

TCSS 422: OPERATING SYSTEMS


Memory Virtualization IV: Translation Lookaside Buffer (TLB), Smaller Tables, Multi-Level Page Tables, Beyond Physical Memory

Wes J. Lloyd
School of Engineering and Technology
University of Washington - Tacoma

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington

Tacoma



1

OBJECTIVES – 5/29

■ Questions from 5/27

■ Memory Segmentation Activity + answers (available in Canvas)

■ Assignment 2 - June 5 AOE

■ Assignment 3 (as a Tutorial) - June 10 AOE

■ Final exam – Thursday June 12 @ 3:40pm

■ Quiz 4 – Page Tables - Due June 8 AOE

■ Chapter 19: Translation Lookaside Buffer (TLB)

■ TLB Algorithm, Hit-to-Miss Ratios

■ Chapter 20: Paging: Smaller Tables

■ Smaller Tables, Multi-level Page Tables, N-level Page Tables

■ Chapter 21/22: Beyond Physical Memory

■ Swapping Mechanisms, Swapping Policies

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.2

2

ONLINE DAILY FEEDBACK SURVEY

■ Daily Feedback Quiz in Canvas – Available After Each Class

■ Extra credit available for completing surveys **ON TIME**

■ Tuesday surveys: due by ~ Wed @ 11:59p

■ Thursday surveys: due ~ Mon @ 11:59p

TCSS 422 A > Assignments

Spring 2021

Search for Assignment

Home

Announcements

Zoom

Syllabus

Assignments

Discussions

Upcoming Assignments

TCSS 422 - Online Daily Feedback Survey - 4/1

Available until Apr 5 at 11:59pm | Due Apr 5 at 10pm | -/1 pts

May 29, 2025

TCSS422: Computer Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.3

3

TCSS 422 - Online Daily Feedback Survey - 4/1

Quiz Instructions

Question 1

0.5 pts

On a scale of 1 to 10, please classify your perspective on material covered in today's class:

12345678910

Mostly Review To MeEqual New and ReviewMostly New to Me

Question 2

0.5 pts

Please rate the pace of today's class:

12345678910

SlowJust RightFast

May 29, 2025

TCSS422: Computer Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.4

4

Slides by Wes J. Lloyd

L17.2

MATERIAL / PACE

- Please classify your perspective on material covered in today’s class (40 of 63 respondents – 63.5 %):
 - 1-mostly review, 5-equal new/review, 10-mostly new
 - **Average – 6.00** (↓ - **previous 6.27**)
- Please rate the pace of today’s class:
 - 1-slow, 5-just right, 10-fast
 - **Average – 5.28** (↑ - **previous 5.17**)

May 29, 2025	TCSS422: Computer Operating Systems [Spring 2025] School of Engineering and Technology, University of Washington - Tacoma	L17.5
--------------	--	-------

5

FEEDBACK FROM 5/29

May 29, 2025	TCSS422: Operating Systems [Spring 2025] School of Engineering and Technology, University of Washington - Tacoma	L17.6
--------------	---	-------

6

FEEDBACK - 2

- Some tips for problems with exponential math and bits:
- >>> It can be helpful to review charts and patterns:
- 8 bits = 1 byte
- 16 bits = 2 bytes
- 32 bits = 4 bytes
- 64 bits = 8 bytes
- 1,024 bytes = 1 kilobyte (2^{10})
- 1,024 kilobytes = 1 megabyte (2^{20})
- 1,024 megabytes = 1 gigabyte (2^{30})
- 1,024 gigabytes = 1 terabyte (2^{40})
- 1,024 terrabytes = 1 petabyte (2^{50})

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.7

7

FEEDBACK - 3

- For simplicity rounding is often acceptable:
- 1 kilobyte (2^{10}) = 1,024 bytes → **1,000 bytes**
- 1,024 kilobytes (2^{20}) = 1 megabyte → **1,000,000 bytes**
- 1,024 megabytes = 1 gigabyte (2^{30})→1,000,000,000 bytes
- 1,024 gigabytes = 1 terabyte (2^{40})→1,000,000,000,000 bytes
- 1,024 terrabytes = 1 petabyte (2^{50})→1,000,000,000,000,000 bytes

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.8

8

2 ¹	2	2 ¹⁷	131,072	2 ³³	8,589,934,592	2 ⁴⁹	562,949,953,421,312
2 ²	4	2 ¹⁸	262,144	2 ³⁴	17,179,869,184	2 ⁵⁰	1,125,899,906,842,624
2 ³	8	2 ¹⁹	524,288	2 ³⁵	34,359,738,368	2 ⁵¹	2,251,799,813,685,248
2 ⁴	16	2 ²⁰ megabyte	1,048,576	2 ³⁶	68,719,476,736	2 ⁵²	4,503,599,627,370,496
2 ⁵	32	2 ²¹	2,097,152	2 ³⁷	137,438,953,472	2 ⁵³	9,007,199,254,740,992
2 ⁶	64	2 ²²	4,194,304	2 ³⁸	274,877,906,944	2 ⁵⁴	18,014,398,509,481,984
2 ⁷	128	2 ²³	8,388,608	2 ³⁹	549,755,813,888	2 ⁵⁵	36,028,797,018,963,968
2 ⁸	256	2 ²⁴	16,777,216	2 ⁴⁰ terabyte	1,099,511,627,776	2 ⁵⁶	72,057,594,037,927,936
2 ⁹	512	2 ²⁵	33,554,432	2 ⁴¹	2,199,023,255,552	2 ⁵⁷	144,115,188,075,855,872
2 ¹⁰ kilobyte	1,024	2 ²⁶	67,108,864	2 ⁴²	4,398,046,511,104	2 ⁵⁸	288,230,376,151,711,744
2 ¹¹	2,048	2 ²⁷	134,217,728	2 ⁴³	8,796,093,022,208	2 ⁵⁹	576,460,752,303,423,488
2 ¹²	4,096	2 ²⁸	268,435,456	2 ⁴⁴	17,592,186,044,416	2 ⁶⁰	1,152,921,504,606,846,976
2 ¹³	8,192	2 ²⁹	536,870,912	2 ⁴⁵	35,184,372,088,832	2 ⁶¹	2,305,843,009,213,693,952
2 ¹⁴	16,384	2 ³⁰ gigabyte	1,073,741,824	2 ⁴⁶	70,368,744,177,664	2 ⁶²	4,611,686,018,427,387,904
2 ¹⁵	32,768	2 ³¹	2,147,483,648	2 ⁴⁷	140,737,488,355,328	2 ⁶³	9,223,372,036,854,775,808
2 ¹⁶	65,536	2 ³²	4,294,967,296	2 ⁴⁸	281,474,976,710,656	2 ⁶⁴ hubbabyte	18,446,744,073,709,551,616

