

# 15-213

## VirtualMemory

### April3,2001

#### Topics

- MotivationsforVM
- Addresstranslation
- Acceleratingtranslationwith TLBs

# Motivations for Virtual Memory

- **Use Physical DRAM as a Cache for the Disk**
  - Address space of a process can exceed physical memory size
  - Sum of address spaces of multiple processes can exceed physical memory
- **Simplify Memory Management**
  - Multiple processes resident in main memory.
    - Each process with its own address space
  - Only “active” code and data is actually in memory
    - Allocate more memory to process as needed.

## Provide Protection

- One process can't interfere with another.
  - because they operate in different address spaces.
- User process cannot access privileged information
  - different sections of address spaces have different permissions.

# Motivation#1: DRAM a “Cache” for Disk


Full address space is quite large:

- 32-bit addresses: ~4,000,000,000 (4 billion) bytes
- 64-bit addresses: ~16,000,000,000,000,000,000 (16 quintillion) bytes


Disk storage is ~170X cheaper than DRAM storage

- 20GB of DRAM: ~\$20,000
- 20GB of disk: ~\$120

To access large amounts of data in a cost-effective manner, the bulk of the data must be stored on disk



# Levels in Memory Hierarchy



	Register	Cache	Memory	DiskMemory
size:	32B	32KB -4MB	128MB	20GB
speed:	1ns	2ns	50ns	8ms
\$/Mbyte:		\$100/MB	\$1.00/MB	\$0.006/MB
linesize:	8B	32B	4KB	


larger, slower, cheaper



# DRAM vs. SRAM as a “Cache”

## DRAM vs. disk is more extreme than SRAM vs. DRAM

- **Access latencies:**
  - DRAM ~ 10X slower than SRAM
  - Disk ~ **100,000X** slower than DRAM
- **Importance of exploiting spatial locality:**
  - First byte is ~ **100,000X** slower than successive bytes on disk
    - » vs. ~ 4X improvement for page -mode vs. regular accesses to DRAM
- **Bottomline:**
  - Design decisions made for DRAM caches driven by enormous cost of misses



# Impact of These Properties on Design

If DRAM was to be organized similar to an SRAM cache, how would we set the following design parameters?

- Line size?
  -
- Associativity?
  -
- Write through or write back?
  -

What should the impact of these choices be on:

- miss rate
  -
- hit time
  -
- miss latency
  -
- tag storage overhead
  -

# Locating an Object in a “Cache”

## SRAM Cache

- Tag stored with cache line
- Maps from cache block to memory blocks
  - From cached to uncached form
- Not a tag for block not in cache
- Hardware retrieves information
  - can quickly match against multiple tags

*ObjectName*

X

=X?


**“Cache”**

	Tag	Data
0:	D	243
1:	X	17
	⋮	⋮
N-1:	J	105

# Locating an Object in a “Cache” (cont.)

## DRAMCache

- Each allocated page of virtual memory has an entry in the page table
- Mapping from virtual pages to physical pages
  - From uncached form to cached form
- Page table entry even if page not in memory
  - Specifies disk address
- OS retrieves information






# A System with Physical Memory Only

## Examples:

- most Cray machines, early PCs, nearly all embedded systems, etc.

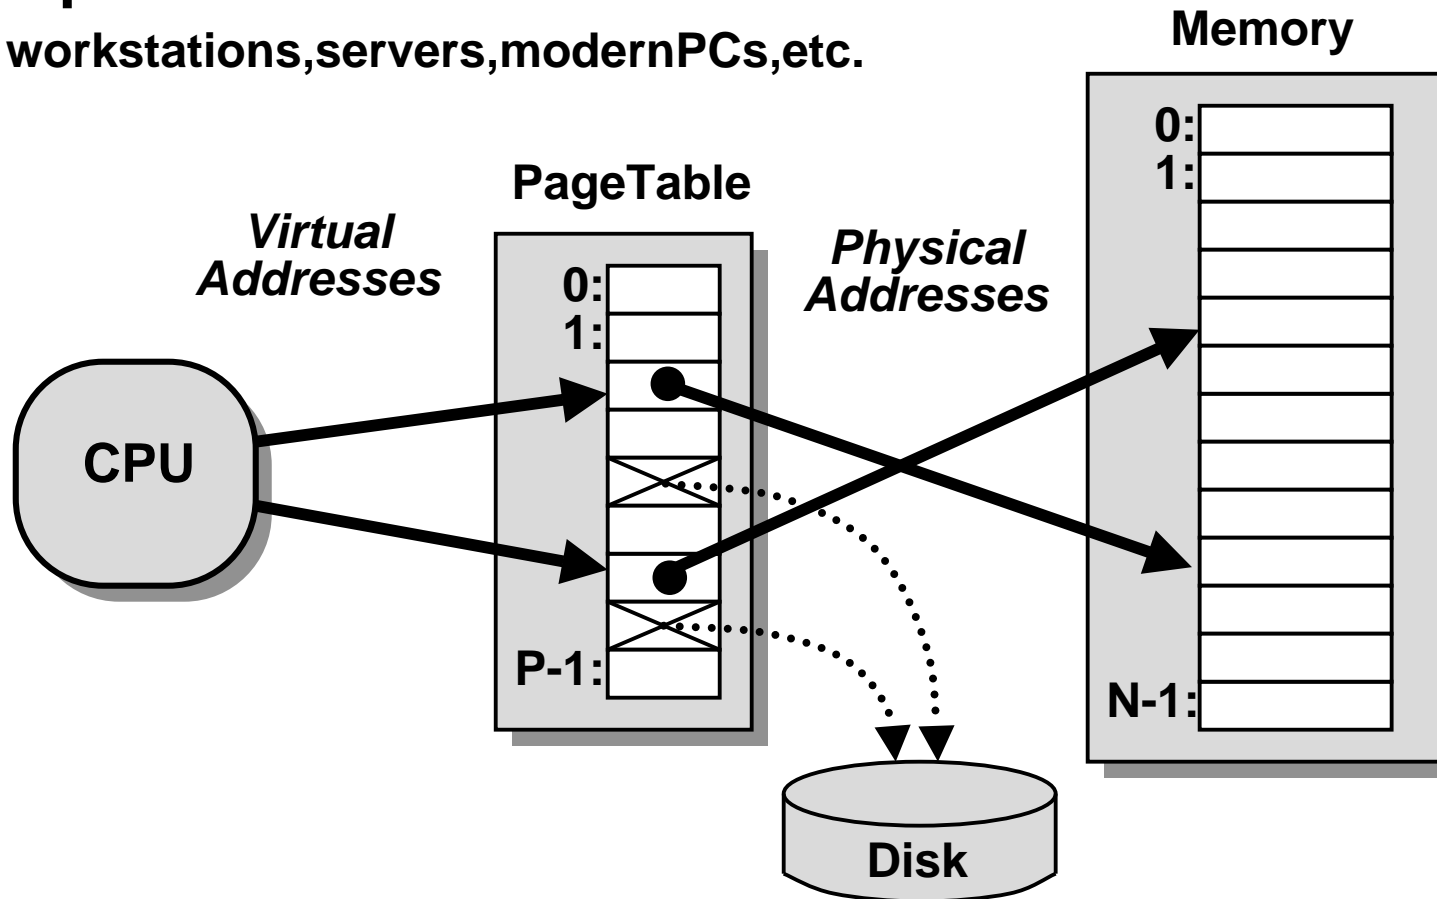


Addresses generated by the CPU point directly to bytes in physical memory

# A System with Virtual Memory

## Examples:

- workstations, servers, modern PCs, etc.




