


# TCSS 422: OPERATING SYSTEMS

## Introduction to Paging

Wes J. Lloyd  
Institute of Technology  
University of Washington - Tacoma



## OBJECTIVES

- Paging
- Address translation
- Paging questions

November 14, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L16.2

## PAGING

- Split up address space of process into fixed sized pieces called **pages**
- Alternative to variable sized pieces (Segmentation) which suffers from significant fragmentation
- Physical memory is split up into an array of fixed-size slots called **page frames**.
- Each process has a **page table** which translates virtual addresses to physical addresses

November 14, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L16.3

## ADVANTAGES OF PAGING

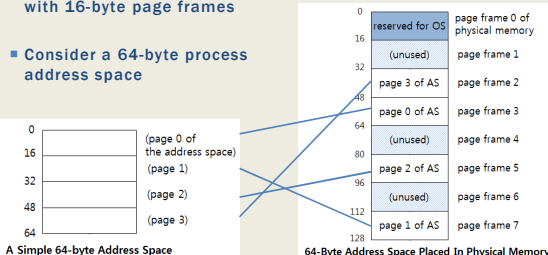
- Flexibility
  - Abstracts the process address space into pages
  - No need to track direction of HEAP / STACK growth
  - No need to store unused space
- Simplicity
  - Pages and page frames are the same size
  - Easy to allocate and keep a free list of pages

November 14, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L16.4

## PAGING: EXAMPLE

**Page Table:**  
VP0 → PF3  
VP1 → PF7  
VP2 → PF5  
VP3 → PF2

- Consider a 128 byte address space with 16-byte page frames
- Consider a 64-byte process address space



A Simple 64-byte Address Space

64-Byte Address Space Placed In Physical Memory

November 14, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L16.5

## PAGING: ADDRESS TRANSLATION

- Two address components
  - VPN: Virtual Page Number
  - Offset: Offset within a Page

VPN				offset			
Va5	Va4	Va3	Va2	Va1	Va0		

■ Example:  
Page Size: 16-bytes, Address Space: 64-bytes

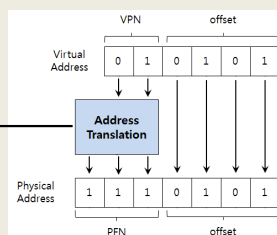
VPN				offset			
0	1	0	1	0	1		

November 14, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L16.6

## EXAMPLE: PAGING ADDRESS TRANSLATION

- Consider a 64-byte address space
- Stored in 128-byte physical memory
- Offset is preserved
- VPN is looked up

**Page Table:**  
VP0 → PF3  
VP1 → PF7  
VP2 → PF5  
VP3 → PF2



November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.7

## PAGING QUESTIONS

- Where are page tables stored?
- What are the typical contents of the page table?
- How big are page tables?
- Does paging make the system too slow?

November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.8

## WHERE ARE PAGE TABLES STORED?

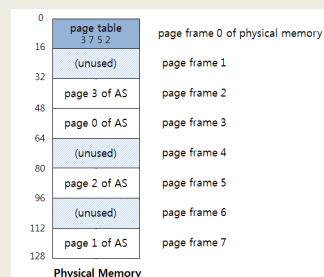
- Real world size example:
- Consider a 32-bit process address space (up to 4GB)
- With 4 KB pages
- 20 bits for VPN
- 12 bits for the page offset
- Page tables for each process are stored in RAM
  - Support potential storage of  $2^{20}$  translations
  - = 1,048,576 pages per process
  - Each page has a page table entry size of 4 bytes
- Consider 100+ OS processes
  - Requires 400+ MB of RAM to store process information

November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.9

## PAGE TABLE WITH KERNEL PHYSICAL MEMORY



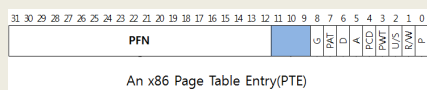
November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.10

## WHAT'S ACTUALLY IN THE PAGE TABLE

- Page table is data structure used to map virtual page numbers (VPN) to the physical address (Physical Frame Number PFN)
  - Linear page table → simple array
- Page-table entry
  - 32 bits for capturing state



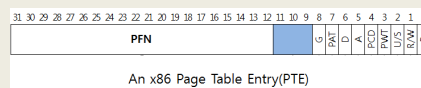
November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.11

## PAGE TABLE ENTRY

- P: present
- R/W: read/write bit
- U/S: supervisor
- A: accessed bit
- D: dirty bit
- PFN: the page frame number



November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.12

## PAGE TABLE ENTRY - 2

- Common flags:
- Valid Bit:** Indicating whether the particular translation is valid.
- Protection Bit:** Indicating whether the page could be read from, written to, or executed from
- Present Bit:** Indicating whether this page is in physical memory or on disk(swapped out)
- Dirty Bit:** Indicating whether the page has been modified since it was brought into memory
- Reference Bit(Accessed Bit):** Indicating that a page has been accessed

November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.13

## HOW BIG ARE PAGE TABLES?

- Page tables are too big to store on the CPU
- Page tables are stored using physical memory
- Paging supports efficiently storing a sparsely populated address space
  - Reduced memory requirement  
Compared to base and bounds, and segments

November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.14

## DOES PAGING MAKE THE SYSTEM TOO SLOW?

- Translation
- Issue #1:** Starting location of the page table is needed
  - HW Support: Page-table base register
    - stores active process
    - Facilitates translation
- Issue #2:** Each memory address translation for paging requires an extra memory reference
  - HW Support: TLBs (Chapter 19)

**Page Table:**  
VP0 → PF3  
VP1 → PF7  
VP2 → PF5  
VP3 → PF2

Stored in RAM →

November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.15

## PAGING MEMORY ACCESS

```

1. // Extract the VPN from the virtual address
2. VPN = (VirtualAddress & VPN_MASK) >> SHIFT
3.
4. // Form the address of the page-table entry (PTE)
5. PTEAddr = PTBR + (VPN * sizeof(PTE))
6.
7. // Fetch the PTE
8. PTE = AccessMemory(PTEAddr)
9.
10. // Check if process can access the page
11. if (PTE.Valid == False)
12.     RaiseException(SEGMENTATION_FAULT)
13. else if (CanAccess(PTE.ProtectBits) == False)
14.     RaiseException(PROTECTION_FAULT)
15. else
16.     // Access is OK: form physical address and fetch it
17.     offset = VirtualAddress & OFFSET_MASK
18.     PhysAddr = (PTE.PFN << PFN_SHIFT) | offset
19.     Register = AccessMemory(PhysAddr)
    
```

November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.16

## COUNTING MEMORY ACCESSES

- Example: Use this Array initialization Code

```

int array[1000];
...
for (i = 0; i < 1000; i++)
    array[i] = 0;
    
```

- Assembly equivalent:

```

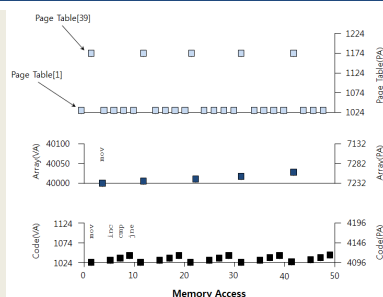
0x1024 movl $0x0, (%edi,%eax,4)
0x1028 incl %eax
0x102c cmpl $0x03e8, %eax
0x1030 jne 0x1024
    
```

November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.17

## MEMORY ACCESSES: FOR THE FIRST 5 LOOP ITERATIONS



November 14, 2016

TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma

L16.18

