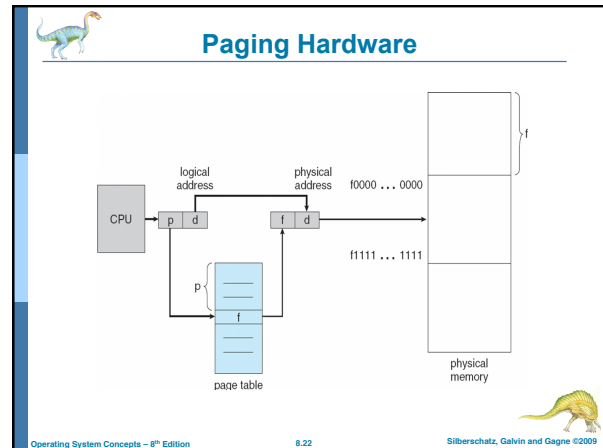
| page number | page offset |
|-------------|-------------|
| p | d |
| $m - n$ | n |

- For given logical address space 2^m and page size 2^n

Operating System Concepts – 8th Edition

8.21

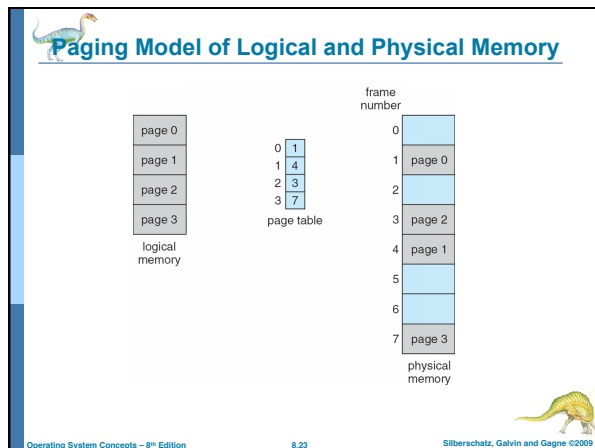
Silberschatz, Galvin and Gagne ©2009



Operating System Concepts – 8th Edition

8.22

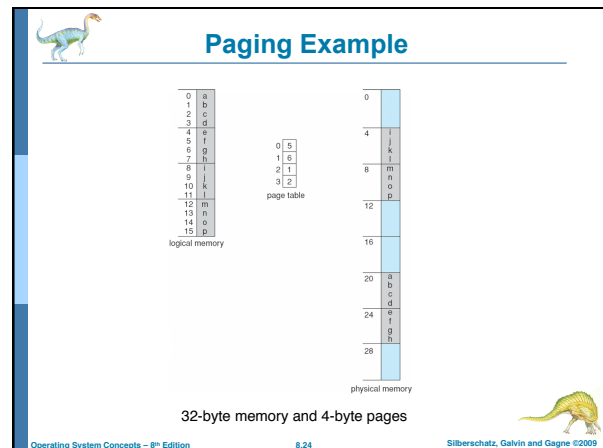
Silberschatz, Galvin and Gagne ©2009



Operating System Concepts – 8th Edition

8.23

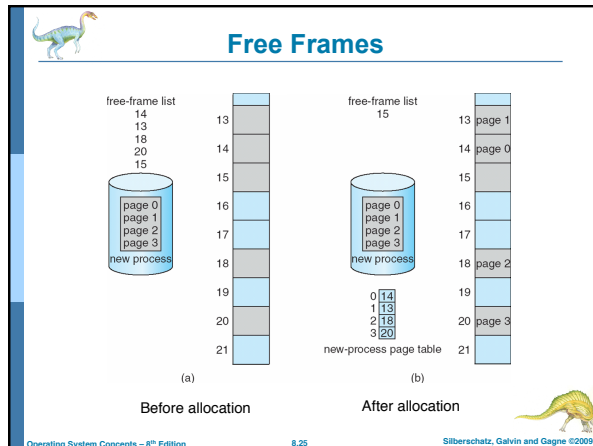
Silberschatz, Galvin and Gagne ©2009



Operating System Concepts – 8th Edition

8.24

Silberschatz, Galvin and Gagne ©2009



Implementation of Page Table

- Page table is kept in main memory
- Page-table base register (PTBR)** points to the page table
- Page-table length register (PRLR)** indicates size of the page table
- In this scheme every data/instruction access requires two memory accesses. One for the page table and one for the data/instruction.
- The two memory access problem can be solved by the use of a special fast-lookup hardware cache called **associative memory** or **translation look-aside buffers (TLBs)**
- Some TLBs store **address-space identifiers (ASIDs)** in each TLB entry – uniquely identifies each process to provide address-space protection for that process

Operating System Concepts – 8th Edition 8.26 Silberschatz, Galvin and Gagne ©2009

Associative Memory

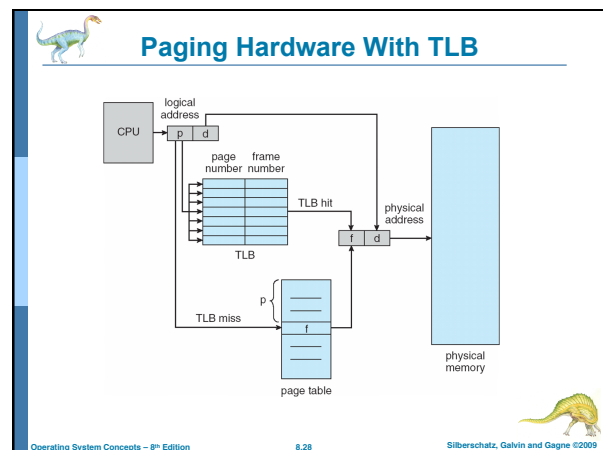
- Associative memory – parallel search


| Page # | Frame # |
|--------|---------|
| | |
| | |
| | |
| | |

Address translation (p, d)

