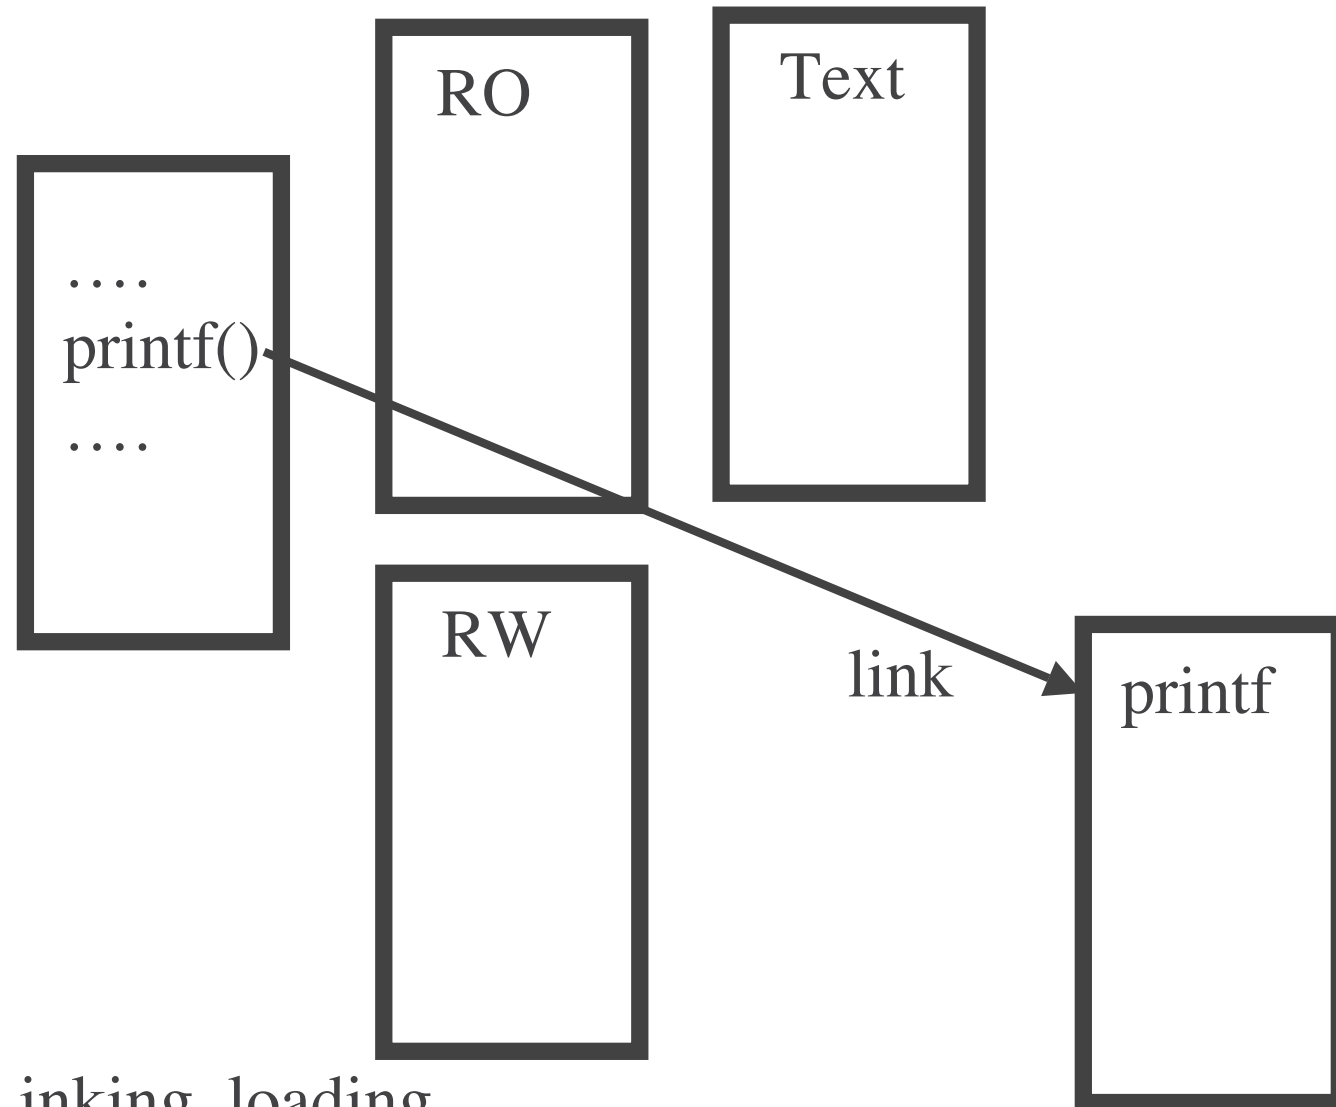


So far...



