

CS61CL : Machine Structures

Lecture #12 – Virtual Memory

2009-08-03



Jeremy Huddleston

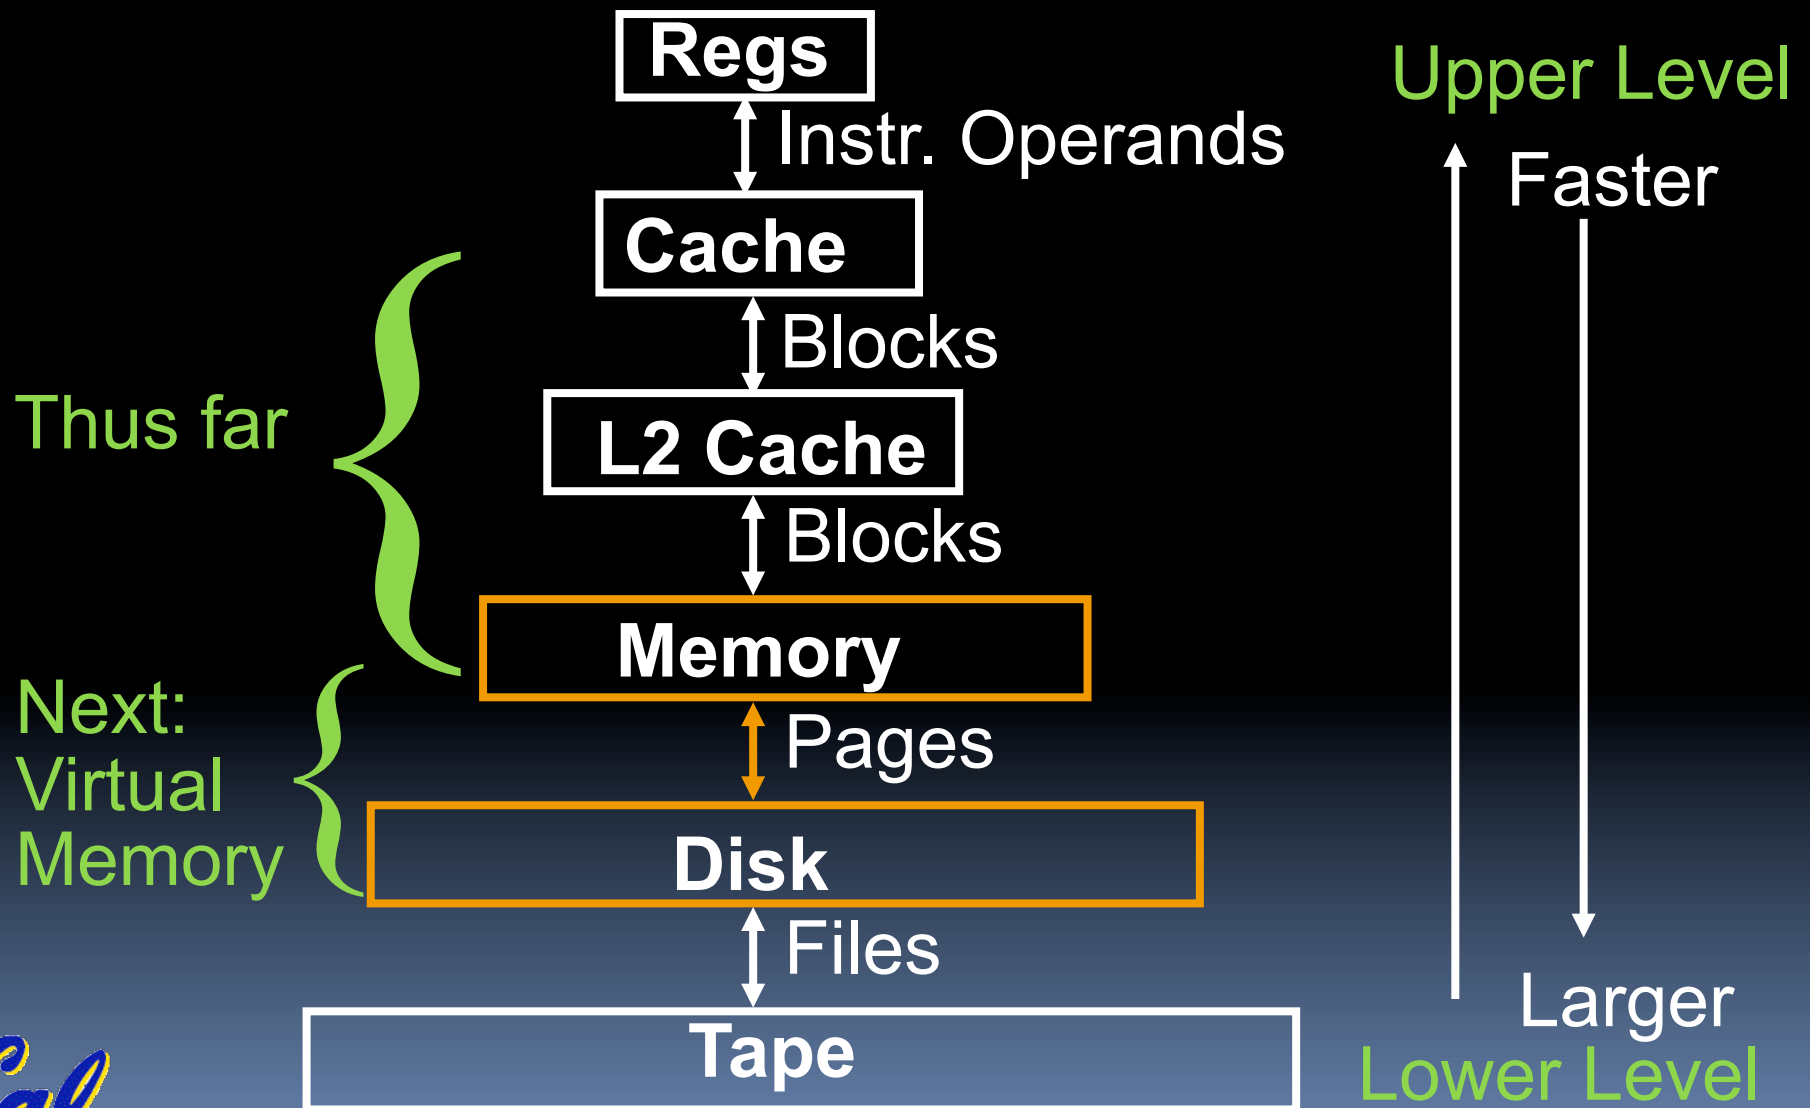


Review

- **Cache design choices:**
 - **Size of cache: speed v. capacity**
 - **Block size (i.e., cache aspect ratio)**
 - **Write Policy (Write through v. write back)**
 - **Associativity choice of N (direct-mapped v. set v. fully associative)**
 - **Block replacement policy**
 - **2nd level cache?**
 - **3rd level cache?**
- **Use performance model to pick between choices, depending on programs, technology, budget, ...**



Another View of the Memory Hierarchy



Memory Hierarchy Requirements

- **If Principle of Locality allows caches to offer (close to) speed of cache memory with size of DRAM memory, then recursively why not use at next level to give speed of DRAM memory, size of Disk memory?**
- **While we're at it, what other things do we need from our memory system?**



Memory Hierarchy Requirements

- Allow multiple **processes** to simultaneously occupy memory and provide protection – don't let one program read/write memory from another
- Address space – give each program the **illusion** that it has its own private memory
 - Suppose code starts at address **0x40000000**. But different processes have different code, both residing at the same address. So each program has a different view of memory.