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.9

9

FEEDBACK - 4

- How many bits are required to index the following amounts of memory?
 - 1,024 bytes = 1 kilobyte
 - (2¹⁰) → 10 bits
 - 1,024 kilobytes = 1 megabyte
 - (2²⁰) → 20 bits
 - 1,024 megabytes = 1 gigabyte
 - (2³⁰) → 30 bits
 - 1,024 gigabytes = 1 terabyte
 - (2⁴⁰) → 40 bits
 - 1,024 terrabytes = 1 petabyte
 - (2⁵⁰) → 50 bits

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.10

10

FEEDBACK - 5

- *With paging, we divide an address space in fixed sized pieces (known as the page size)*
- *Assuming a computer indexes memory using 1 kilobyte memory pages (2^{10})*
- *How many unique pages are required to manage/index memory?*
- 1 kilobyte (2^{10}) of memory
 - 1 page
- 1 megabyte (2^{20}) of memory
 - 1024 pages (2^{10})
- 1 gigabyte (2^{30}) of memory
 - 1,048,576 pages (2^{20})
- 1 terabyte (2^{40}) of memory
 - 1,073,741,824 pages (2^{30})
- 1 petabyte (2^{50}) of memory
 - 1,099,511,627,776 pages (2^{40})

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.11

11

OBJECTIVES – 5/29

- Questions from 5/27
- **Memory Segmentation Activity + answers (available in Canvas)**
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.12

12

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- **Assignment 2 - June 5 AOE**
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.13

13

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- **A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE**
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.14

14

ASSIGNMENT 3: INTRODUCTION TO LINUX KERNEL MODULES

- Assignment 3 provides an introduction to kernel programming by demonstrating how to create a Linux Kernel Module
- Kernel modules are commonly used to write device drivers and can access protected operating system data structures
 - For example: Linux `task_struct` process data structure

May 29, 2025	TCSS422: Operating Systems [Spring 2025] School of Engineering and Technology, University of Washington - Tacoma	L17.15
--------------	---	--------

15

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- **Final exam – Thursday June 12 @ 3:40pm**
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies

May 29, 2025	TCSS422: Operating Systems [Spring 2025] School of Engineering and Technology, University of Washington - Tacoma	L17.16
--------------	---	--------

16

FINAL EXAM – THURSDAY JUNE 12 @
3:40PMTH

- Thursday June 12 from 3:40 to 5:40 pm
 - Final (100 points)
 - Similar number of questions as the midterm
 - 2-hours
 - Focus on new content - since the midterm (~70% new, 30% before)
- Final Exam Review -
 - Complete Memory Segmentation Activity
 - Complete Quiz 4
 - Practice Final Exam Questions – 2nd hour of June 5th class session
 - Individual work
 - 3 pages of notes (any sized paper), double sided
 - Basic calculators allowed
 - NO smartphones, laptop, book, Internet, group work

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.17

17

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.18

18

CATCH UP FROM LECTURE 16

- Switch to Lecture 16 Slides
- Slides L18.54 to L18.61
(Chapter 18 – Introduction to Paging)

May 29, 2025	TCSS422: Operating Systems [Spring 2025] School of Engineering and Technology, University of Washington - Tacoma	L17.19
--------------	---	--------


19

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies

May 29, 2025	TCSS422: Operating Systems [Spring 2025] School of Engineering and Technology, University of Washington - Tacoma	L17.20
--------------	---	--------

20



CHAPTER 19: TRANSLATION LOOKASIDE BUFFER (TLB)

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.21

21

TRANSLATION LOOKASIDE BUFFER

- Legacy name...
- Better name, “Address Translation Cache”
- TLB is an on CPU cache of address translations
 - virtual → physical memory

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.22

22

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
```

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.23

23

VISUALIZING MEMORY ACCESSES:
FOR THE FIRST 5 LOOP ITERATIONS

■ Locations:

■ Page table

■ Array

■ Code

■ 50 accesses
for 5 loop
iterations

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.24

24

TRANSLATION LOOKASIDE BUFFER - 2

■ Goal:
Reduce access
to the page
tables

■ Example:
50 RAM accesses
for first 5 for-loop
iterations

■ Move lookups
from RAM to TLB
by caching page
table entries

The top graph shows Page Table access (PA) vs Memory Access (0-50). It highlights two specific entries: Page Table[1] and Page Table[39]. The middle graph shows Array access (VA) vs Memory Access (0-50) with labels 'mov' at indices 0, 10, 20, 30, and 40. The bottom graph shows Code access (VA) vs Memory Access (0-50) with labels 'mov', 'inc', 'cmp', and 'jne' at indices 0, 1, 2, and 3 respectively.

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.25

25

TRANSLATION LOOKASIDE BUFFER (TLB)

■ Part of the CPU's Memory Management Unit (MMU)

■ Address translation cache

The diagram illustrates the address translation process. A CPU provides a Logical Address. This address is used for a TLB Lookup in the MMU. If it's a TLB Hit, the Physical Address is retrieved directly from the MMU. If it's a TLB Miss, the MMU consults the Page Table (which contains all v to p entries) to find the Physical Address. This Physical Address is then used to access Physical Memory, which is organized into pages (Page 0, Page 1, Page 2, ..., Page n).

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

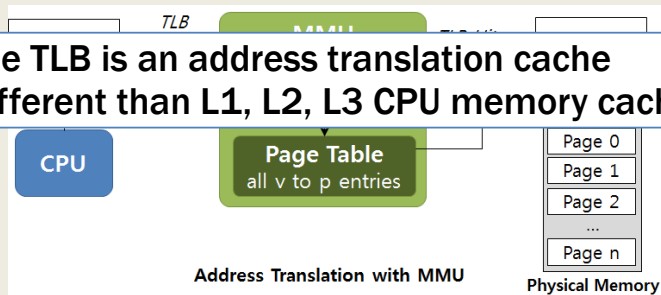
L17.26

26

TRANSLATION LOOKASIDE BUFFER (TLB)

- Part of the CPU's Memory Management Unit (MMU)
- Address translation cache

The TLB is an address translation cache
Different than L1, L2, L3 CPU memory caches



May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.27

27

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - **TLB Algorithm** Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.28

28

TLB BASIC ALGORITHM

- For: array based page table
- Hardware managed TLB