**Address Translation:** Hardware converts *virtual addresses* to *physical addresses* via an OS -managed lookup table( *pagetable* )

# PageFaults(Similar to “CacheMisses”)

## What if an object is on disk rather than in memory?


- Page table entry indicates virtual address not in memory
- OS exception handler invoked to move data from disk into memory
  - current process suspends, others can resume
  - OS has full control over placement, etc.

### Before fault



class20.ppt

### After fault



CS213S'01

# Servicing a Page Fault

## Processor Signals Controller


- Read block of length  $P$  starting at disk address  $X$  and store starting at memory address  $Y$

## Read Occurs

- Direct Memory Access (DMA)
- Under control of I/O controller

## I/O Controller Signals Completion

- Interrupt processor
- OS resumes suspended process



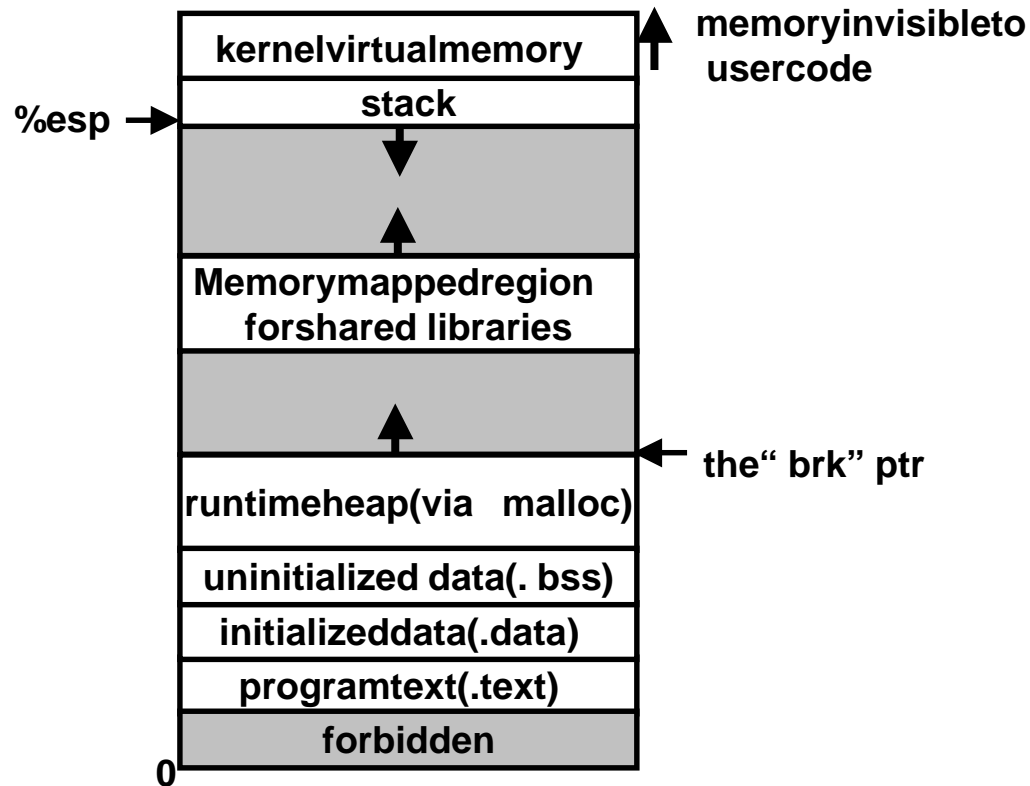
# Motivation#2:MemoryManagement

Multipleprocessescanresideinphysicalmemory.

Howdoweresolveaddressconflicts?


- whatiftwoprocessesaccesssomethingatthesameaddress?

Linux/x86  
process  
memory  
image



# Solution: Separate Virtual Addr. Spaces

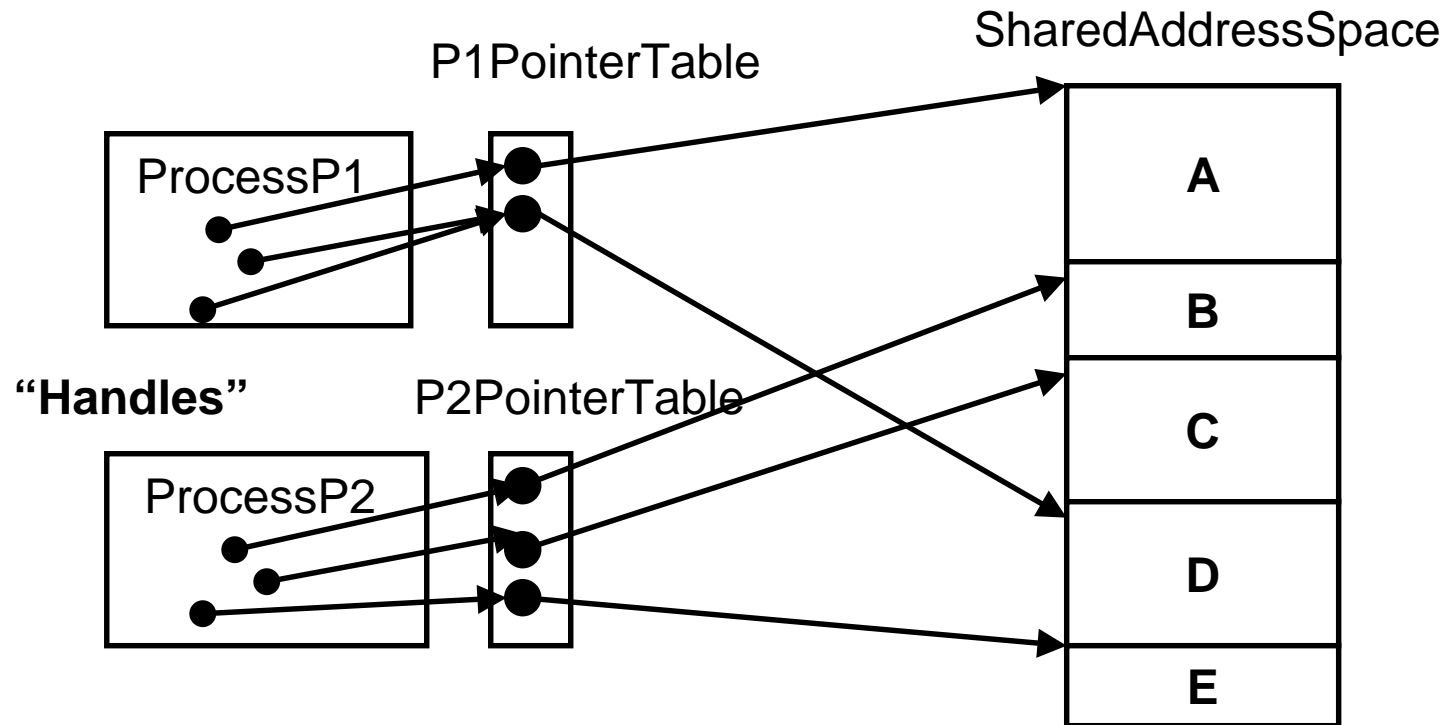
- Virtual and physical address spaces divided into equal -sized blocks
  - blocks are called “pages” (both virtual and physical)
- Each process has its own virtual address space
  - operating system controls how virtual pages are assigned to physical memory