Virtual Memory

- **Next level in the memory hierarchy:**
 - Provides program with illusion of a very large main memory:
 - Working set of “pages” reside in main memory - others reside on disk.
- Also allows OS to share memory, protect programs from each other
- Today, more important for **protection** vs. just another level of memory hierarchy
- Each process thinks it has all the memory to itself
- (Historically, it predates caches)



Three Advantages of Virtual Memory

- **1) Translation:**
 - **Program can be given consistent view of memory, even though physical memory is scrambled**
 - **Makes multiple processes reasonable**
 - **Only the most important part of program (“Working Set”) must be in physical memory**
 - **Contiguous structures (like stacks) use only as much physical memory as necessary yet still grow later**



Three Advantages of Virtual Memory

▪ 2) Protection:

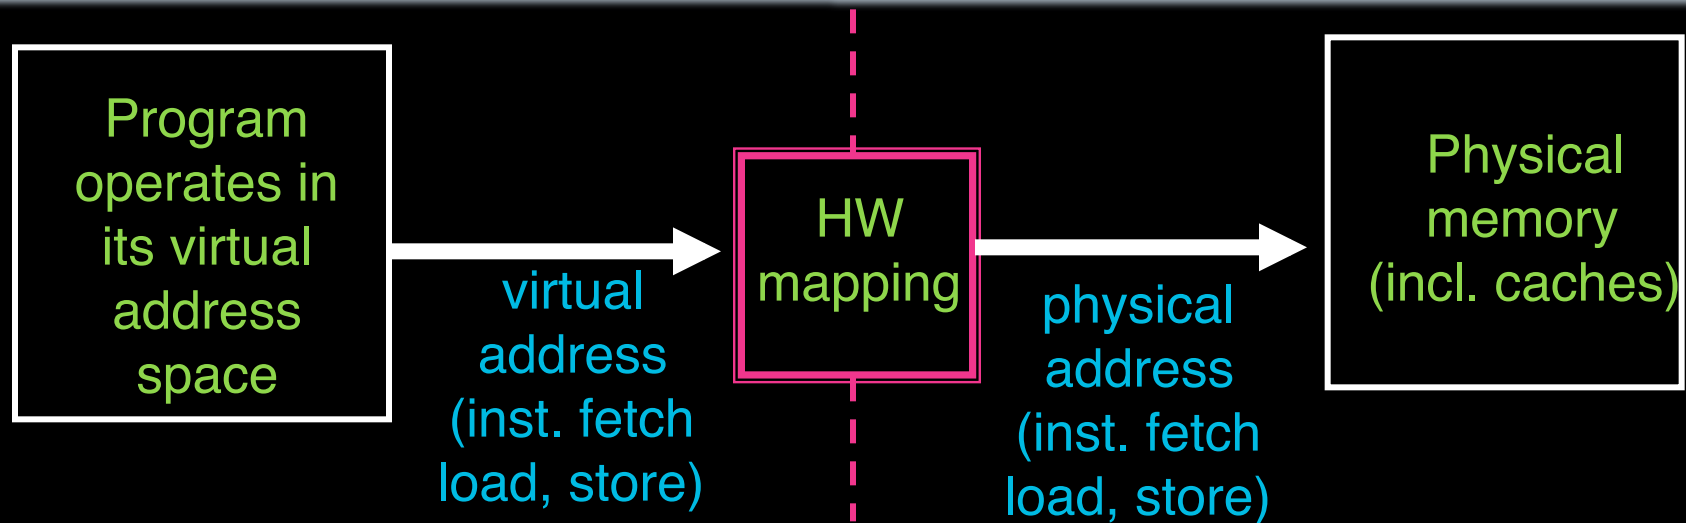
- Different processes protected from each other
- Different pages can be given special behavior
 - (Read Only, Invisible to user programs, etc).
- Kernel data protected from User programs
- Very important for protection from malicious programs ⇒ Far more “viruses” under Microsoft Windows
- Special Mode in processor (“Kernel mode”) allows processor to change page table/TLB

