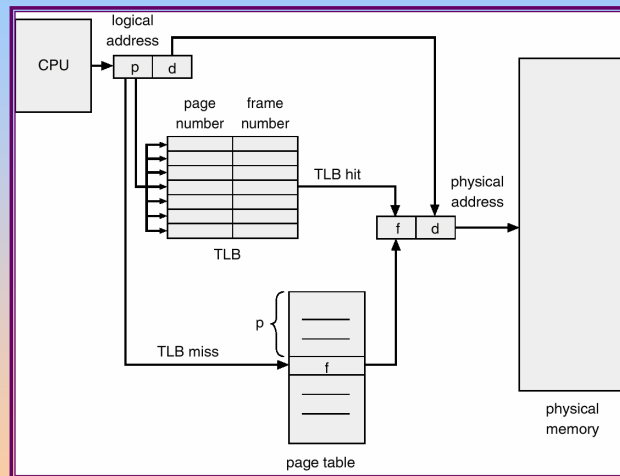


Paging Hardware With TLB



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.

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

00000	page 0
	page 1
	page 2
	page 3
	page 4
10,468	page 5
12,287	

0	2	v
1	3	v
2	4	v
3	7	v
4	8	v
5	9	v
6	0	i
7	0	i

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

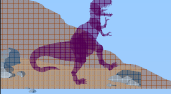
page table

Operating System Concepts
9.3
Silberschatz, Galvin and Gagne ©2002

Page Table Structure


- Hierarchical Paging
- Hashed Page Tables
- Inverted Page Tables

Operating System Concepts
9.4
Silberschatz, Galvin and Gagne ©2002

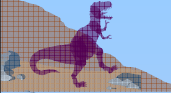


Hierarchical Page Tables

- Break up the logical address space into multiple page tables.
- A simple technique is a two-level page table.



Operating System Concepts 9.5 Silberschatz, Galvin and Gagne ©2002




Two-Level Paging Example

- A logical address (on 32-bit machine with 4K page size) is divided into:
 - a page number consisting of 20 bits.
 - a page offset consisting of 12 bits.
- Since the page table is paged, the page number is further divided into:
 - a 10-bit page number.
 - a 10-bit page offset.
- Thus, a logical address is as follows:

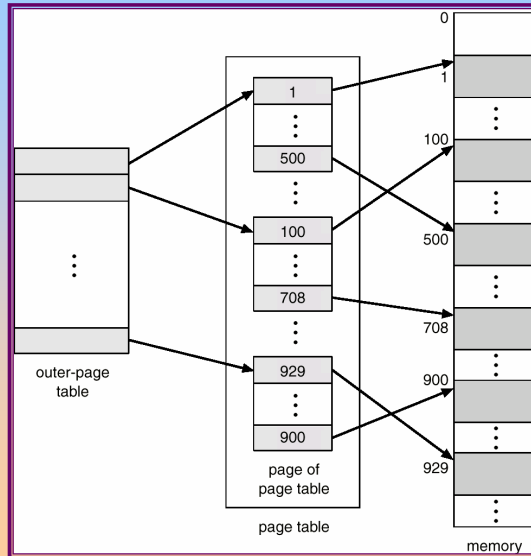
page number		page offset
p_1	p_2	d
10	10	12

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 9.6 Silberschatz, Galvin and Gagne ©2002

Two-Level Page-Table Scheme



Address-Translation Scheme

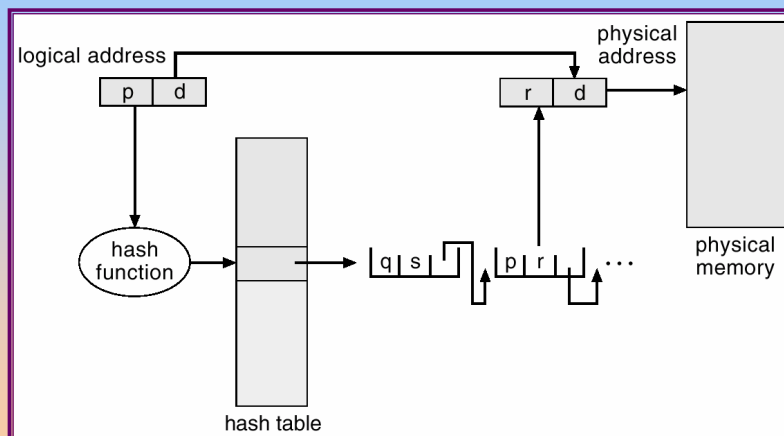
- Address-translation scheme for a two-level 32-bit paging architecture



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.