# Contrast:MacintoshMemoryModel

## MACOS1 -9

- Doesnotusetraditionalvirtualmemory



**Allprogramobjectsaccessedthrough“handles”**

- Indirectreferencethroughpointertable
- Objectsstoredinsharedglobaladdressspace

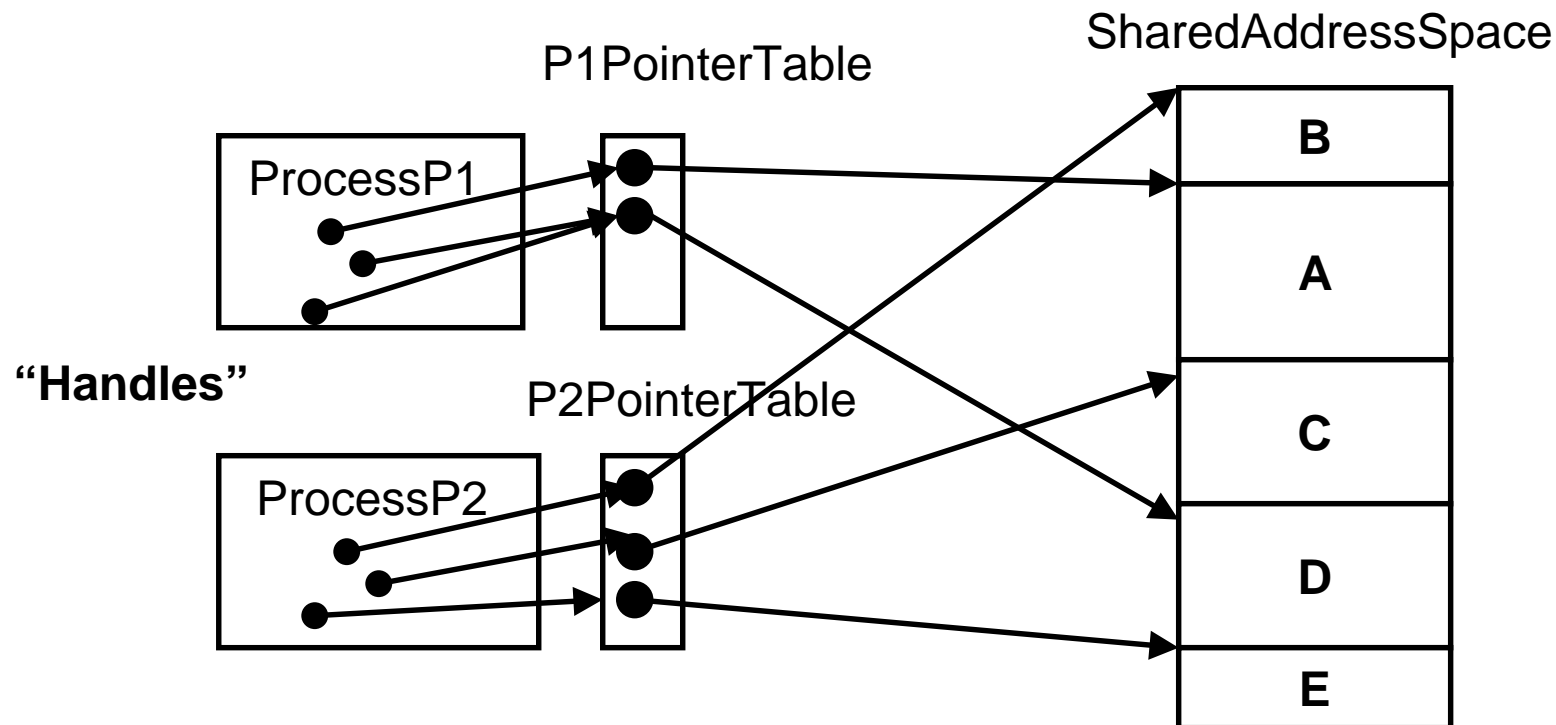
# MacintoshMemoryManagement

## Allocation/ Deallocation

- Similar to free-list management of malloc/free

## Compaction

- Can move any object and just update the (unique) pointer in pointer table





# Mac vs. VM - Based Memory Mgmt

## Allocating, deallocating, and moving memory:

- can be accomplished by both techniques

## Block sizes:

- **Mac: variable -sized**
  - may be very small or very large
- **VM: fixed -size**
  - size is equal to *one page* (4KB on x86 Linux systems)

## Allocating contiguous chunks of memory:

- **Mac: contiguous allocation is *required***
- **VM: can map contiguous range of virtual addresses to disjoint ranges of physical addresses**

## Protection

- **Mac: “wildwrite” by one process can corrupt another’s data**

# MACOSX

## “Modern” Operating System

- Virtual memory with protection
- Preemptive multitasking
  - Other versions of MACOS require processes to voluntarily relinquish control