▪ 3) Sharing:

- Can map same physical page to multiple users (“Shared memory”)



Virtual to Physical Address Translation



- **Each program operates in its own virtual address space; ~only program running**
- **Each is protected from the other**
- **OS can decide where each goes in memory**
- **Hardware gives virtual \Rightarrow physical mapping**




Analogy

- Book title like **virtual address**
- Library of Congress call number like **physical address**
- Card catalogue like **page table**, mapping from book title to call #
- On card for book, in local library vs. in another branch like **valid bit** indicating in main memory vs. on disk
- On card, available for 2-hour in library use (vs. 2-week checkout) like **access rights**



Simple Example: Base and Bound Reg



Enough space for User D,
but discontinuous
("fragmentation problem")

- Want:
 - discontinuous mapping
 - Process size \gg mem
 - Addition not enough!
- \Rightarrow use Indirection!




Mapping Virtual Memory to Physical Memory

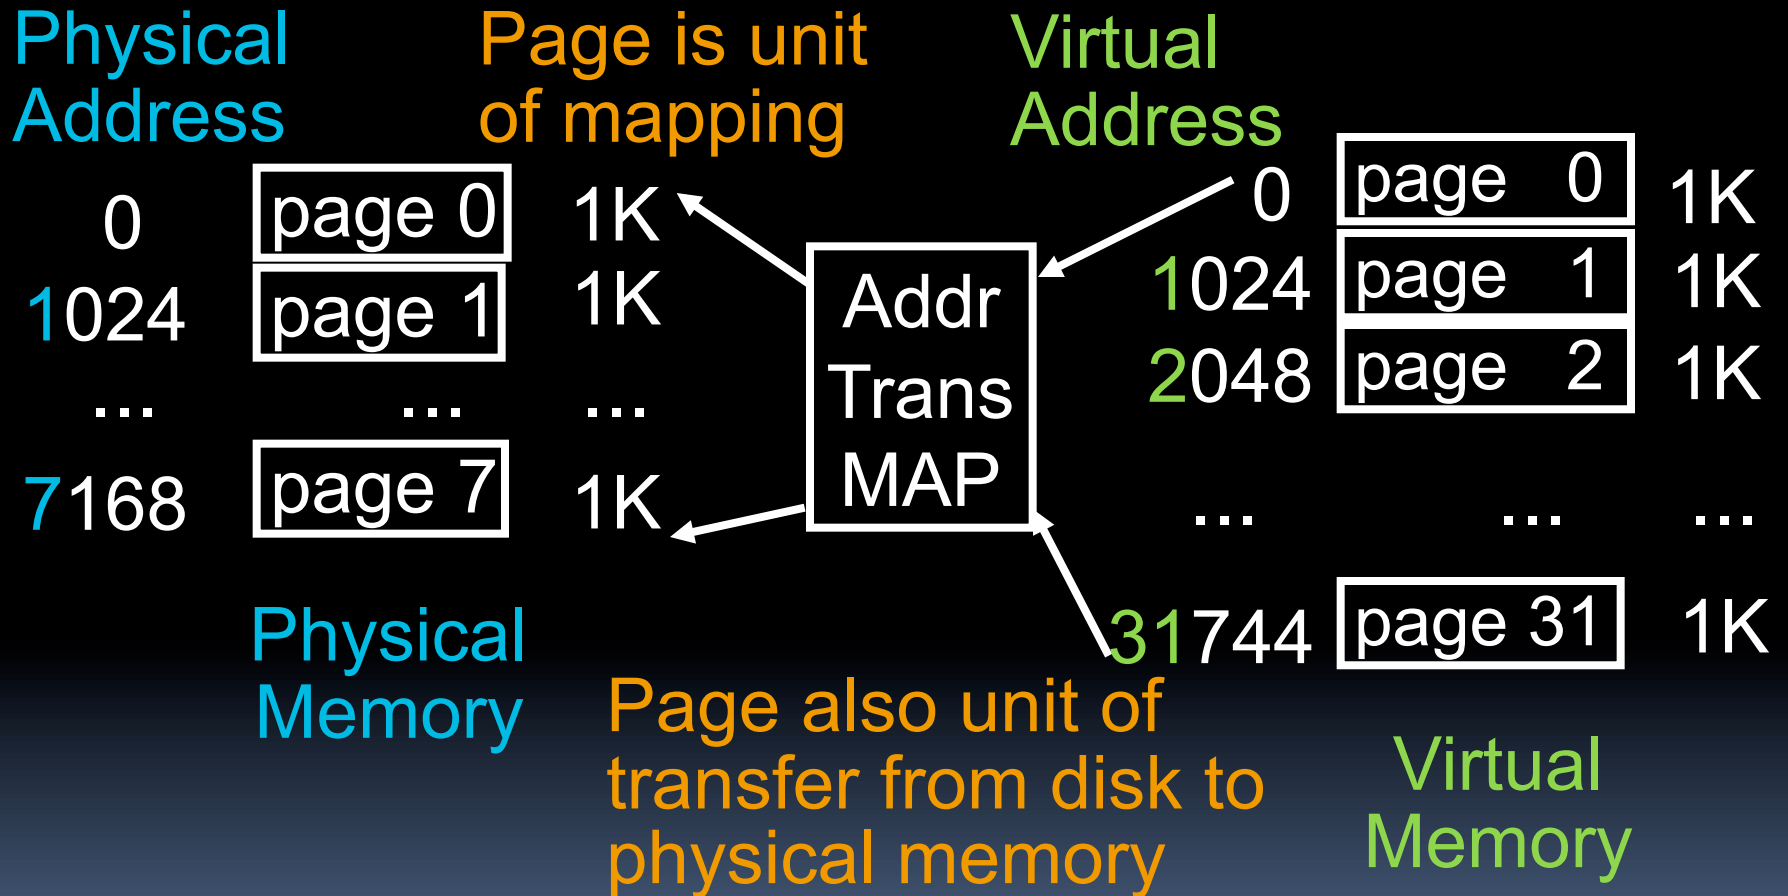
- Divide into equal sized chunks (about 4 KB - 8 KB)
- Any chunk of Virtual Memory assigned to any chunk of Physical Memory (“page”)