- If p is in associative register, get frame # out
- Otherwise get frame # from page table in memory

Operating System Concepts – 8th Edition 8.27 Silberschatz, Galvin and Gagne ©2009






Effective Access Time


- Associative Lookup = ϵ time unit
- Assume memory cycle time is 1 microsecond
- Hit ratio – percentage of times that a page number is found in the associative registers; ratio related to number of associative registers
- Hit ratio = α
- **Effective Access Time (EAT)**

$$\text{EAT} = (1 + \epsilon) \alpha + (2 + \epsilon)(1 - \alpha)$$

$$= 2 + \epsilon - \alpha$$




Operating System Concepts – 8th Edition 8.29 Silberschatz, Galvin and Gagne ©2009




Memory Protection

- Memory protection implemented by associating protection bit with each frame
- **Valid-invalid** bit attached to each entry in the page table:
 - "valid" indicates that the associated page is in the process' logical address space, and is thus a legal page
 - "invalid" indicates that the page is not in the process' logical address space



Operating System Concepts – 8th Edition 8.30 Silberschatz, Galvin and Gagne ©2009



Valid (v) or Invalid (i) Bit In A Page Table

00000

| |
|--------|
| page 0 |
| page 1 |
| page 2 |
| page 3 |
| page 4 |
| page 5 |


10,468
12,287

frame number


| | | |
|---|---|---|
| 0 | 2 | V |
| 1 | 3 | V |
| 2 | 4 | V |
| 3 | 7 | V |
| 4 | 8 | V |
| 5 | 9 | V |
| 6 | 0 | I |
| 7 | 0 | I |

page table

| |
|--------|
| 0 |
| 1 |
| 2 |
| 3 |
| 4 |
| 5 |
| 6 |
| 7 |
| 8 |
| 9 |
| ⋮ |
| page n |




Operating System Concepts – 8th Edition 8.31 Silberschatz, Galvin and Gagne ©2009

