

CS 162 Operating Systems and Systems Programming

Professor: Anthony D. Joseph

Spring 2004

Lecture 13: Address Translation

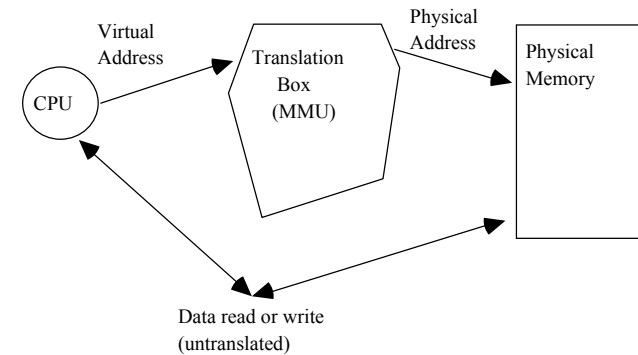
13.0 Main Points

Options for managing memory:

- Paging
- Segmentation
- Multi-level translation
- Paged page tables
- Inverted page tables

Comparison among options

13.1 Hardware Translation Overview




Think of memory in two ways:

- View from the CPU – what program sees, virtual memory
- View from memory – physical memory

Translation implemented in hardware; controlled in software. There are many kinds of hardware translation schemes. Start with the simplest!

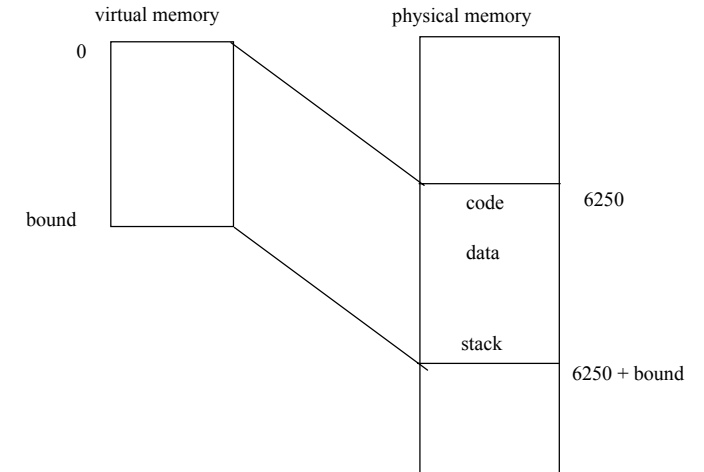
13.2 Base and Bounds

Base and bounds: Each program loaded into contiguous regions of physical memory, but with protection between programs. First built in the Cray-1.



Hardware Implementation of Base and Bounds Translation

Program has illusion it is running on its own dedicated machine, with memory starting at 0 and going up to size = bounds. Like linker-loader, program gets contiguous region of memory. But unlike linker-loader, protection: program can only touch locations in physical memory between base and base + bounds.



Virtual and Physical Memory Views in Base and Bounds System

Provides level of indirection: OS can move bits around behind the program's back, for instance, if program needs to grow beyond its bounds, or if need to coalesce fragments of memory. Stop program, copy bits, change base and bounds registers, restart.

Only the OS gets to change the base and bounds! Clearly, user program can't, or else lose protection.

With base&bounds system, what gets saved/restored on a context switch?

Hardware cost:

- 2 registers
- Adder
- Comparator

Plus, slows down hardware because need to take time to do add/compare on every memory reference.

Base and bounds, Pros:

+ Simple, fast

Cons:

1. Hard to share between programs

For example, suppose two copies of “vi”

- Want to share code
- Want data and stack to be different

Can't do this with base and bounds!

2. Complex memory allocation

First fit, best fit, buddy system. Particularly bad if want address space to grow dynamically (e.g., the heap).


In worst case, have to shuffle large chunks of memory to fit new program.

3. Doesn't allow heap, stack to grow dynamically – want to put these as far apart as possible in virtual memory, so that they can grow to whatever size is needed.

13.3 Segmentation

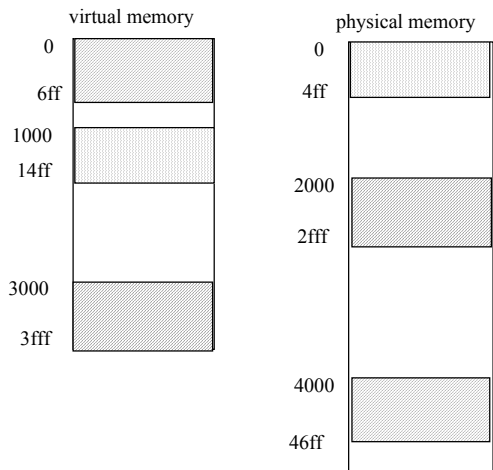
A **segment** is a region of logically contiguous memory.

Idea is to generalize base and bounds, by allowing a **table** of base&bound pairs.



For example, what does it look like with this segment table, in virtual memory and physical memory? Assume a 14 bit addresses divided up as:
2 bit segment ID, and a 12 bit segment offset.

