


Chapter 7: Memory Management





Chapter 7: Memory Management

- Background
- Swapping
- Contiguous Memory Allocation
- Paging
- Structure of the Page Table
- Segmentation
- Example: The Intel 32 and 64-bit Architectures
- Example: ARMv8 Architecture





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

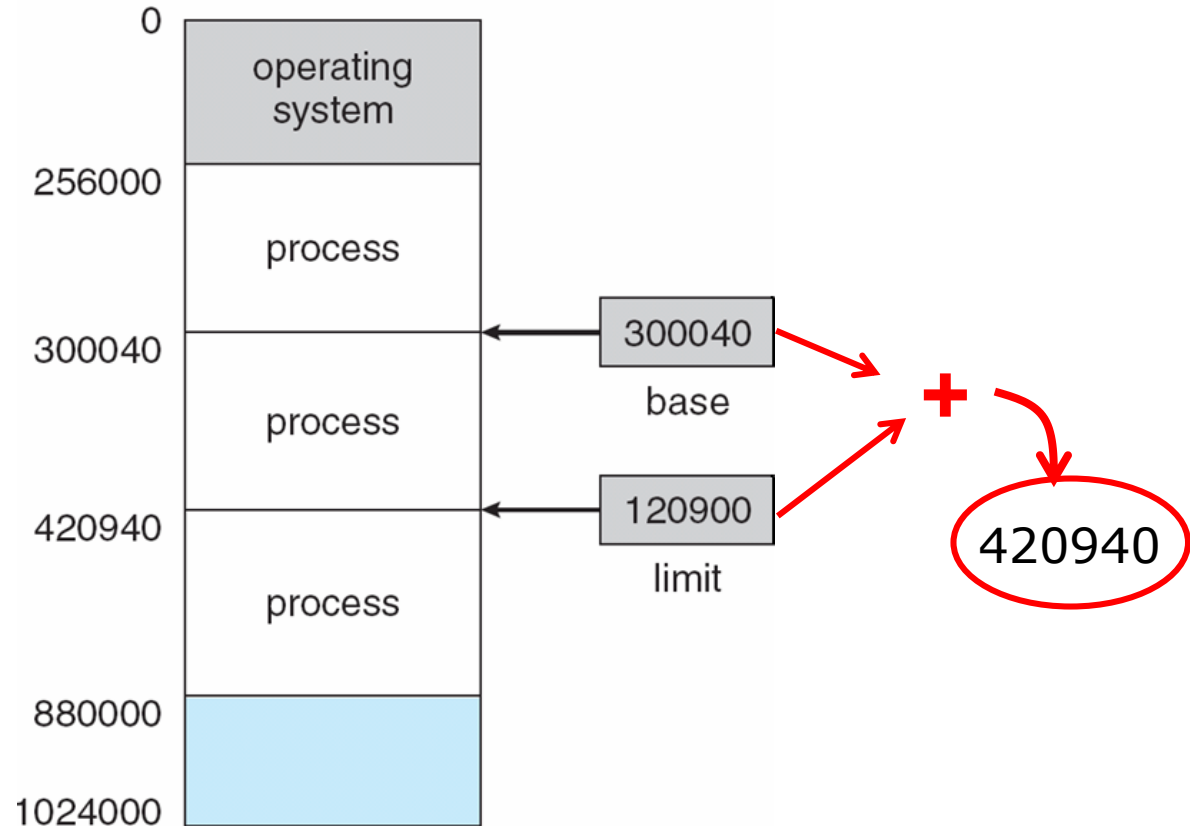
brought : ถูกนำพา





Base and Limit Registers

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





Binding of Instructions and Data to Memory

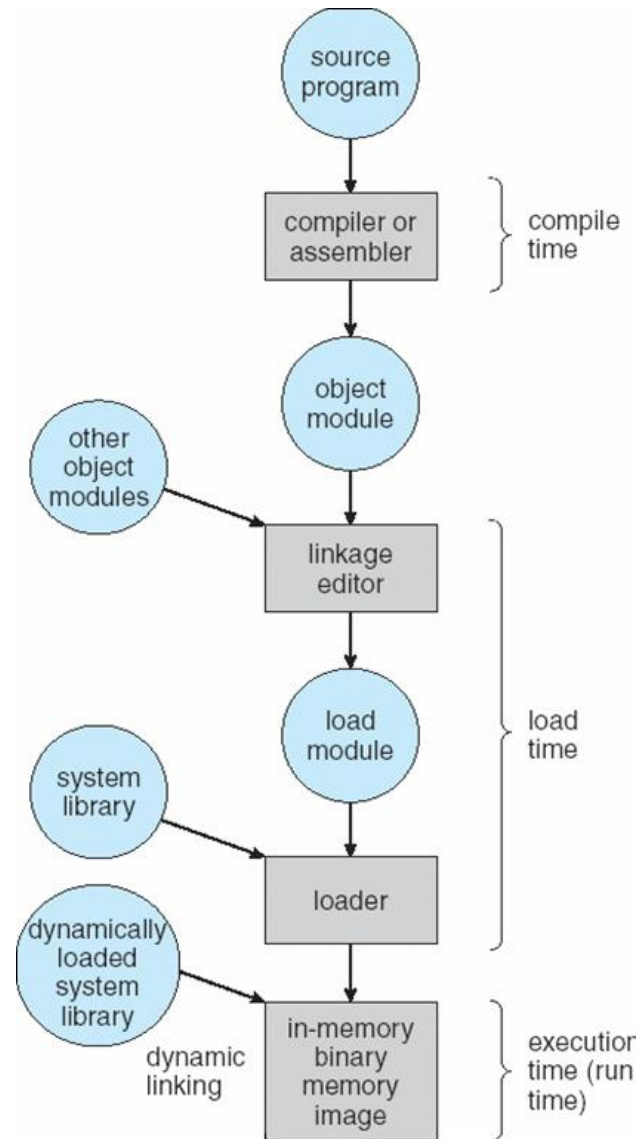
- Address binding of instructions and data to memory addresses can happen at three different stages
 - **Compile time:** If memory location known a priori, **absolute code** can be generated; must recompile code if starting location changes
 - **Load time:** Must generate **relocatable code** if memory 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)

Binding: การกำหนดค่า





Multistep Processing of a User Program





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





Memory-Management Unit (MMU)