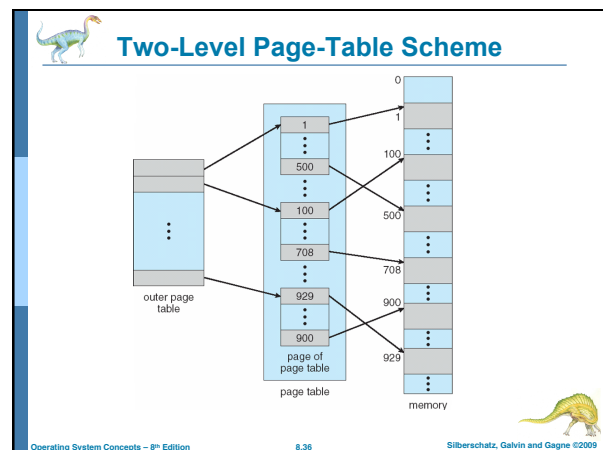
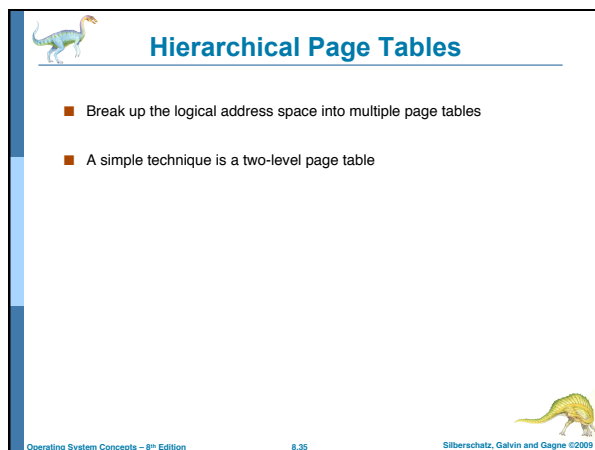
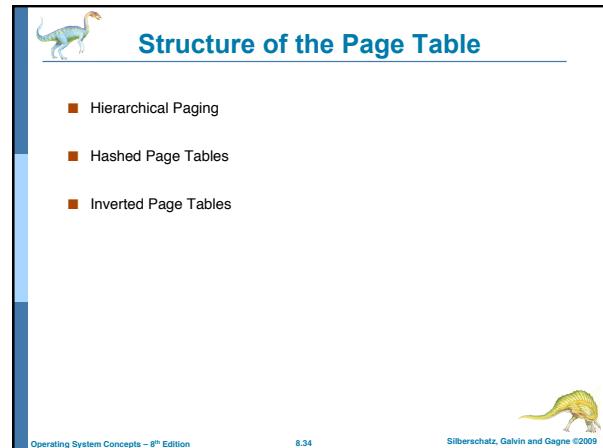
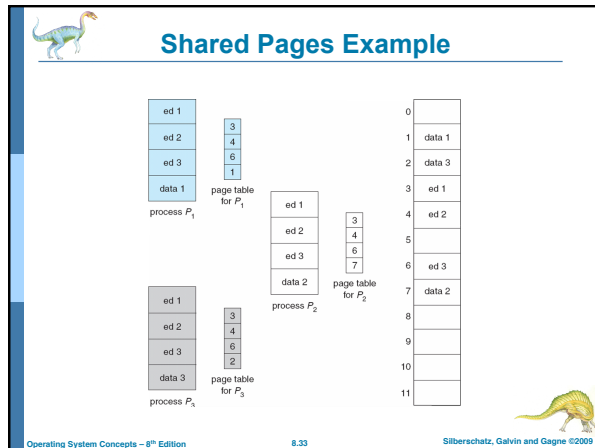


Shared Pages

- **Shared code**
 - One copy of read-only (reentrant) code shared among processes (i.e., text editors, compilers, window systems).
 - Shared code must appear in same location in the logical address space of all processes
- **Private code and data**
 - Each process keeps a separate copy of the code and data
 - The pages for the private code and data can appear anywhere in the logical address space