```

1: VPN = (VirtualAddress & VPN_MASK ) >> SHIFT
2: (Success , TlbEntry) = TLB_Lookup(VPN)
3: if(Success == True){ // TLB Hit
4:   if(CanAccess(TlbEntry.ProtectBits) == True ){
5:     Offset = VirtualAddress & OFFSET_MASK
6:     PhysAddr = (TlbEntry.PFN << SHIFT) | Offset
7:     AccessMemory( PhysAddr )
8:   }else RaiseException( PROTECTION_ERROR)

```

Generate the physical address to access memory

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.29

29

TLB BASIC ALGORITHM - 2

```

11: else{ //TLB Miss
12:   PTEAddr = PTBR + (VPN * sizeof(PTE))
13:   PTE = AccessMemory(PTEAddr)
14:   (...) // Check for, and raise exceptions...
15:
16:   TLB_Insert( VPN , PTE.PFN , PTE.ProtectBits)
17:   RetryInstruction ()
18: }
19:}

```

Retry the instruction... (requery the TLB)

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.30

30

TLB – ADDRESS TRANSLATION CACHE

- Key detail:
- For a TLB miss, we first access the page table in RAM to populate the TLB... we then requery the TLB
- All address translations go through the TLB

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.31

31

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, **Hit-to-Miss Ratios**
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.32

32

TLB EXAMPLE

```
0:  int sum = 0 ;
1:  for( i=0; i<10; i++){
2:      sum+=a[i];
3:  }
```

- Example:
- Program address space: 256-byte
 - Addressable using 8 total bits (2^8)
 - 4 bits for the VPN (16 total pages)
- Page size: 16 bytes
 - Offset is addressable using 4-bits
- Store an array: of (10) 4-byte integers

	00	04	08	12	16
VPN = 00					
VPN = 01					
VPN = 03					
VPN = 04					
VPN = 05					
VPN = 06		a[0]	a[1]	a[2]	
VPN = 07	a[3]	a[4]	a[5]	a[6]	
VPN = 08	a[7]	a[8]	a[9]		
VPN = 09					
VPN = 10					
VPN = 11					
VPN = 12					
VPN = 13					
VPN = 14					
VPN = 15					

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.33

33

TLB EXAMPLE - 2

```
0:  int sum = 0 ;
1:  for( i=0; i<10; i++){
2:      sum+=a[i];
3:  }
```

- Consider the code above:
- Initially the TLB does not know where a[] is
- Consider the accesses:
 - a[0], a[1], a[2], a[3], a[4], a[5], a[6], a[7], a[8], a[9]
- How many pages are accessed?
- What happens when accessing a page not in the TLB?

	00	04	08	12	16
VPN = 00					
VPN = 01					
VPN = 03					
VPN = 04					
VPN = 05					
VPN = 06		a[0]	a[1]	a[2]	
VPN = 07	a[3]	a[4]	a[5]	a[6]	
VPN = 08	a[7]	a[8]	a[9]		
VPN = 09					
VPN = 10					
VPN = 11					
VPN = 12					
VPN = 13					
VPN = 14					
VPN = 15					

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.34

34

TLB EXAMPLE - 3

```
0:  int sum = 0 ;
1:  for( i=0; i<10; i++){
2:      sum+=a[i];
3:  }
```

- For the accesses: a[0], a[1], a[2], a[3], a[4], a[5], a[6], a[7], a[8], a[9]
- How many are hits?
- How many are misses?
- What is the hit rate? (%)
 - 70% (3 misses one for each VP, 7 hits)

	00	04	08	12	16
VPN = 00					
VPN = 01					
VPN = 03					
VPN = 04					
VPN = 05					
VPN = 06		a[0]	a[1]	a[2]	
VPN = 07	a[3]	a[4]	a[5]	a[6]	
VPN = 08	a[7]	a[8]	a[9]		
VPN = 09					
VPN = 10					
VPN = 11					
VPN = 12					
VPN = 13					
VPN = 14					
VPN = 15					

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.35

35

TLB EXAMPLE - 4