Logical Address




- 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**





Dynamic relocation using a relocation register





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

Routine: โปรแกรมย่อย





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**





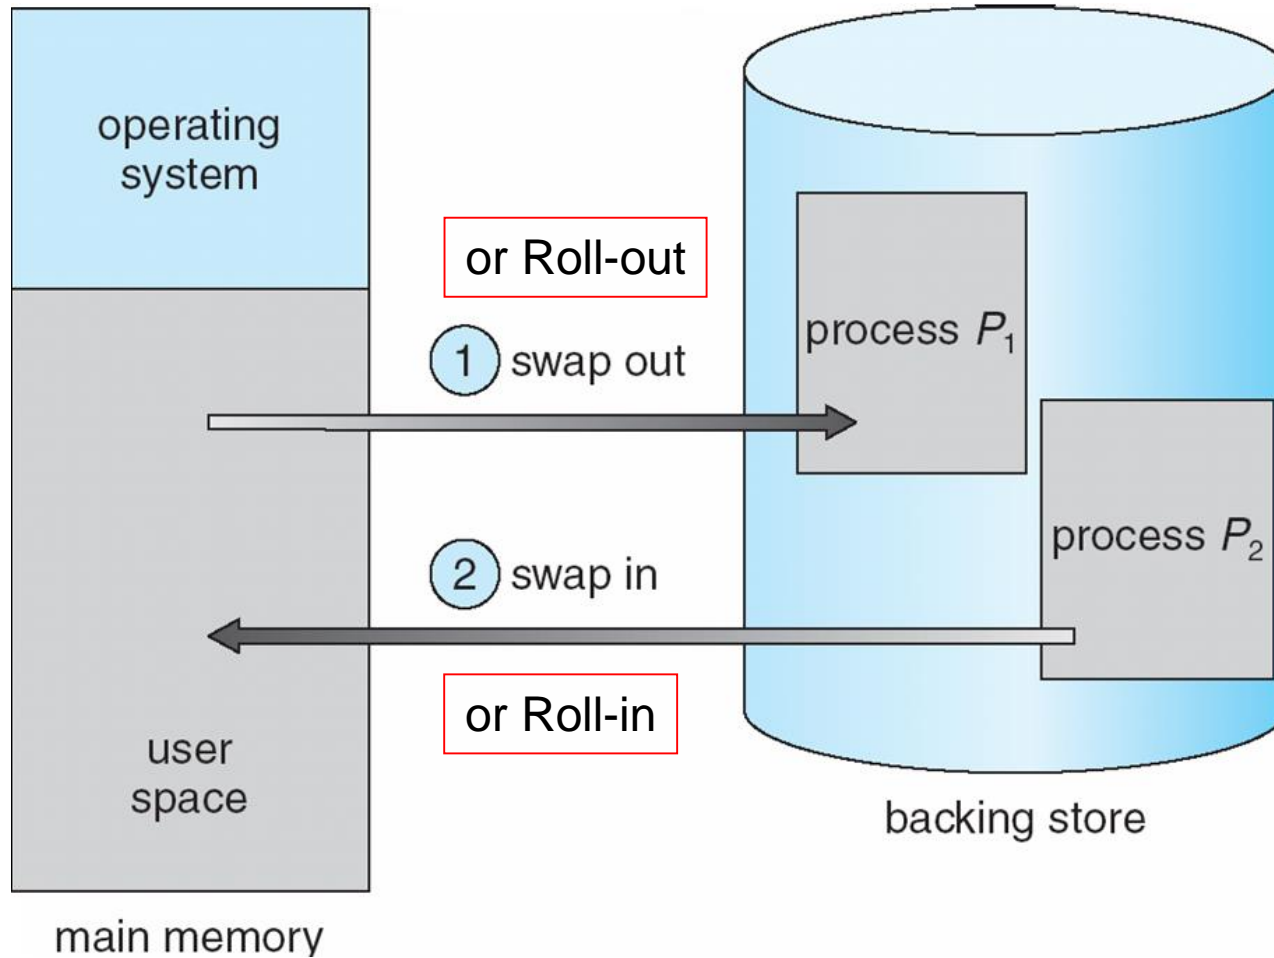
Swapping

- 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
- System maintains a **ready queue** of ready-to-run processes which have memory images on disk
- Modified versions of swapping are found on many systems (i.e., **UNIX, Linux, and Windows**)
 - Swapping normally disabled
 - Started if more than threshold amount of memory allocated
 - Disabled again once memory demand reduced below threshold





Schematic View of Swapping





Contiguous Allocation

(การจัดสรรพื้นที่ที่ติดกัน)

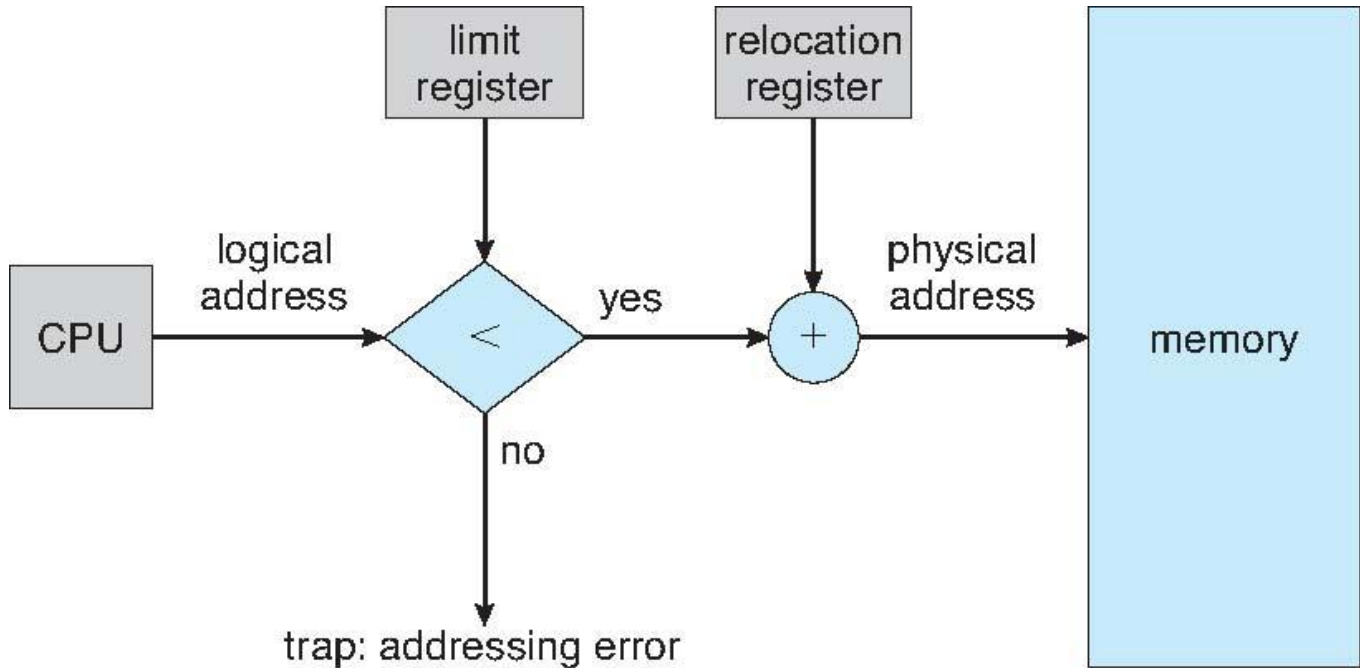
- Main memory usually into **two partitions**:
 - Resident operating system, usually held in low memory with interrupt vector (โปรแกรมของระบบปฏิบัติการเอง)
 - User processes then held in high memory
- **Relocation registers used to protect user processes from each other**, and from changing operating-system 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*

contiguous : ติดกัน





Hardware Support for Relocation and Limit Registers

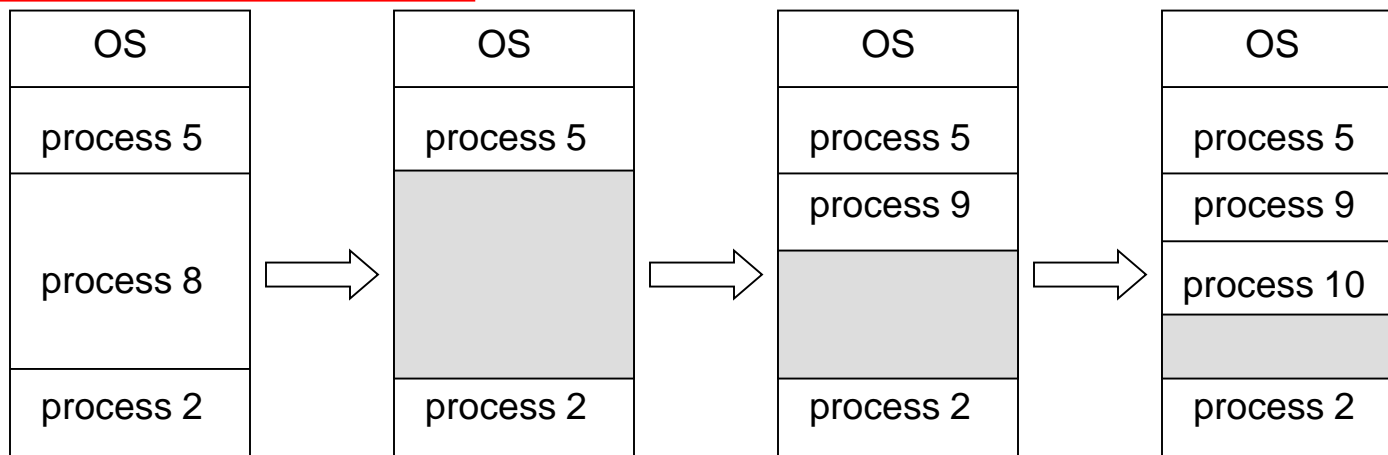




Contiguous Allocation (Cont)

- Multiple-partition allocation (การจัดสรรเนื้อที่แบบหลายส่วน)
 - Degree of multiprogramming limited by number of partitions
 - **Variable-partition** sizes for efficiency (sized to a given process' needs)
 - **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)**