## Based on MACHOS

- Developed at CMU in late 1980's

# Motivation#3:Protection

## Protectiongoals:

- Cannotread/writememoryfromanotherprocess
- Cannotwriteintosharedlibraries

## Processescanonlyseevirtualaddresses

- Cannotgettophysicaladdressesdirectly
- Canonlygothroughthepagetable
- Ifaphysicalpageisnotinaprocess'pagetable,itis“invisible”

## Pagetableentrycontainsaccessrightsinformation

- hardwareenforcesthisprotection(trapintoOSifviolationoccurs)
- Thepagetableitselfisinprotectedmemory

## Whenallocatinganewphysicalpage,itiscleared

- Importantthattheprocesscannotseethepreviouscontents

# Protection: Example

PageTables


Memory

Processi:

	Read?	Write?	Physical Addr
VP0:	Yes	No	PP9
VP1:	Yes	Yes	PP4
VP2:	No	No	XXXXXXXX
	⋮	⋮	⋮

Processj:

	Read?	Write?	Physical Addr
VP0:	Yes	Yes	PP6
VP1:	Yes	No	PP9
VP2:	No	No	XXXXXXXX
	⋮	⋮	⋮



- Processi and j can only read physical page 9
- Process i cannot even see page 6

# Motivations for Virtual Memory

- **Use Physical DRAM as a Cache for the Disk**
  - Address space of a process can exceed physical memory size
  - Sum of address spaces of multiple processes can exceed physical memory
- **Simplify Memory Management**
  - Multiple processes resident in main memory.
    - Each process with its own address space
  - Only “active” code and data is actually in memory
    - Allocate more memory to process as needed.

## Provide Protection

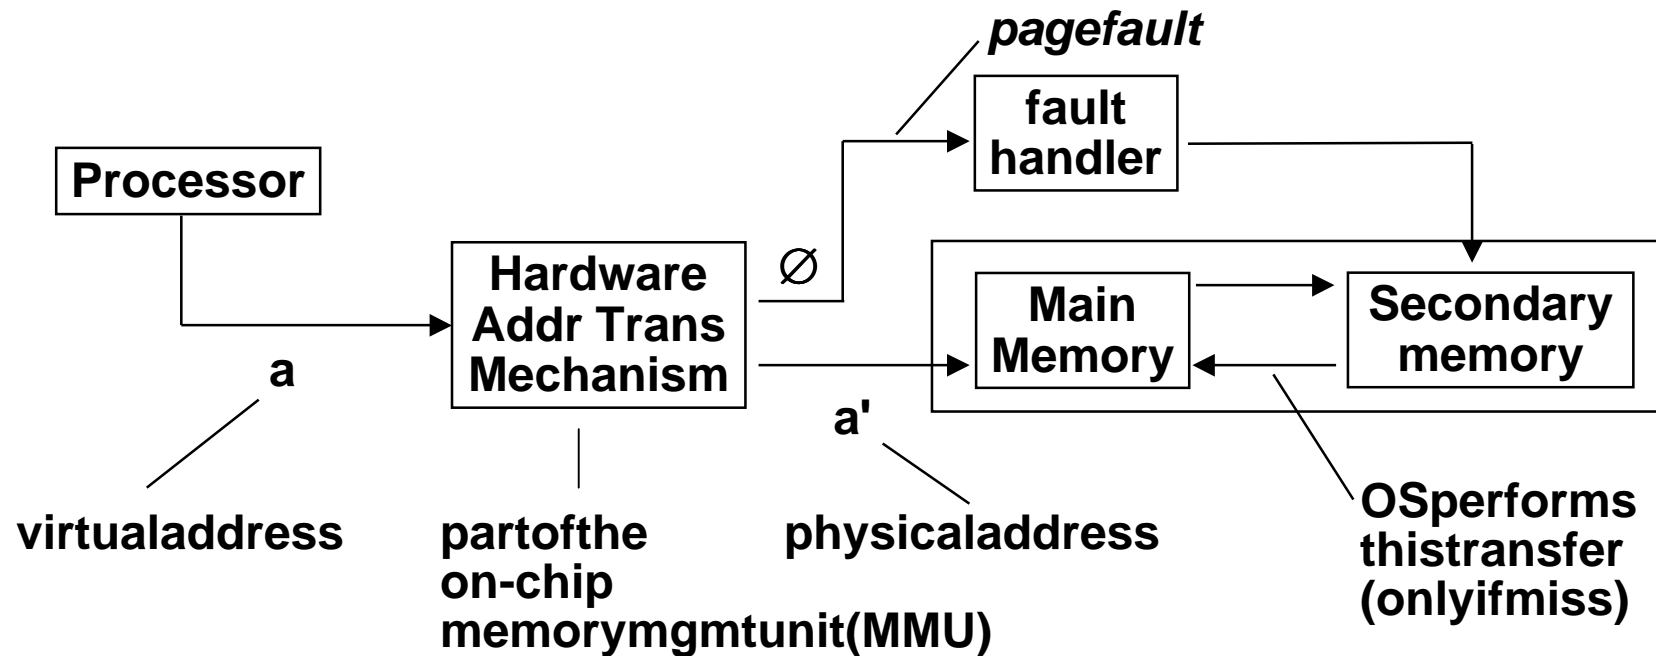
- One process can't interfere with another.
  - because they operate in different address spaces.
- User process cannot access privileged information
  - different sections of address spaces have different permissions.

# VMAddressTranslation

$V = \{0, 1, \dots, N-1\}$  virtual address space  $N > M$   
 $P = \{0, 1, \dots, M-1\}$  physical address space

