
5DV118

Computer Organization and Architecture

Umeå University

Department of Computing Science

Stephen J. Hegner

Topic 5: The Memory Hierarchy

Part B: Address Translation

These slides are mostly taken verbatim, or with minor changes, from those prepared by

Mary Jane Irwin (www.cse.psu.edu/~mji)

of The Pennsylvania State University


[Adapted from *Computer Organization and Design, 4th Edition*,

Patterson & Hennessy, © 2008, MK]

Key to the Slides

- ❑ The source of each slide is coded in the footer on the right side:
 - [Irwin CSE331](#) = slide by Mary Jane Irwin from the course CSE331 (Computer Organization and Design) at Pennsylvania State University.
 - [Irwin CSE431](#) = slide by Mary Jane Irwin from the course CSE431 (Computer Architecture) at Pennsylvania State University.
 - [Hegner UU](#) = slide by Stephen J. Hegner at Umeå University.

Review: Major Components of a Computer

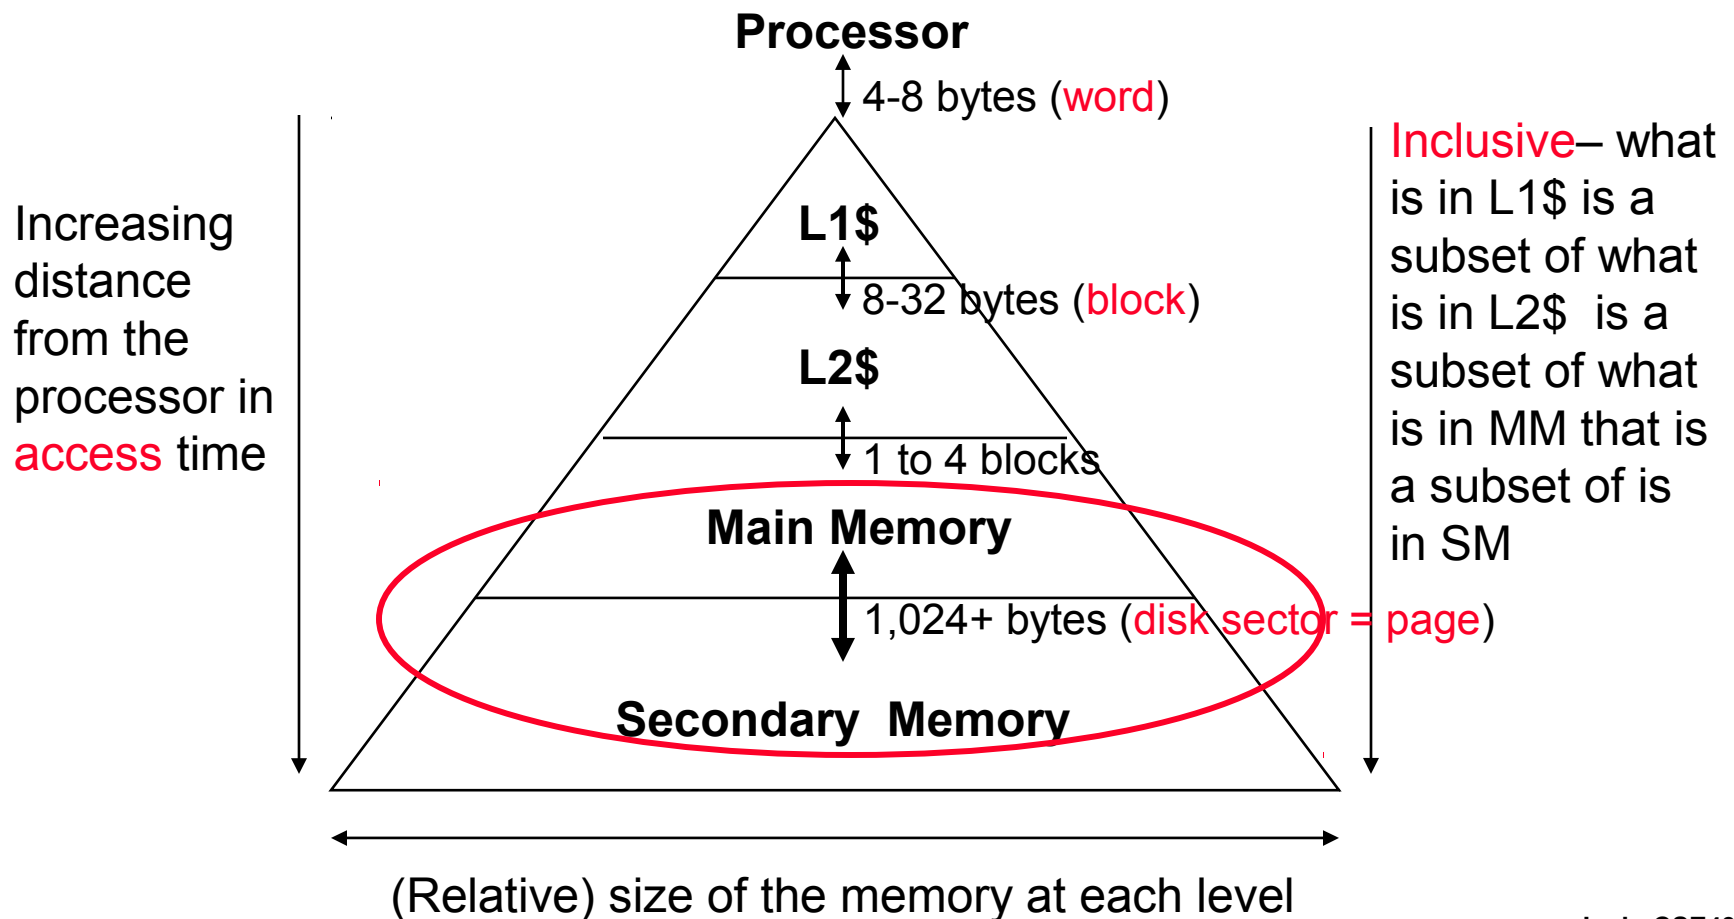


How is the Hierarchy Managed?