อาจใช้การจัด Schedule แบบ FCFS





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

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





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

compaction: การบีบอัด

satisfy: ปฏิบัติตาม

shuffle: สับเปลี่ยน





Paging


- Physical 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 16 Mbytes)
- 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
- Still have **Internal fragmentation**





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

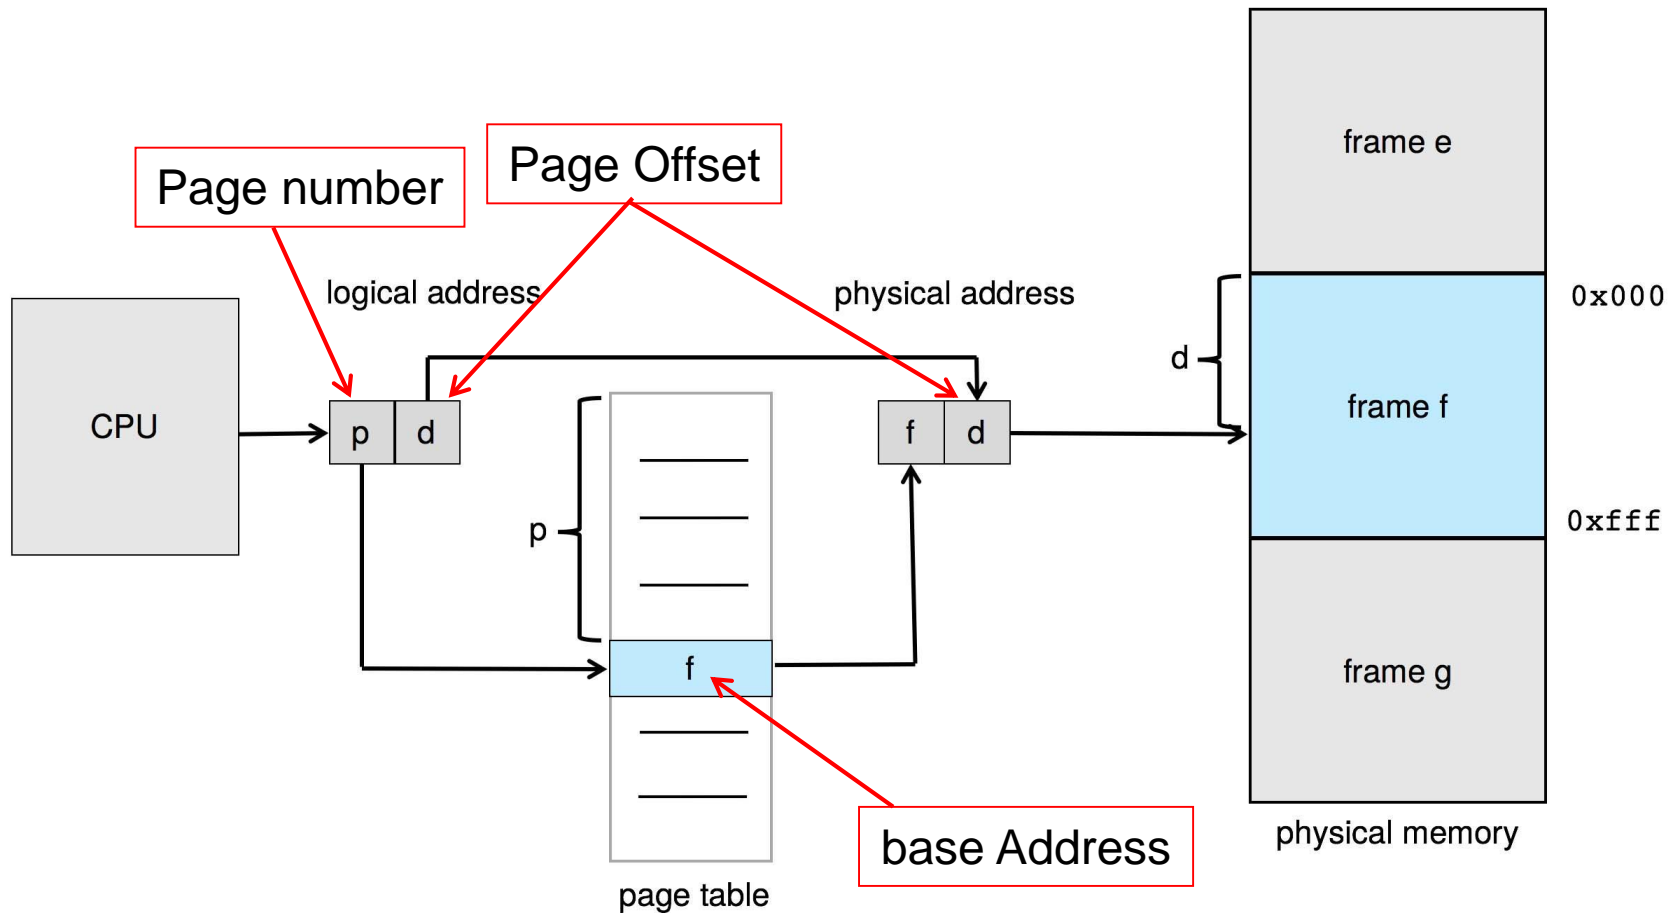


- For given logical address space 2^m and page size 2^n



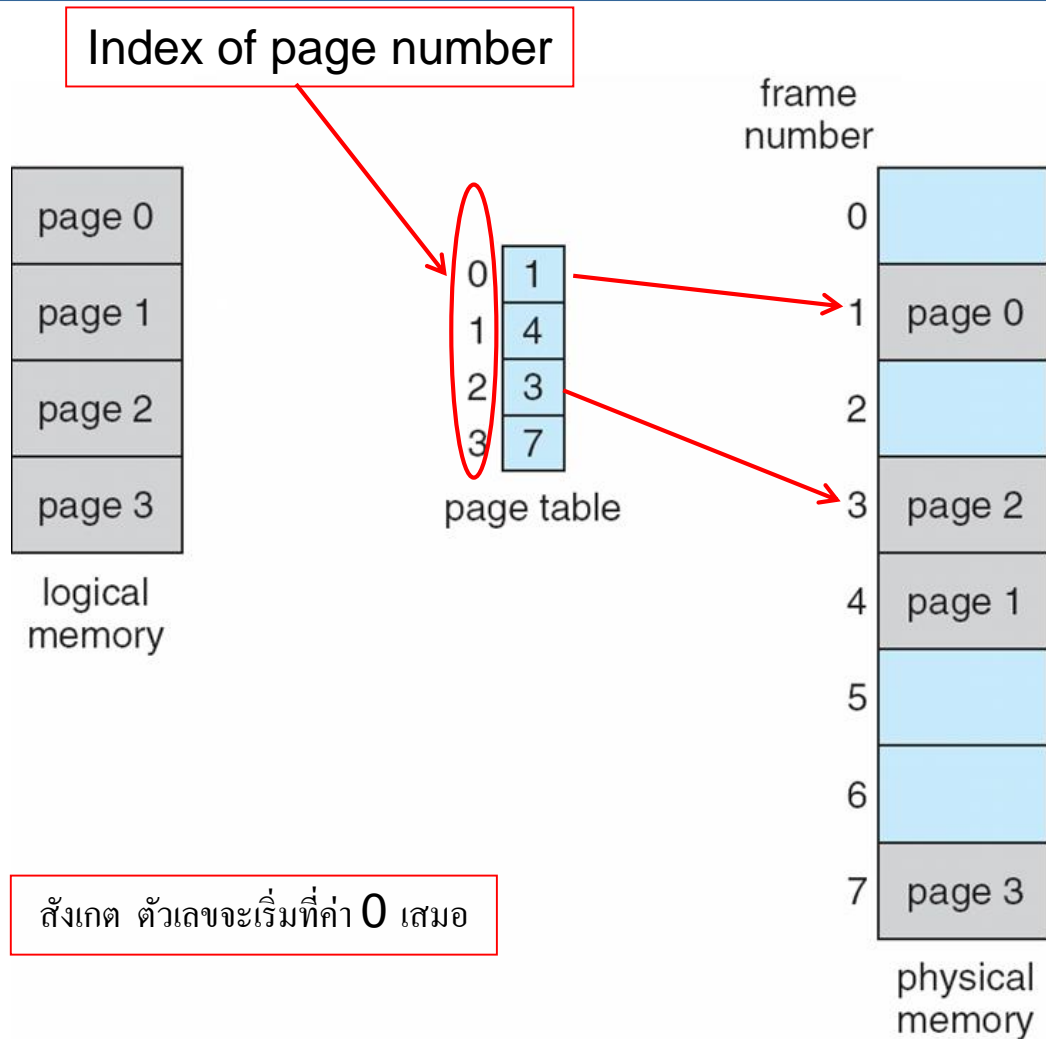


Paging Hardware





Paging Model of Logical and Physical Memory



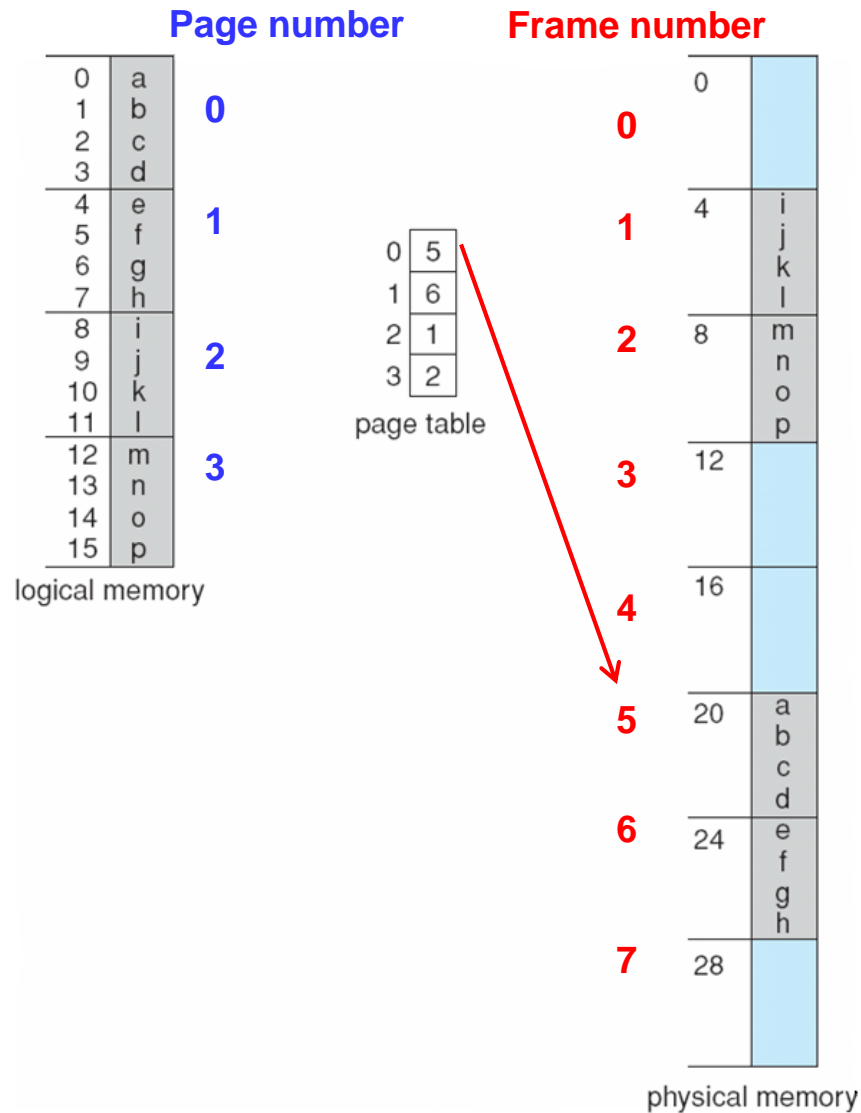