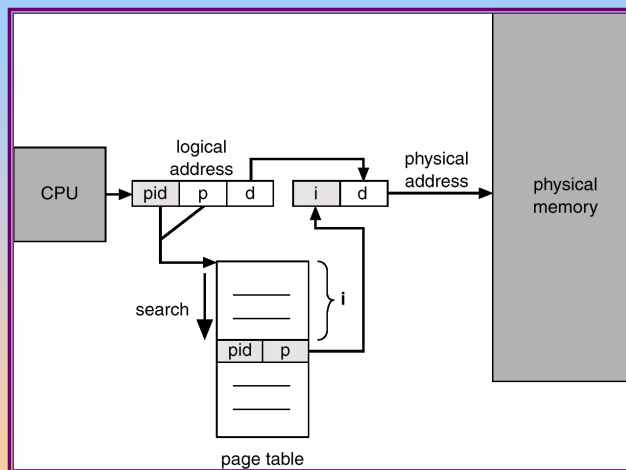
Hashed Page Table



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.

Inverted Page Table Architecture

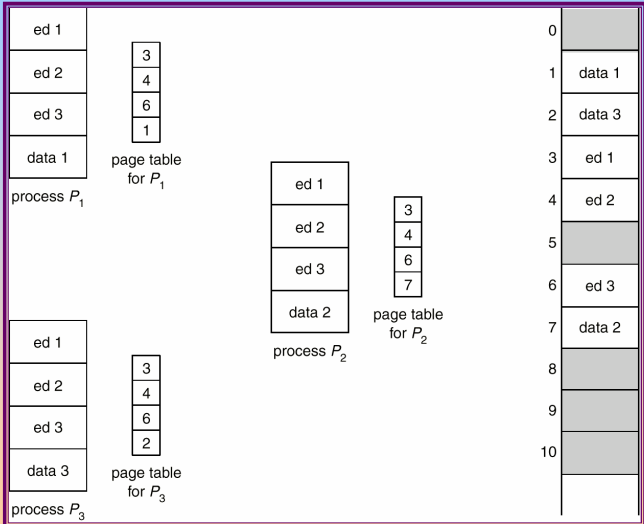


Shared Pages

- 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.

Shared Pages Example




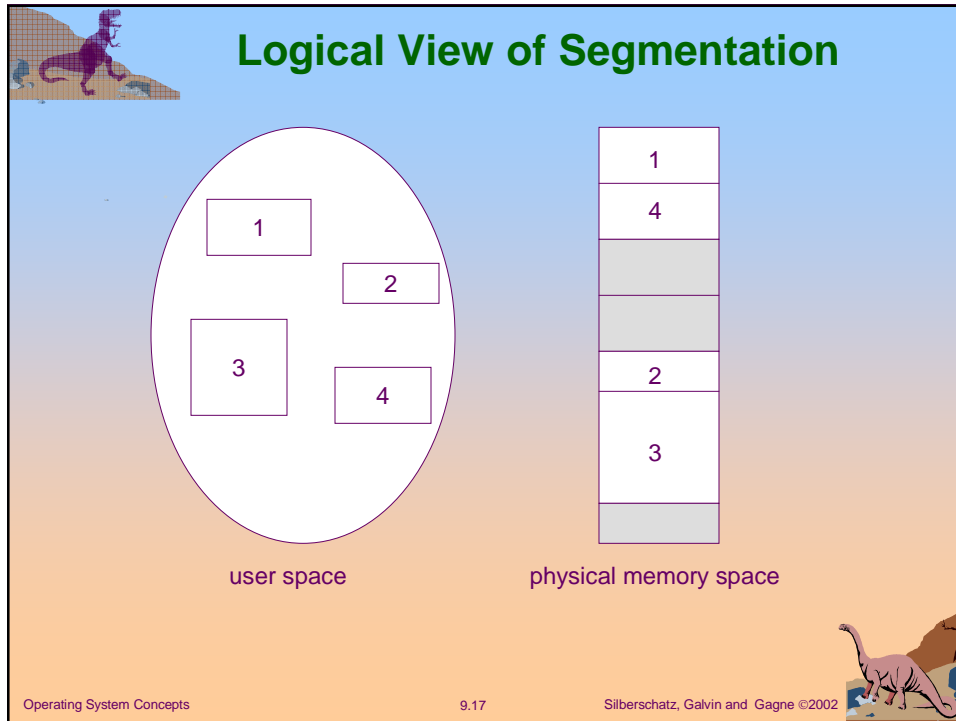
Segmentation

- Memory-management scheme that supports user view of memory.
- A program is a collection of segments. A segment is a logical unit such as:

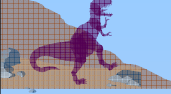
main program,
procedure,
function,
method,
object,
local variables, global variables,
common block,
stack,
symbol table, arrays

User's View of a Program






- ## Segmentation Architecture
- Logical address consists of a two tuple:
 $\langle \text{segment-number, offset} \rangle$,
 - *Segment table* – maps two-dimensional physical addresses; each table entry has:
 - ☞ *base* – contains the starting physical address where the segments reside in memory.
 - ☞ *limit* – specifies the length of the segment.
 - *Segment-table base register (STBR)* points to the segment table's location in memory.
 - *Segment-table length register (STLR)* indicates number of segments used by a program;
segment number s is legal if $s < \text{STLR}$.
- Operating System Concepts 9.18 Silberschatz, Galvin and Gagne ©2002

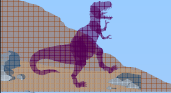


Segmentation Architecture (Cont.)