- ❑ registers ↔ memory
 - by compiler (programmer?)
- ❑ cache ↔ main memory
 - by the cache controller hardware
- ❑ main memory ↔ disks
 - by the operating system (virtual memory)
 - virtual to physical address mapping assisted by the hardware (TLB)
 - by the programmer (files)

Review: The Memory Hierarchy

- Take advantage of the principle of locality to present the user with as much memory as is available in the cheapest technology at the speed offered by the fastest technology

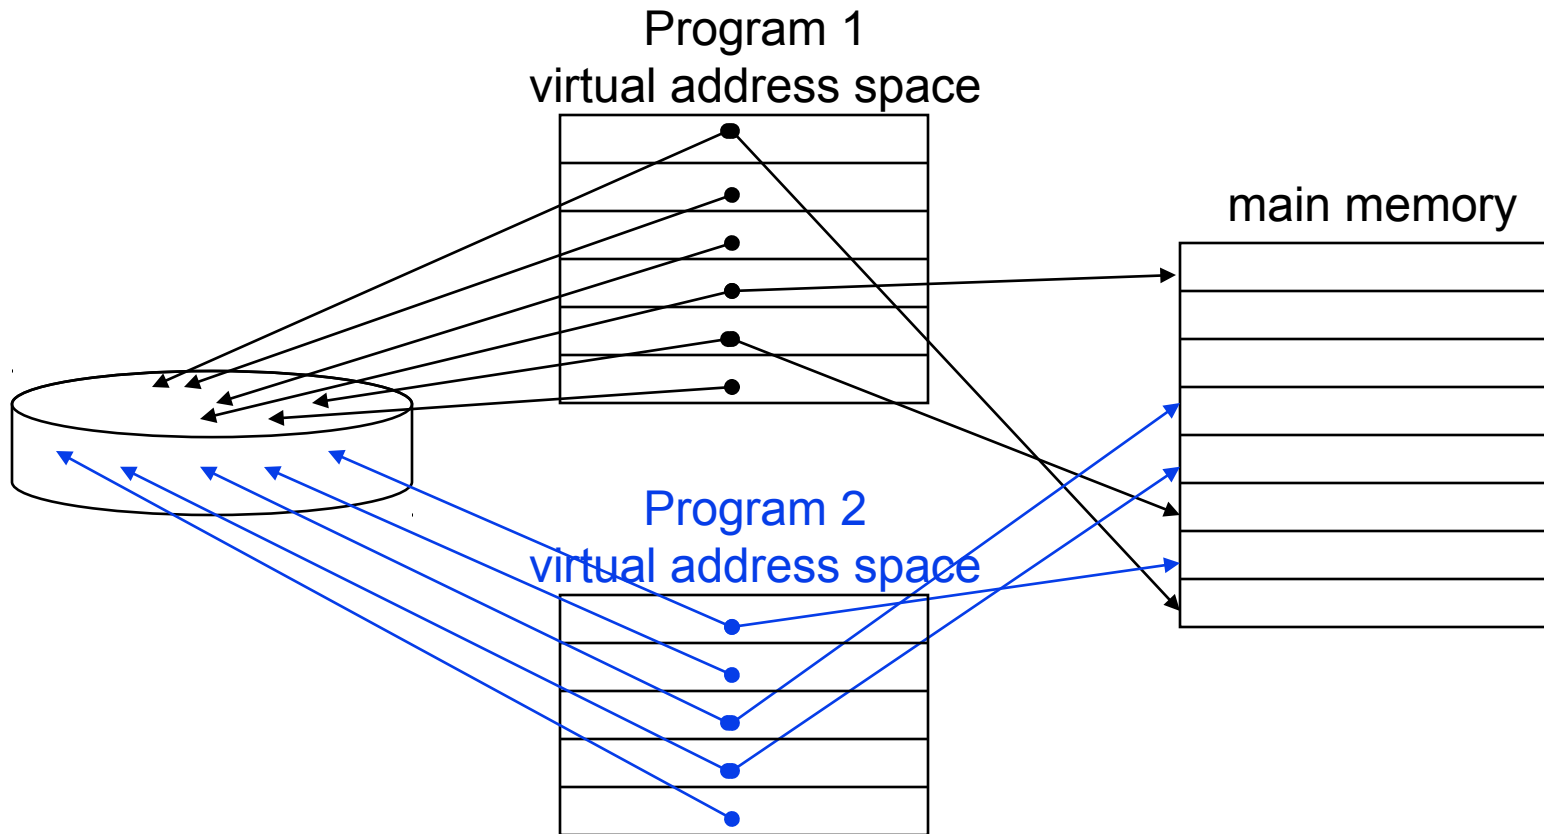


Virtual Memory

- ❑ Use main memory as a “cache” for secondary memory
 - Allows efficient and **safe** sharing of memory among multiple programs
 - Provides the ability to easily run programs larger than the size of physical memory
 - Simplifies loading a program for execution by providing for code relocation (i.e., the code can be loaded anywhere in main memory)
- ❑ What makes it work? – again the Principle of Locality
 - A program is likely to access a relatively small portion of its address space during any period of time
- ❑ Each program is compiled into its own address space – a “virtual” address space
 - During run-time each **virtual** address must be translated to a **physical** address (an address in main memory)