Operating System Concepts – 8th Edition 8.32 Silberschatz, Galvin and Gagne ©2009



Two-Level Paging Example

- A logical address (on 32-bit machine with 1K page size) is divided into:
 - a page number consisting of 22 bits
 - a page offset consisting of 10 bits
- Since the page table is paged, the page number is further divided into:
 - a 12-bit page number
 - a 10-bit page offset
- Thus, a logical address is as follows:

| page number | | page offset |
|-------------|-------|-------------|
| p_1 | p_2 | d |
| 12 | 10 | 10 |

where p_1 is an index into the outer page table, and p_2 is the displacement within the page of the outer page table

Operating System Concepts – 8th Edition 8.37 Silberschatz, Galvin and Gagne ©2009

Address-Translation Scheme

Operating System Concepts – 8th Edition 8.38 Silberschatz, Galvin and Gagne ©2009

Three-level Paging Scheme

| outer page | inner page | offset |
|------------|------------|--------|
| p_1 | p_2 | d |
| 42 | 10 | 12 |

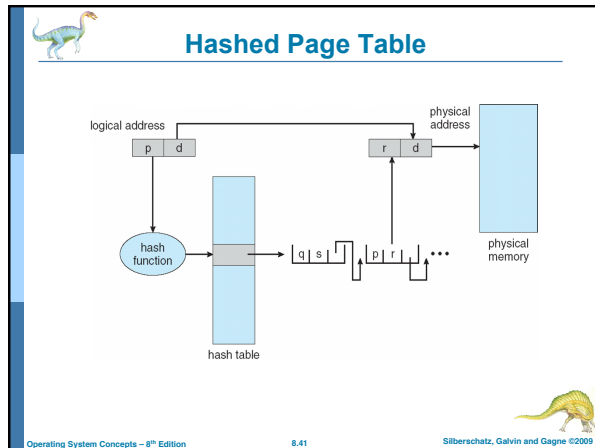
| 2nd outer page | outer page | inner page | offset |
|----------------|------------|------------|--------|
| p_1 | p_2 | p_3 | d |
| 32 | 10 | 10 | 12 |

Operating System Concepts – 8th Edition 8.39 Silberschatz, Galvin and Gagne ©2009

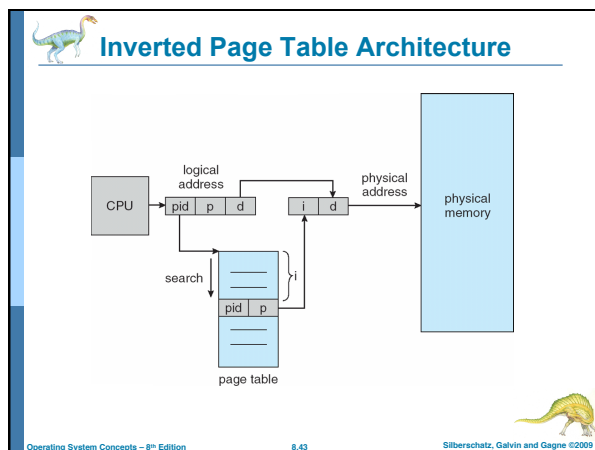
Hashed Page Tables

- Common in address spaces > 32 bits
- The virtual page number is hashed into a page table
 - This page table contains a chain of elements hashing to the same location
- Virtual page numbers are compared in this chain searching for a match
 - If a match is found, the corresponding physical frame is extracted