$MAP: V \rightarrow P \cup \{\emptyset\}$  address mapping function

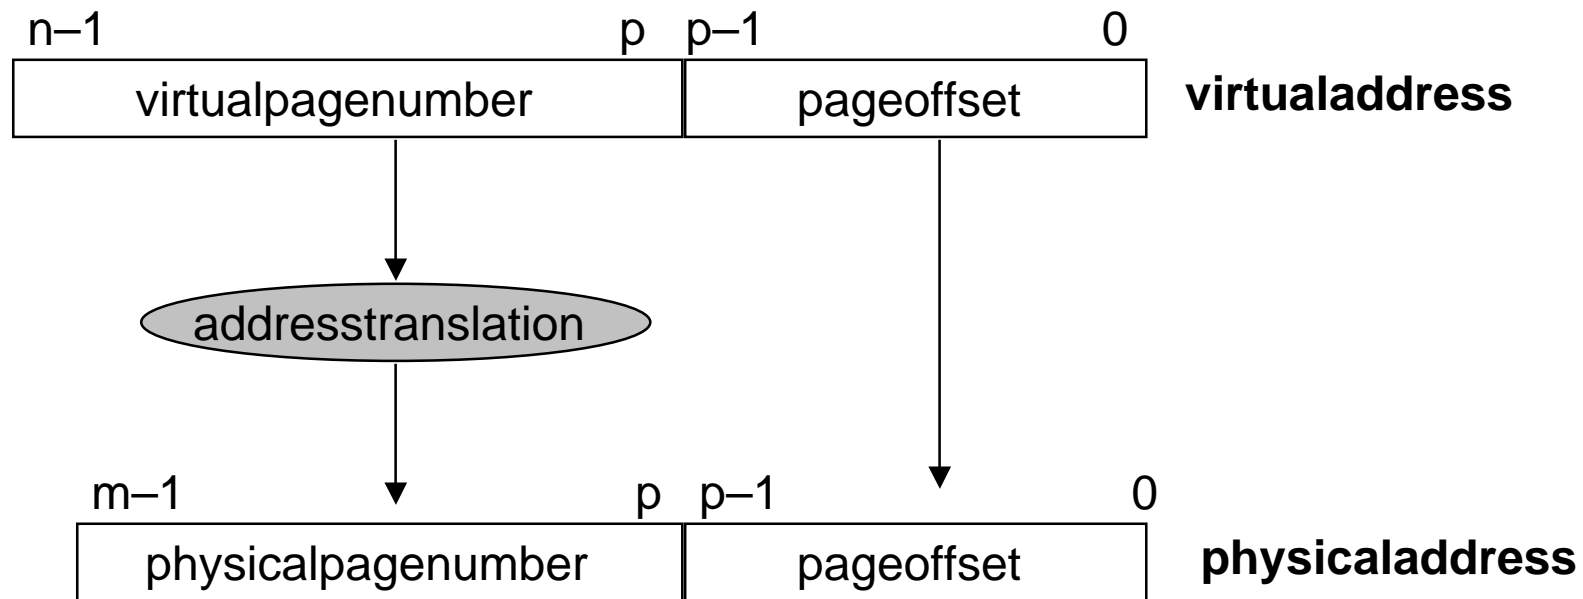
$MAP(a) = a'$  if data at virtual address  $a$  is present at physical address  $a'$  in  $P$   
 $= \emptyset$  if data at virtual address  $a$  is not present in  $P$



# VMAddressTranslation

## Parameters

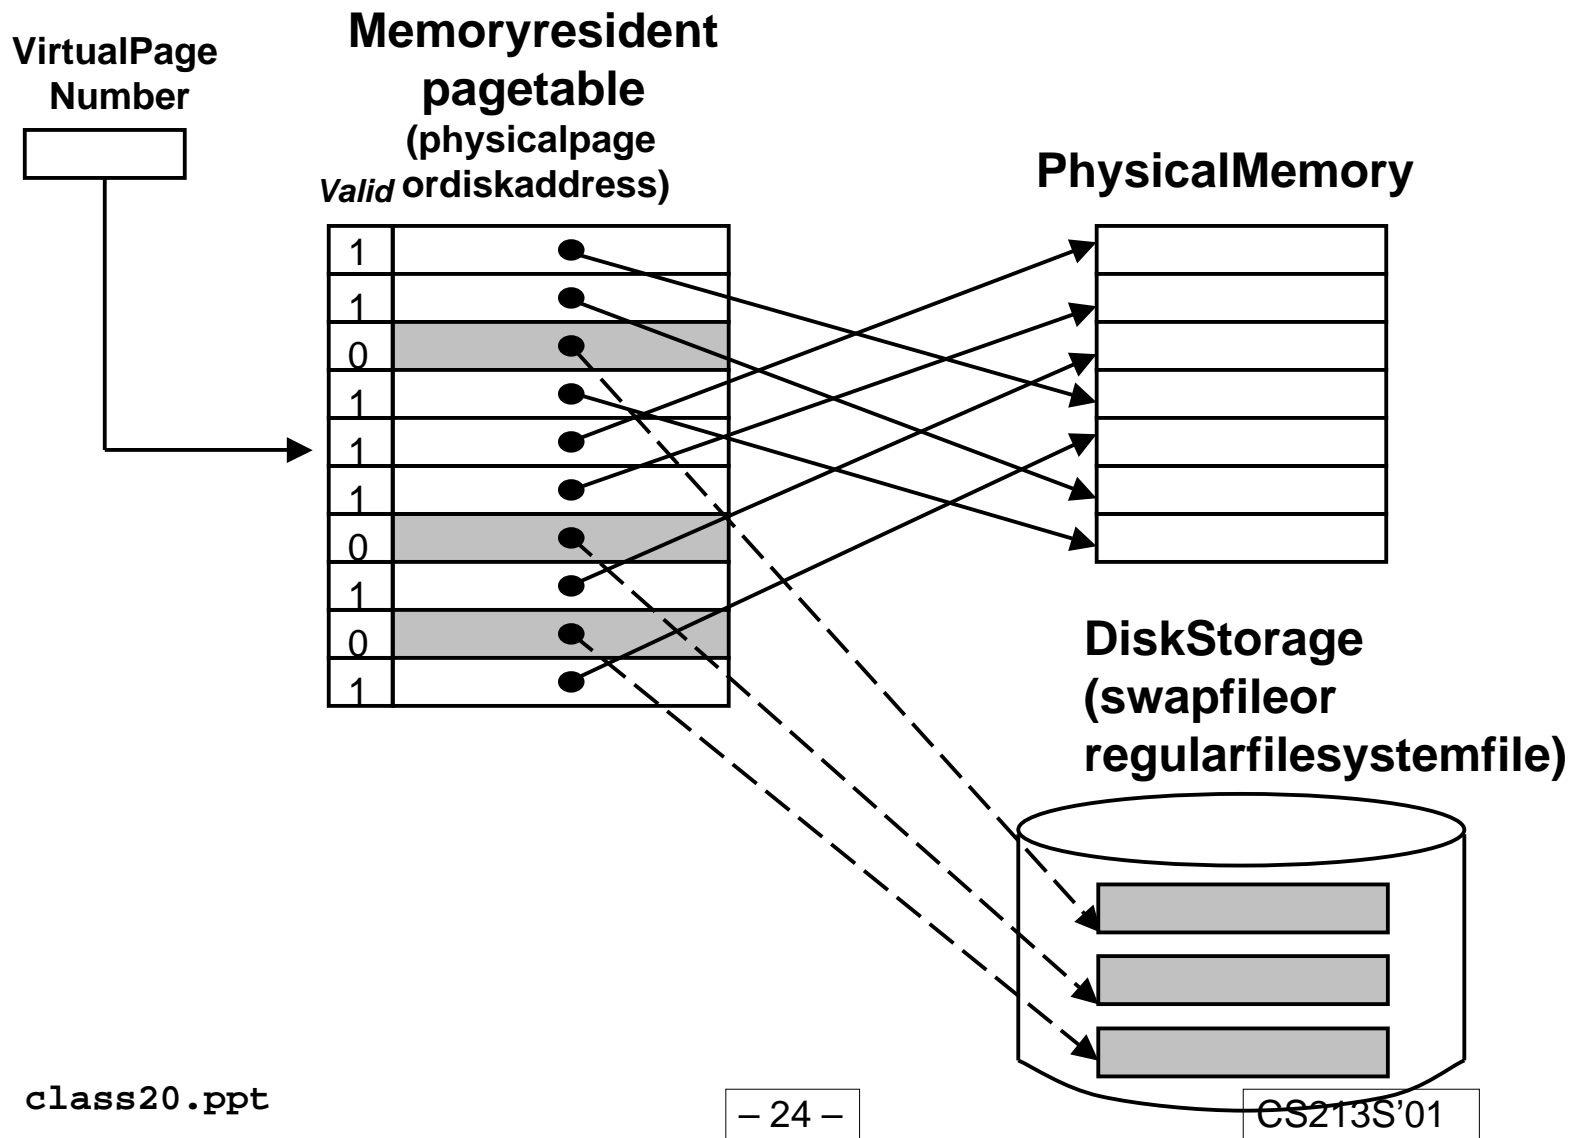
- $P=2^p$  =pagesize(bytes).
- $N=2^n$  =Virtualaddresslimit
- $M=2^m$  =Physicaladdresslimit