Two Programs Sharing Physical Memory



- ❑ A program's address space is divided into **pages** (all one fixed size) or segments (variable sizes)
 - The starting location of each page (either in main memory or in secondary memory) is contained in the program's **page table**



Address Translation

- A **virtual address** is translated to a **physical address** by a combination of hardware and software

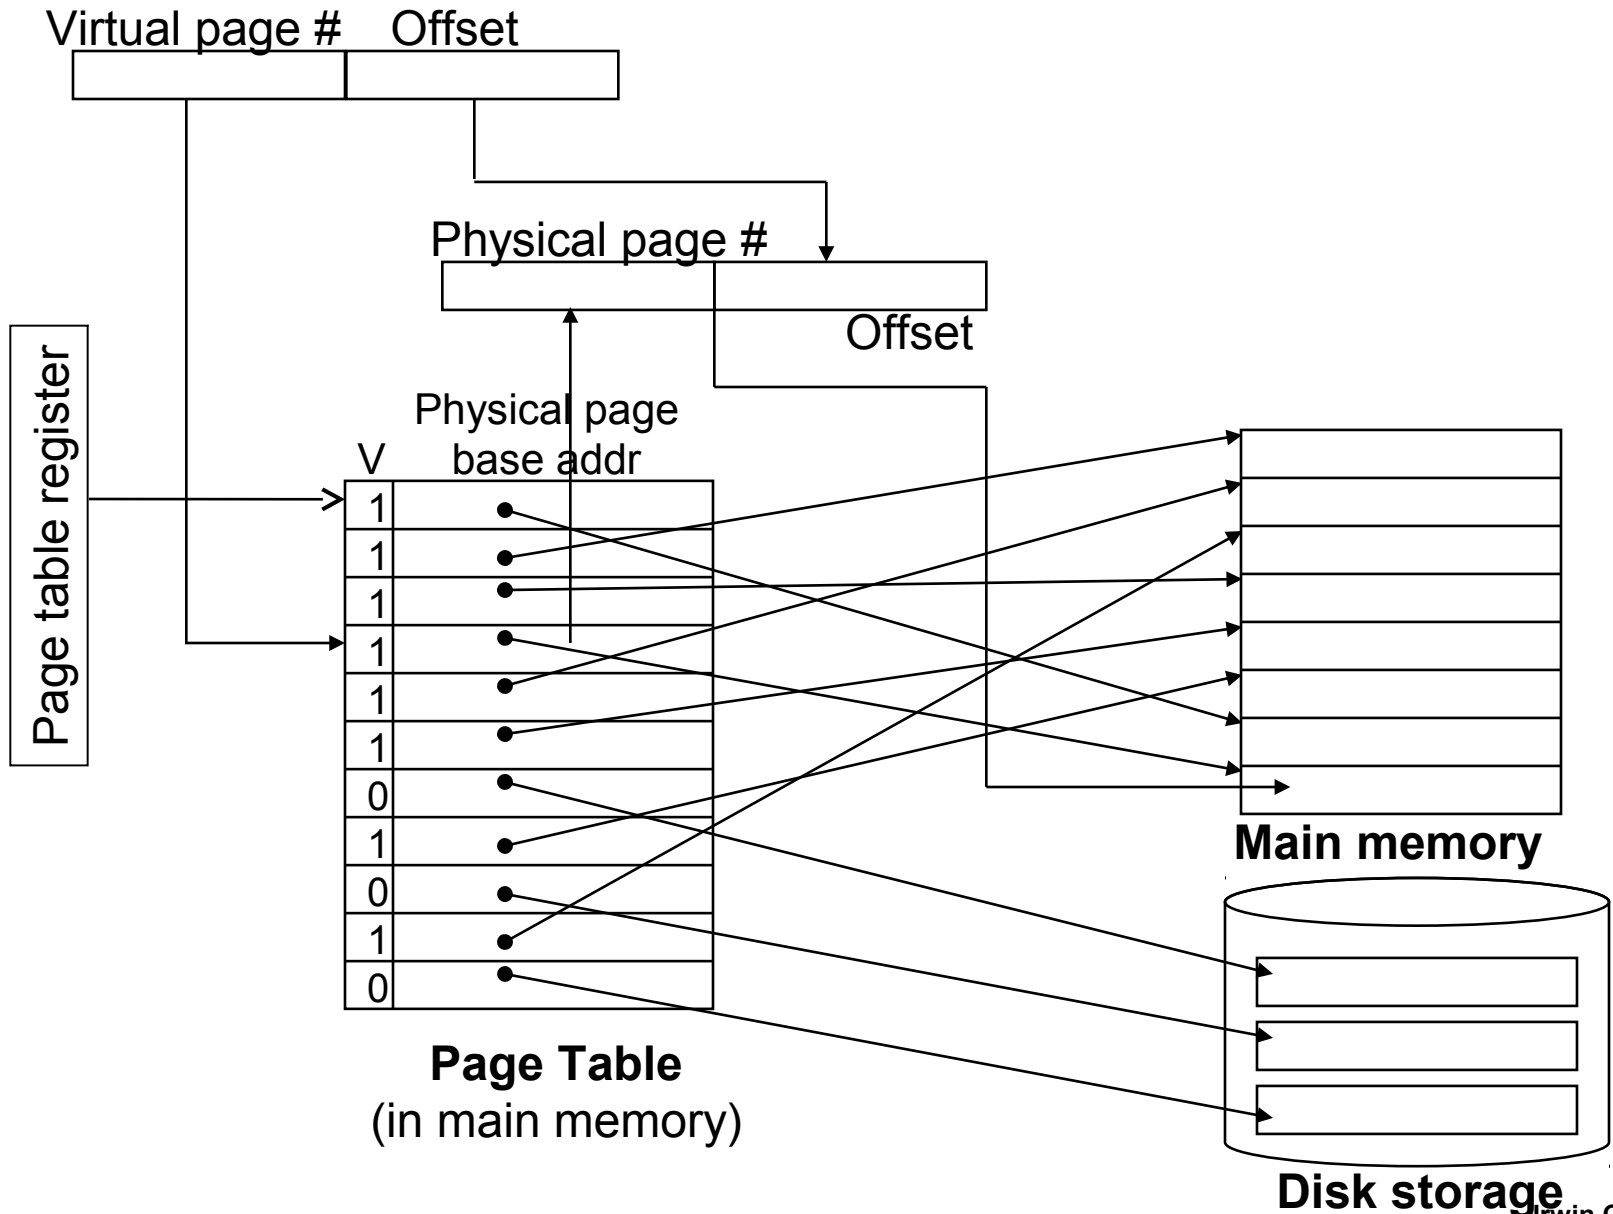
Virtual Address (VA)