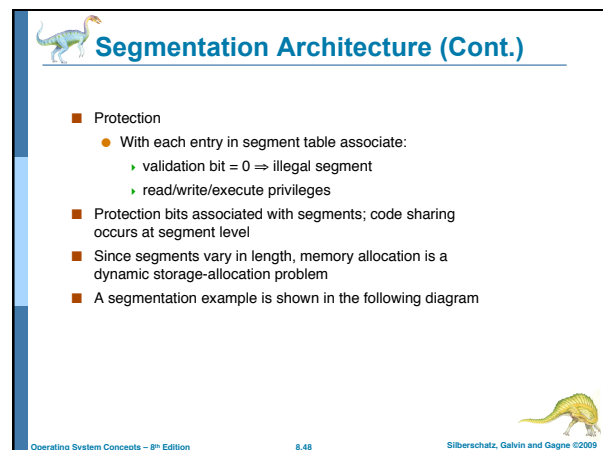
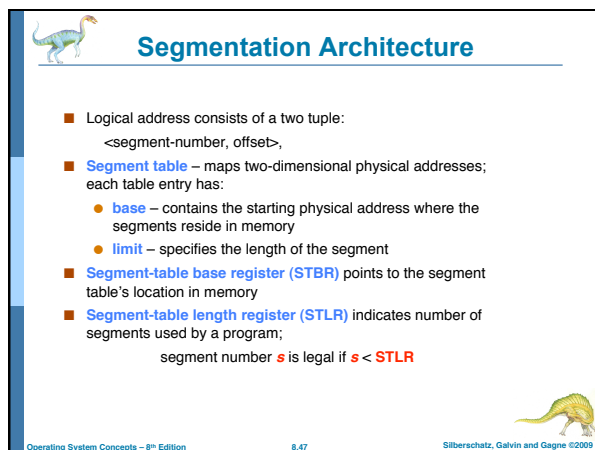
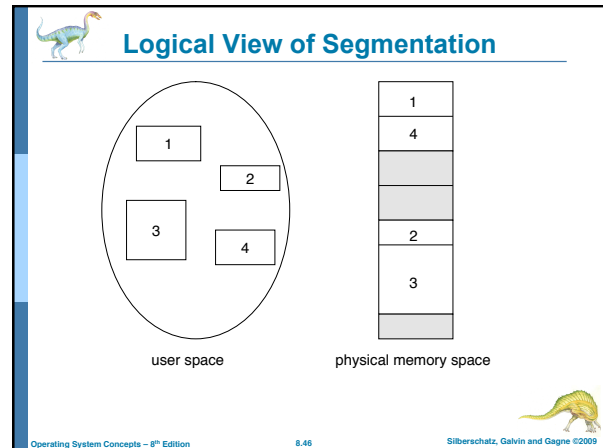
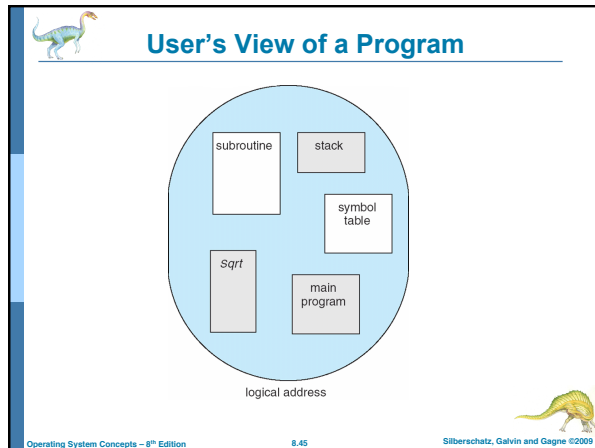
Operating System Concepts – 8th Edition 8.40 Silberschatz, Galvin and Gagne ©2009

