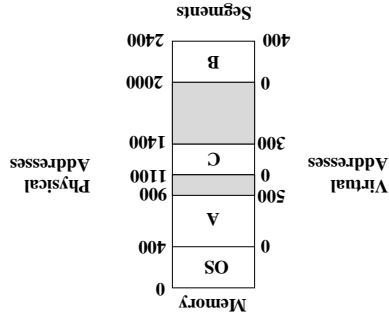


- Processes typically do not use their entire space in memory all the time.
- Paging
1. divides and assigns processes to fixed sized *pages*,
 2. then selectively allocates pages to *frames* in memory, and
 3. manages (moves, removes, reallocates) pages in memory.

Today: Paging



- Unprogramming
- Static Relocation
- Dynamic Relocation

Last Class: Memory Management

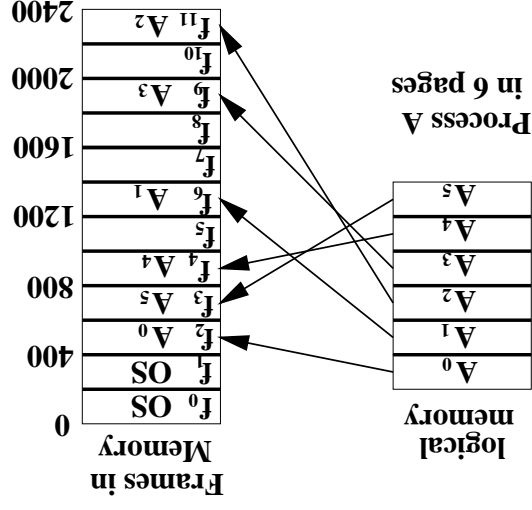
Paging: Motivation & Features

90/10 rule: Processes spend 90% of their time accessing 10% of their space in memory.

⇒ Keep only those parts of a process in memory that are actually being used

- Pages greatly simplify the hole fitting problem
- The logical memory of the process is contiguous, but pages need not be allocated contiguously in memory.
- By dividing memory into fixed size pages, we can eliminate external fragmentation.
- Paging does not eliminate internal fragmentation (1/2 page per process)

Paging: Example



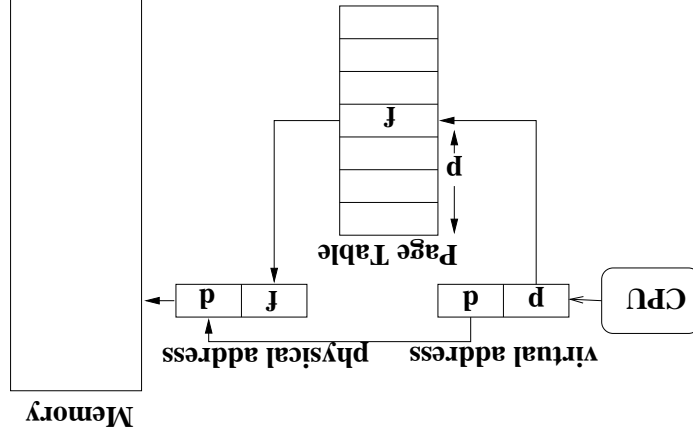
Paging Hardware

- **Problem:** How do we find addresses when pages are not allocated contiguously in memory?

- **Virtual Address:**

- Processes use a virtual (logical) address to name memory locations.
- Process generates contiguous, virtual addresses from 0 to size of the process.
- The OS lays the process down on pages and the paging hardware translates virtual addresses to actual physical addresses in memory.
- In paging, the virtual address identifies the page and the page offset.
- A *page table* keeps track of the page frame in memory in which the page is located.

Paging Hardware

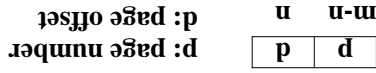


Paging Hardware

- Paging is a form of dynamic relocation, where each virtual address is bound by the paging hardware to a physical address.
- Think of the page table as a set of relocation registers, one for each frame.
- Mapping is invisible to the process; the OS maintains the mapping and the hardware does the translation.
- Protection is provided with the same mechanisms as used in dynamic relocation.