Physical Address (PA)


- So each memory request *first* requires an address **translation** from the virtual space to the physical space
 - A virtual memory miss (i.e., when the page is not in physical memory) is called a **page fault**

Address Translation Mechanisms



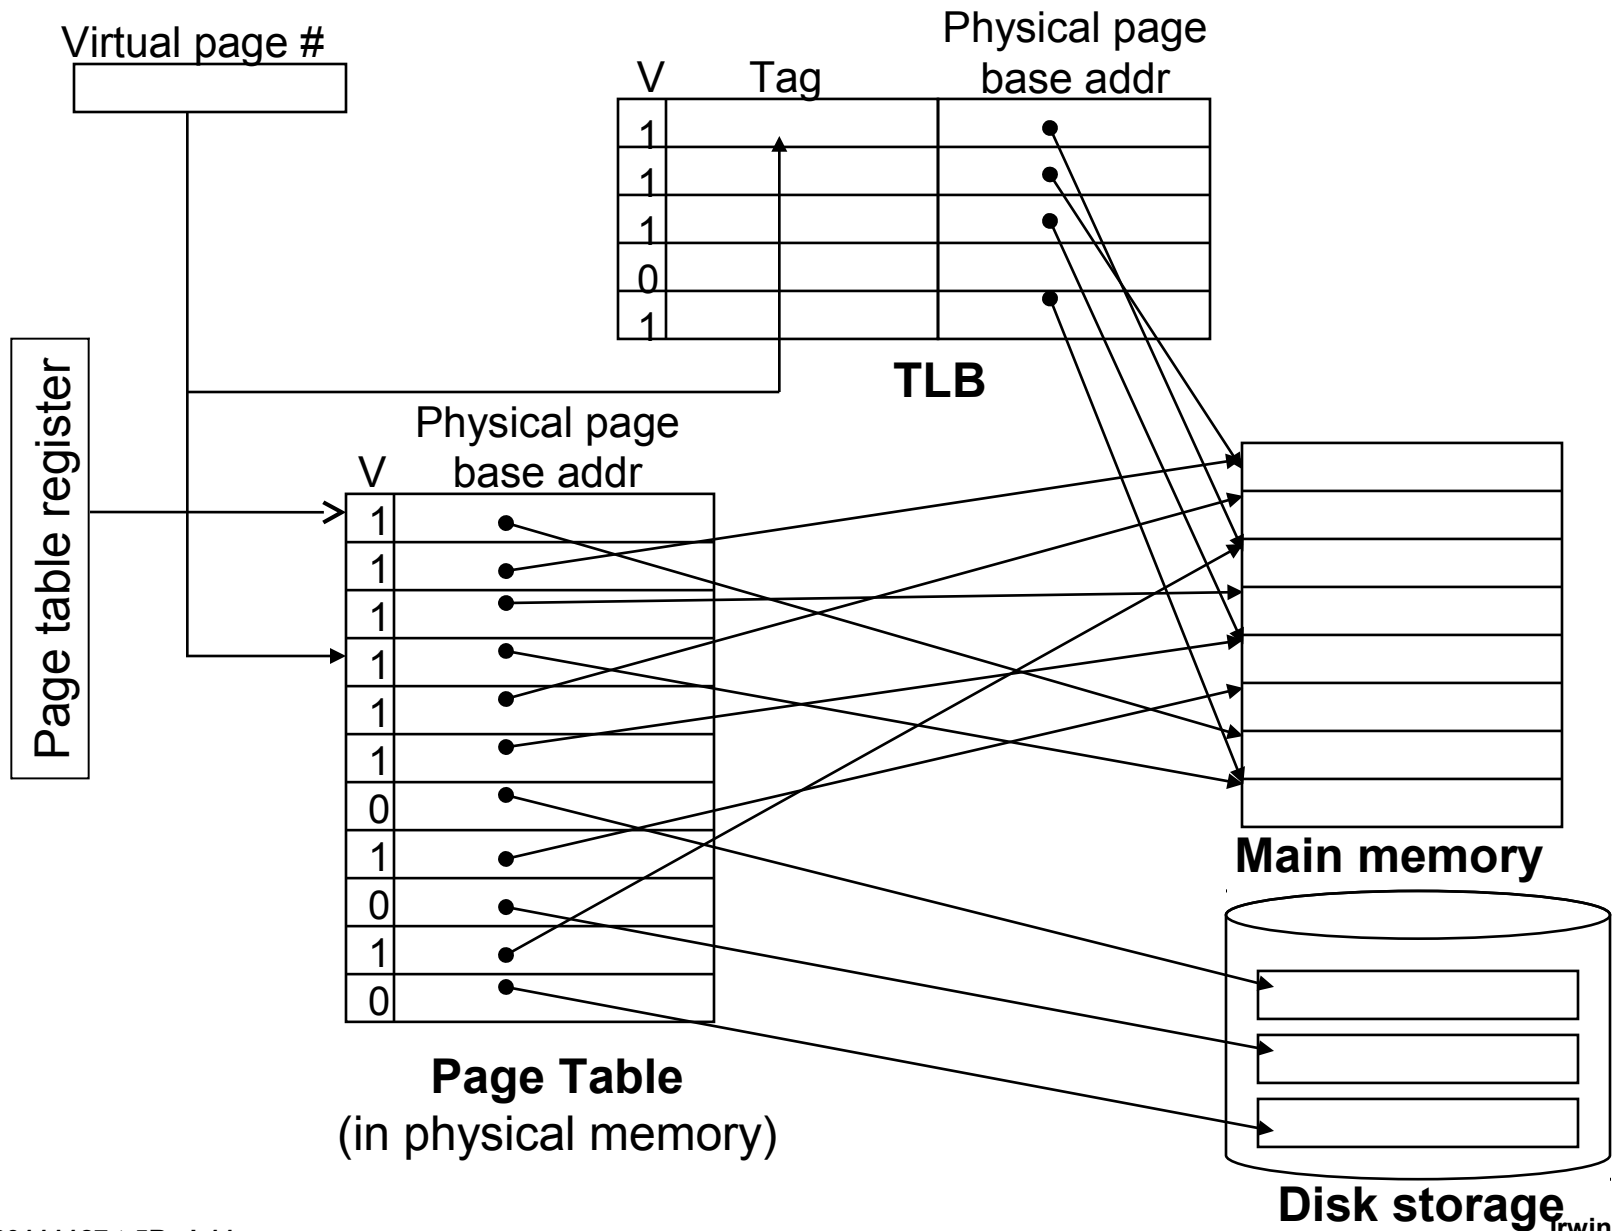
Virtual Addressing with a Cache

- ❑ Thus it takes an *extra* memory access to translate a VA to a PA