Virtual segment #	Physical segment start	Segment size
0 code	0x4000	0x700
1 data	0	0x500
2 -	0	0
3 stack	0x2000	0x1000



This should seem a bit strange: the virtual address space has gaps in it! Each segment gets mapped to contiguous locations in physical memory, but may be gaps between segments.

But a correct program will never address gaps; if it does, trap to kernel and then core dump. Minor exception: stack, heap can grow. In UNIX, `sbzrk()` increases size of heap segment. For stack, just take fault, system automatically increases size of stack.

Detail: Need protection mode in segmentation table. For example, code segment would be read-only (only execution and loads are allowed). Data and stack segment would be read-write (stores allowed).

What must be saved/restored on context switch? Typically, segment table stored in CPU, not in memory, because it's small.

Example: What happens with the above segment table, with the following as virtual memory contents? Code does:

```
strlen(x);
```

Virtual memory	
Main: 240	store 1108, r2
244	store pc +8, r31
248	jump 360
24c	...
...	
Strlen: 360	loadbyte (r2), r3
...	
420	jump (r31)
...	
x: 1108	a b c \0
...	

Initially, pc = 240.

Physical Memory	
x: 108	666
...	
Main: 4240	store 1108, r2
4244	store pc +8, r31
4248	jump 360
424c	...
...	
Strlen: 4360	loadbyte (r2), r3
...	
420	jump (r31)

Segmentation Pros & Cons:

- + Efficient for sparse address spaces
- + Easy to share whole segments (for example, code segment)
- Complex memory allocation

Still need first fit, best fit, etc., and re-shuffling to coalesce free fragments, if no single free space is big enough for a new segment.

How do we make memory allocation simple and easy?

13.4 Paging

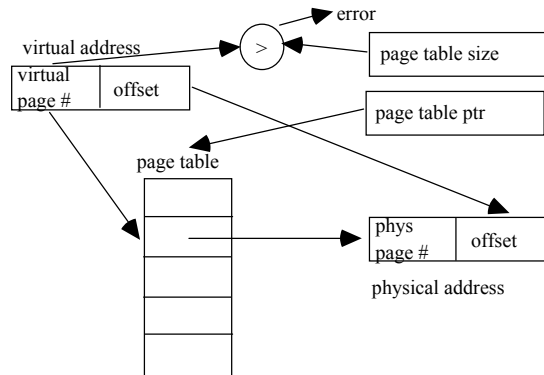
Allocate physical memory in terms of fixed size chunks of memory, or **pages**.

Simpler, because allows use of a bitmap. What's a bitmap?

00111110000001100

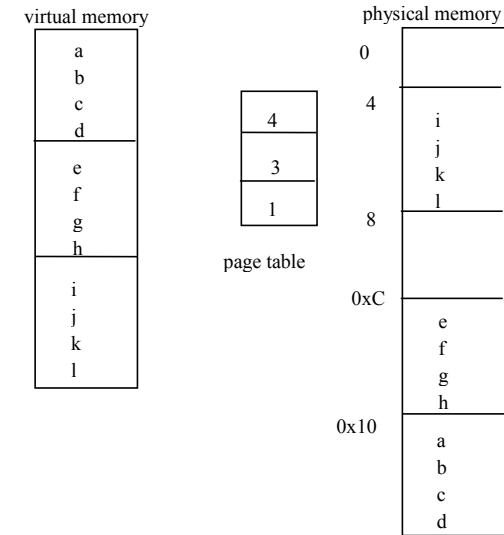
Each bit represents one page of physical memory – 1 means allocated, 0 means unallocated. Lots simpler than base&bounds or segmentation

Operating system controls mapping: any page of virtual memory can go anywhere in physical memory.



Each address space has its own page table, in physical memory. Hardware needs two special registers – pointer to physical location of page table, and page table size.

Example: suppose page size is 4 bytes.



Where is virtual address 6? 9?

Questions:

1. What must be saved and restored on a context switch?
2. What if page size is very small? For example, VAX had a page size of 512 bytes.
3. What if page size is really big? Why not use an infinite page size?

Fragmentation: wasted space

- **External** – free gaps between allocated chunks
- **Internal** – free gaps because don't need all of allocated chunk

With segmentation need to re-shuffle segments to avoid external fragmentation. Paging suffers from internal fragmentation.

4. What if address space is sparse? For example: on UNIX, code starts at 0, stack starts at $2^{31} - 1$. With 1KB pages, 2 million page table entries!