```
0:  int sum = 0 ;
1:  for( i=0; i<10; i++){
2:      sum+=a[i];
3:  }
```

- What factors affect the hit/miss rate?
 - Page size
 - Data/Access locality (how is data accessed?)
 - Sequential array access vs. random array access
 - Temporal locality
 - Size of the TLB cache (how much history can you store?)

	00	04	08	12	16
VPN = 00					
VPN = 01					
VPN = 03					
VPN = 04					
VPN = 05					
VPN = 06		a[0]	a[1]	a[2]	
VPN = 07	a[3]	a[4]	a[5]	a[6]	
VPN = 08	a[7]	a[8]	a[9]		
VPN = 09					
VPN = 10					
VPN = 11					
VPN = 12					
VPN = 13					
VPN = 14					
VPN = 15					

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.36

36

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies


May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.37

37

CHAPTER 20:
PAGING:
SMALLER TABLES



May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.38

38

LINEAR PAGE TABLES

- Consider array-based page tables:
 - Each process has its own page table
 - 32-bit process address space (up to 4GB)
 - With 4 KB pages
 - 20 bits for VPN
 - 12 bits for the page offset

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.39

39

LINEAR PAGE TABLES - 2

- Page tables stored in RAM
- Support potential storage of 2^{20} translations
= 1,048,576 pages per process @ 4 bytes/page
- Page table size 4MB / process

$$\text{Page table size} = \frac{2^{32}}{2^{12}} * 4\text{Byte} = 4\text{MByte}$$

- Consider 100+ OS processes
 - Requires 400+ MB of RAM to store process information

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.40

40

LINEAR PAGE TABLES - 2

- Page tables stored in RAM
- Support potential storage of 2^{20} translations
= 1,048,576 pages per process @ 4 bytes/page
- Page table size 4MB / process

Page tables are too big and
consume too much memory.

Need Solutions ...

- Consider 100+ OS processes
 - Requires 400+ MB of RAM to store process information

May 29, 2025	TCSS422: Operating Systems [Spring 2025] School of Engineering and Technology, University of Washington - Tacoma	L17.41
--------------	---	--------

41

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - **Smaller Tables** Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies

May 29, 2025	TCSS422: Operating Systems [Spring 2025] School of Engineering and Technology, University of Washington - Tacoma	L17.42
--------------	---	--------

42

PAGING: USE LARGER PAGES

- **Larger pages** = 16KB = 2^{14}
- 32-bit address space: 2^{32}
- 2^{18} = 262,144 pages

$$\frac{2^{32}}{2^{14}} * 4 = 1MB \text{ per page table}$$

- Memory requirement cut to $\frac{1}{4}$
- However pages are huge
- Internal fragmentation results
- 16KB page(s) allocated for small programs with only a few variables

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.43

43

PAGE TABLES: WASTED SPACE

- **Process: 16KB Address Space w/ 1KB pages**

code

heap

stack

0

1

2

3

4

5

6

7

8

9

10

11

12

13

14

Allocate

A 16KB Address Space with 1KB Pages

PFN	valid	prot	present	dirty
10	1	r-x	1	0
-	0	-	-	-
-	0	-	-	-
-	0	-	-	-
15	1	rw-	1	1
...
-	0	-	-	-
3	1	rw-	1	1
23	1	rw-	1	1

A Page Table For 16KB Address Space

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.44

44

PAGE TABLES: WASTED SPACE

■ Process: 16KB Address Space w/ 1KB pages

Page Table

Virtual Address Space

code

0

1

2

heap

8

9

10

11

12

13

stack

14

Physical Memory

Allocate

PFN	valid	prot	present	dirty
				0
	-			-
	-			-
15	1	rw-	1	1
...
-	0	-	-	-
3	1	rw-	1	1
23	1	rw-	1	1

A Page Table For 16KB Address Space

Most of the page table is unused and full of wasted space. (73%)

A 16KB Address Space with 1KB Pages

May 29, 2025

TCCS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.45

45

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies

May 29, 2025

TCCS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.46

46

MULTI-LEVEL PAGE TABLES

- Consider a page table:
- 32-bit addressing, 4KB pages
- 2²⁰ page table entries
- Even if memory is sparsely populated the *per process* page table requires:

Page table size = $\frac{2^{32}}{2^{12}} * 4Byte = 4MByte$

- Often most of the 4MB *per process* page table is empty
- Page table must be placed in 4MB contiguous block of RAM
- MUST SAVE MEMORY!

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.47

47

MULTI-LEVEL PAGE TABLES - 2

- Add level of indirection, the “page directory”

Linear Page Table

PBTR201

valid	prot	PFN
1	rx	12
1	rx	13
0	-	-
1	rw	100
0	-	-
0	-	-
0	-	-
0	-	-
0	-	-
1	rw	86
1	rw	15

PFN201

Multi-level Page Table

PBTR200

valid	PFN
1	201
0	-
0	-
1	203

PFN200

The Page Directory

[Page 1 of PT: Not Allocated]

valid	prot	PFN
1	rx	12
1	rx	13
0	-	-
1	rw	100

PFN201

valid	prot	PFN
0	-	-
0	-	-
1	rw	86
1	rw	15

PFN204

Linear (Left) And Multi-Level (Right) Page Tables

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.48

48

MULTI-LEVEL PAGE TABLES - 2

■ Add level of indirection, the “page directory”

Linear Page Table

PBTR 201

PFN203

0	-	-
0	-	-
1	rw	86
1	rw	15

Multi-level Page Table

PBTR 200

PFN204

0	-	-
0	-	-
1	rw	86
1	rw	15

Two level page table:
2²⁰ pages addressed with
two level-indexing
(page directory index, page table index)

Linear (Left) And Multi-Level (Right) Page Tables

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.49

49

4 GB computer (2³²) and 4KB pages (2¹²)

1. How much space is required for a 2-level page table with one page directory (PD) and one page table (PT)?

2. How much memory can a single PD pointing to a single PT address?

50

4 GB computer (2^{32}) and 4KB pages (2^{12})

1. How much space is required for a 2-level page table with one page directory (PD) and one page table (PT)?

2. How much memory can a single PD pointing to a single PT address?

1023
0
1
10 bits
1023
1024 x 4
4096
8192 bytes
1023
10 bits
1024 x 4
4096
1024 x 4096
4096 b
1024 x 4 KB
 $2^{10} \quad 2^{12}$
2. $2^{22} \rightarrow 4\text{MB} \rightarrow 8\text{MB}$
WITH A
SECOND PT


51

May 23, 2023

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.52

52



53

MULTI-LEVEL PAGE TABLES - 3

- Advantages

- Only allocates page table space in proportion to the address space actually used
- Can easily grab next free page to expand page table

- Disadvantages

- Multi-level page tables are an example of a time-space tradeoff
- Sacrifice address translation time (now 2-level) for space
- Complexity: multi-level schemes are more complex

54

EXAMPLE

- 16KB address space, 64byte pages
- How large would a one-level page table need to be?
- 2^{14} (address space) / 2^6 (page size) = 2^8 = 256 (pages)

0000 0000
0000 0001
—

1111 1111

code
code
(free)
(free)
heap
heap
(free)
(free)
stack
stack

Flag	Detail
Address space	16 KB
Page size	64 byte
Virtual address	14 bit
VPN	8 bit
Offset	6 bit
Page table entry	2^8 (256)

A 16-KB Address Space With 64-byte Pages

13 12 11 10 9 8 7 6 5 4 3 2 1 0

Offset

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.55

55

EXAMPLE - 2

- 256 total page table entries (64 bytes each)
- 1,024 bytes page table size, stored using 64-byte pages
= (1024/64) = 16 page directory entries (PDEs)
- Each page directory entry (PDE) can hold 16 page table entries (PTEs) *e.g. lookups*
- 16 page directory entries (PDE) x 16 page table entries (PTE)
= 256 total PTEs
- Key idea: the page table is stored using pages too!**

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.56

56

PAGE DIRECTORY INDEX

- Now, let's split the page table into two:
 - 8 bit VPN to map 256 pages
 - 4 bits for page directory index (PDI – 1st level page table)
 - 6 bits offset into 64-byte page