- ❑ This makes memory (cache) accesses **very expensive** (if every access was really *two* accesses)
- ❑ The hardware fix is to use a Translation Lookaside Buffer (TLB) – a small cache that keeps track of recently used address mappings to avoid having to do a page table lookup

Making Address Translation Fast



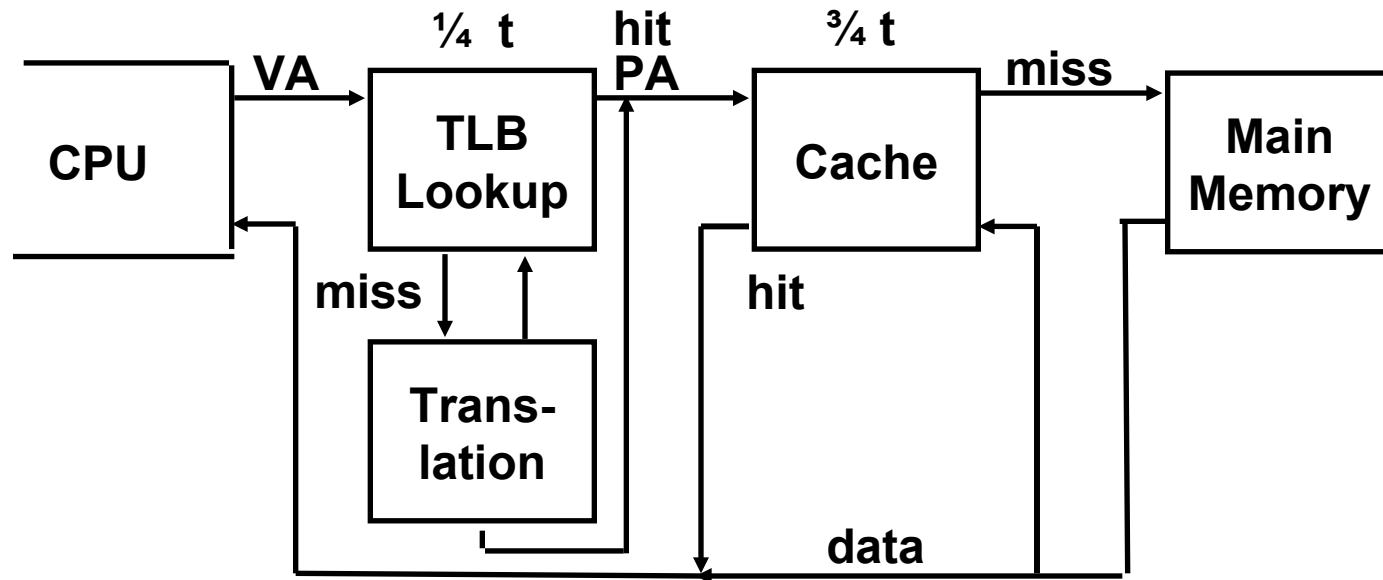
Translation Lookaside Buffers (TLBs)