Paging Hardware: Practical Details

- Page size (frame sizes) are typically a power of 2 between 512 bytes and 8192 bytes per page.
- Powers of 2 make the translation of virtual addresses into physical addresses easier. For example, given
 1. virtual address space of size 2^m bytes and a page of size 2^n , then
 2. the high order $m - n$ bits of a virtual address select the page,
 3. the low order n bits select the offset in the page



- How should we store the page table?
- **Registers:** Advantages? Disadvantages?
 - **Memory:** Advantages? Disadvantages?
 - **TLB:** a fast fully associative memory that stores page numbers (key) and the frame (value) in which they are stored.
 - if memory accesses have locality, address translation has locality too.
 - typical TLB sizes range from 8 to 2048 entries.

Making Paging Efficient

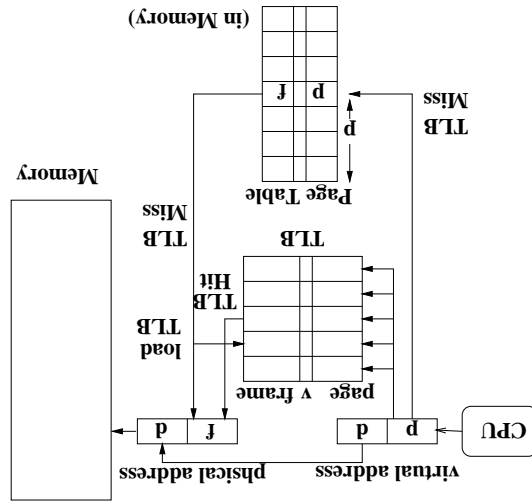
- How many bits for an address? Assume we can address only 1 word (4 byte) increments?
- What part is p, and d?
- Given virtual address 13, do the virtual to physical translation.
- What needs to happen on a context switch?

Address Translation Example

1. What is the effective memory access cost if the page table is in memory?
 2. What is the effective memory access cost with a TLB?
- A large TLB improves hit ratio, decreases average memory cost.

Costs of Using The TLB

v: valid bit that says the entry is up-to-date



The Translation Look-aside Buffer (TLB)

- **Multilevel Paging:** If the virtual address space is huge, page tables get too big, and many systems use a multilevel paging scheme (refer OSC for details)
 1. Copy the page table base register value to the PCB.
 2. Copy the TLB to the PCB (optionally).
 3. Flush the TLB.
 4. Restore the page table base register.
 5. Restore the TLB if it was saved.
- On a context switch:
 - The page table
 - Possibly a copy of the TLB
- The Process Control Block (PCB) must be extended to contain:

Saving/Restoring Memory on a Context Switch

1. Process needing k pages arrives.
2. If k page frames are free, then allocate these frames to pages. Else free frames that are no longer needed.
3. The OS puts each page in a frame and then puts the frame number in the corresponding entry in the page table.
4. OS marks all TLB entries as invalid (flushes the TLB).
5. OS starts process.
6. As process executes, OS loads TLB entries as each page is accessed, replacing an existing entry if the TLB is full.

Initializing Memory when Starting a Process

Sharing

Paging allows sharing of memory across processes, since memory used by a process no longer needs to be contiguous.

- Shared code must be reentrant, that means the processes that are using it cannot change it (e.g., no data in reentrant code).
- Sharing of pages is similar to the way threads share text and memory with each other.
- A shared page may exist in different parts of the virtual address space of each process, but the virtual addresses map to the same physical address.
- The user program (e.g., emacs) marks text segment of a program as reentrant with a system call.
- The OS keeps track of available reentrant code in memory and reuses them if a new process requests the same program.
- Can greatly reduce overall memory requirements for commonly used applications.

Summary

- Paging is a big improvement over segmentation:
 - They eliminate the problem of external fragmentation and therefore the need for compaction.
 - They allow sharing of code pages among processes, reducing overall memory requirements.
 - They enable processes to run when they are only partially loaded in main memory.
- However, paging has its costs:
 - Translating from a virtual address to a physical address is more time-consuming.
 - Paging requires hardware support in the form of a TLB to be efficient enough.
 - Paging requires more complex OS to maintain the page table.