Memory Management — Last Time

Uniprogramming:

Static Relocation:

Dynamic Relocation:

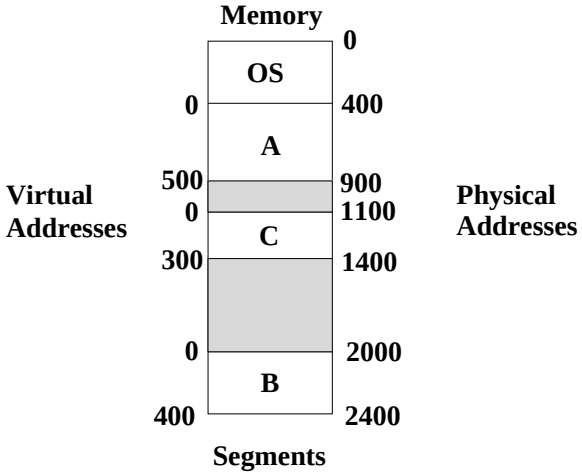

Paging — Today

A process typically does not use all their memory all the time. Paging

1. divides and assigns a process to fixed sized *pages*,
2. stores a copy of the *process* on disk,
3. then demand allocates pages to *frames* in memory, and
4. manages (moves, removes, reallocates) pages in memory.

Motivation & Features of Paging

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

$\Longrightarrow$ Implication: we only want to keep 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.

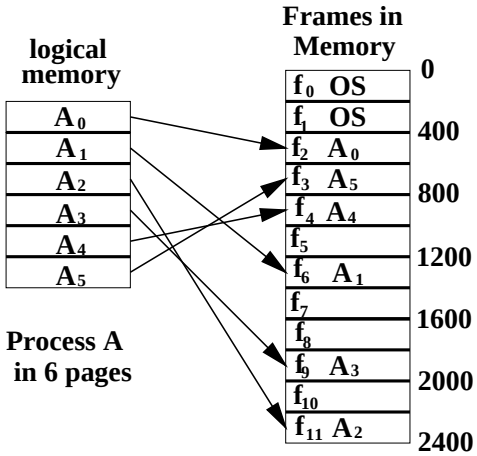

- By dividing memory into fixed size pages, we can eliminate external fragmentation. Of course, we introduce what? and how much of it?

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.

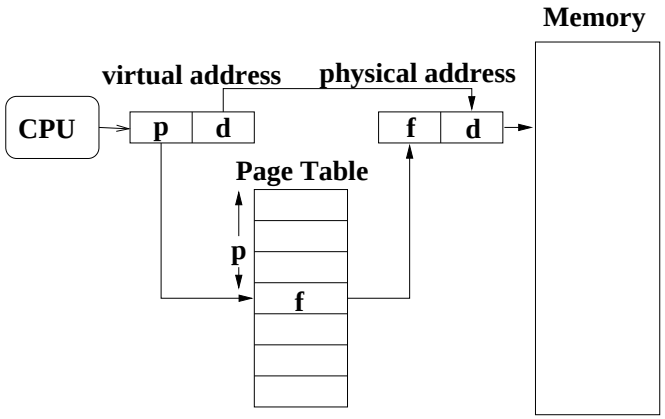

Paging Hardware

Paging is a form of dynamic relocation, where each virtual address is bound by the paging hardware to a physical address.

$\Longrightarrow$ 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.

Some details

- The page size (frame size), like memory 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 or words)
  2. a page of size 2^n, then
  3. the high order $m - n$ bits of a virtual address select the page,
  4. the low order n bits select the offset in the page.

p	d
m-n	n

p: page number
d: page offset

Address Translation Example

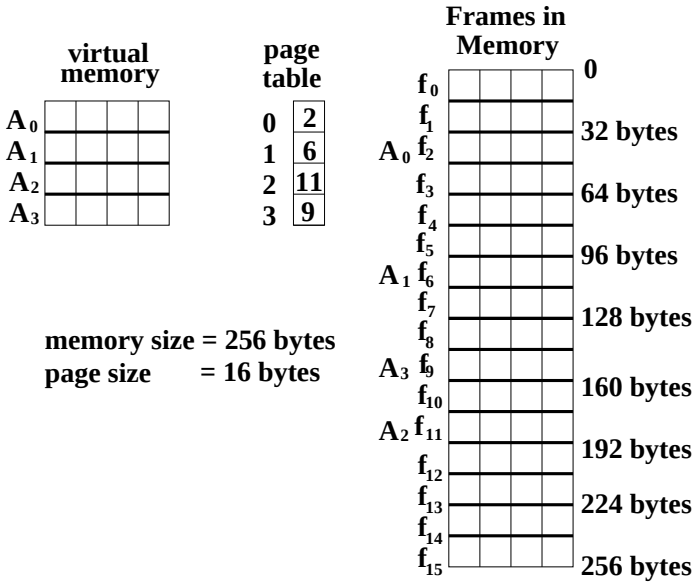

- How big is the page table?
- How many bits for an address. Assume we can address 1 byte increments?
- What part is p, and d?
- Given virtual address 24, do the virtual to physical translation.