Linking, loading
Logical vs Physical addresses

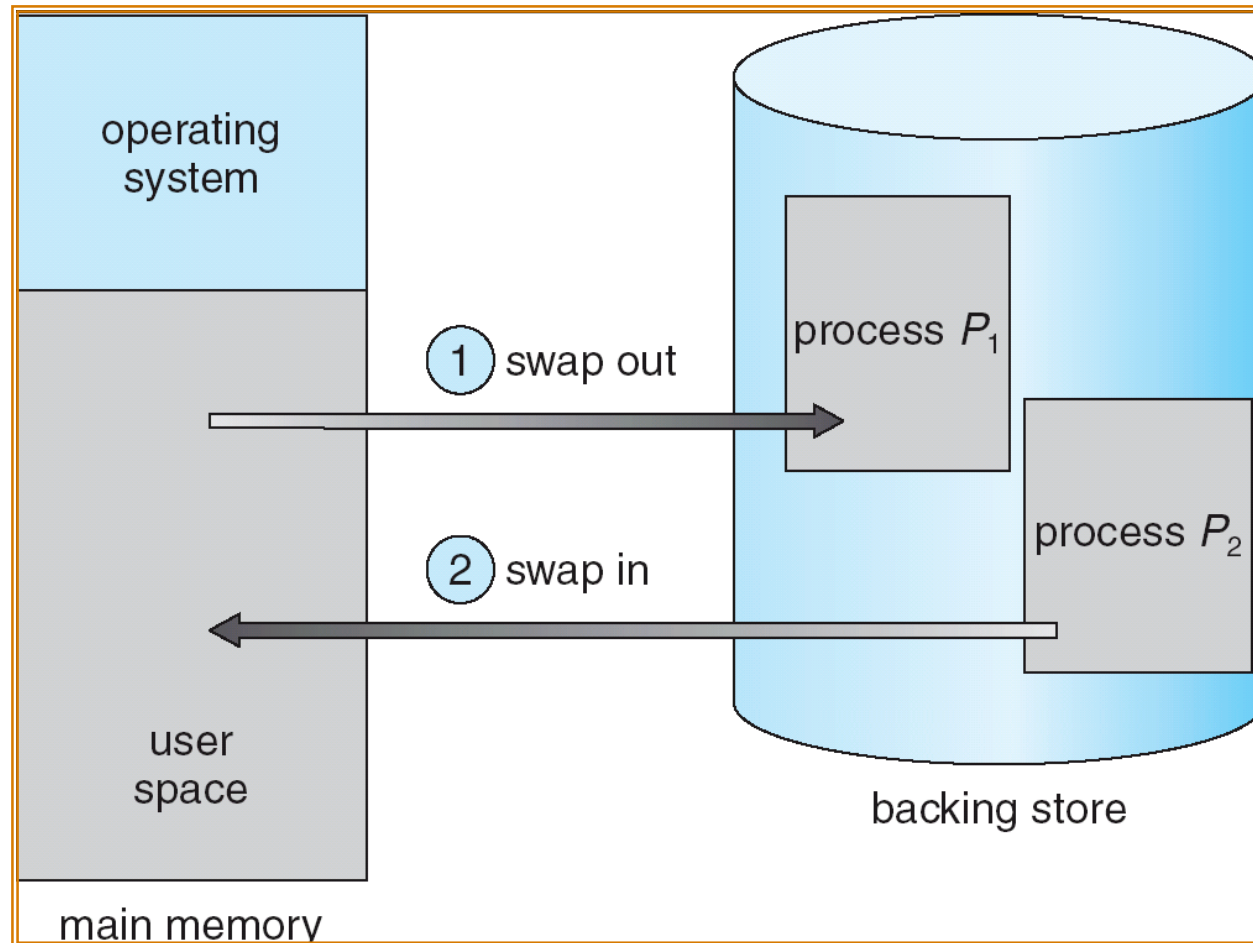


Swapping

- ▶ 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)