- ❑ Just like any other cache, the TLB can be organized as fully associative, set associative, or direct mapped

V	Virtual Page #	Physical Page #	Dirty	Ref	Access

- ❑ TLB access time is typically smaller than cache access time (because TLBs are much smaller than caches)
 - TLBs are typically not more than 512 entries even on high end machines

A TLB in the Memory Hierarchy



- ❑ A TLB miss – is it a page fault or merely a TLB miss?
 - If the page is loaded into main memory, then the TLB miss can be handled (in hardware or software) by loading the translation information from the page table into the TLB
 - Takes 10's of cycles to find and load the translation info into the TLB
 - If the page is not in main memory, then it's a true page fault
 - Takes 1,000,000's of cycles to service a page fault
- ❑ TLB misses are much more frequent than true page faults

TLB Event Combinations

TLB	Page Table	Cache	Possible? Under what circumstances?
Hit	Hit	Hit	
Hit	Hit	Miss	
Miss	Hit	Hit	
Miss	Hit	Miss	
Miss	Miss	Miss	
Hit	Miss	Miss/ Hit	
Miss	Miss	Hit	

TLB Event Combinations

TLB	Page Table	Cache	Possible? Under what circumstances?
Hit	Hit	Hit	Yes – what we want!
Hit	Hit	Miss	Yes – although the page table is not checked if the TLB hits
Miss	Hit	Hit	Yes – TLB miss, PA in page table
Miss	Hit	Miss	Yes – TLB miss, PA in page table, but data not in cache
Miss	Miss	Miss	Yes – page fault
Hit	Miss	Miss/ Hit	Impossible – TLB translation not possible if page is not present in memory
Miss	Miss	Hit	Impossible – data not allowed in cache if page is not in memory

Handling a TLB Miss

- ❑ Consider a TLB miss for a page that is present in memory (i.e., the Valid bit in the page table is set)
 - A TLB miss (or a page fault exception) must be asserted by the end of the same clock cycle that the memory access occurs so that the next clock cycle will begin exception processing

Register	CP0 Reg #	Description
EPC	14	Where to restart after exception
Cause	13	Cause of exception
BadVAddr	8	Address that caused exception
Index	0	Location in TLB to be read/written
Random	1	Pseudorandom location in TLB
EntryLo	2	Physical page address and flags
EntryHi	10	Virtual page address
Context	4	Page table address & page number

A MIPS Software TLB Miss Handler

- ❑ When a TLB miss occurs, the hardware saves the address that caused the miss in `BadVAddr` and transfers control to `8000 0000hex`, the location of the TLB miss handler

TLBmiss:

```
mfc0    $k1, Context    #copy addr of PTE into $k1
lw      $k1, 0($k1)     #put PTE into $k1
mtc0    $k1, EntryLo   #put PTE into EntryLo
tlbwr                                       #put EntryLo into TLB
                                                #    at Random
eret                                         #return from exception
```

- ❑ `tlbwr` copies from `EntryLo` into the TLB entry selected by the control register `Random`
- ❑ A TLB miss takes about a dozen clock cycles to handle

Some Virtual Memory Design Parameters

	Paged VM	TLBs
Total size	16,000 to 250,000 words	16 to 512 entries
Total size (KB)	250,000 to 1,000,000,000	0.25 to 16
Block size (B)	4000 to 64,000	4 to 8
Hit time		0.5 to 1 clock cycle
Miss penalty (clocks)	10,000,000 to 100,000,000	10 to 100
Miss rates	0.00001% to 0.0001%	0.01% to 1%