64 MB Physical Memory

Virtual Memory



Paging Organization (assume 1 KB pages)



Virtual Memory Mapping Function

- Cannot have simple function to predict arbitrary mapping
- Use table lookup of mappings

Page Number

Offset

- Use table lookup (“Page Table”) for mappings: Page number is index
- Virtual Memory Mapping Function
 - Physical Offset = Virtual Offset
 - Physical Page Number = PageTable[Virtual Page Number]
(P.P.N. also called “Page Frame”)



Page Table

- A page table is an operating system structure which contains the mapping of virtual addresses to physical locations
 - There are several different ways, all up to the operating system, to keep this data around
- Each process running in the operating system has its own page table
 - “**State**” of process is PC, all registers, plus page table
 - OS changes page tables by changing contents of **Page Table Base Register**



Requirements revisited

- Remember the motivation for VM:
- **Sharing memory with protection**
 - Different physical pages can be allocated to different processes (sharing)
 - A process can only touch pages in its own page table (protection)
- **Separate address spaces**
 - Since programs work only with virtual addresses, different programs can have different data/code at the same address!
- What about the memory hierarchy?



Page Table Entry (PTE) Format

- Contains either Physical Page Number or indication not in Main Memory
- OS maps to disk if Not Valid ($V = 0$)

Page Table

...		
V	A.R.	P. P.N.
Val	Access	Physical
-id	Rights	Page
		Number
V	A.R.	P. P. N.
...		

P.T.E.

- If valid, also check if have permission to use page: **Access Rights** (A.R.) may be Read Only, Read/Write, Executable



Paging/Virtual Memory Multiple Processes


User A:

User B:

Virtual Memory

Virtual Memory

Physical Memory



Comparing the 2 levels of hierarchy

Cache version

Block or Line

Miss

Block Size: 32-64B

Placement:

Direct Mapped,
N-way Set Associative

Replacement:

LRU or Random

Write Thru or Back

Virtual Memory vers.

Page

Page Fault

Page Size: 4K-8KB

Fully Associative

Least Recently Used
(LRU)

Write Back



Notes on Page Table

- **Solves Fragmentation problem: all chunks same size, so all holes can be used**
- **OS must reserve “Swap Space” on disk for each process**
- **To grow a process, ask Operating System**
 - **If unused pages, OS uses them first**
 - **If not, OS swaps some old pages to disk**
 - **(Least Recently Used to pick pages to swap)**
- **Each process has own Page Table**
- **Will add details, but Page Table is essence of Virtual Memory**



Why would a process need to “grow”?

- A program’s **address space** contains 4 regions:
 - **stack**: local variables, grows downward
 - **heap**: space requested for pointers via `malloc()` ; resizes dynamically, grows upward
 - **static data**: variables declared outside main, does not grow or shrink
 - **code**: loaded when program starts, does not change

~ FFFF FFFF_{hex}



~ 0_{hex}

For now, OS somehow prevents accesses between stack and heap (gray hash lines).