Schematic View of Swapping

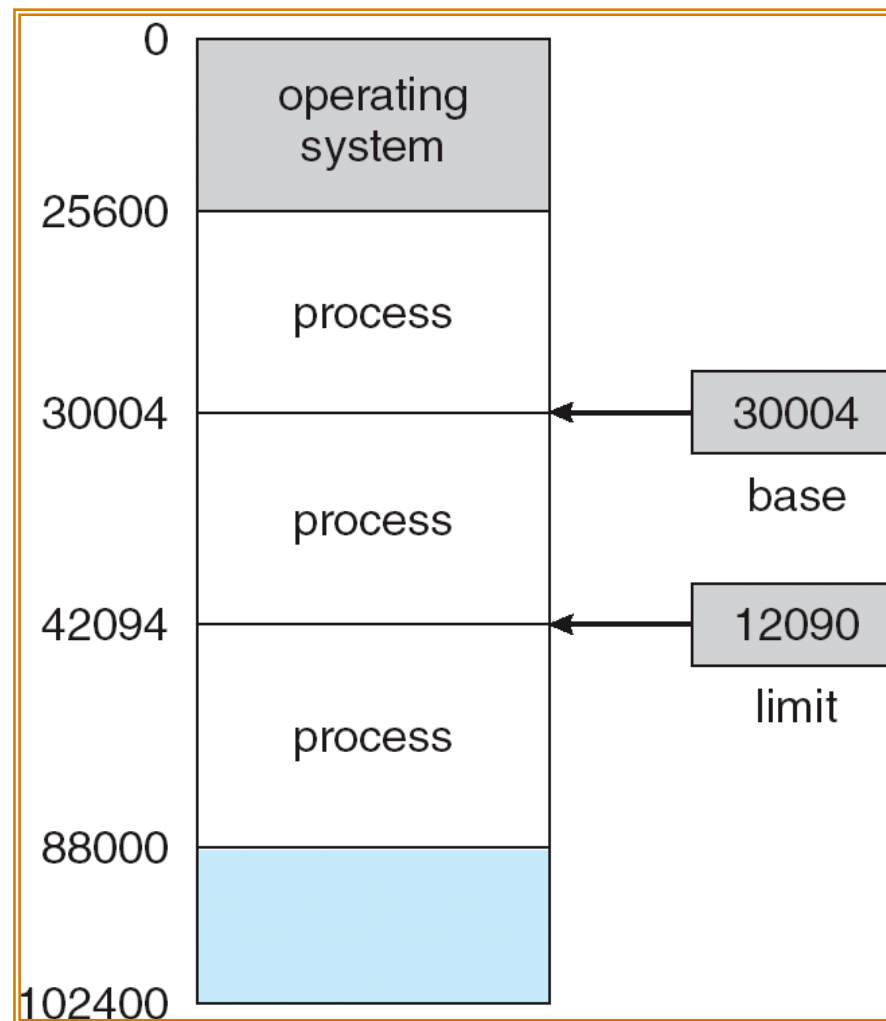


Contiguous Allocation

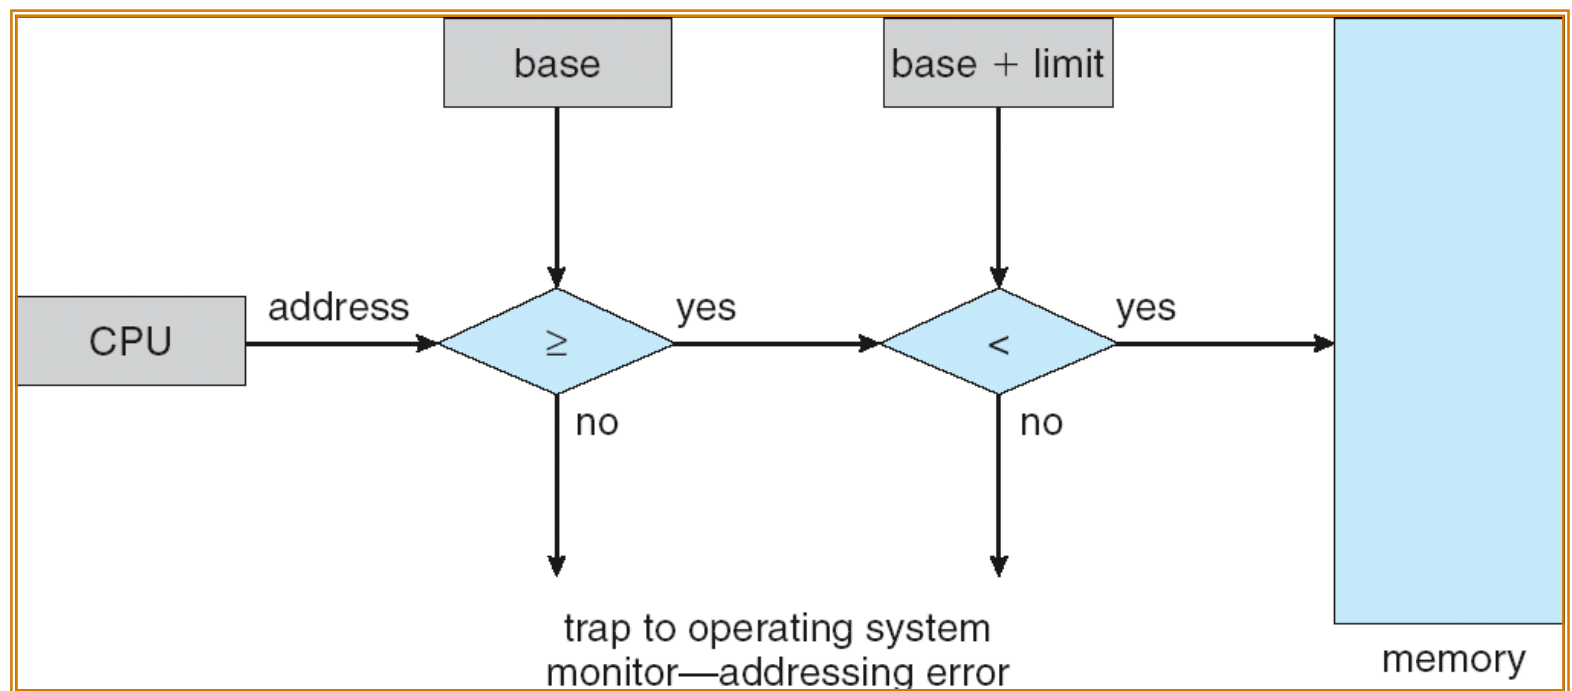
- ▶ Main memory usually into two partitions:
 - Resident operating system, usually held in low memory with interrupt vector
 - User processes then held in high memory
- ▶ Single-partition allocation
 - Relocation-register scheme used to protect user processes from