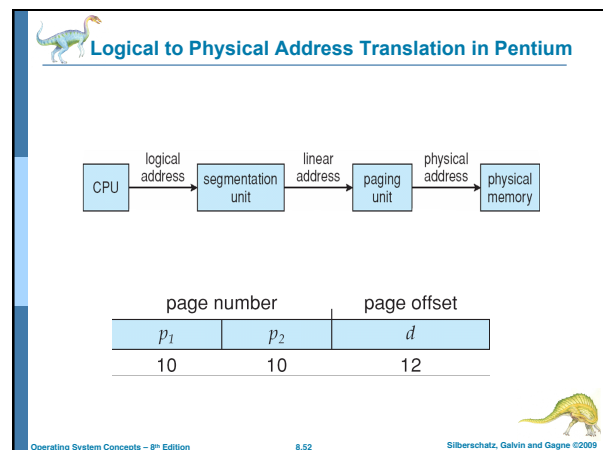
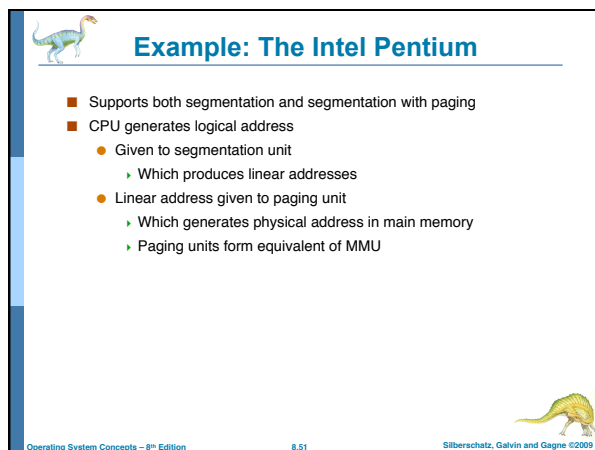
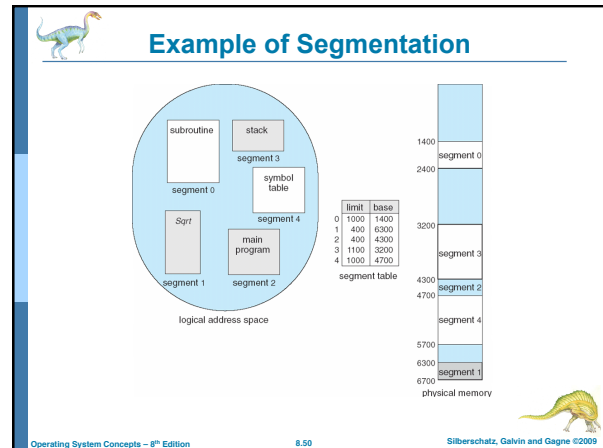
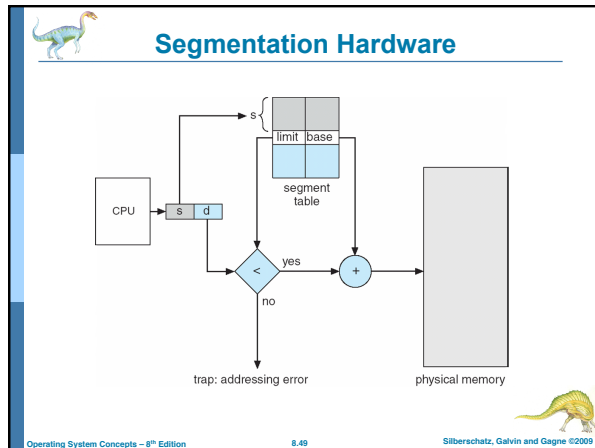