Virtual Memory Problem #1

- Map every address \Rightarrow 1 indirection via Page Table in memory per virtual address \Rightarrow 1 virtual memory accesses = 2 physical memory accesses \Rightarrow SLOW!
- Observation: since locality in pages of data, there must be locality in **virtual address translations** of those pages
- Since small is fast, why not use a small cache of virtual to physical address translations to make translation fast?
- For historical reasons, this cache is called a **Translation Lookaside Buffer, or TLB**



Translation Look-Aside Buffers (TLBs)

- TLBs usually small, typically 128 - 256 entries
- Like any other cache, the TLB can be direct mapped, set associative, or fully associative




On TLB miss, get page table entry from main memory

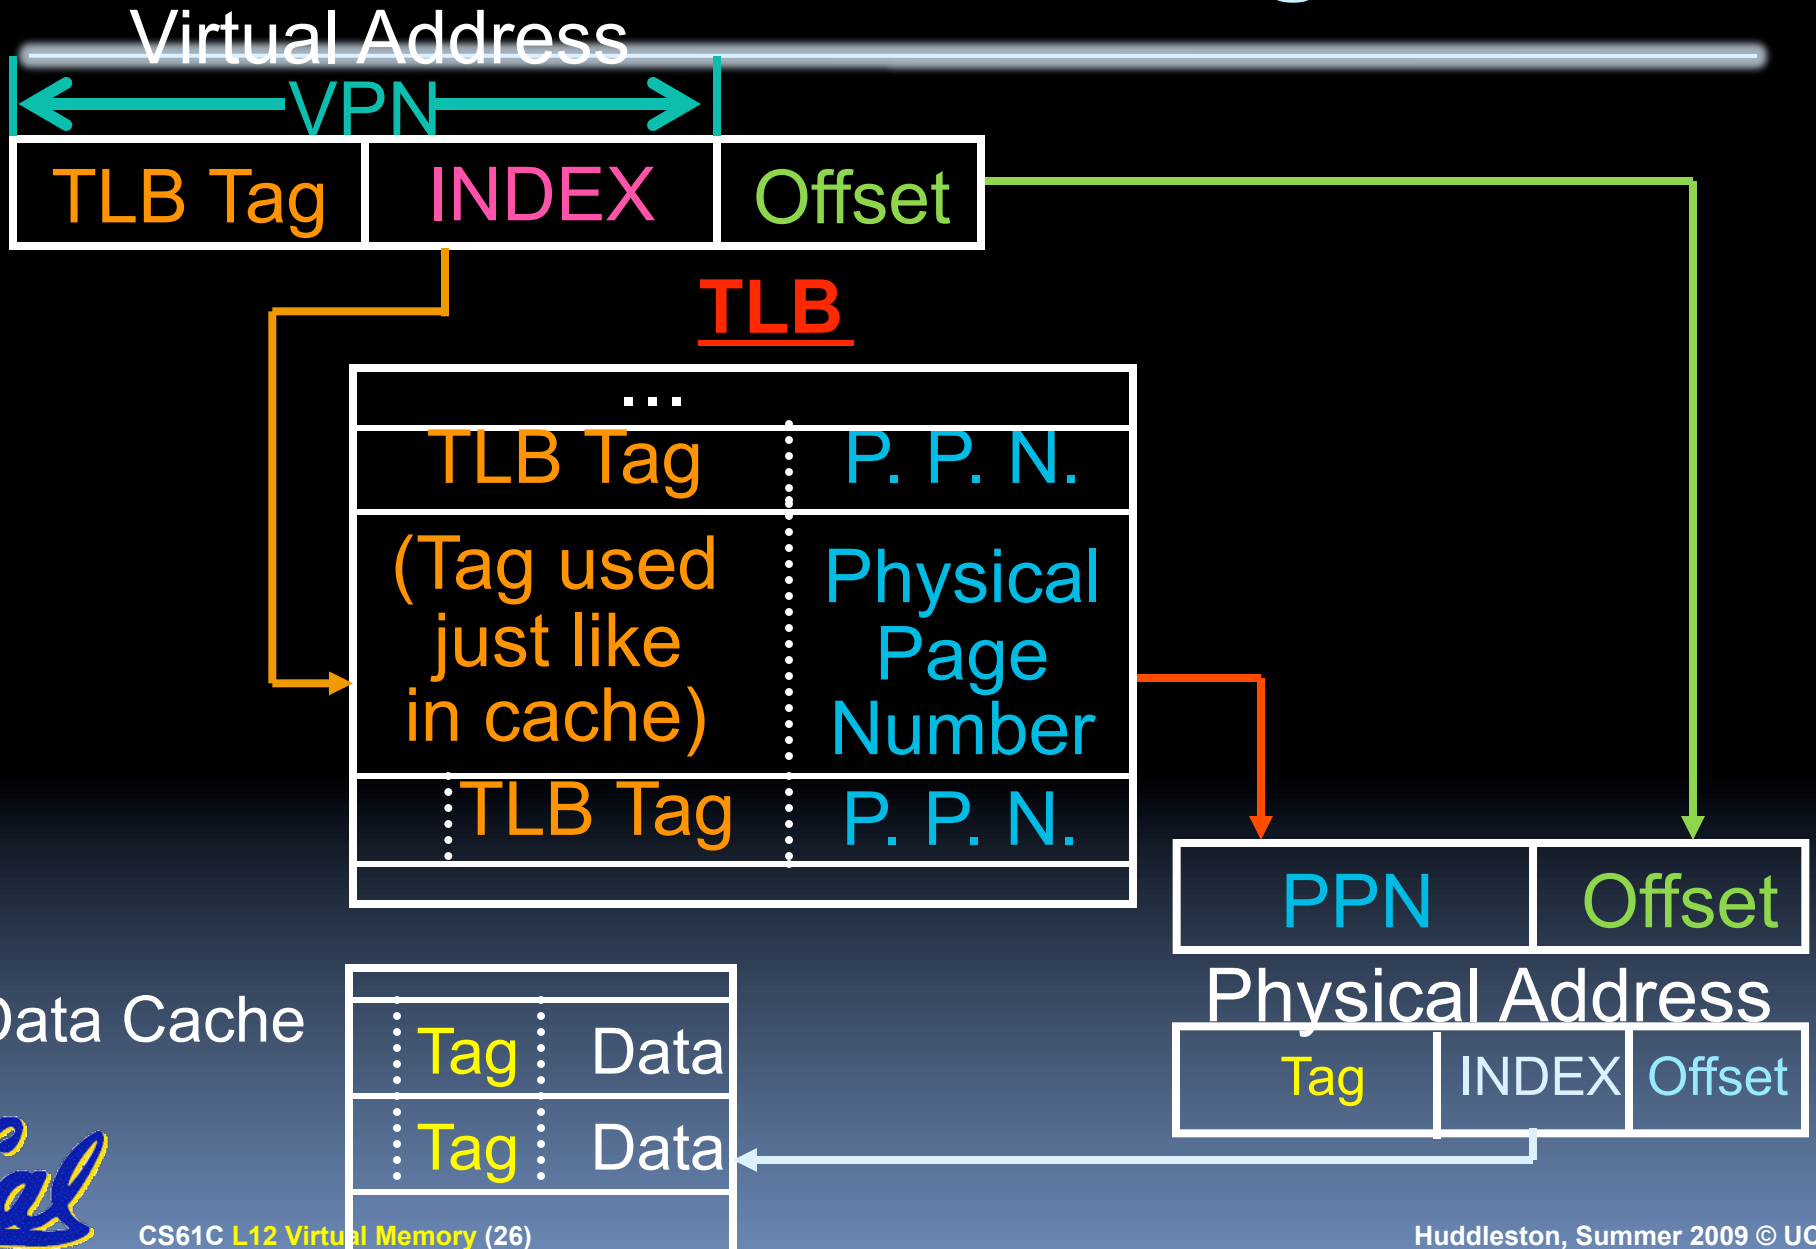


Fetching data on a memory read

- Check TLB (input: VPN, output: PPN)
 - hit: fetch translation
 - miss: check page table (in memory)
 - Page table hit: fetch translation
 - Page table miss: page fault, fetch page from disk to memory, return translation to TLB
- Check cache (input: PPN, output: data)
 - hit: return value
 - miss: fetch value from memory, remember it in cache, return value



Address Translation using TLB



Typical TLB Format

Two ways of looking at it:

Valid	Dirty	Ref	Tag	Cached Data (What is in the PT)		
Valid TLB	Dirty TLB	Ref	Tag TLB	Dirty Page	Access Rights	Physical Page #

- TLB just a cache on the page table mappings
- Dirty Page: since use write back, need to know whether or not to write page to disk when replaced
- Dirty TLB: since we may update the “cached” data (like the dirty page bit), need to know if we need to write back to the page table.
- Ref: Used to help calculate LRU on replacement
 - Cleared by OS periodically, then checked to see if page was referenced



What if not in TLB?

- **Option 1: Hardware checks page table and loads new Page Table Entry into TLB**
