


## TCSS 422: OPERATING SYSTEMS

### Translation Lookaside Buffer (TLB)



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

## OBJECTIVES

- TLB Algorithm
- TLB Tradeoffs
- TLB Context Switch

November 18, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L17.2

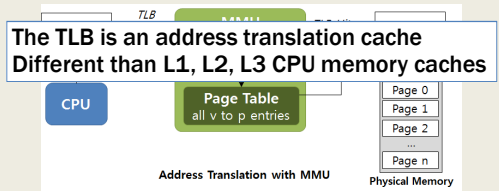
## TRANSLATION LOOKASIDE BUFFER

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

November 18, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L17.3

## TRANSLATION LOOKASIDE BUFFER (TLB)

- Part of the CPU's Memory Management Unit (MMU)
- Goal: Reduce accesses to the page tables



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

Address Translation with MMU

November 18, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L17.4

## 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

November 18, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L17.5

## 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: } else {
16:   TLB_Insert(VPN, PTE.PFN, PTE.ProtectBits)
17:   RetryInstruction()
18: }
19: }

```

Retry the instruction... (requery the TLB)

November 18, 2016 TCSS422: Operating Systems [Fall 2016]  
Institute of Technology, University of Washington - Tacoma L17.6

## TLB - ADDRESS TRANSLATION CACHE

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

November 18, 2016

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

L17.7

## TLB EXAMPLE

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

### Consider:

- Address space: 256-byte
  - Addressable using 8 total bits
  - 4 bits for the VPN
- Page size: 16 bytes
  - Offset is addressable using 4-bits
- Array: 10 x 4-byte integers

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

November 18, 2016

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

L17.8

## 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					
VPN = 07	a[0]	a[1]	a[2]		
VPN = 08	a[3]	a[4]	a[5]	a[6]	
VPN = 09	a[7]	a[8]	a[9]		
VPN = 10					
VPN = 11					
VPN = 12					
VPN = 13					
VPN = 14					
VPN = 15					

November 18, 2016

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

L17.9

## 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? (%)

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

November 18, 2016

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

L17.10

## 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?

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

November 18, 2016

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

L17.11

## TLB TRADEOFFS

- Page size
  - Larger page sizes increase the probability of a TLB hit
  - Example: 16-bytes (very small), 4096-bytes (common)
  - Larger sizes increase memory requirement of offset

November 18, 2016

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

L17.12

## TLB TRADEOFFS - 2

- **Spatial locality**
  - Accessing addresses local to each improves the hit rate.
  - Intel example - nehalem architecture: two-level TLB
  - First level TLB: split into data and code instruction cache
  - Second level TLB: shared TLB for code and data
- Each level may have a small and large page cache

November 18, 2016

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

L17.13

## TLB TRADEOFFS - 3

- **Temporal locality**
  - Higher cache hit ratios are expected for repeated memory accesses close in time
- Can dramatically improve performance for "second iteration"

November 18, 2016

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

L17.14

## HW CACHE TRADEOFF

- **Speed vs. size**
- Speed  $\longleftrightarrow$  Size
- In order to be fast, caches must be small
  - Too large of a cache will mimic physical memory
  - Limitations for on chip memory



November 18, 2016

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

L17.15

## HANDLING TLB MISS

- **Historical view**
- **CISC – Complex instruction set computer**
  - Intel x86 CPUs
- Traditionally have provided on CPU HW instructions and handling of TLB misses
- HW has a page table register to store location of page table

November 18, 2016

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

L17.16

## HANDLING TLB MISS - 2

- **RISC – Reduced instruction set computer**
  - ARM CPUs
- Traditionally the OS handles TLB misses
- HW raises an exception
- Trap handler is executed to handle the miss
- **Advantages**
  - HW Simplicity: simply needs to raise an exception
  - Flexibility: OS provided page table implementations can use different data structures, etc.

November 18, 2016

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

L17.17

## TLB CONTENTS

- TLB typically may have 32, 64, or 128 entries
- HW searches the entire TLB in parallel to find the translation
- **Other bits**
  - Valid bit: valid translation?
  - Protection bit: read/execute, read/write
  - Address-space identifier: identify entries by process
  - Dirty bit

VPN	PFN	other bits

Typical TLB entry look like this

November 18, 2016

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

L17.18

## TLB: CONTEXT SWITCH

- TLB stores address translations for current running process
- A context/switch to a new process invalidates the TLB
- Must "switch" out the TLB
- TLB flush
  - Flush TLB on context switches, set all entries to 0
  - Requires time to flush
  - TLB must be reloaded for each C/S
  - If process doesn't reside in the CPU for long, the TLB may not get reloaded
- Share TLB across processes during C/S
  - Use address space identifier (ASID) to tag TLB entries by process


November 18, 2016

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

L17.19

## TLB: CONTEXT SWITCH - 2

- Address space identifier (ASID) enables data state to persist during context switches



November 18, 2016

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

L17.20

## SHARED MEMORY SPACE

VPN	PFN	valid	prot	ASID
10	101	1	rwX	1
-	-	-	-	-
50	101	1	rwX	2
-	-	-	-	-

- When processes share a code page
  - Shared libraries ok
  - Code pages typically are RX, not RWX

Sharing of pages is useful as it reduces the number of physical pages in use.

November 18, 2016

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

L17.21

## CACHE REPLACEMENT


- When TLB cache is full, how add a new address translation to the TLB?
  - Least Recently Used (LRU)
    - Evict the oldest entry
  - Note how the TLB is loaded / unloaded...
  - Goal minimize miss rate, increase hit rate
  - Random policy
    - Pick a candidate at random to free-up space in the TLB

November 18, 2016

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

L17.22

## LEAST RECENTLY USED



- RED - miss
- WHITE - hit
- For 3-page TLB, observe replacement

Total 11 TLB miss

November 18, 2016

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

L17.23

## EXAMPLE TLB ENTRY - MIPS R4000

- Early 64-bit RISC processor

All 64 bits of this TLB entry(example of MIPS R4000)

0	1	2	3	4	5	6	7	8	9	10	11	...	19	...	31									
VPN												G	ASID											
PFN												C	D	V										

Flag	Content
19-bit VPN	The rest reserved for the kernel.
24-bit PFN	Systems can support with up to 64GB of main memory( $2^{24} \times 4KB$ pages ).
Global bit(G)	Used for pages that are globally-shared among processes.
ASID	OS can use to distinguish between address spaces.

Coherence bit(C)	determine how a page is cached by the hardware.
Dirty bit(D)	marking when the page has been written.
Valid bit(V)	tells the hardware if there is a valid translation present in the entry.



November 18, 2016

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

L17.24