สังเกต ตัวเลขจะเริ่มที่ค่า 0 เสมอ





Paging Example



- Logical address: $n = 2$ and $m = 4$. Using a page size of 4 bytes and a physical memory of 32 bytes (8 pages)





Free Frames


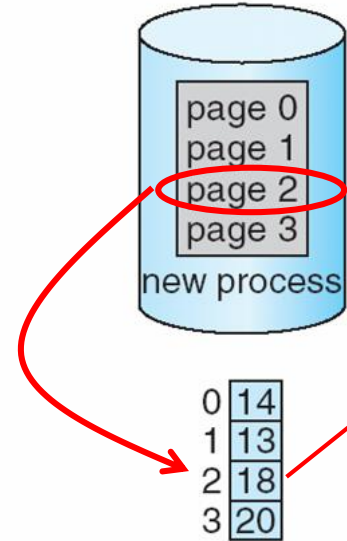
เรียงตามลำดับ
Free-frame
ที่ว่าง



(a)

Before allocation

free-frame list
15



(b)

After allocation







Exercise: Free Frames

free-frame list



15
18
13
20
14




(a)

Before allocation

free-frame list



new-process page table



(b)

After allocation





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
 - TLBs typically small (64 to 1,024 entries)





Associative Memory

- Associative memory – parallel search

Page #	Frame #

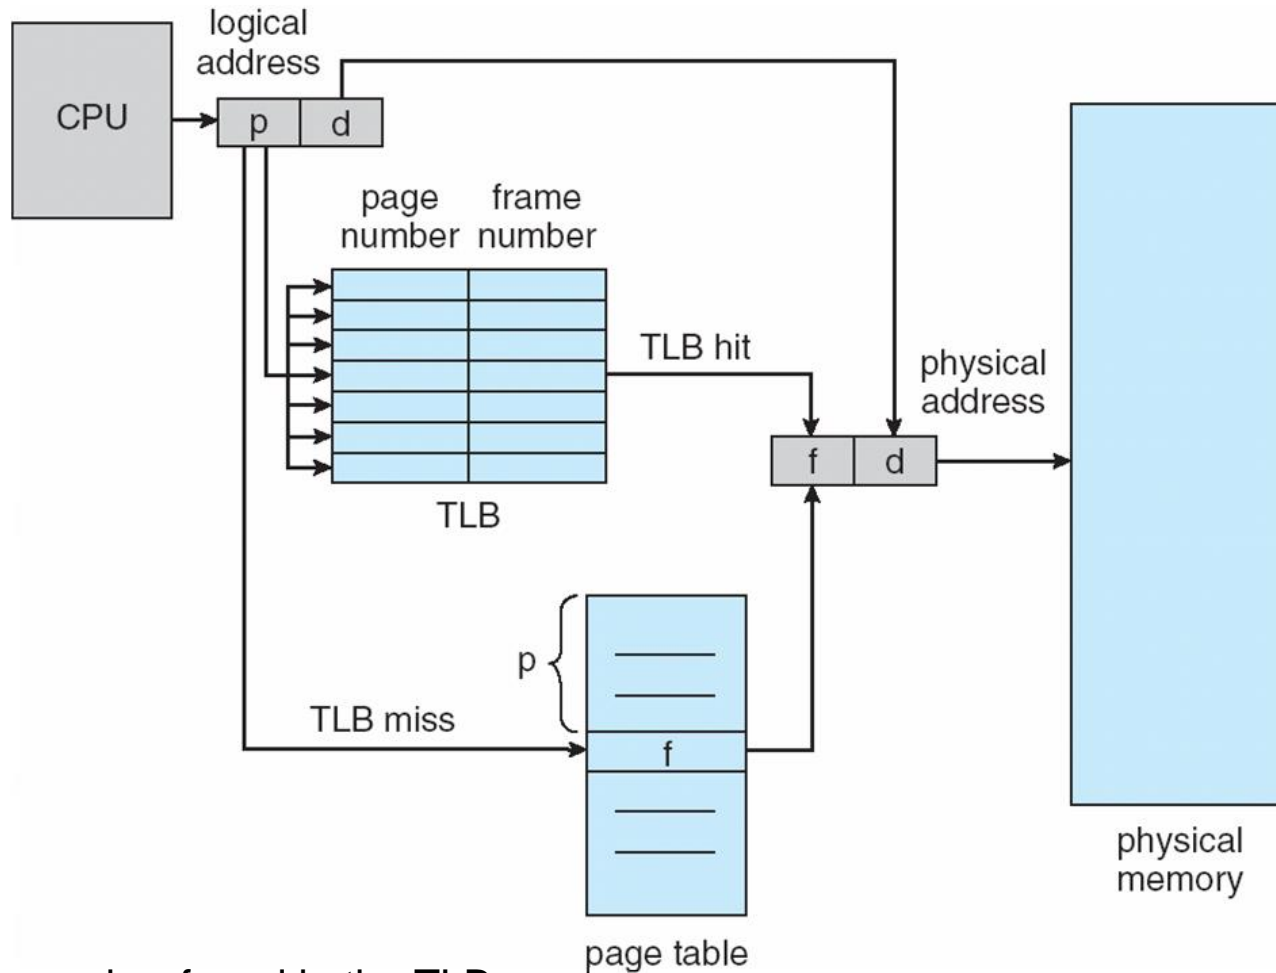
Address translation (p , d)