- **Option 2: Hardware traps to kernel (OS), up to kernel to decide what to do**
 - **MIPS follows Option 2: Hardware knows nothing about page table**
 - **A trap is a synchronous exception in a user process, often resulting in the kernel taking over and performing some action before returning to the program.**
 - **More about exceptions next lecture**



What if the data is on disk?

- We load the page off the disk into a free block of memory, using a **DMA transfer** (Direct Memory Access – special hardware support to avoid processor)
 - Meantime we switch to some other process waiting to be run
- When the DMA is complete, we get an interrupt and update the process's page table
 - So when we switch back to the task, the desired data will be in memory



What if we don't have enough memory?

- **We chose some other page belonging to a program and transfer it onto the disk if it is dirty**
 - **If clean (disk copy is up-to-date), just overwrite that data in memory**
 - **We chose the page to evict based on replacement policy (e.g., LRU)**
- **And update that program's page table to reflect the fact that its memory moved somewhere else**
- **If continuously swap between disk and memory, called **Thrashing****



Question (1/3)

- 40-bit virtual address, 16 KB page

Virtual Page Number (? bits)

Page Offset (? bits)

- 36-bit physical address

Physical Page Number (? bits)

Page Offset (? bits)

- Number of bits in
Virtual Page Number/Page offset,
Physical Page Number/Page offset?

1: 22/18 (VPN/PO), 22/14 (PPN/PO)

2: 24/16, 20/16

3: 26/14, 22/14

4: 26/14, 26/10

5: 28/12, 24/12



(1/3) Answer

- 40-bit virtual address, 16 KbiB page

Virtual Page Number (26 bits)

Page Offset (14 bits)

- 36-bit physical address

Physical Page Number (22 bits)

Page Offset (14 bits)

- Number of bits in
Virtual Page Number/Page offset,
Physical Page Number/Page offset?