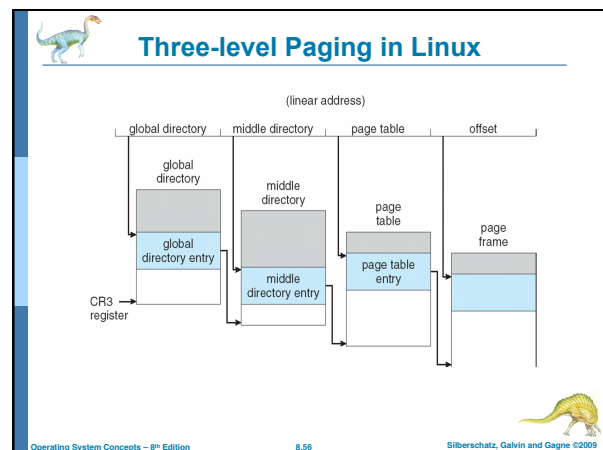
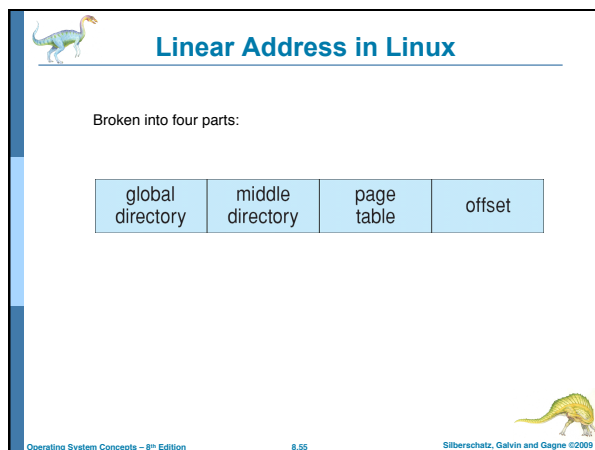
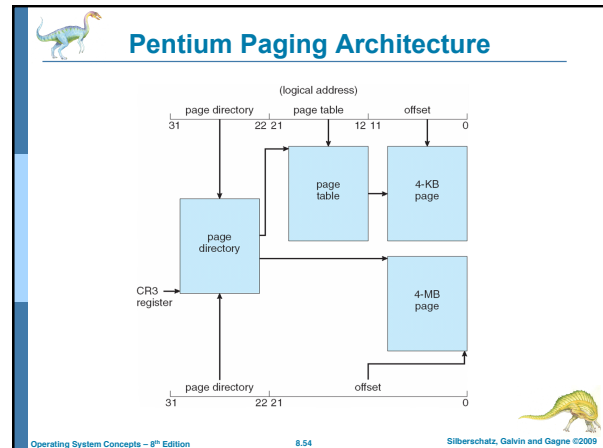
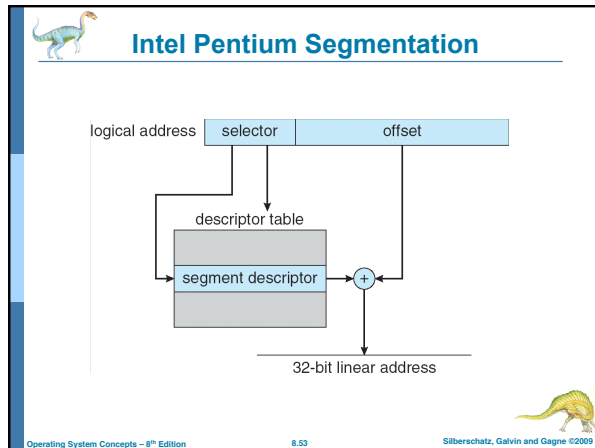
- ### Inverted Page Table
- One entry for each real page of memory
 - Entry consists of the virtual address of the page stored in that real memory location, with information about the process that owns that page
 - Decreases memory needed to store each page table, but increases time needed to search the table when a page reference occurs
 - Use hash table to limit the search to one — or at most a few — page-table entries
- Operating System Concepts – 8th Edition 8.42 Silberschatz, Galvin and Gagne ©2009



- ### Segmentation
- Memory-management scheme that supports user view of memory
 - A program is a collection of segments
 - A segment is a logical unit such as:
 - main program
 - procedure
 - function
 - method
 - object
 - local variables, global variables
 - common block
 - stack
 - symbol table
 - arrays
- Operating System Concepts – 8th Edition 8.44 Silberschatz, Galvin and Gagne ©2009







End of Chapter 8