each other, and from changing operating-system code and data
 - Relocation register contains value of smallest physical address; limit register contains range of logical addresses – each logical address must be less than the limit register



A base and a limit register define a logical address space



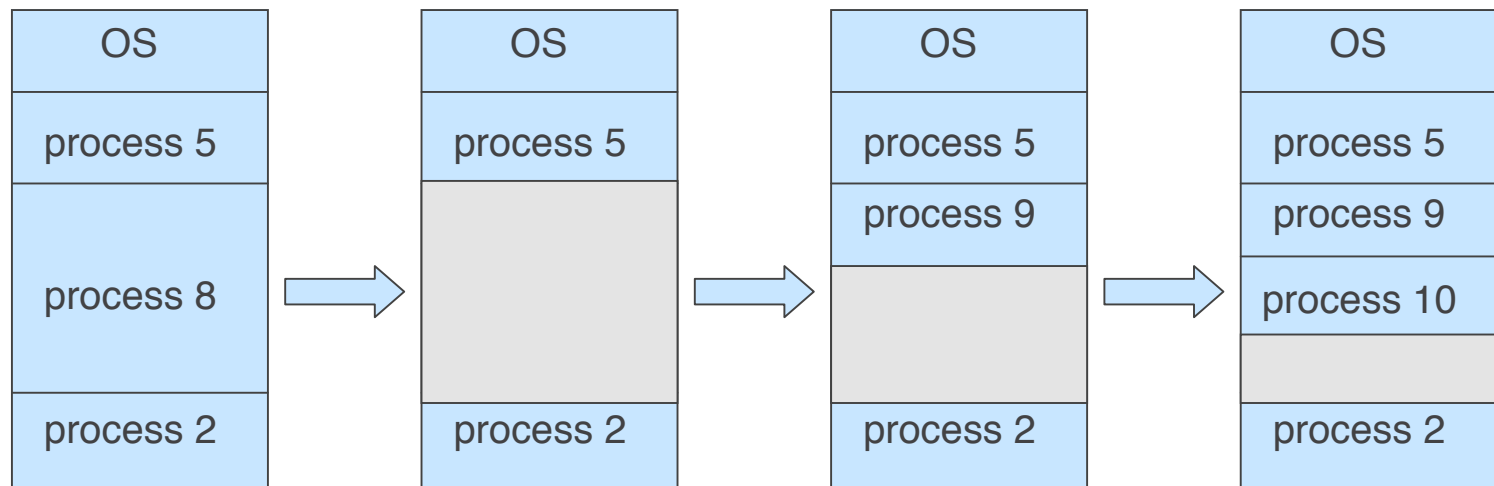
HW address protection with base and limit registers