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





Paging Hardware With TLB



hit : page number found in the TLB
miss : page number not in the TLB





Effective Access Time

(เวลาแท้จริงที่ใช้ในการเข้าถึงหน่วยความจำหลัก)

- Hit ratio – percentage of times that a page number is found in the TLB
- An 80% hit ratio means that we find the desired page number in the TLB 80% of the time.
- Suppose that 10 nanoseconds to access memory.
 - If we find the desired page in TLB then a mapped-memory access take 10 ns
 - Otherwise we need two memory access so it is 20 ns
- **Effective Access Time (EAT)**

$$\text{EAT} = 0.80 \times 10 + 0.20 \times 20 = 12 \text{ nanoseconds}$$

implying 20% slowdown in access time

- Consider amore realistic hit ratio of 99%,
$$\text{EAT} = 0.99 \times 10 + 0.01 \times 20 = 10.1\text{ns}$$
implying only 1% slowdown in access time.

hit ratio: อัตราส่วนการพบ page number ต่อจำนวน page number ที่อ้างอิงทั้งหมด

ratio : อัตราส่วน





Memory Protection

- Memory protection implemented by **associating protection bit with each frame**

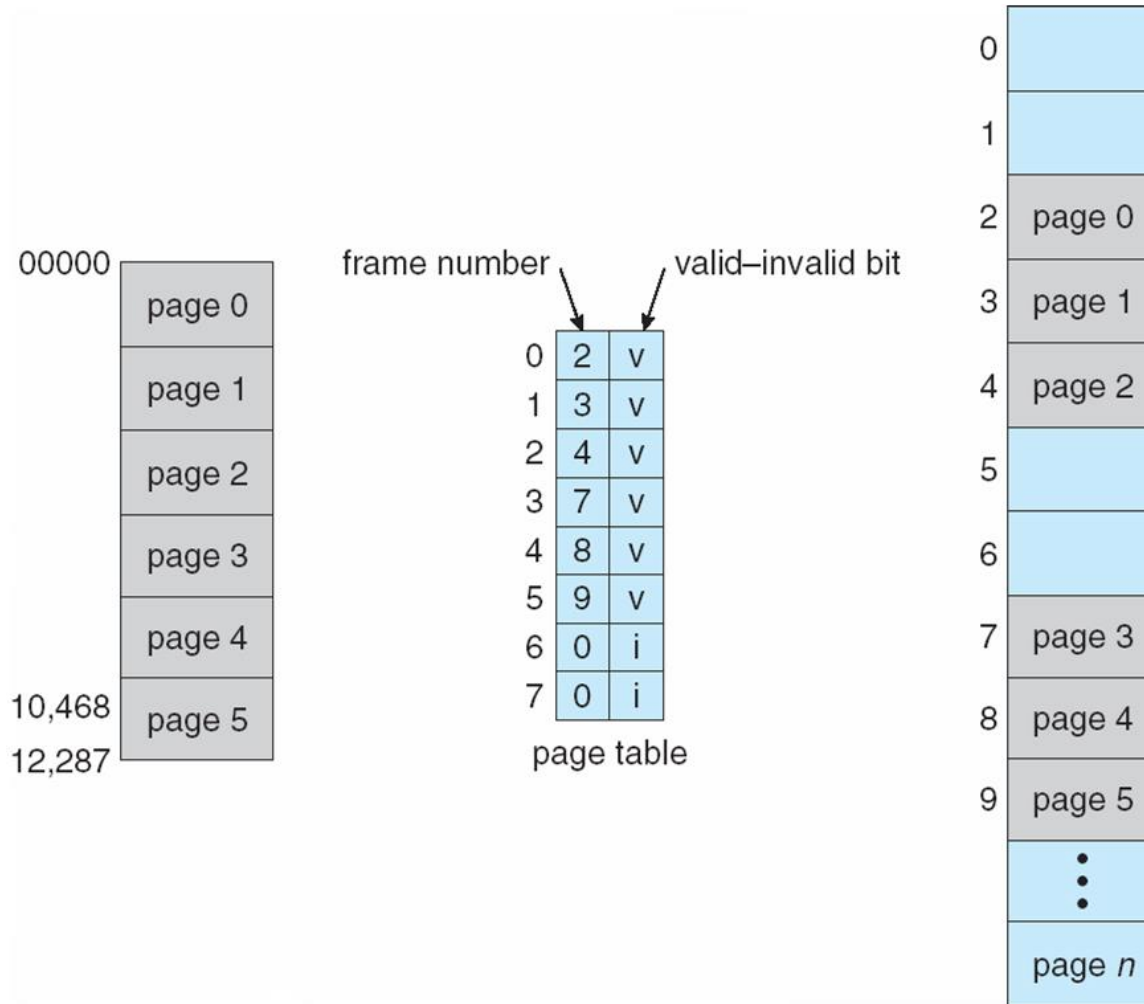
- **Valid-invalid** bit attached to each entry in the page table:
 - “**valid**” indicates that the associated **page is in the process**’ logical address space, and is thus a legal page
 - “**invalid**” indicates that the **page is not in the process**’ logical address space
 - Or use **page-table length register (PTLR)**

