# Memory Management

## References:

- https://en.wikipedia.org/wiki/Page (computer memory)
- <u>https://en.wikipedia.org/wiki/Page\_table</u>
- https://en.wikipedia.org/wiki/Virtual memory

Kartik Gopalan

• Chapter 3, Modern Operating Systems, Andrew S. Tanenbaum

## Memory Management

## • Ideally programmers want memory that is

- large
- fast
- persistent (non-volatile)

# Memory Hierarchy

- Registers & Cache
  - small amount of fast, expensive, volatile memory
- Main memory
  - some medium-speed, medium price, volatile/persistent memory
- Disk & Tape
  - Lots of slow, cheap, persistent, storage

## Typical access time



- Consider an instruction that reads from a memory location
  - load reg memory address
  - But programmer doesn't know the memory address where data will be stored when the process runs!
- **Solution: Relocation** 
  - Programmer assumes a "relative" address, which is converted to a "physical" address by the OS+hardware when the process runs.

## Protection

• Access to address locations larger than **limit value** results in an error

## **Relocation and Protection**



BASE

**Relocated Addresses** in Executing Binary





## Swapping and External Fragmentation



- Memory allocation changes as
  - processes come into memory
  - leave memory and are *swapped out* to disk
  - Re-enter memory by getting *swapped-in* from disk
- Shaded regions are unused memory
- memory allocations.

• Physical memory may not be enough to accommodate the needs of all processes

• **External Fragmentation** is when we have free memory that is too small for

## Virtual Memory

- Virtual memory: means that <u>each process gets</u> <u>an illusion that it has its own memory space</u> <u>whose size is independent of the size of</u> <u>physical RAM in the system</u>.
- How? Break up the memory space of a process into equal-sized <u>PAGES</u>.
  - Typically, a page = 4KiB
- Memory is allocated to processes at the granularity of pages
  - E.g. 4KiB, 8KiB, 12KiB etc.
- OS then decides which pages stay in memory and which get paged (moved) out to disk.



## Internal Fragmentation

- <u>Internal Fragmentation</u> occurs when some of part of allocated memory is wasted.
- E.g. malloc() of 100 bytes might fetch a 4KiB page from OS.
  - Then 4KiB-100 bytes might be wasted, unless used by future malloc() operations
- <u>Virtual Memory introduces internal fragmentation</u> • Larger the page size, more internal fragmentation
- <u>Virtual Memory eliminates external fragmentation</u> • All memory allocations occur at the granularity of page size, so no small unused memory fragments are left
- around.

## Memory Management Unit (MMU)



The MMU sends physical addresses to the memory

• MMU is a hardware module that accompanies the CPU • It translates the Virtual Address used by executing instructions to Physical Addresses in the main memory.

# Size of address space (in bytes) as a function of address size (in bits)

Number of bits in address	Maximum address space size (bytes)	
0	2 <sup>0</sup> = 1 byte	
1	2 <sup>1</sup> = 2 bytes	
2	2 <sup>2</sup> = 4 bytes	
10	$2^{10} = 1024 = 1$ KiB	
12	$2^{12} = 4$ KiB	
16	2 <sup>16</sup> = 64 KiB	
32	2 <sup>32</sup> = 4GiB (Gibibytes)	
64	2 <sup>64</sup> = 16 EiB (Exbibytes)	

# Page Table

- An array that stores the mapping from virtual page numbers to physical numbers
- The OS maintains
  - <u>One page table per userspace process</u>.
  - And usually another page table for kernel memory.



10



Translating Virtual address (VA) to physical address (PA)

## Virtual Address Translation For Small Address Space



Internal operation of MMU with 16 4 KB pages

12

## Virtual Address Translation For Large Address Space





- 32 bit address with 2 page table fields
- Two-level page tables
- PT too Big for MMU
  - Keep it in main memory
- But how does MMU know where to find PT?
  - Registers (CR2 on Intel)

# Typical Page Table Entry (PTE)



- represented by the PTE
- bit was reset.
- since the last time the bit was reset.
- Protection bits: Whether the page is readable? writeable? executable? contains higher privilege code/data?
- number. Used for marking swapped/unallocated pages.

Present/absent

Page frame number

• Page Frame number = physical page number for the virtual page

Referenced bit: Whether the page was accessed since last time the

Modified bit: Also called "Dirty" bit. Whether the page was written to,

• Present/Absent bit: Whether the PTE contains a valid page frame

## TLB – Translation Lookaside Buffer

Valid	Virtual page	Modified	Protection	Page frame
1	140	1	RW	31
1	20	0	RX	38
1	130	1	RW	29
1	129	1	RW	62
1	19	0	RX	50
1	21	0	RX	45
1	860	1	RW	14
1	861	1	RW	75

- need to access the page-table in memory too often.
- TLB is part of the MMU hardware (comes with CPU)
- It is not a Data Cache or anInstruction Cache. Those are separate.
- field in TLB to identify the process context.
  - Tagged TLB can reduce this overhead

• TLB is a small cache that speeds up the translation of virtual addresses to physical addresses.

• <u>TLB caches translations from virtual page number to physical page number so that the MMU doesn't</u>

• On older x86 processors, TLB had to be "flushed" upon every context switch because there was no

## Impact of Page Size on Page tables

Small page size

- Advantages
  - less internal fragmentation
  - page-in/page-out less expensive

- Disadvantages

  - Smaller "TLB Coverage" (discussed later)

• process that needs more pages has larger page table

## Bit distribution in a memory address

- For a 32 bit address and ...
  - 4KiB page
    - 12 bit offset and 20 bit page number
  - 8KiB page
    - 13 bit offset and 19 bit page number
  - 64KiB page
    - 16 bit offset and 16 bit page number

- Consider a machine that has a 32-bit virtual address space and 8KiByte page size.
- 1. What is the total size (in bytes) of the virtual address space for each process?
- 2. How many bits in a 32-bit address are needed to determine the page number of the address?
- 3. How many bits in a 32-bit address represent the byte offset into a page?
- 4. How many page-table entries are present in the page table?



# Quiz Answers

- page size.
- = 4 \* 1024 \* 1024 \* 1024 bytes = 4 GiB
- $2^{9}*2^{10} = 2^{19}$
- - $\log 2(8 \text{KiB}) = \log 2(2^{13}) = 13$  bits
  - Also, 32 19 = 13 bits
- 4. How many page-table entries are present in the page table?
  - pages

Consider a machine that has a 32-bit virtual address space and 8KiByte

1. Total size (in bytes) of the virtual address space for each process =  $2^{32}$ 

2. Number of pages in virtual address space = 4GiB/8KiB = 512\*1024 =

So the number of bits in a 32-bit address are needed to determine the page number of the address =  $log2(4GiB/8KiB) = log2(2^19) = 19$  bits

3. How many bits in a 32-bit address represent the byte offset into a page?

Number of PTEs = Number of pages in virtual address =  $4GiB/8KiB = 2^{19}$ 

## References

- Chapter 3: Modern Operating Systems, Andrew S. Tanenbaum
- X86 architecture http://en.wikipedia.org/wiki/X86
- Memory segment http://en.wikipedia.org/wiki/Memory segment
- Memory model http://en.wikipedia.org/wiki/Memory model
- Architecture

• IA-32 Intel Architecture Software Developer's Manual, Volume 1: Basic