Notice that the pageoffset bits don't change as a result of tra

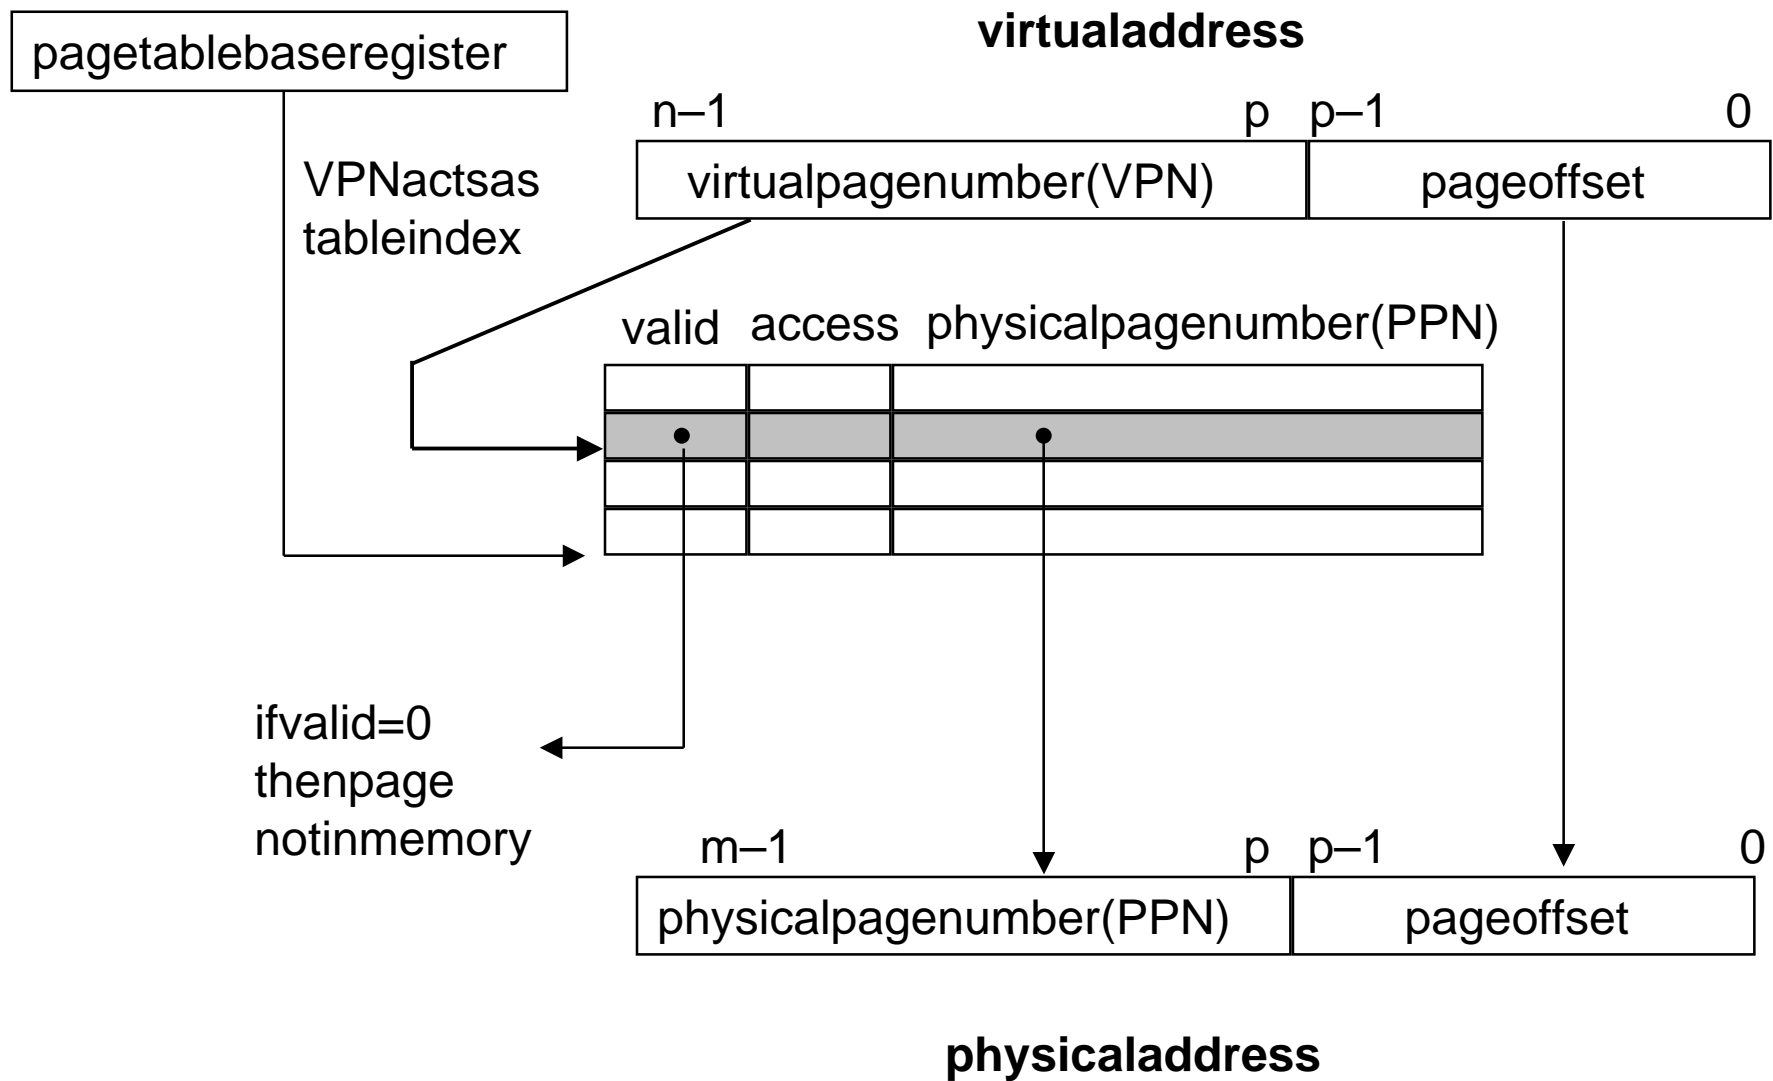
nslation

# PageTables





# Address Translation via Page Table



# PageTableOperation

## Translation

- Separate(set of) pagetable(s) per process
- VPN forms index into pagetable (points to a pagetable entry)


## Computing Physical Address

- PageTableEntry (PTE) provides information about page
  - if (valid bit = 1) then the page is in memory.
    - » Use physical page number (PPN) to construct address
  - if (valid bit = 0) then the page is on disk
    - » Page fault
    - » Must load page from disk into main memory before continuing

## Checking Protection

- Access rights field indicate allowable access
