A decorative vertical bar on the left side of the slide. It consists of a dark teal background with a white dotted vertical line running through its center. To the right of this bar, there are several orange circles of varying sizes, arranged in a cluster. The largest circle is at the top, with several smaller ones below and to its right. The entire slide is framed by thin orange vertical lines on the far left and far right.

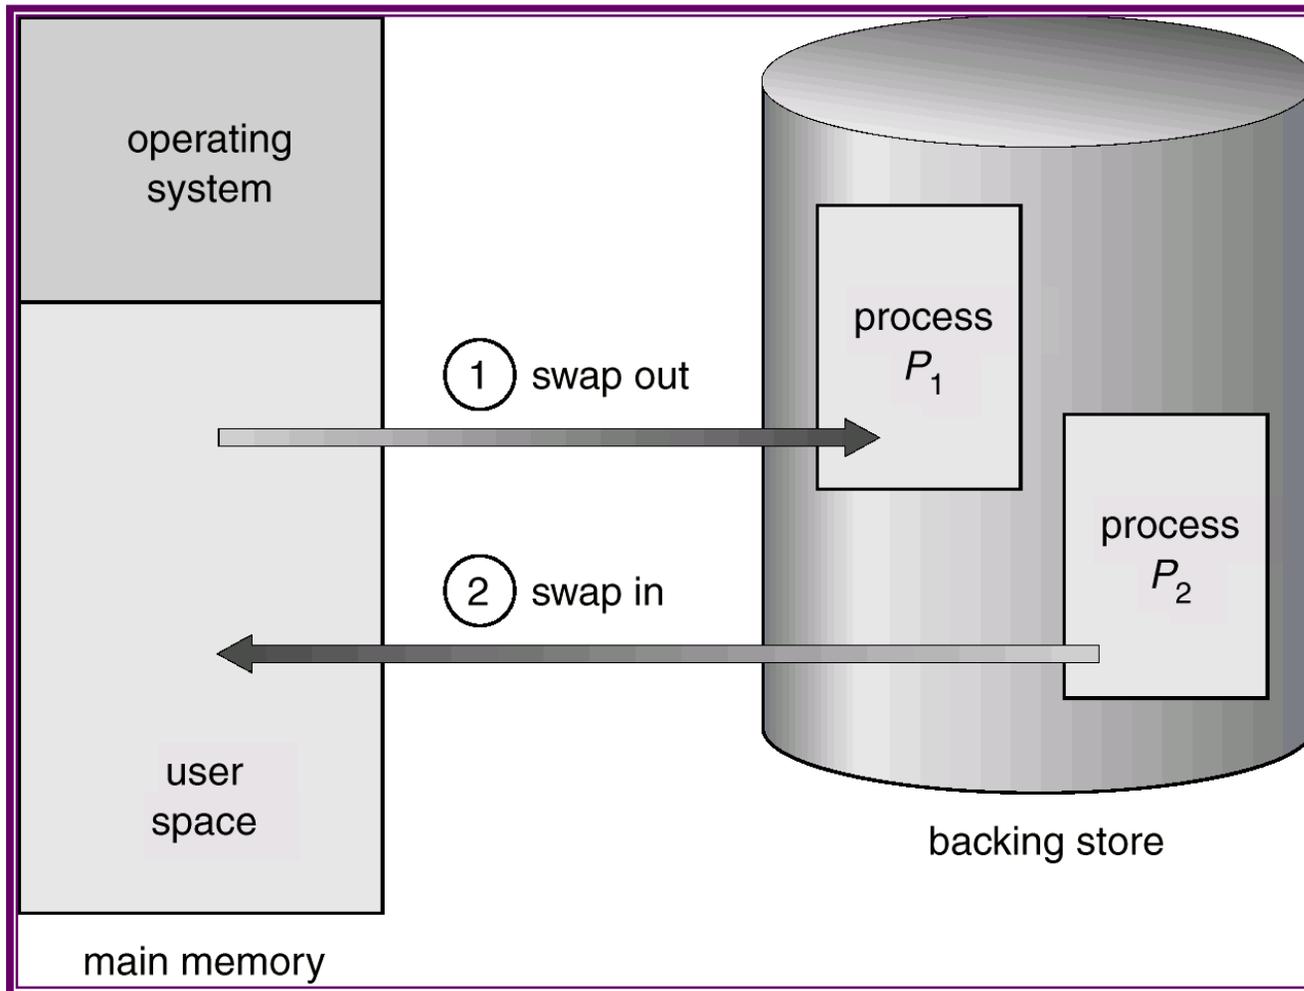