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)



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



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 - I/O may be performed by a DMA controller
 - Latch job in memory while it is involved in I/O
 - Do I/O only into OS buffers



Paging for noncontiguous allocation

- ▶ 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 8192 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
- ▶ This scheme will create internal fragmentation

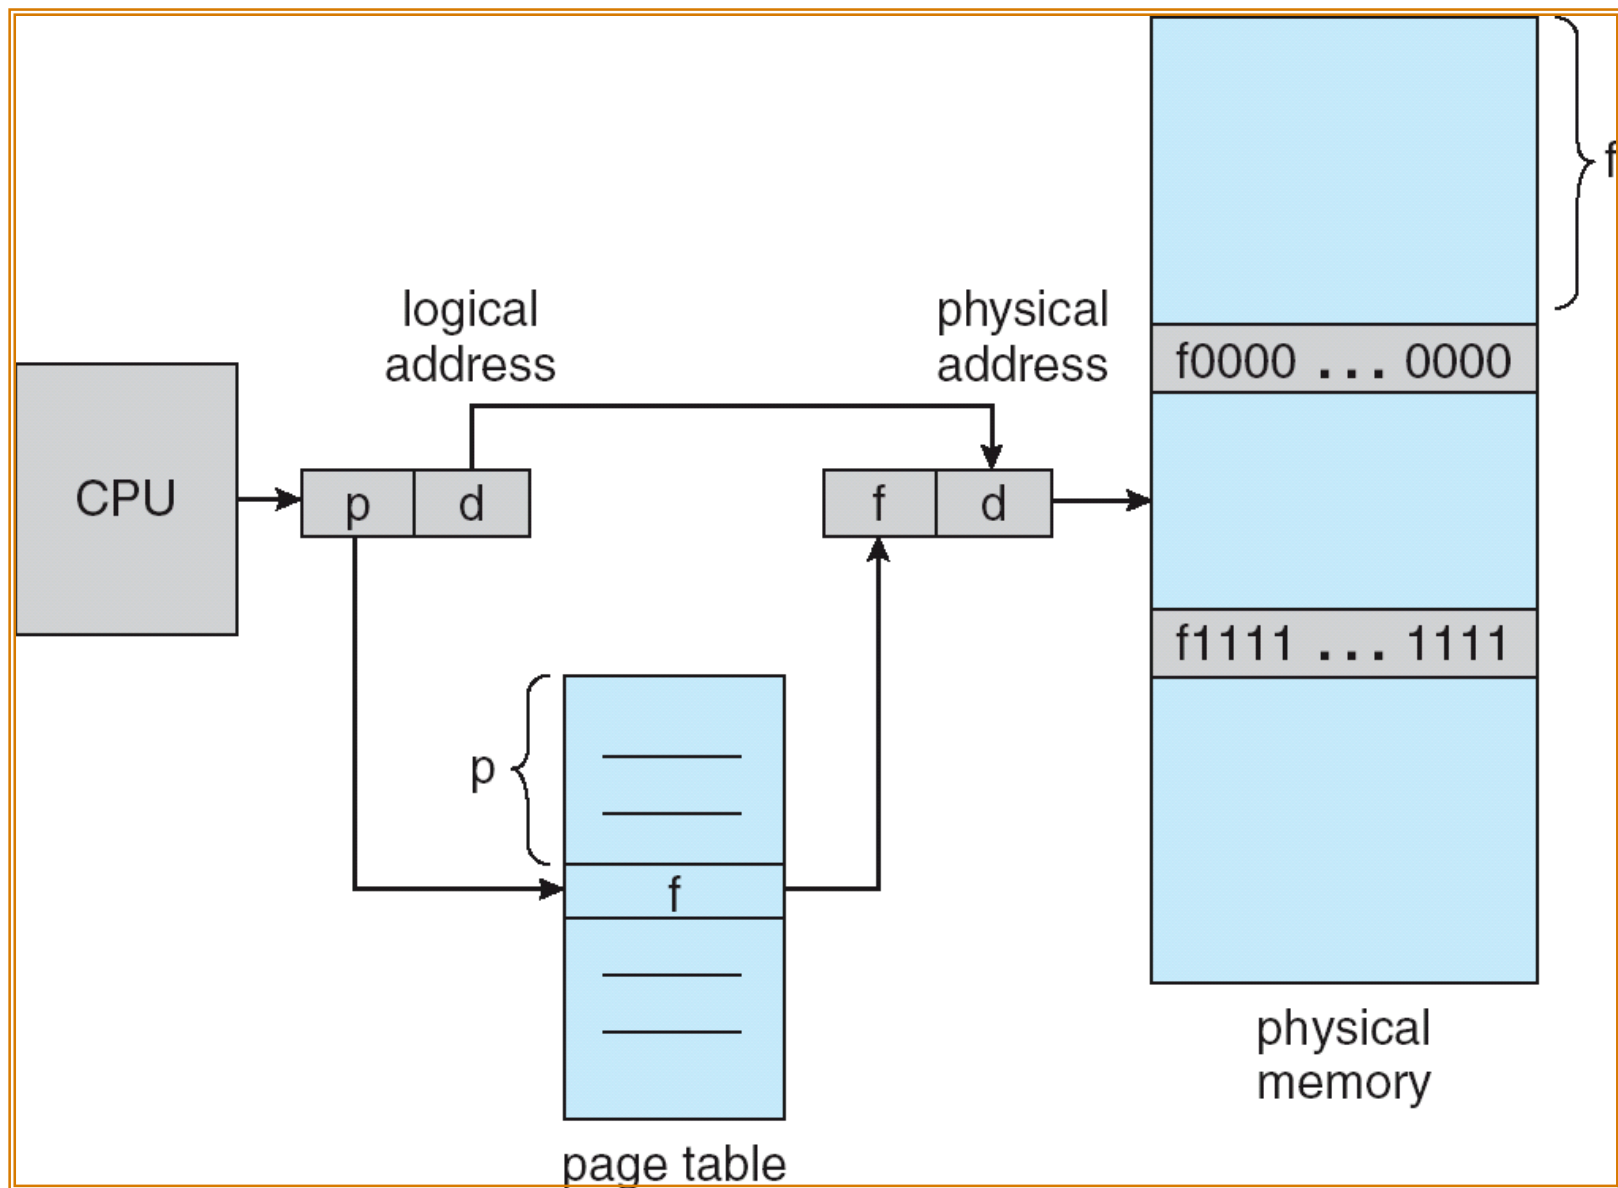


Address Translation Scheme

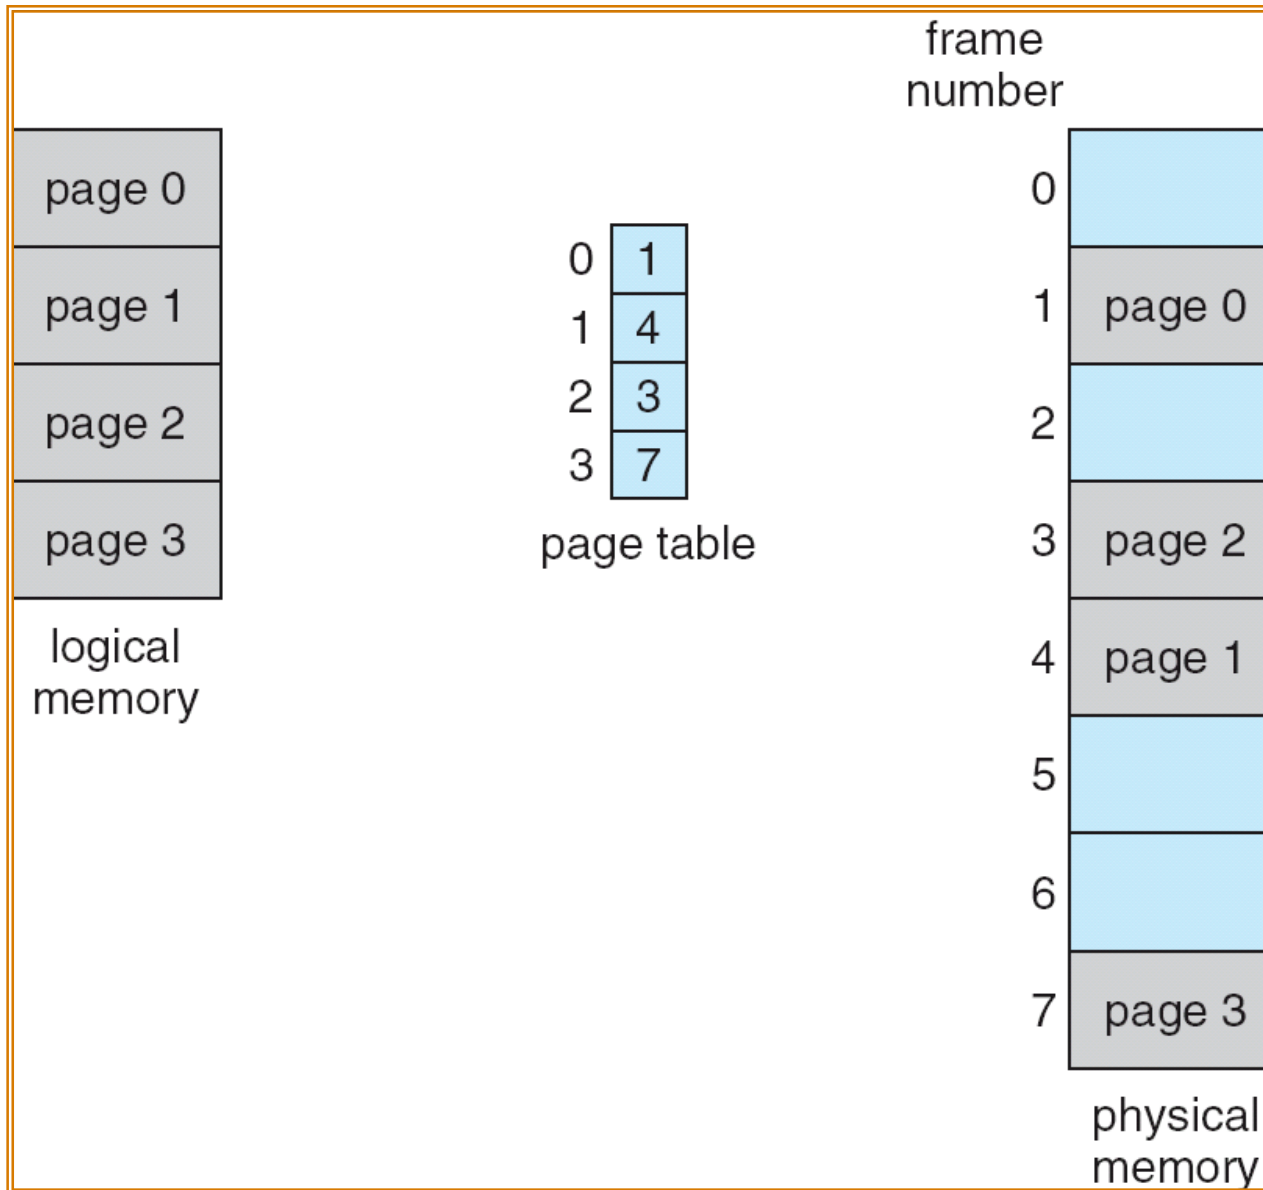
- ▶ Address generated by CPU is divided into:
 - *Page number (p)* – used as an index into a *page table* which contains base address of each page in physical memory
 - *Page offset (d)* – combined with base address to define the physical memory address that is sent to the memory unit