1: 22/18 (VPN/PO), 22/14 (PPN/PO)

2: 24/16, 20/16

3: 26/14, 22/14

4: 26/14, 26/10

5: 28/12, 24/12



Question (2/3): 40b VA, 36b PA

- 2-way assoc. TLB, 512 entries, 40b VA:

TLB Tag (? bits)	TLB Index (? bits)	Page Offset (14 bits)
------------------	--------------------	-----------------------

- TLB Entry: Valid bit, TLB Dirty bit, TLB LRU, Access Control (say 2 bits), Virtual Page Tag, Page Dirty bit, Physical Page Number

V	D	lru	TLB Tag (? bits)	D	Access (2 bits)	Physical Page No. (? bits)
---	---	-----	------------------	---	-----------------	----------------------------

- Number of bits in TLB Tag / Index / Entry?

1: 12 / 14 / 40 (TLB Tag / Index / Entry)
2: 14 / 12 / 42
3: 18 / 8 / 46
4: 18 / 8 / 60



(2/3) Answer

- 2-way set-assoc data cache, 256 (2^8) “sets”, 2 TLB entries per set => 8 bit index

TLB Tag (18 bits)	TLB Index (8 bits)	Page Offset (14 bits)
-------------------	--------------------	-----------------------

Virtual Page Number (26 bits)

- TLB Entry: Valid bit, TLB Dirty bit, LRU bit, Page Dirty bit, Access Control (2 bits), Virtual Page Number, Physical Page Number

V	D	lru	TLB Tag (18 bits)	D	Access (2 bits)	Physical Page No. (22 bits)
---	---	-----	-------------------	---	-----------------	-----------------------------

1: 12 / 14 / 40 (TLB Tag / Index / Entry)

2: 14 / 12 / 42

3: 18 / 8 / 46

4: 18 / 8 / 60



Question (3/3)

- 2-way set-assoc, 64KbiB data cache, 64B block

Cache Tag (? bits)	Cache Index (? bits)	Block Offset (? bits)
--------------------	----------------------	-----------------------

Physical Page Address (36 bits)

- Data Cache Entry: Valid bit, Dirty bit, Cache tag + ? bits of Data

V	D	Cache Tag (? bits)
---	---	--------------------

Cache Data (? bits)

- Number of bits in Data cache Tag / Index / Offset / Entry?

1: 12 / 9 / 14 / 87 (Tag/Index/Offset/Entry)

2: 20 / 10 / 6 / 86

3: 20 / 10 / 6 / 534

4: 21 / 9 / 6 / 87

5: 21 / 9 / 6 / 535



(3/3) Answer

- 2-way set-associative data cache, $64\text{KiB} / 64\text{B} = 1\text{Ki}$ (2^{10}) “sets”, 2 entries per sets \Rightarrow 9 bit index

Cache Tag (21 bits) | Cache Index (9 bits) | Block Offset (6 bits)

Physical Page Address (36 bits)

- Data Cache Entry: Valid bit, Dirty bit, Cache tag + 64 Bytes of Data

V | D | Cache Tag (21 bits)

Cache Data (64 Bytes = 512 bits)

1: 12 / 9 / 14 / 87 (Tag/Index/Offset/Entry)
2: 20 / 10 / 6 / 86
3: 20 / 10 / 6 / 534
4: 21 / 9 / 6 / 87
5: 21 / 9 / 6 / 535



And in conclusion...

- **Manage memory to disk? Treat as cache**
 - **Included protection as bonus, now critical**
 - **Use Page Table of mappings for each user vs. tag/data in cache**
 - **TLB is cache of Virtual \Rightarrow Physical addr trans**
- **Virtual Memory allows protected sharing of memory between processes**
- **Spatial Locality means Working Set of Pages is all that must be in memory for process to run fairly well**



And in conclusion...

- **Virtual memory to Physical Memory Translation too slow?**
 - **Add a cache of Virtual to Physical Address Translations, called a TLB**
- **Spatial Locality means Working Set of Pages is all that must be in memory for process to run fairly well**
- **Virtual Memory allows protected sharing of memory between processes with less swapping to disk**



Bonus slides

- These are extra slides that used to be included in lecture notes, but have been moved to this, the “bonus” area to serve as a supplement.
- The slides will appear in the order they would have in the normal presentation

BONUS




4 Qs for any Memory Hierarchy

- **Q1: Where can a block be placed?**
 - One place (direct mapped)
 - A few places (set associative)
 - Any place (fully associative)
- **Q2: How is a block found?**
 - Indexing (as in a direct-mapped cache)
 - Limited search (as in a set-associative cache)
 - Full search (as in a fully associative cache)
 - Separate lookup table (as in a page table)
- **Q3: Which block is replaced on a miss?**
 - Least recently used (LRU)
 - Random
- **Q4: How are writes handled?**
 - Write through (Level never inconsistent w/lower)
 - Write back (Could be “dirty”, must have dirty bit)




Q1: Where block placed in upper level?