- Any violations result in a trap to the kernel





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





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

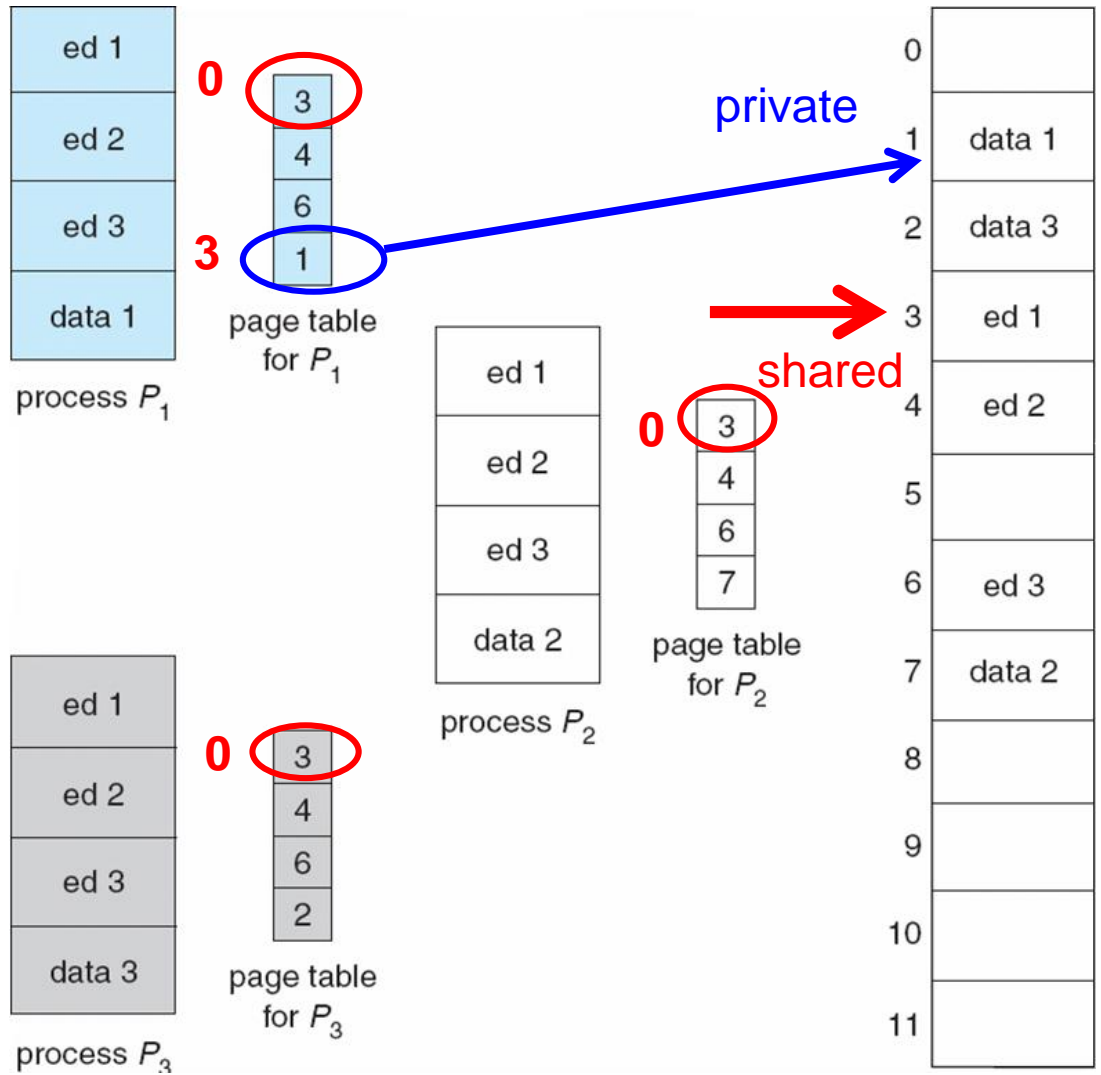
■ Private code and data

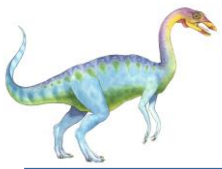
- Each process keeps a separate copy of the code and data
- The pages for the private code and data can appear anywhere in the logical address space





Shared Pages Example





Structure of the Page Table


- Hierarchical Paging
- Hashed Page Tables
- Inverted Page Tables






Hierarchical Page Tables

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





Two-Level Page-Table Scheme



Outer page table : หน้าอื่นๆ ที่อยู่คนละตาราง





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 10-bit page number
 - a 12-bit page offset
- Thus, a logical address is as follows:

page number		page offset
p_1	p_2	d
10	10	12

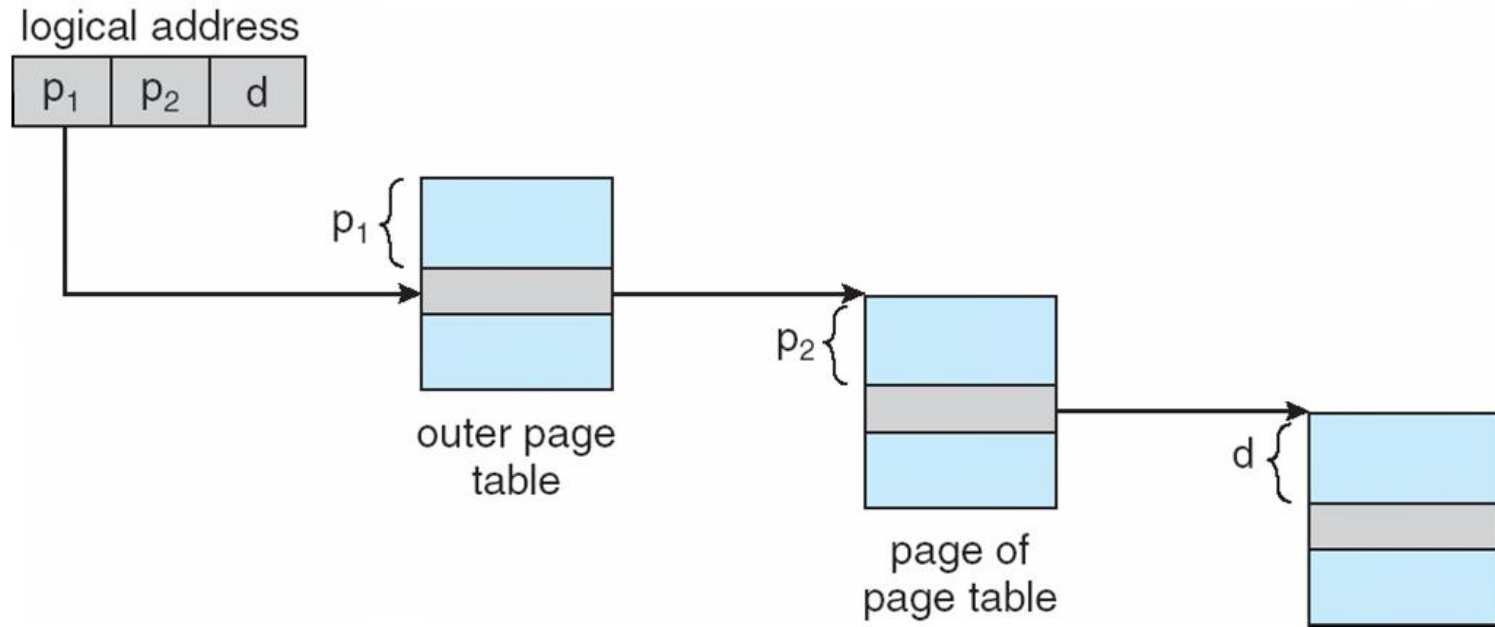
where p_1 is an index into the outer page table, and p_2 is the displacement within the page of the outer page table

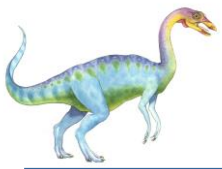
Outer page table : หน้าอื่นๆ ที่อยู่คนละตาราง