Address Translation Architecture



Paging Example



Paging Example

0	a
1	b
2	c
3	d
4	e
5	f
6	g
7	h
8	i
9	j
10	k
11	l
12	m
13	n
14	o
15	p

logical memory

0	5
1	6
2	1
3	2

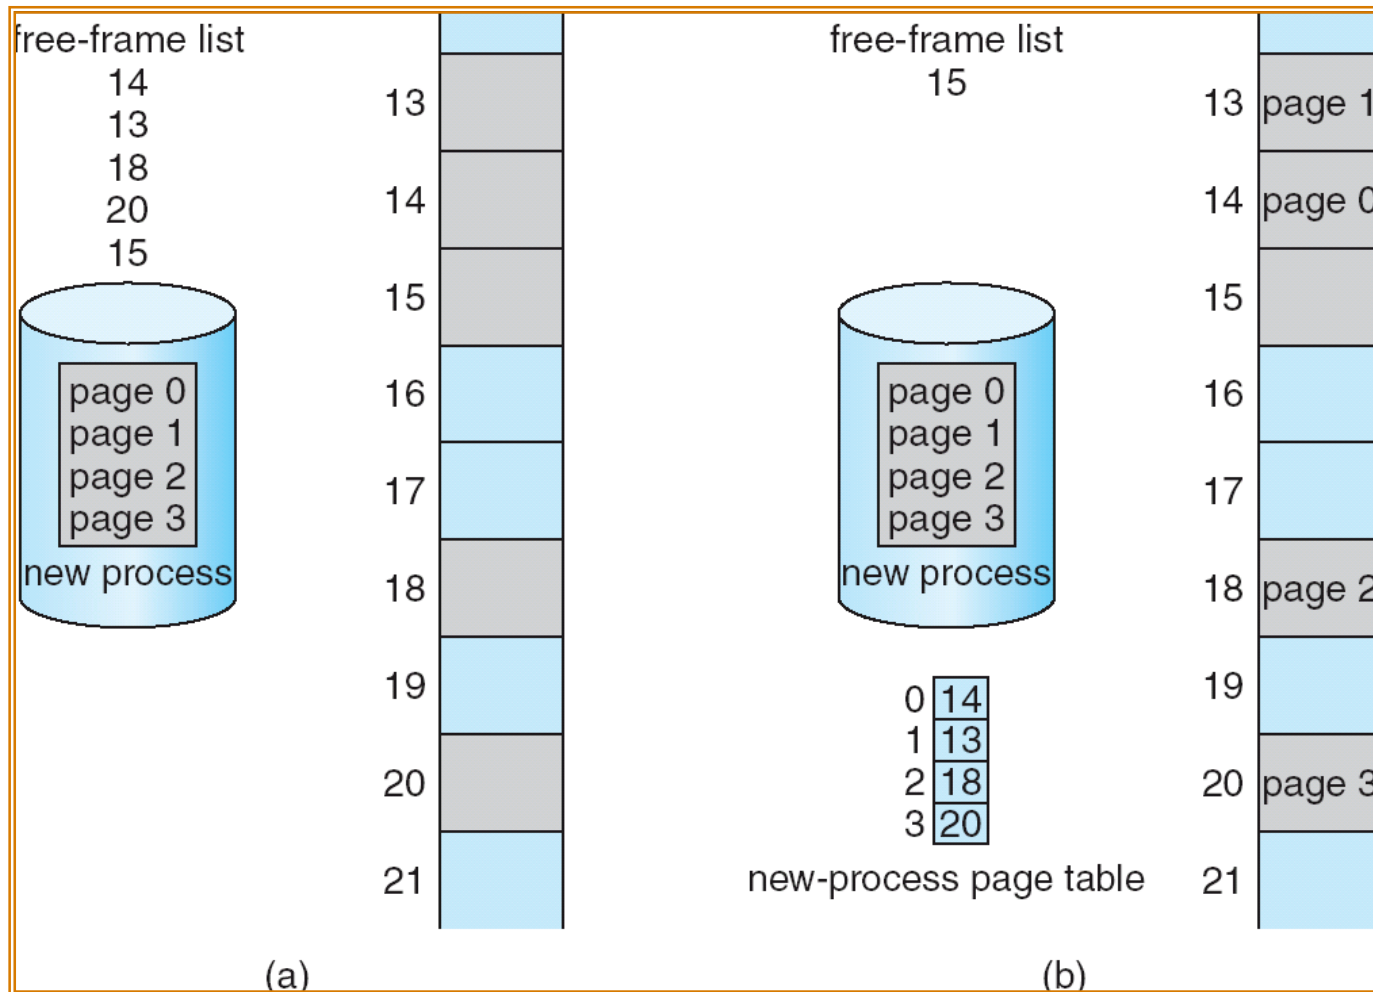
page table

0	
4	i j k l
8	m n o p
12	
16	
20	a b c d
24	e f g h
28	

physical memory



Free Frames



Before allocation

After allocation