Two Machines' TLB Parameters

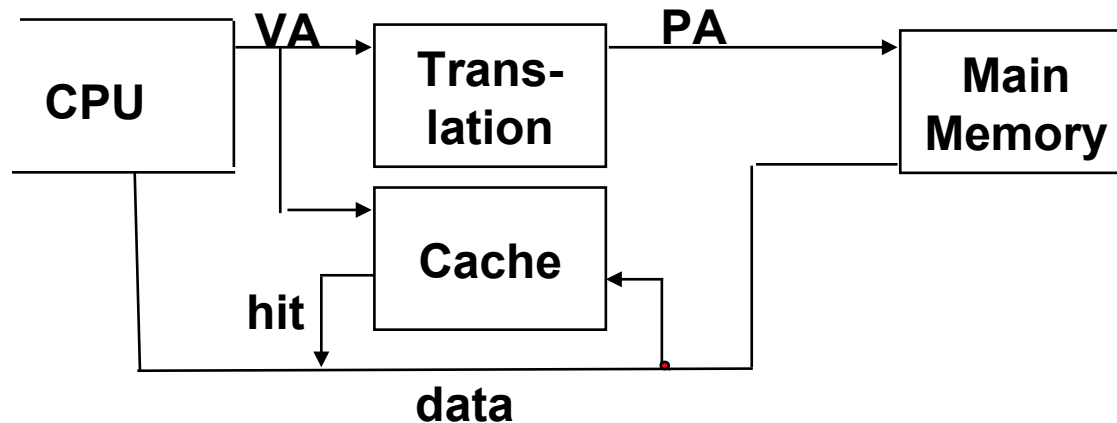
	Intel Nehalem	AMD Barcelona
Address sizes	48 bits (vir); 44 bits (phy)	48 bits (vir); 48 bits (phy)
Page size	4KB	4KB
TLB organization	<p>L1 TLB for instructions and L1 TLB for data per core; both are 4-way set assoc.; LRU</p> <p>L1 ITLB has 128 entries, L2 DTLB has 64 entries</p> <p>L2 TLB (unified) is 4-way set assoc.; LRU</p> <p>L2 TLB has 512 entries</p> <p>TLB misses handled in hardware</p>	<p>L1 TLB for instructions and L1 TLB for data per core; both are fully assoc.; LRU</p> <p>L1 ITLB and DTLB each have 48 entries</p> <p>L2 TLB for instructions and L2 TLB for data per core; each are 4-way set assoc.; round robin LRU</p> <p>Both L2 TLBs have 512 entries</p> <p>TLB misses handled in hardware</p>

Two Machines' TLB Parameters

	Intel P4	AMD Opteron
TLB organization	<p>1 TLB for instructions and 1 TLB for data</p> <p>Both 4-way set associative</p> <p>Both use ~LRU replacement</p> <p>Both have 128 entries</p> <p>TLB misses handled in hardware</p>	<p>2 TLBs for instructions and 2 TLBs for data</p> <p>Both L1 TLBs fully associative with ~LRU replacement</p> <p>Both L2 TLBs are 4-way set associative with round-robin LRU</p> <p>Both L1 TLBs have 40 entries</p> <p>Both L2 TLBs have 512 entries</p> <p>TBL misses handled in hardware</p>

Why Not a Virtually Addressed Cache?

- ❑ A virtually addressed cache would only require address translation on cache misses