Address-Translation Scheme





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






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





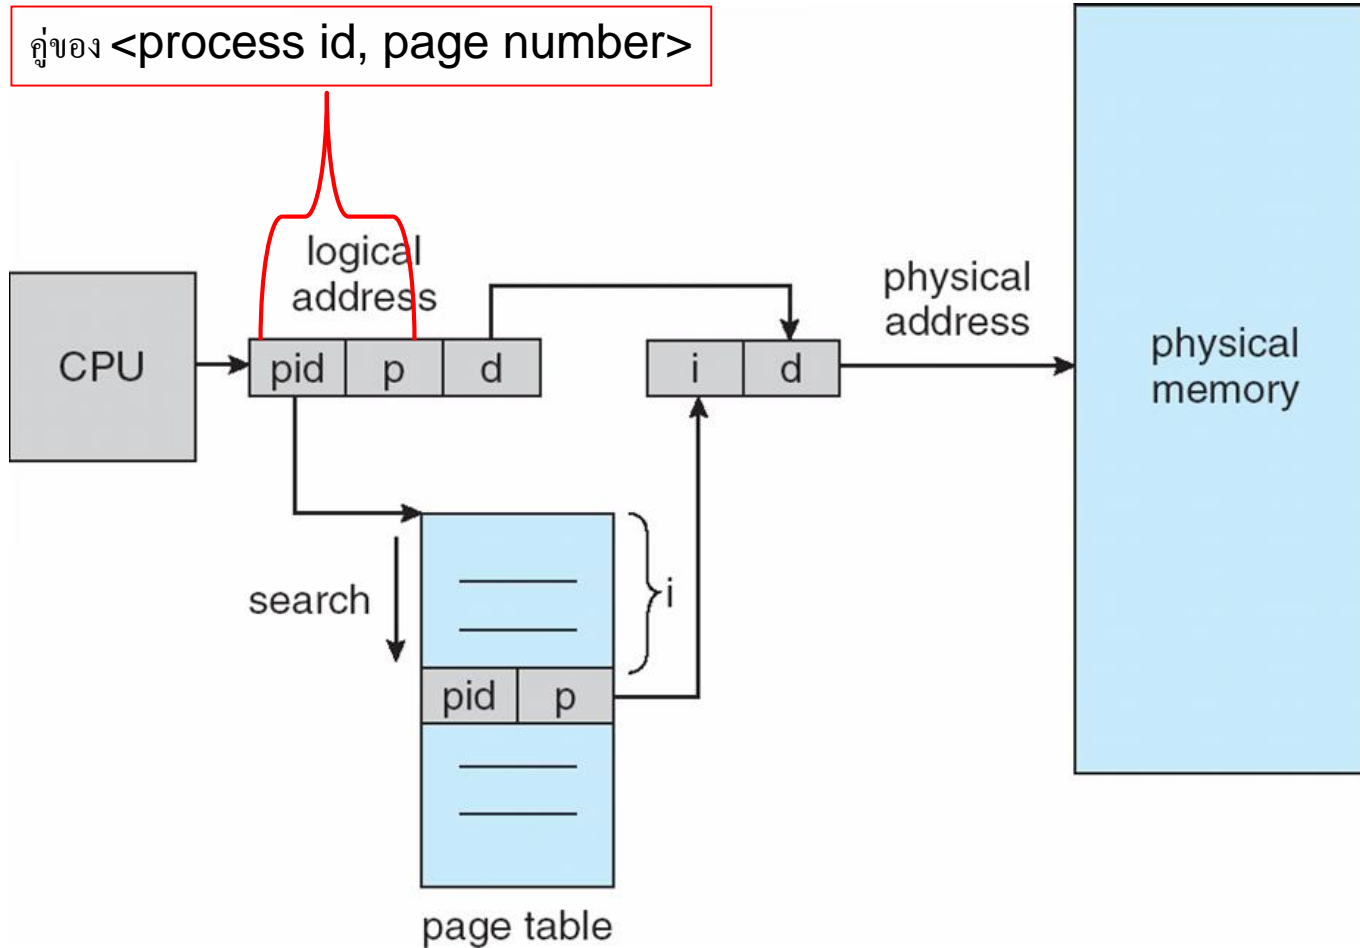
Inverted Page Table

- 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



pid : Process id (หมายเลข process)

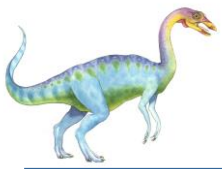





Segmentation (แบ่งเป็นตอน)

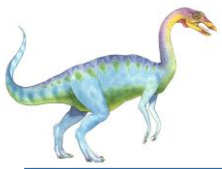
- Memory-management scheme that supports user view of memory
- A **program is a collection of segments**
 - A segment is a logical unit such as:
 - main program
 - procedure
 - function
 - method
 - object
 - local variables, global variables
 - common block
 - stack
 - symbol table
 - arrays






User's View of a Program






Logical View of Segmentation



user space



physical memory space






Segmentation Architecture

อ้างอิงความเข้าใจในเรื่อง paging

- Logical address consists of a two tuple:
 <segment-number, offset>,
- **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** < **STLR**





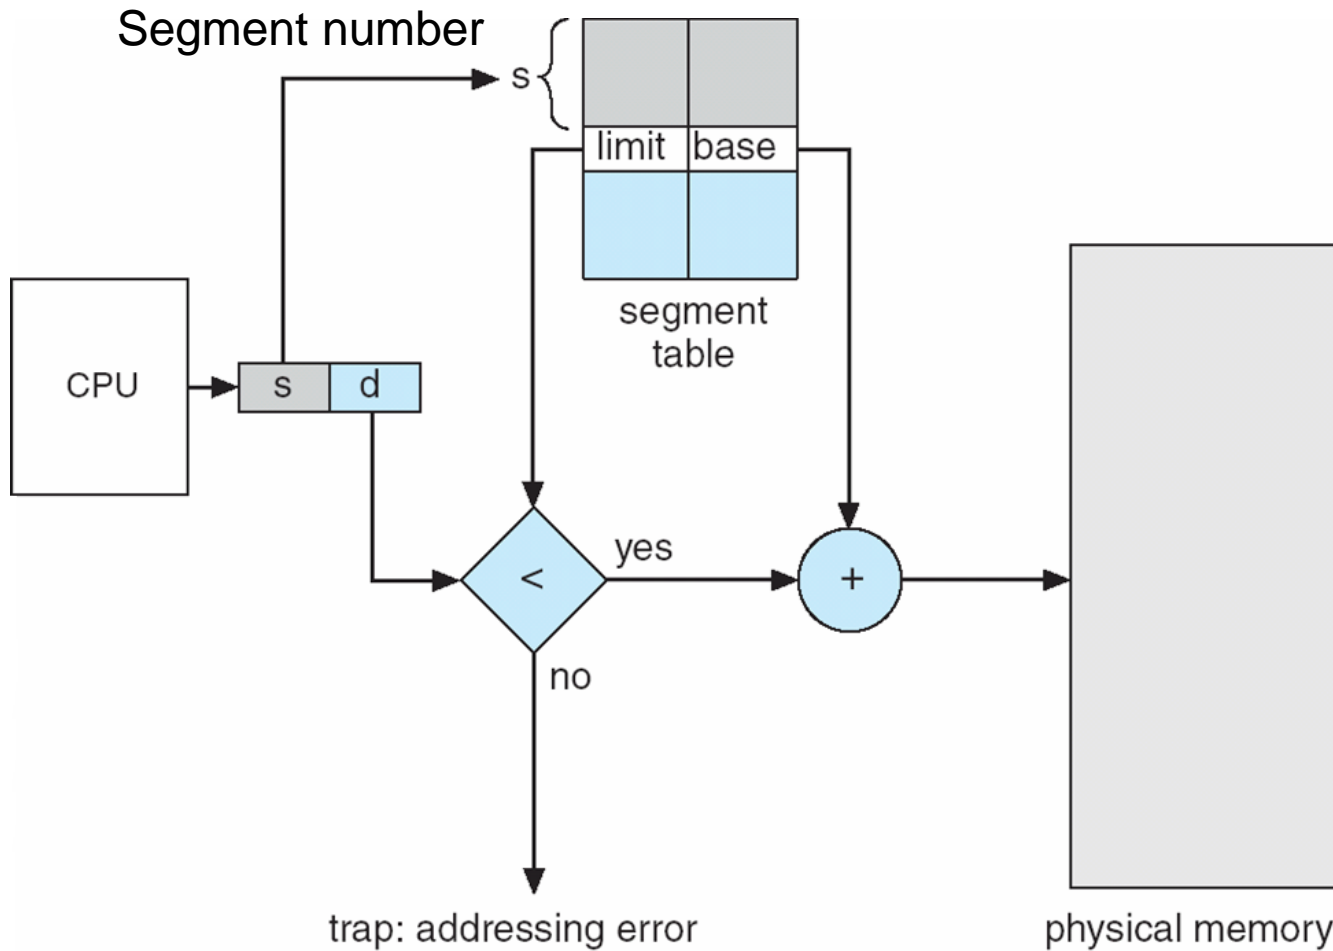
Segmentation Architecture (Cont.)