The diagram shows a 14-bit virtual address represented as a row of 14 boxes numbered 13 down to 0. Boxes 13, 12, 11, and 10 are orange and grouped under a bracket labeled 'Page Directory Index'. Boxes 9, 8, 7, and 6 are orange and grouped under a bracket labeled 'VPN'. Boxes 5, 4, 3, 2, 1, and 0 are blue and grouped under a bracket labeled 'Offset'. A large bracket below the entire row is labeled '14-bits Virtual address'.

May 29, 2025

TCCS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.57

57

PAGE TABLE INDEX

- 4 bits page directory index (PDI – 1st level)
- 4 bits page table index (PTI – 2nd level)

The diagram shows a 14-bit virtual address represented as a row of 14 boxes numbered 13 down to 0. Boxes 13, 12, 11, and 10 are orange and grouped under a bracket labeled 'Page Directory Index'. Boxes 9, 8, 7, and 6 are orange and grouped under a bracket labeled 'Page Table Index'. Boxes 5, 4, 3, 2, 1, and 0 are blue and grouped under a bracket labeled 'Offset'. A large bracket below the entire row is labeled '14-bits Virtual address'.

- To dereference one 64-byte memory page,
 - We need one page directory entry (PDE)
 - One page table Index (PTI) – can address 16 pages

May 29, 2025

TCCS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.58

58

EXAMPLE - 3

- For this example, how much space is required to store as a single-level page table with any number of PTEs?
 - 16KB address space, 64 byte pages
 - 256 page frames, 4 byte page size
 - 1,024 bytes required (*single level*)
- How much space is required for a two-level page table with only 4 page table entries (PTEs) ?
 - Page directory = 16 entries x 4 bytes (1 x 64 byte page)
 - Page table = 16 entries (4 used) x 4 bytes (1 x 64 byte page)
 - 128 bytes required (2 x 64 byte pages)
 - Savings = using just 12.5% the space !!!

May 29, 2025

TCCS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.59

59

For this example, how much space is required to store as a single-level page table with any number of PTEs?

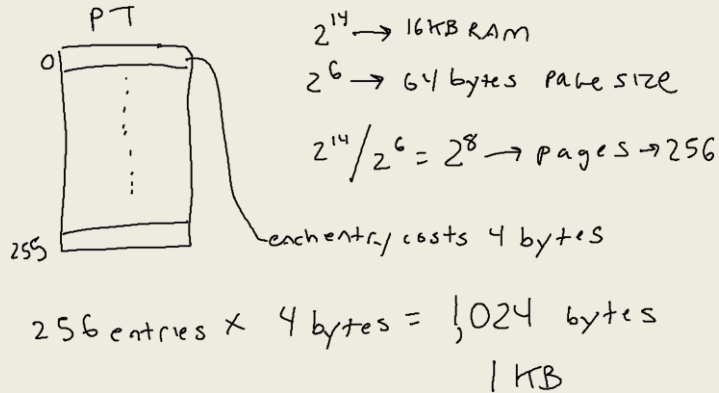
16KB address space, 64 byte pages, 256 page frames, 4 byte page size

Storage requirement: bytes required (*single level*)

60

For this example, how much space is required to store as a single-level page table with any number of PTEs?

16KB address space, 64 byte pages, 256 page frames, 4 byte page size



Storage requirement: bytes required (single level)

61

How much space is required for a two-level page table with only 4 page table entries (PTEs) ? (one page each for code segment, stack segment, heap segment, data segment)

16KB address space, 64 byte pages, 256 page frames, 4 byte page size

Page directory = 16 entries x 4 bytes (1 x 64 byte page)
Page table = 16 entries (4 used) x 4 bytes (1 x 64 byte page)
Store requirement = 128 bytes required (2 x 64 byte pages)
Savings =

62

How much space is required for a two-level page table with only 4 page table entries (PTEs) ? (one page each for code segment, stack segment, heap segment, data segment)
16KB address space, 64 byte pages, 256 page frames, 4 byte page size

8 bit VPN

PD

0

...

15

NULL

PT

0

code

1

stack

2

heap

3

data

...

15

64 byte page size

16×64

$2^4 \times 2^6$

2^{10}

1,024 bytes

$16 \times 4 \text{ bytes}$

64 bytes

$+$

$16 \times 4 \text{ bytes}$

64 bytes

128 bytes

$1,024$

$- 128$

$896 \text{ savings in bytes}$

$128 / 1,024 = 12.5\%$

Page directory = 16 entries x 4 bytes (1 x 64 byte page)
Page table = 16 entries (4 used) x 4 bytes (1 x 64 byte page)
Store requirement = 128 bytes required (2 x 64 byte pages)
Savings =

63

32-BIT EXAMPLE

- Consider: 32-bit address space, 4KB pages, 2^{20} pages
- Only 4 mapped pages
- Single level: 4 MB (we've done this before) 12.5%
- Two level: (old VPN was 20 bits, split in half)
 - Page directory = 2^{10} entries x 4 bytes = 1 x 4 KB page
 - Page table = 4 entries x 4 bytes (mapped to 1 4KB page)
 - 8KB (8,192 bytes) required
 - Savings = using just .78 % the space !!!
- 100 sparse processes now require < 1MB for page tables
 - $8\text{KB} \times 100 = 800\text{KB}$

May 29, 2025

TCCS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.64

64

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables **N-level Page Tables**
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.65

65

WE WILL RETURN AT
5:00 PM

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma



com

L17.66

66

MORE THAN TWO LEVELS - 2

- Page table entries per page = $512 / 4 = 128$
- 7 bytes – for page table index (PTI)

Flag	Detail
Virtual address	30 bit
Page size	512 byte
VPN	21 bit
Offset	9 bit
Page entry per page	128 PTEs

$\log_2 128 = 7$

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.67

67

MORE THAN TWO LEVELS - 3

- Consider 1 GB computer: 2^{30} =1GB RAM, 512-byte (2^9 pages)
- $2^{14} = 16,384$ page directory entries (PDEs) are required
- When using 2^7 (128 entry) page tables...
- Page size = 512 bytes * 4 bytes per addr

Flag	Detail
Virtual address	30 bit
Page size	512 byte
VPN	21 bit
Offset	9 bit
Page entry per page	128 PTEs

$\log_2 128 = 7$

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.68

68

MORE THAN TWO LEVELS - 3

- Consider 1 GB computer: 2^{30} =1GB RAM, 512-byte (2^9 pages)
- 2^{14} = 16,384 page directory entries (PDEs) are required
- When using 2^7 (128 entry) page tables...
- Page size = 512 bytes / 4 bytes per addr

Can't Store Page Directory with 16K pages, using 512 bytes pages. Pages only dereference 128 addresses (512 bytes / 32 bytes)

Virtual address	30 bit
Page size	512 byte
VPN	21 bit
Offset	9 bit
Page entry per page	128 PTEs

→ $\log_2 128 = 7$

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.69

69

MORE THAN TWO LEVELS - 3

- Consider 1 GB computer: 2^{30} =1GB RAM, 512-byte (2^9 pages)
- 2^{14} = 16,384 page directory entries (PDEs) are required
- When using 2^7 (128 entry) page tables...
- Page size = 512 bytes / 4 bytes per addr

Need three level page table:
Page directory 0 (PD Index 0)
Page directory 1 (PD Index 1)
Page Table Index

Virtual address	30 bit
Page size	512 byte
VPN	21 bit
Offset	9 bit
Page entry per page	128 PTEs

→ $\log_2 128 = 7$

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.70

70

MORE THAN TWO LEVELS - 4

- We can now address 1GB with “fine grained” 512 byte pages
- Using multiple levels of indirection

The diagram shows a 32-bit Virtual Page Number (VPN) divided into four sections. The first section, labeled 'PD Index 0', covers bits 30 down to 23. The second section, labeled 'PD Index 1', covers bits 22 down to 16. The third section, labeled 'Page Table Index', covers bits 15 down to 9. The final section, covering bits 8 down to 0, is shaded blue and represents the remaining bits of the VPN. Arrows indicate the bit ranges for each index.

- Consider the implications for address translation!
- How much space is required for a virtual address space with 4 entries on a 512-byte page? (let's say 4 32-bit integers)
- PD0 1 page, PD1 1 page, PT 1 page = 1,536 bytes
- Memory Usage= 1,536 (3-level) / 8,388,608 (1-level) = .0183% !!!

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.71

71

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.7
2

72

$2^{30} = 1\text{GB}$ $2^9 = 512\text{ byte PAGE SIZE}$
 $2^{30} / 2^9 \rightarrow 2^{21} \text{ PAGES} \rightarrow 2 \text{ million PAGES}$
 $2^{21} \times 4 \text{ bytes} \rightarrow 2^{23} \rightarrow 8 \text{ MB}$

SINGLE LEVEL PG-TABLE

VPN - 21 bits
 PD0
 0
 ...
 127 NULL
 128×4
 512

PD1
 0
 ...
 127 NULL
 128×4
 512

PT
 0
 CODE
 STACK
 heap
 DATA
 NULL 127
 128×4
 512

Three LEVEL PG-TABLE

$512 + 512 + 512$
1,536 bytes

May 29, 2025 TCSS422: Operating Systems [Spring 2025]
 School of Engineering and Technology, University of Washington - Tacoma L17.7
 3

73

ADDRESS TRANSLATION CODE