- Relocation.
 - ☞ dynamic
 - ☞ by segment table
- Sharing.
 - ☞ shared segments
 - ☞ same segment number
- Allocation.
 - ☞ first fit/best fit
 - ☞ external fragmentation




Operating System Concepts 9.19 Silberschatz, Galvin and Gagne ©2002



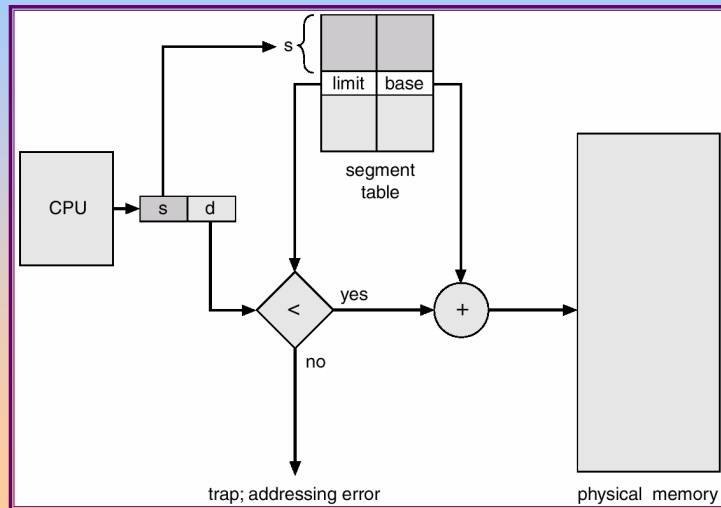
Segmentation Architecture (Cont.)

- Protection. With each entry in segment table associate:
 - ☞ validation bit = 0 \Rightarrow illegal segment
 - ☞ read/write/execute privileges
- Protection bits associated with segments; code sharing occurs at segment level.
- Since segments vary in length, memory allocation is a dynamic storage-allocation problem.
- A segmentation example is shown in the following diagram



Operating System Concepts 9.20 Silberschatz, Galvin and Gagne ©2002

Segmentation Hardware




Chapter 10: Virtual Memory

- Background
- Demand Paging
- Process Creation
- Page Replacement
- Allocation of Frames
- Thrashing
- Operating System Examples

Background


- **Virtual memory** – separation of user logical memory from physical memory.
 - ☞ Only part of the program needs to be in memory for execution.
 - ☞ Logical address space can therefore be much larger than physical address space.
 - ☞ Allows address spaces to be shared by several processes.
 - ☞ Allows for more efficient process creation.

- Virtual memory can be implemented via:
 - ☞ Demand paging
 - ☞ Demand segmentation



Operating System Concepts 10.23 Silberschatz, Galvin and Gagne ©2002

Virtual Memory That is Larger Than Physical Memory



Operating System Concepts 10.24 Silberschatz, Galvin and Gagne ©2002

Demand Paging

- Bring a page into memory only when it is needed.
 - ☞ Less I/O needed
 - ☞ Less memory needed
 - ☞ Faster response
 - ☞ More users

- Page is needed ⇒ reference to it
 - ☞ invalid reference ⇒ abort
 - ☞ not-in-memory ⇒ bring to memory

Operating System Concepts 10.25 Silberschatz, Galvin and Gagne ©2002

Transfer of a Paged Memory to Contiguous Disk Space

The diagram illustrates the transfer of paged memory to contiguous disk space. On the left, a vertical stack of boxes represents 'main memory'. The top four boxes are labeled 'program A' and contain pages 4, 5, 6, and 7. The next four boxes are labeled 'program B' and contain pages 16, 17, 18, and 19. On the right, a cylindrical disk represents contiguous disk space, divided into 24 pages numbered 0 through 23. Arrows labeled 'swap out' point from the pages of program A in main memory to the corresponding pages (4, 5, 6, 7) on the disk. An arrow labeled 'swap in' points from the pages of program B on the disk back to the corresponding pages (16, 17, 18, 19) in main memory.