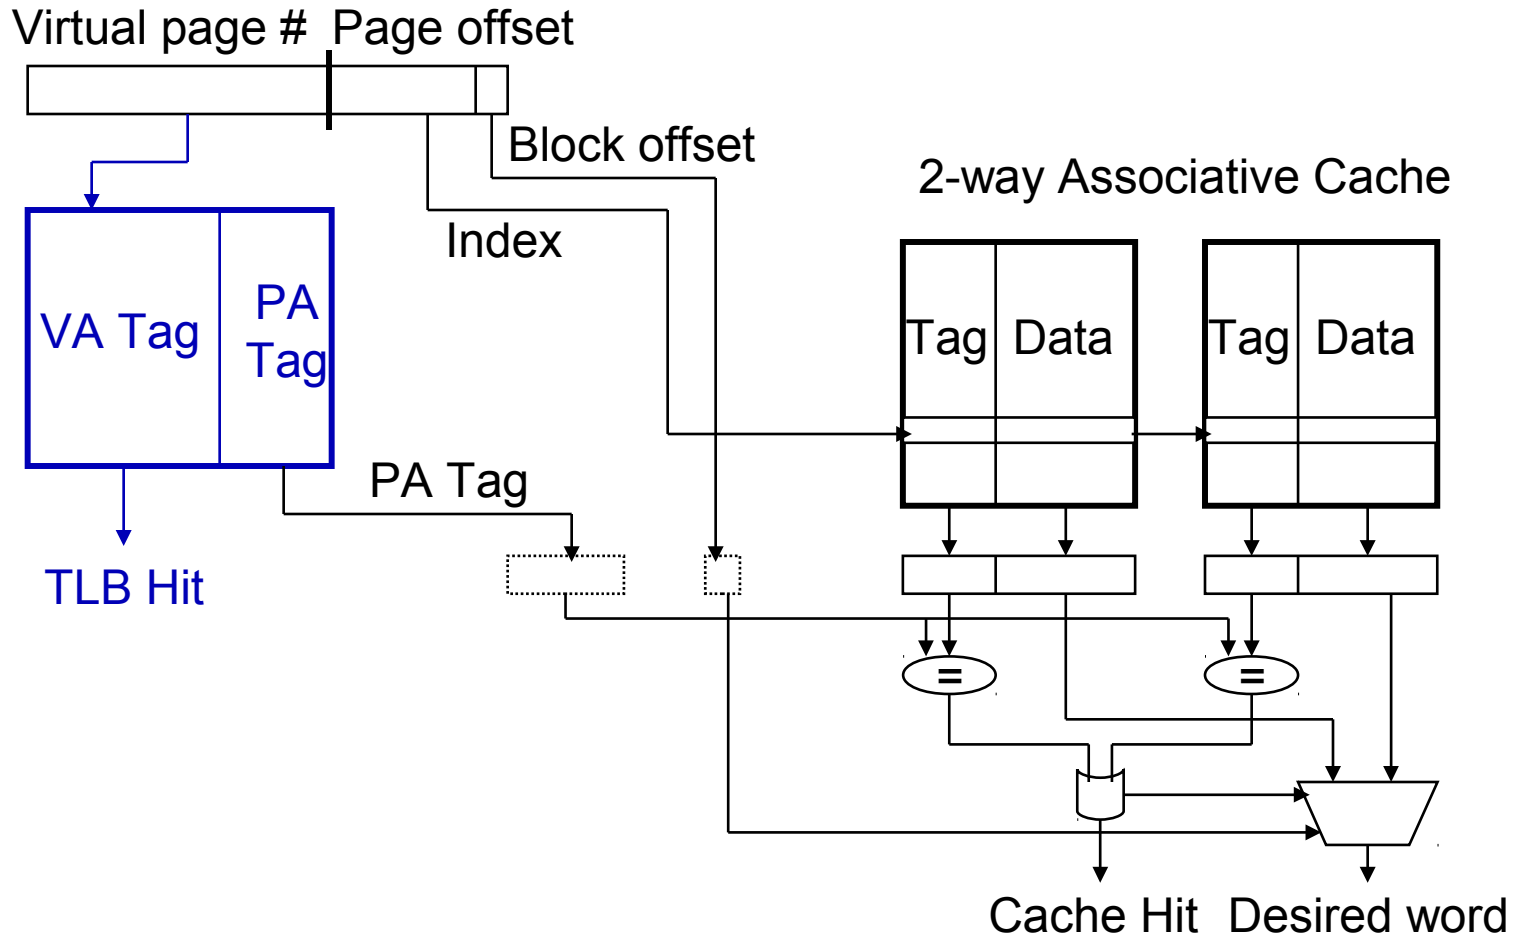


but

- Two programs which are sharing data will have two different virtual addresses for the same physical address – **aliasing** – so have two copies of the shared data in the cache and two entries in the TBL which would lead to coherence issues
 - Must update all cache entries with the same physical address or the memory becomes inconsistent

Reducing Translation Time

- ❑ Can **overlap** the cache access with the TLB access
 - Works when the high order bits of the VA are used to access the TLB while the low order bits are used as index into cache



The Hardware/Software Boundary

- ❑ What parts of the virtual to physical address translation is done by or assisted by the hardware?
 - Translation Lookaside Buffer (TLB) that caches the recent translations
 - TLB access time is part of the cache hit time
 - May allot an extra stage in the pipeline for TLB access
 - Page table storage, fault detection and updating
 - Page faults result in interrupts (precise) that are then handled by the OS
 - Hardware must support (i.e., update appropriately) Dirty and Reference bits (e.g., ~LRU) in the Page Tables
 - Disk placement
 - Bootstrap (e.g., out of disk sector 0) so the system can service a limited number of page faults before the OS is even loaded

4 Questions for the Memory Hierarchy

- ❑ Q1: Where can an entry be placed in the upper level?
(Entry placement)
- ❑ Q2: How is an entry found if it is in the upper level?
(Entry identification)
- ❑ Q3: Which entry should be replaced on a miss?
(Entry replacement)
- ❑ Q4: What happens on a write?
(Write strategy)

Q1&Q2: Where can an entry be placed/found?

	# of sets	Entries per set
Direct mapped	# of entries	1
Set associative	(# of entries)/ associativity	Associativity (typically 2 to 16)
Fully associative	1	# of entries

	Location method	# of comparisons
Direct mapped	Index	1
Set associative	Index the set; compare set's tags	Degree of associativity
Fully associative	Compare all entries' tags Separate lookup (page) table	# of entries 0

Q3: Which entry should be replaced on a miss?

- ❑ Easy for direct mapped – only one choice
- ❑ Set associative or fully associative
 - Random
 - LRU (Least Recently Used)
- ❑ For a 2-way set associative, random replacement has a miss rate about 1.1 times higher than LRU
- ❑ LRU is too costly to implement for high levels of associativity (> 4 -way) since tracking the usage information is costly

Q4: What happens on a write?

- ❑ Write-through – The information is written to the entry in the current memory level *and* to the entry in the next level of the memory hierarchy
 - Always combined with a write buffer so write waits to next level memory can be eliminated (as long as the write buffer doesn't fill)
- ❑ Write-back – The information is written only to the entry in the current memory level. The modified entry is written to next level of memory only when it is replaced.
 - Need a dirty bit to keep track of whether the entry is clean or dirty
 - Virtual memory systems always use write-back of dirty pages to disk
- ❑ Pros and cons of each?
 - Write-through: read misses don't result in writes (so are simpler and cheaper), easier to implement
 - Write-back: writes run at the speed of the cache; repeated writes require only one write to lower level

Summary

- ❑ The Principle of Locality:
 - Program likely to access a relatively small portion of the address space at any instant of time.
 - **Temporal Locality**: Locality in Time
 - **Spatial Locality**: Locality in Space

- ❑ Caches, TLBs, Virtual Memory all understood by examining how they deal with the four questions
 1. Where can entry be placed?
 2. How is entry found?
 3. What entry is replaced on miss?
 4. How are writes handled?

- ❑ Page tables map virtual address to physical address
 - TLBs are important for fast translation