  - e.g., read -only, read -write, execute -only
  - typically support multiple protection modes (e.g., kernel vs. user)
- Protection violation fault if user doesn't have necessary permission

# Integrating VM and Cache



## Most Caches “Physically Addressed”

- Accessed by physical addresses
- Allows multiple processes to have blocks in cache at same time
- Allows multiple processes to share pages
- Cache doesn't need to be concerned with protection issues
  - Access rights checked as part of address translation

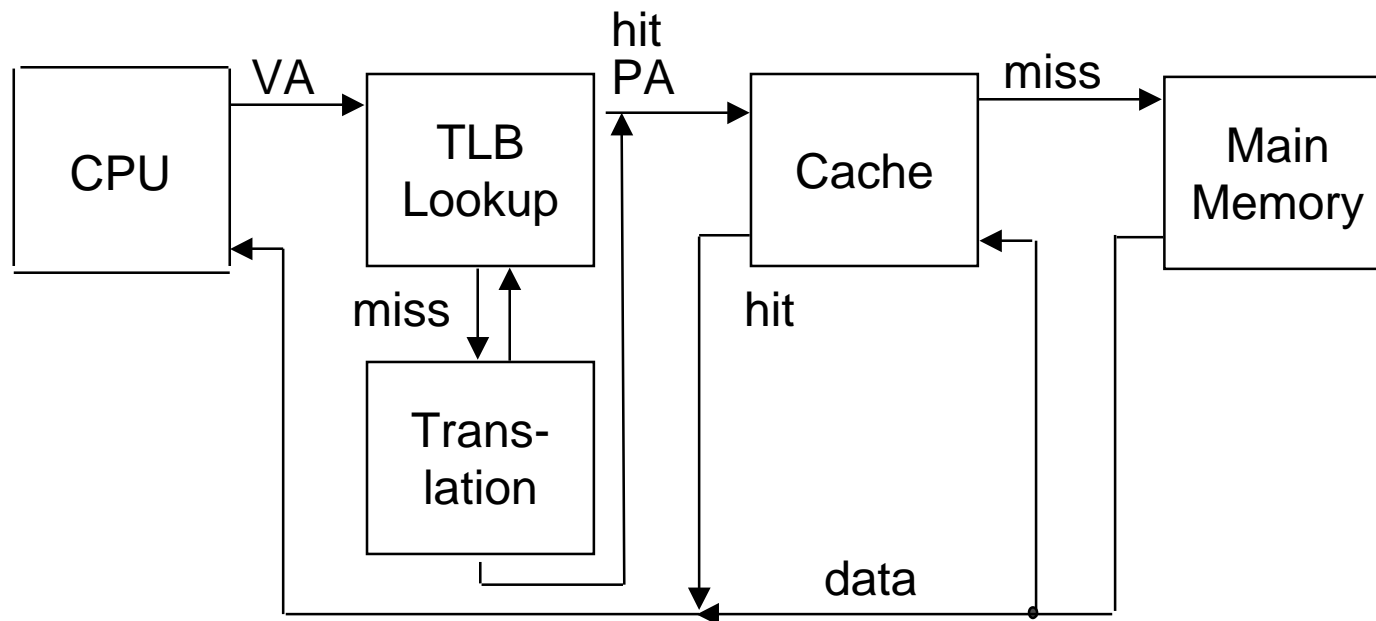
## Perform Address Translation Before Cache Lookup

- But this could involve a memory access itself (of the PTE)
- Of course, page table entries can also become cached