Paging Pros&Cons:

- + simple memory allocation
- + easy to share
- big page tables if sparse address space

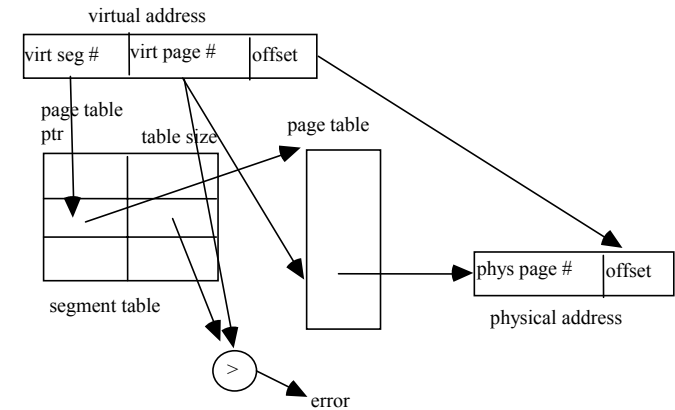
Is there a solution that allows simple memory allocation, easy to share memory, **and** is efficient for sparse address spaces?

How about combining paging and segmentation?

13.5 Multi-level translation

Multi-level translation. Use tree of tables. Lowest level is page table, so that physical memory can be allocated using a bitmap. Higher levels are typically segmented.

For example, here's a picture of a 2-level translation scheme:



Just like recursion – could have any number of levels. Most architectures today do this.

Questions:

1. What must be saved/restored on context switch?
2. How do we share memory? Can share entire segment, or a single page.

Example: Suppose we have 24 bit virtual addresses partitioned as 4 bits of segment #, 8 bits of virtual page #, and 12 bits of offset.

Segment Table	
Page table pointer	Page table size
0x2000	0x14
-	-
0x1000	0xD
-	-

Physical Memory	
0x1000	0x6
	0xb
	0x4
	...
0x2000	0x13
	0x2a
	0x3
	...

(the above are portions of the page tables for the segments)

What do the following addresses translate to?

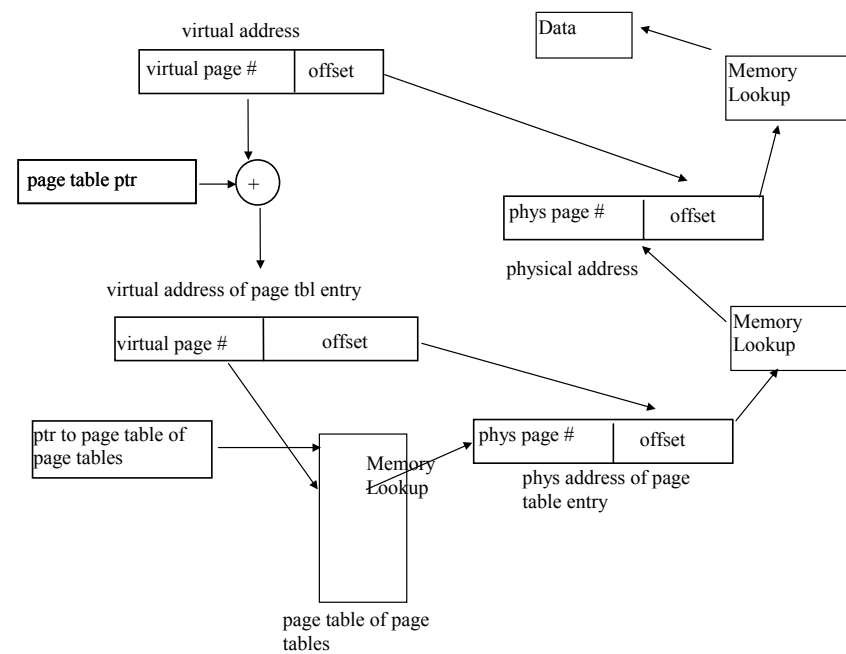
- 0x002070?
- 0x201016 ?
- 0x14c684 ?
- 0x210014 ?

Multilevel translation:

- + Only need to allocate as many page table entries as we need.
In other words, sparse address spaces are easy.
- + Easy memory allocation
- + Share at segment or page level (need additional reference counting)
- Pointer per page (typically 4KB - 16KB pages today)
- Page tables need to be contiguous
- Two (or more, if > 2 levels) lookups per memory reference

13.6 Paged page tables

A different solution to sparse address spaces is to allow the page tables to be paged – only need to allocate physical memory for page table entries you really use. Top level page table is in physical memory, all lower levels of hierarchy are in virtual memory (and therefore can be allocated in fixed size page frames in physical memory).



This means that potentially, each memory reference involves three memory references (one for the system page table, one for the user page table, and one for the real data).

How do we reduce the overhead of translation?

Caching in Translation Lookaside Buffers (TLB's).

Relative to multilevel translation, paged page tables are more efficient if using a TLB. If virtual address of page table entry is in TLB, can skip one or more levels of translation.

13.7 Inverted page tables

What is an efficient data structure for doing lookups? Hash table. Why not use a hash table to translate from virtual address to a physical address.

This is called an inverted page table for historical reasons.

Take virtual page #, run hash function on it, index into hash table to find page table entry with physical page frame #.

Independent of size of address space,

Advantages:

- + $O(1)$ lookup to do translation
- + Requires page table space proportional to how many pages are actually being used, not proportional to size of address space – with 64 bit address spaces, this is a big win!
- Overhead of managing hash chains, etc.