```

// 5-level Linux page table address lookup
//
// Inputs:
// mm_struct - process's memory map struct
// vpage - virtual page address

// Define page struct pointers
pgd_t *pgd;
p4d_t *p4d;
pud_t *pud;
pmd_t *pmd;
pte_t *pte;
struct page *page;
    
```

May 29, 2025 TCSS422: Operating Systems [Spring 2025]
 School of Engineering and Technology, University of Washington - Tacoma L17.74

74

ADDRESS TRANSLATION - 2

```
pgd = pgd_offset(mm, vpage);
if (pgd_none(*pgd) || pgd_bad(*pgd))
    return 0;
p4d = p4d_offset(pgd, vpage);
if (p4d_none(*p4d) || p4d_bad(*p4d))
    return 0;
pud = pud_offset(p4d, vpage);
if (pud_none(*pud) || pud_bad(*pud))
    return 0;
pmd = pmd_offset(pud, vpage);
if (pmd_none(*pmd) || pmd_bad(*pmd))
    return 0;
if (!(pte = pte_offset_map(pmd, vpage)))
    return 0;
if (!(page = pte_page(*pte)))
    return 0;
physical_page_addr = page_to_phys(page);
pte_unmap(pte);
return physical_page_addr; // param to send back
```

pgd_offset():
Takes a vpage address and the mm_struct for the process, returns the PGD entry that covers the requested address...

p4d/pud/pmd_offset():
Takes a vpage address and the pgd/p4d/pud entry and returns the relevant p4d/pud/pmd.

pte_unmap()
release temporary kernel mapping for the page table entry


May 29, 2025

TCCS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.75

75

INVERTED PAGE TABLES



- Keep a single page table for each physical page of memory
- Consider 4GB physical memory
- Using 4KB pages, page table requires 4MB to map all of RAM
- Page table stores
 - Which process uses each page
 - Which process virtual page (from process virtual address space) maps to the physical page
- All processes share the same page table for memory mapping, kernel must isolate all use of the shared structure
- Finding process memory pages requires search of 2^{20} pages
- Hash table: can index memory and speed lookups

May 29, 2025

TCCS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.76

76

MULTI-LEVEL PAGE TABLE EXAMPLE

- Consider a 16 MB computer which indexes memory using 4KB pages
- (#1) For a single level page table, how many pages are required to index memory?
- (#2) How many bits are required for the VPN?
- (#3) Assuming each page table entry (PTE) can index any byte on a 4KB page, how many offset bits are required?
- (#4) Assuming there are 8 status bits, how many bytes are required for each page table entry?

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.77

77

MULTI LEVEL PAGE TABLE EXAMPLE - 2

- (#5) How many bytes (or KB) are required for a single level page table?
- Let's assume a simple HelloWorld.c program.
- HelloWorld.c requires virtual address translation for 4 pages:
 - 1 – code page 1 – stack page
 - 1 – heap page 1 – data segment page
- (#6) Assuming a two-level page table scheme, how many bits are required for the Page Directory Index (PDI)?
- (#7) How many bits are required for the Page Table Index (PTI)?

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.78

78

MULTI LEVEL PAGE TABLE EXAMPLE - 3

- Assume each page directory entry (PDE) and page table entry (PTE) requires 4 bytes:
 - 6 bits for the Page Directory Index (PDI)
 - 6 bits for the Page Table Index (PTI)
 - 12 offset bits
 - 8 status bits
- (#8)** How much **total** memory is required to index the HelloWorld.c program using a two-level page table when we only need to translate 4 total pages?
- HINT:** we need to allocate one Page Directory and one Page Table...
- HINT:** how many entries are in the PD and PT

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.79

79

MULTI LEVEL PAGE TABLE EXAMPLE - 4

- (#9)** Using a single page directory entry (PDE) pointing to a single page table (PT), if all of the slots of the page table (PT) are in use, what is the total amount of memory a two-level page table scheme can address?
- (#10)** And finally, for this example, as a percentage (%), how much memory does the 2-level page table scheme consume compared to the 1-level scheme?
- HINT:** two-level memory use / one-level memory use

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.80

80

ANSWERS

- #1 – 4096 pages
- #2 – 12 bits
- #3 – 12 bits
- #4 – 4 bytes
- #5 – $4096 \times 4 = 16,384$ bytes (16KB)
- #6 – 6 bits
- #7 – 6 bits
- #8 – 256 bytes for Page Directory (PD) (64 entries x 4 bytes)
256 bytes for Page Table (PT) **TOTAL = 512 bytes**
- #9 – 64 entries, where each entry maps a 4,096 byte page
With 12 offset bits, can address 262,144 bytes (256 KB)
- #10- $512/16384 = .03125 \rightarrow 3.125\%$

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.81

81

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, Swapping Policies


May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.82

82

CHAPTER 21/22: BEYOND PHYSICAL MEMORY



May 29, 2025

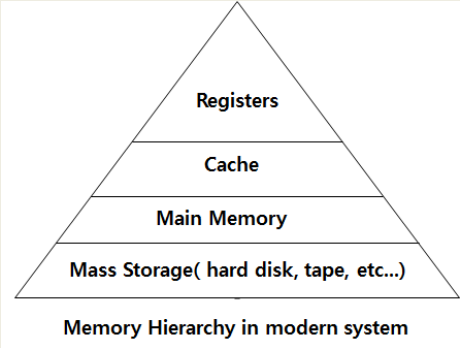
TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.83

83

MEMORY HIERARCHY

- Disks (HDD, SSD) provide another level of storage in the memory hierarchy



May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.84

84

MOTIVATION FOR EXPANDING THE ADDRESS SPACE

- Provide the illusion of an address space larger than physical RAM
- For a single process
 - Convenience
 - Ease of use
- For multiple processes
 - Large virtual memory space supports running *many concurrent processes*. . .

May 29, 2025

TCCS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.85

85

LATENCY TIMES

- Design considerations:
 - SSDs 4x the time of DRAM
 - HDDs 80x the time of DRAM

Action	Latency (ns)	(μs)	
L1 cache reference	0.5ns		
L2 cache reference	7 ns		14x L1 cache
Mutex lock/unlock	25 ns		
Main memory reference	100 ns		20x L2 cache, 200x L1
Read 4K randomly from SSD*	150,000 ns	150 μs	~1GB/sec SSD
Read 1 MB sequentially from memory	250,000 ns	250 μs	
Read 1 MB sequentially from SSD*	1,000,000 ns	1,000 μs	1 ms ~1GB/sec SSD, 4X memory
Read 1 MB sequentially from disk	20,000,000 ns	20,000 μs	20 ms 80x memory, 20X SSD

- Latency numbers every programmer should know
- From: <https://gist.github.com/jboner/2841832#file-latency-txt>

May 29, 2025

TCCS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.86

86

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - **Swapping Mechanisms** Swapping Policies

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.87

87

SWAP SPACE

- Disk space for storing memory pages
- “Swap” them in and out of memory to disk as needed

	PFN 0	PFN 1	PFN 2	PFN 3
Physical Memory	Proc 0 [VPN 0]	Proc 1 [VPN 2]	Proc 1 [VPN 3]	Proc 2 [VPN 0]

	Block 0	Block 1	Block 2	Block 3	Block 4	Block 5	Block 6	Block 7
Swap Space	Proc 0 [VPN 1]	Proc 0 [VPN 2]	[Free]	Proc 1 [VPN 0]	Proc 1 [VPN 1]	Proc 3 [VPN 0]	Proc 2 [VPN 1]	Proc 3 [VPN 1]

Physical Memory and Swap Space

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.88

88

SWAP SPACE - 2

- The size of the swap space can be seen using the Linux free command: “free -h”