- 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



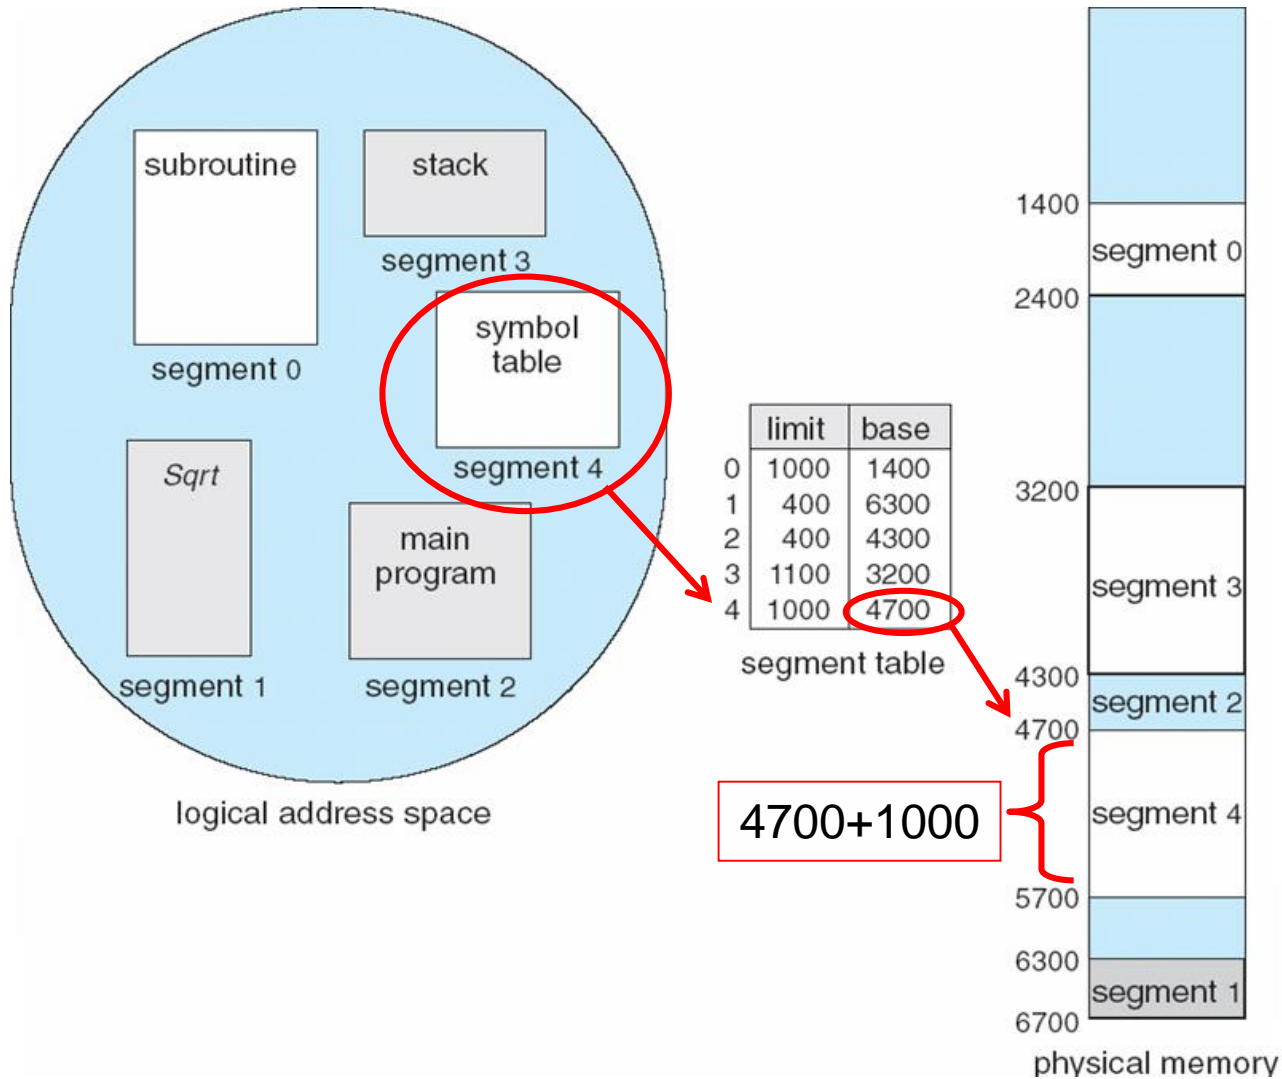


Segmentation Hardware





Example of Segmentation





Example: The Intel IA-32 Architecture


- Supports both segmentation and segmentation with paging
 - Each segment can be 4 GB
 - Up to 16 K segments per process
 - Divided into two partitions
 - ▶ First partition of up to 8 K segments are private to process (kept in **local descriptor table (LDT)**)
 - ▶ Second partition of up to 8K segments shared among all processes (kept in **global descriptor table (GDT)**)





Example: The Intel IA-32 Architecture (Cont.)

- CPU generates logical address
 - Selector given to segmentation unit
 - ▶ Which produces linear addresses





- Linear address given to paging unit
 - ▶ Which generates physical address in main memory
 - ▶ Paging units form equivalent of MMU
 - ▶ Pages sizes can be 4 KB or 4 MB



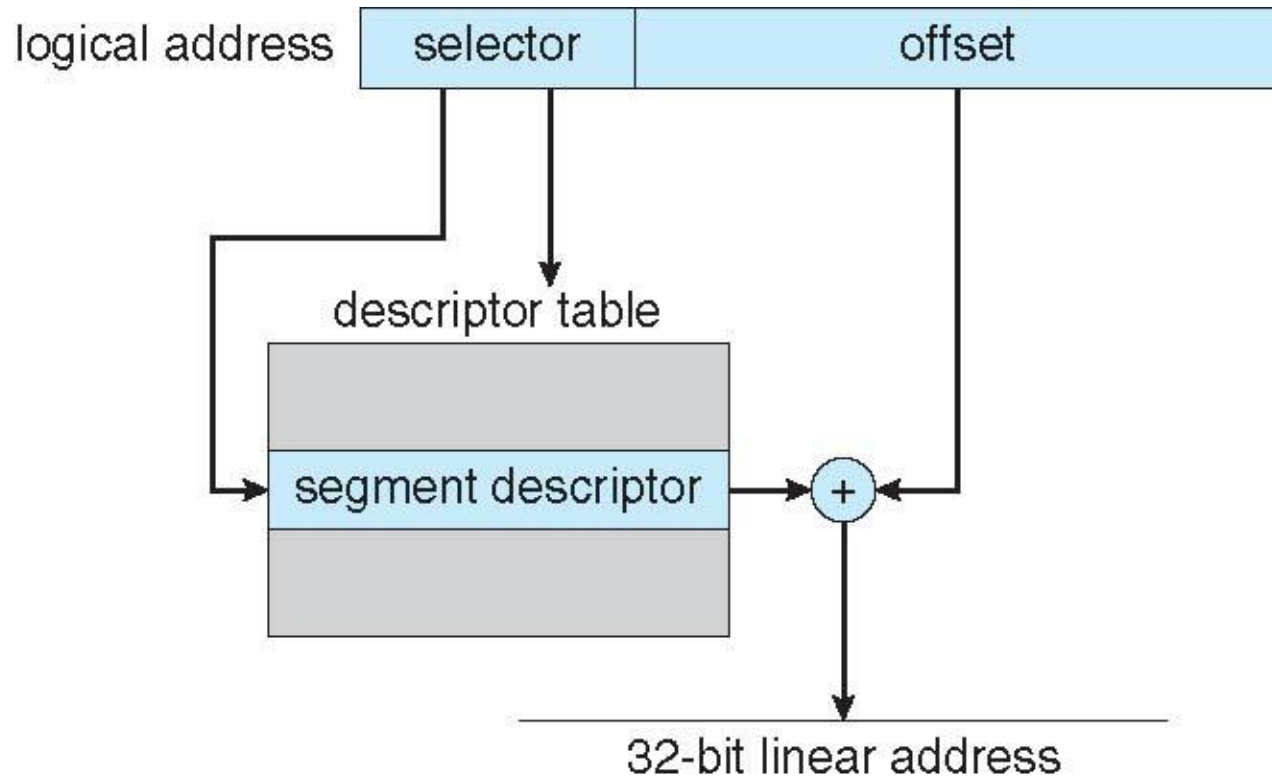


Logical to Physical Address Translation in IA-32






Intel IA-32 Segmentation






Intel IA-32 Paging Architecture





Intel IA-32 Page Address Extensions


- 32-bit address limits led Intel to create **page address extension (PAE)**, allowing 32-bit apps access to more than 4GB of memory space
 - Paging went to a 3-level scheme
 - Top two bits refer to a **page directory pointer table**
 - Page-directory and page-table entries moved to 64-bits in size
 - Net effect is increasing address space to 36 bits – 64GB of physical memory





Intel x86-64


- Current generation Intel x86 architecture
- 64 bits is ginormous (> 16 exabytes)
- In practice only implement 48 bit addressing
 - Page sizes of 4 KB, 2 MB, 1 GB
 - Four levels of paging hierarchy
- Can also use PAE so virtual addresses are 48 bits and physical addresses are 52 bits





Example: ARM Architecture

- Dominant mobile platform chip (Apple iOS and Google Android devices for example)
- Modern, energy efficient, 32-bit CPU
- 4 KB and 16 KB pages
- 1 MB and 16 MB pages (termed **sections**)
- One-level paging for sections, two-level for smaller pages
- Two levels of TLBs
 - Outer level has two micro TLBs (one data, one instruction)
 - Inner is single main TLB
 - First inner is checked, on miss outers are checked, and on miss page table walk performed by CPU



End of Chapter 7