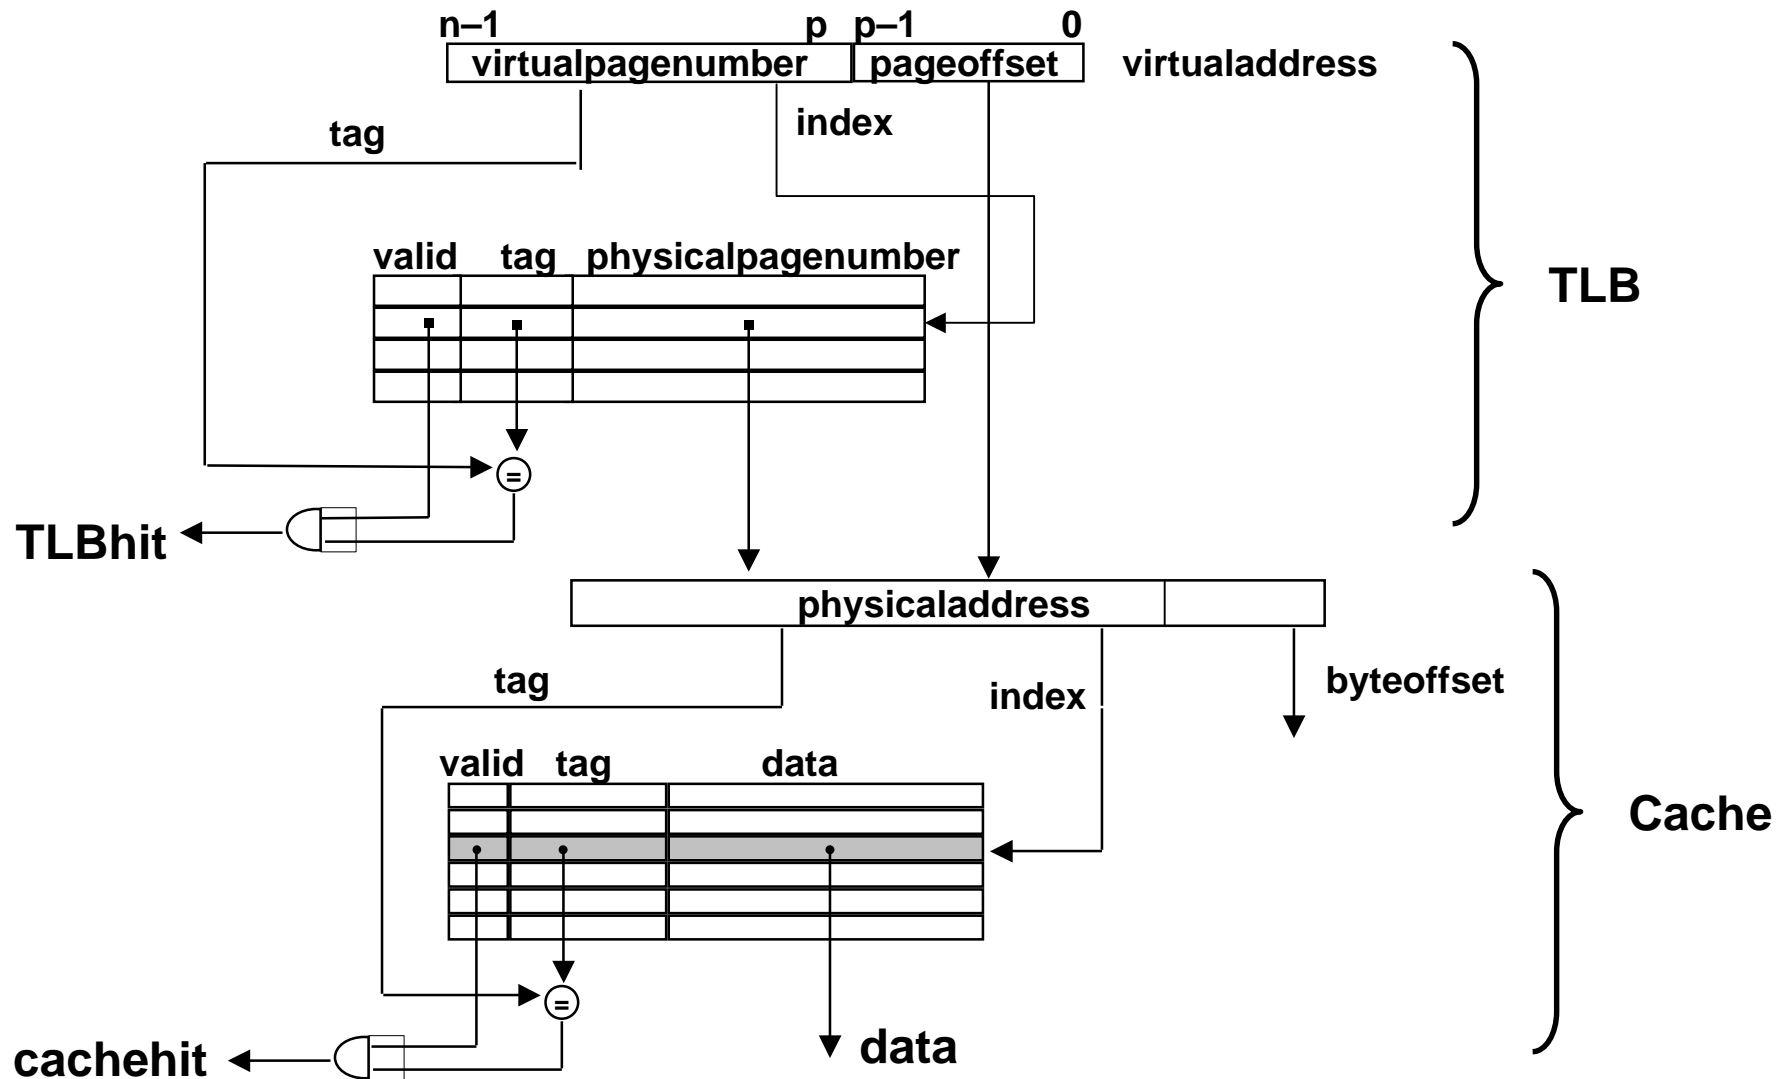
# Speedingup Translation with a TLB

## “Translation Lookaside Buffer”(TLB)

- Small hardware cache in MMU (Memory Management Unit)
- Maps virtual page number to physical page numbers
- Contains complete page table entries for small number of pages



# Address Translation with a TLB



# Example Sizes

## VirtualAddress(32bits)

- 19bitspagenumber
- 13bitspageoffset(8 Kbyte pages)

## TLB

- 128entries
- 4-waysetassociative
- HowmanybitsistheTLBtag?

### Virtualaddress

tag	idx	pageoffset
-----	-----	------------

## L1Cache

- 32Kbytes
- 4-waysetassociative
- 32-bytelinesize
- HowmanybitsintheCacheTag?

### physicaladdress

tag	idx	offst
-----	-----	-------

# Multi-Level Page Tables

## Given:


- 4KB( $2^{12}$ ) page size
- 32-bit address space
- 4-byte PTE

## Problem:

- Would need a 4MB page table!
  - $2^{20} * 4 \text{ bytes}$

## Common solution

- multi-level page tables
- e.g., 2-level table (P6)
  - Level 1 table: 1024 entries, each of which points to a Level 2 page table.
  - Level 2 table: 1024 entries, each of which points to a page



# Main Themes

## Programmer's View

- **Large “flat” address space**
  - Can allocate large blocks of contiguous addresses
- **Processor “owns” machine**
  - Has private address space
  - Unaffected by behavior of other processes

## System View

- **User virtual address space created by mapping to set of pages**
  - Need not be contiguous
  - Allocated dynamically
  - Enforce protection during address translation
- **OS manages many processes simultaneously**
  - Continually switching among processes
  - Especially when one must wait for resource
    - » E.g., disk I/O to handle page fault