PRINCIPLES OF OPERATING SYSTEMS

LECTURE 12: SWAPPING

Introduction

- 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

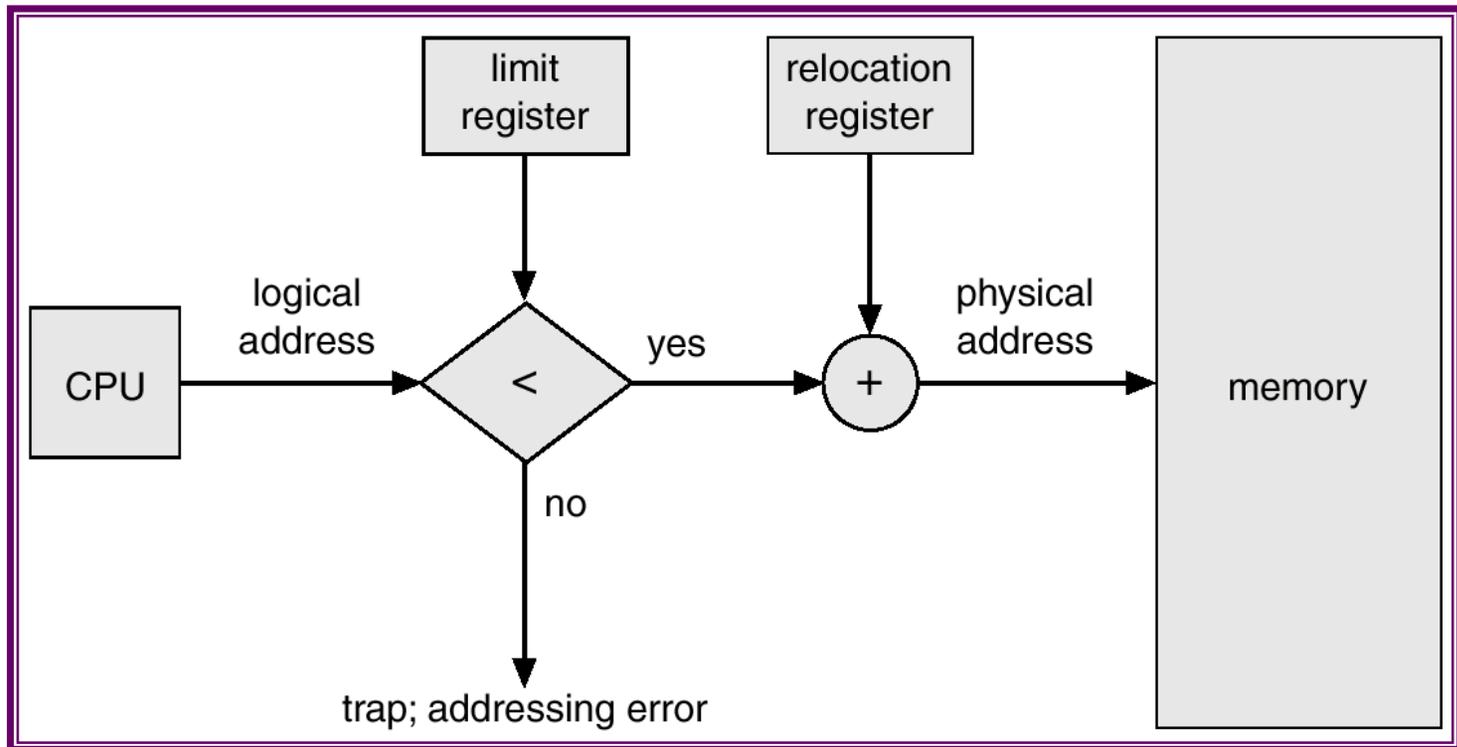


Contiguous Allocation

- 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.

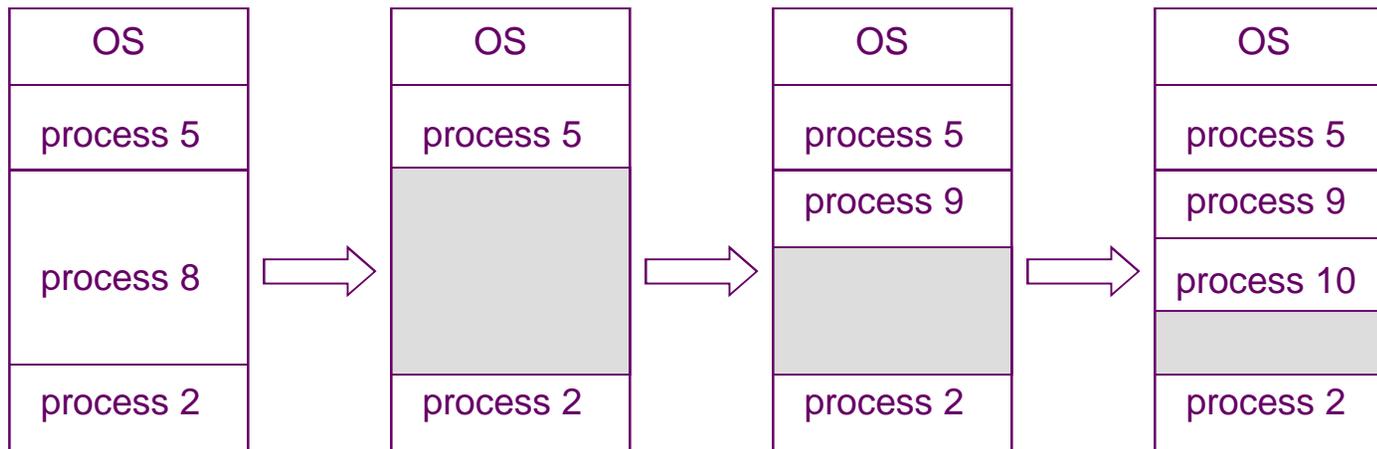
Hardware Support for Relocation and Limit Registers



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
 - 📄 Latch job in memory while it is involved in I/O.
 - 📄 Do I/O only into OS buffers.