```
wlloyd@dlone:~$ free -h
```

	total	used	free	shared	buff/cache	available
Mem:	30G	11G	14G	1.3G	4.4G	17G
Swap:	31G	0B	31G			

- With sufficient disk space, a common allocation is to create Swap space greater than or equal to physical RAM

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.89

89

SWAP SPACE - 3

- Swap space lives on a separate logical volume in Ubuntu Linux that is managed separately from the root file system
- Check logical volumes with “sudo lvs” command:

```
--- Logical volume ---
LV Path                /dev/ubuntu-vg/swap_1
LV Name                 swap_1
VG Name                ubuntu-vg
LV UUID                G1ovj6-4M33-2YXY-YETH-wF7V-93vF-QRQytG
LV Write Access        read/write
LV Creation host, time ubuntu, 2018-09-30 15:44:16 -0700
LV Status              available
# open                 2
LV Size                976.00 MiB
Current LE             244
Segments              1
Allocation              inherit
Read ahead sectors     auto
- currently set to    256
Block device           253:1
```

- See also “lvm lvs” command

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.90

90

PAGE LOCATION

- Memory pages are:
 - Stored in memory
 - Swapped to disk
- Present bit
 - In the page table entry (PTE) indicates if page is present
- Page fault
 - Memory page is accessed, but has been swapped to disk

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.91

91

PAGE FAULT

- OS steps in to handle the page fault
- Loading page from disk requires a free memory page
- Page-Fault Algorithm