Address Translation Example 2

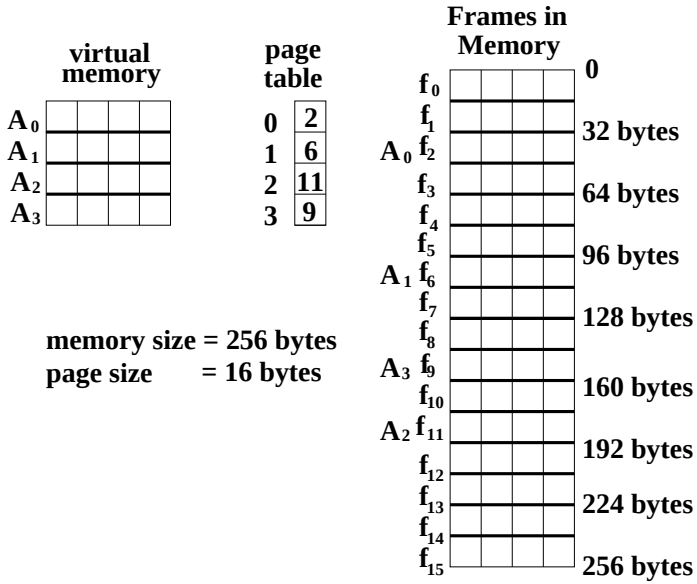

- How many bits for an address? Assume we can address only 1 word (4 byte) increments?
- What part is p, and d?
- Given virtual address 13, do the virtual to physical translation.

What needs to happen on a context switch?

Making Paging Go Fast

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.

The Translation Look-aside Buffer (TLB)

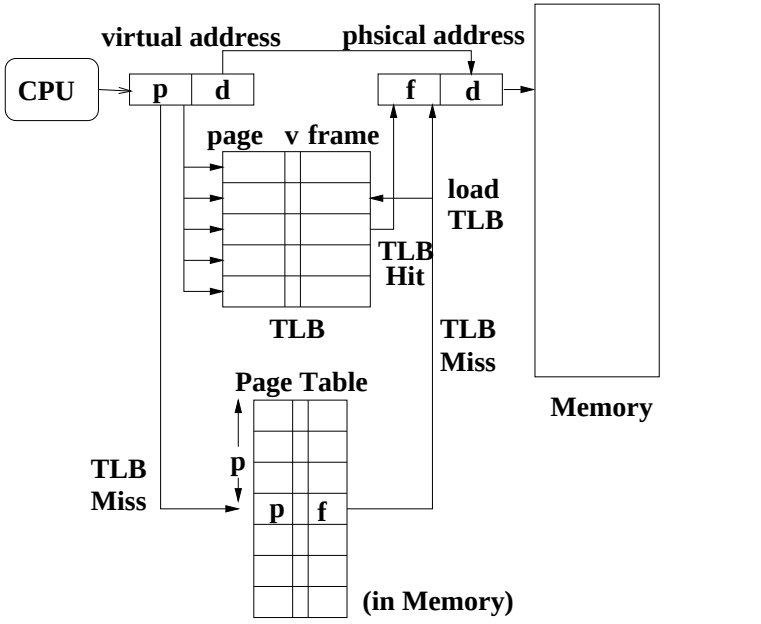

v: valid bit that says the entry is up-to-date

Computing Relative Costs

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?

Initializing Memory when Starting a Process

1. Process needing k pages arrives.
2. If k page frames are free, the OS allocates all k pages to the free frames. Otherwise, we need to free a frame that is no longer needed.
3. The OS puts the first page in a frame, and puts the frame number in the first entry in the page table; it then puts the second page in the frame and its frame number in the second entry of the page table and so on.
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.

Saving/Restoring Memory on a Context Switch

The Process Control Block (PCB) must be extended to contain:

- The page table
- Possibly a copy of the TLB

1. Copy the page table base register value to the PCB.
2. Copy the TLB to the PCB (optional).
3. Flush the TLB.
4. Restore the page table base register.
5. Restore the TLB, if it was saved.

Multilevel Paging

If the virtual address space is huge, page tables get too big, and many systems use a multilevel paging scheme (read about it in your book!)

Sharing

Paging introduces the possibility of sharing since the 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.

$\Longrightarrow$ The expense of memory accesses and the flexibility of paging, make paging cost effective.

Next Time

Combining segmentation (a logical view of programs as heap, code, and stack) and paging.