Operating System Concepts 10.26 Silberschatz, Galvin and Gagne ©2002

Valid-Invalid Bit

- With each page table entry a valid–invalid bit is associated
(1 ⇒ in-memory, 0 ⇒ not-in-memory)
- Initially valid–invalid but is set to 0 on all entries.
- Example of a page table snapshot.

Frame #	valid-invalid bit
	1
	1
	1
	1
	0
⋮	
	0
	0

page table

- During address translation, if valid–invalid bit in page table entry is 0 ⇒ page fault.

Operating System Concepts
10.27
Silberschatz, Galvin and Gagne ©2002

Page Table When Some Pages Are Not in Main Memory

0	A
1	B
2	C
3	D
4	E
5	F
6	G
7	H

logical memory

frame	valid-invalid bit
0	4 v
1	i
2	6 v
3	i
4	i
5	9 v
6	i
7	i

page table

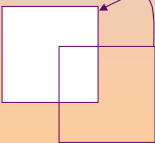
0	
1	
2	
3	
4	A
5	
6	C
7	
8	
9	F
10	
11	
12	
13	
14	
15	

physical memory

Operating System Concepts
10.28
Silberschatz, Galvin and Gagne ©2002

Page Fault

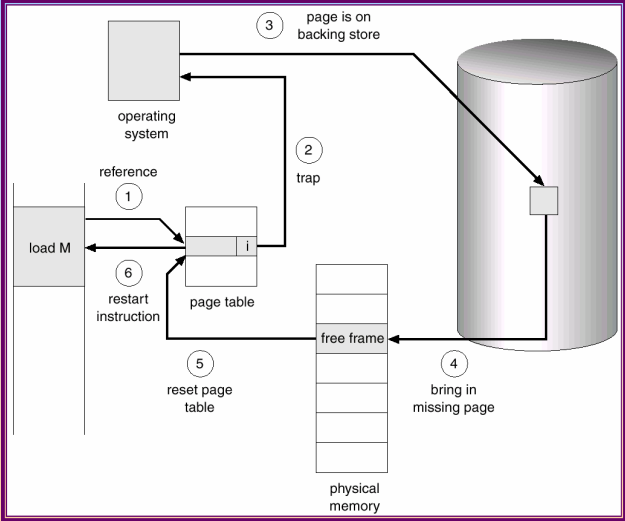
- If there is ever a reference to a page, first reference will trap to OS \Rightarrow page fault
- OS looks at another table to decide:
 - ✦ Invalid reference \Rightarrow abort.
 - ✦ Just not in memory.
- Get empty frame.
- Swap page into frame.
- Reset tables, validation bit = 1.
- Restart instruction: Least Recently Used
 - ✦ block move



✦ auto increment/decrement location

Operating System Concepts 10.29 Silberschatz, Galvin and Gagne ©2002

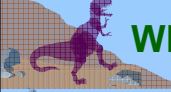
Steps in Handling a Page Fault



The diagram illustrates the following steps:


- reference to page i in load M
- trap to operating system
- page is on backing store
- bring in missing page
- reset page table
- restart instruction

Operating System Concepts 10.30 Silberschatz, Galvin and Gagne ©2002

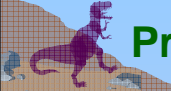


What happens if there is no free frame?

- Page replacement – find some page in memory, but not really in use, swap it out.
 - ↻ algorithm
 - ↻ performance – want an algorithm which will result in minimum number of page faults.
- Same page may be brought into memory several times.




Operating System Concepts 10.31 Silberschatz, Galvin and Gagne ©2002

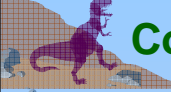


Process Creation

- Virtual memory allows other benefits during process creation:
 - Copy-on-Write
 - Memory-Mapped Files




Operating System Concepts 10.32 Silberschatz, Galvin and Gagne ©2002



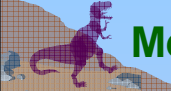
Copy-on-Write

- Copy-on-Write (COW) allows both parent and child processes to initially *share* the same pages in memory.

If either process modifies a shared page, only then is the page copied.
- COW allows more efficient process creation as only modified pages are copied.
- Free pages are allocated from a *pool* of zeroed-out pages.




Operating System Concepts 10.33 Silberschatz, Galvin and Gagne ©2002



Memory-Mapped Files

- Memory-mapped file I/O allows file I/O to be treated as routine memory access by *mapping* a disk block to a page in memory.
- A file is initially read using demand paging. A page-sized portion of the file is read from the file system into a physical page. Subsequent reads/writes to/from the file are treated as ordinary memory accesses.
- Simplifies file access by treating file I/O through memory rather than **read()** **write()** system calls.
- Also allows several processes to map the same file allowing the pages in memory to be shared.



Operating System Concepts 10.34 Silberschatz, Galvin and Gagne ©2002

Memory Mapped Files