```
1:     PFN = FindFreePhysicalPage()
2:     if (PFN == -1)                // no free page found
3:         PFN = EvictPage()         // run replacement algorithm
4:     DiskRead(PTE.DiskAddr, pfn)   // sleep (waiting for I/O)
5:     PTE.present = True            // set PTE bit to present
6:     PTE.PFN = PFN                // reference new loaded page
7:     RetryInstruction()           // retry instruction
```

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.92

92

PAGE REPLACEMENTS

- Page daemon
 - Background threads which monitors swapped pages
- Low watermark (LW)
 - Threshold for when to swap pages to disk
 - Daemon checks: free pages < LW
 - Begin swapping to disk until reaching the highwater mark
- High watermark (HW)
 - Target threshold of free memory pages
 - Daemon free until: free pages >= HW

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.93

93

OBJECTIVES – 5/29

- Questions from 5/27
- Memory Segmentation Activity + answers (available in Canvas)
- Assignment 2 - June 5 AOE
- A3: (Tutorial) Intro to Linux Kernel Modules - June 10 AOE
- Final exam – Thursday June 12 @ 3:40pm
- Quiz 4 – Page Tables - Due June 8 AOE
- Chapter 19: Translation Lookaside Buffer (TLB)
 - TLB Algorithm, Hit-to-Miss Ratios
- Chapter 20: Paging: Smaller Tables
 - Smaller Tables, Multi-level Page Tables, N-level Page Tables
- Chapter 21/22: Beyond Physical Memory
 - Swapping Mechanisms, **Swapping Policies**

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma


L17.94

94

REPLACEMENT POLICIES

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma



POLICY CHANGES

L17.9
5

95

CACHE MANAGEMENT

- Replacement policies apply to “any” cache
- Goal is to minimize the number of misses
- Average memory access time (AMAT) can be estimated:

$$AMAT = (P_{Hit} * T_M) + (P_{Miss} * T_D)$$

Argument	Meaning
T_M	The cost of accessing memory (time)
T_D	The cost of accessing disk (time)
P_{Hit}	The probability of finding the data item in the cache(a hit)
P_{Miss}	The probability of not finding the data in the cache(a miss)

- Consider $T_M = 100\text{ ns}$, $T_D = 10\text{ms}$
- Consider $P_{hit} = .9\text{ (90\%)}$, $P_{miss} = .1$
- Consider $P_{hit} = .999\text{ (99.9\%)}$, $P_{miss} = .001$

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.96

96

OPTIMAL REPLACEMENT POLICY

- What if:
 - We could predict the future (... with a magical oracle)
 - All future page accesses are known
 - Always replace the page in the cache used farthest in the future
- Used for a comparison
- Provides a “best case” replacement policy
- Consider a 3-element empty cache with the following page accesses:

0 1 2 0 1 3 0 3 1 2 1

What is the hit/miss ratio?

6 hits

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.97

97

FIFO REPLACEMENT

- Queue based
- Always replace the oldest element at the back of cache
- Simple to implement
- Doesn't consider importance... just arrival ordering
- Consider a 3-element empty cache with the following page accesses:

0 1 2 0 1 3 0 3 1 2 1

What is the hit/miss ratio?

4 hits

How is FIFO different than LRU?

LRU incorporates history

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.98

98

RANDOM REPLACEMENT

- Pick a page at random to replace
- Simple and fast implementation
- Performance depends on luck of random choices

0 1 2 0 1 3 0 3 1 2 1

Number of Hits	Frequency
1	0
2	2
3	10
4	20
5	40
6	45

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.99

99

HISTORY-BASED POLICIES

- LRU: Least recently used
- Always replace page with oldest access time (front)
- Always move end of cache when element is read again
- LRU requires constant reorganization of the cache
- Considers temporal locality (*when pg was last accessed*)

0 1 2 0 1 3 0 3 1 2 1

What is the hit/miss ratio?

6 hits

- LFU: Least frequently used
- Always replace page with the fewest # of accesses (front)
- Incorporates frequency of use - *must track pg accesses*
- Consider frequency of page accesses

0 1 2 0 1 3 0 3 1 2 1

Hit/miss ratio is=6 hits

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.100

100

2 hits

3 hits

4 hits

5 hits

6 hits

Consider a 3-element cache. With a FIFO replacement policy, how many hits occur with the following page access sequence:
1 2 0 1 3 1 2 0 2 1 3

May 29, 2025

TCSS422: Operating Systems [Spring 2025]

L17.01

101

2 hits

3 hits

4 hits

5 hits

6 hits

Consider a 3-element cache. With an LRU replacement policy, how many hits occur with the following page access sequence:
1 2 0 1 3 1 2 0 2 1 3

May 29, 2025

TCSS422: Operating Systems [Spring 2025]

L17.02

102

WORKLOAD EXAMPLES: NO-LOCALITY

- No-Locality (Random Access) Workload
 - Perform 10,000 random page accesses
 - Across set of 100 memory pages

The No-Locality Workload

When the cache is large enough to fit the entire workload, it doesn't matter which policy you use.

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.103

103

WORKLOAD EXAMPLES: 80/20

- 80/20 Workload
 - Perform 10,000 page accesses, against set of 100 pages
 - 80% of accesses are to 20% of pages (hot pages)
 - 20% of accesses are to 80% of pages (cold pages)

The 80-20 Workload

LRU is more likely to hold onto hot pages
(recalls history)

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.104

104

WORKLOAD EXAMPLES: SEQUENTIAL

- **Looping sequential workload**
 - Refer to 50 pages in sequence: 0, 1, ..., 49
 - Repeat loop

The Looping-Sequential Workload

Cache Size (Blocks)	OPT Hit Rate (%)	LRU Hit Rate (%)	FIFO Hit Rate (%)	RAND Hit Rate (%)
0	0	0	0	0
20	40	0	0	10
40	80	0	0	40
50	100	0	0	100
60	100	0	0	100
80	100	0	0	100
100	100	100	100	100

Random performs better than FIFO and LRU for cache sizes < 50

Algorithms should provide "scan resistance"

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.105

105

With small cache sizes, for the looping sequential workload, why do FIFO and LRU fail to provide cache hits?

Cache hits in this scenario require consideration of how frequently accessed memory is for cache replacement

Memory accesses are unpredictable and too random. Unpredictable accesses require a random cache replacement policy for cache hits

Memory accesses to elements that are accessed repeatedly are too spread apart temporally to benefit from caching

Unlike Random cache replacement, both FIFO and LRU fail to speculate memory accesses in advance to improve caching

None of the above

Start the presentation to see live content. For screen share software, share the entire screen. Get help at pollev.com/app

106

IMPLEMENTING LRU

- Implementing last recently used (LRU) requires tracking access time for all system memory pages
- Times can be tracked with a list
- For cache eviction, we must scan an entire list
- Consider: 4GB memory system (2^{32}), with 4KB pages (2^{12})
- This requires 2^{20} comparisons !!!
- Simplification is needed
 - Consider how to approximate the oldest page access

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.107

107

IMPLEMENTING LRU - 2

- Harness the Page Table Entry (PTE) Use Bit
- HW sets to 1 when page is used
- OS sets to 0
- Clock algorithm (*approximate LRU*)
 - Refer to pages in a circular list
 - Clock hand points to current page
 - Loops around
 - IF USE_BIT=1 set to USE_BIT = 0
 - IF USE_BIT=0 replace page



May 29, 2025


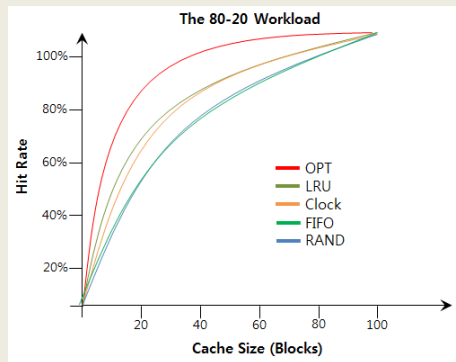
TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.108

108

CLOCK ALGORITHM

- Not as efficient as LRU, but better than other replacement algorithms that do not consider history



May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.109

109

CLOCK ALGORITHM - 2

- Consider dirty pages in cache
- If DIRTY (modified) bit is FALSE
 - No cost to evict page from cache
- If DIRTY (modified) bit is TRUE
 - Cache eviction requires updating memory
 - Contents have changed
- Clock algorithm should favor no cost eviction

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.110

110

WHEN TO LOAD PAGES

- On demand → demand paging
- Prefetching
 - Preload pages based on anticipated demand
 - Prediction based on locality
 - Access page P, suggest page P+1 may be used
- What other techniques might help anticipate required memory pages?
 - Prediction models, historical analysis
 - In general: accuracy vs. effort tradeoff
 - High analysis techniques struggle to respond in real time

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.111

111

OTHER SWAPPING POLICIES

- Page swaps / writes
 - Group/cluster pages together
 - Collect pending writes, perform as batch
 - Grouping disk writes helps amortize latency costs
- Thrashing
 - Occurs when system runs many memory intensive processes and is low in memory
 - Everything is constantly swapped to-and-from disk

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.112

112

OTHER SWAPPING POLICIES - 2

- **Working sets**
 - Groups of related processes
 - When thrashing: prevent one or more working set(s) from running
 - Temporarily reduces memory burden
 - Allows some processes to run, reduces thrashing

May 29, 2025

TCSS422: Operating Systems [Spring 2025]
School of Engineering and Technology, University of Washington - Tacoma

L17.113

113

QUESTIONS



114