- **Block #12 placed in 8 block cache:**
 - **Fully associative**
 - **Direct mapped**
 - **2-way set associative**
 - **Set Associative Mapping = Block # Mod # of**



Fully associative:
block 12 can go
anywhere



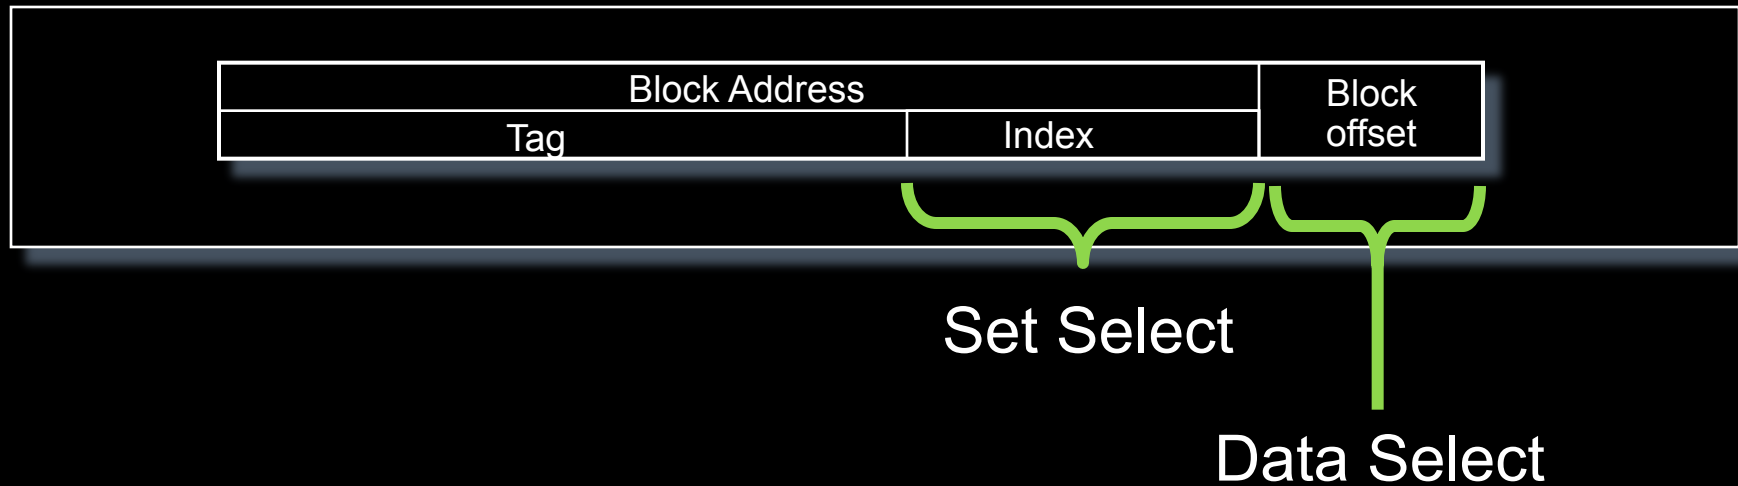
Direct mapped:
block 12 can go
only into block 4
(12 mod 8)



Set Set Set Set
0 1 2 3
Set associative:
block 12 can go
anywhere in set 0
(12 mod 4)



Q2: How is a block found in upper level?



- **Direct indexing (using index and block offset), tag compares, or combination**
- **Increasing associativity shrinks index, expands tag**

Q3: Which block replaced on a miss?

- Easy for Direct Mapped
- Set Associative or Fully Associative:
 - Random
 - LRU (Least Recently Used)

Miss Rates

Associativity: 2-way 4-way 8-way

Size LRU Ran LRU Ran LRU Ran

16 KB 5.2% 5.7% 4.7% 5.3% 4.4% 5.0%

64 KB 1.9% 2.0% 1.5% 1.7% 1.4% 1.5%

256 KB 1.15% 1.17% 1.13% 1.13% 1.12% 1.12%



Q4: What to do on a write hit?

- **Write-through**
 - **update the word in cache block and corresponding word in memory**
- **Write-back**
 - **update word in cache block**
 - **allow memory word to be “stale”**
 - **=> add ‘dirty’ bit to each line indicating that memory be updated when block is replaced**
 - **=> OS flushes cache before I/O !!!**
- **Performance trade-offs?**
 - **WT: read misses cannot result in writes**
 - **WB: no writes of repeated writes**



Why Translation Lookaside Buffer (TLB)?

- **Paging is most popular implementation of virtual memory (vs. base/bounds)**
- **Every paged virtual memory access must be checked against Entry of Page Table in memory to provide protection / indirection**
- **Cache of Page Table Entries (TLB) makes address translation possible without memory access in common case to make fast**



Bonus slide: Virtual Memory Overview (1/3)

- **User program view of memory:**
 - **Contiguous**
 - **Start from some set address**
 - **Infinitely large**
 - **Is the only running program**
- **Reality:**
 - **Non-contiguous**
 - **Start wherever available memory is**
 - **Finite size**
 - **Many programs running at a time**



Bonus slide: Virtual Memory Overview (2/3)

- **Virtual memory provides:**
 - **illusion of contiguous memory**
 - **all programs starting at same set address**
 - **illusion of ~ infinite memory
(2³² or 2⁶⁴ bytes)**
 - **protection**



Bonus slide: Virtual Memory Overview (3/3)

- **Implementation:**
 - **Divide memory into “chunks” (pages)**
 - **Operating system controls page table that maps virtual addresses into physical addresses**
 - **Think of memory as a cache for disk**
 - **TLB is a cache for the page table**



Address Map, Mathematically

$V = \{0, 1, \dots, n - 1\}$ virtual address space ($n > m$)

$M = \{0, 1, \dots, m - 1\}$ physical address space

MAP: $V \rightarrow M \cup \{\theta\}$ address mapping function

MAP(a) = a' if data at virtual address a
is present in physical address a' and a' in M
= θ if data at virtual address a is not present in M

