15-213

Virtual Memory
April 3, 2001

Topics

- Motivations for VM
- Address translation
- Accelerating translation with 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

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

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

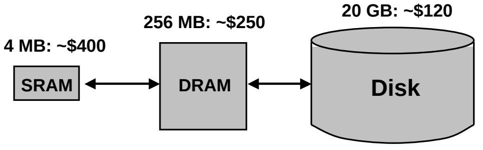

Levels in Memory Hierarchy

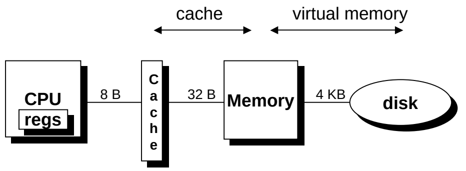

	Register	Cache	Memory	Disk Memory
size:	32 B	32 KB-4MB	128 MB	20 GB
speed:	1 ns	2 ns	50 ns	8 ms
$/Mbyte:		$100/MB	$1.00/MB	$0.006/MB
line size:	8 B	32 B	4 KB	

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
- **Bottom line:**
  - Design decisions made for DRAM caches driven by enormous cost of misses

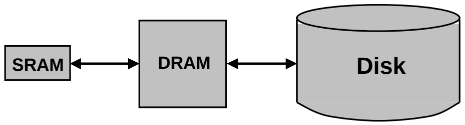

Impact of These Properties on Design

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
- **No tag for block not in cache**
- **Hardware retrieves information**
  - can quickly match against multiple tags

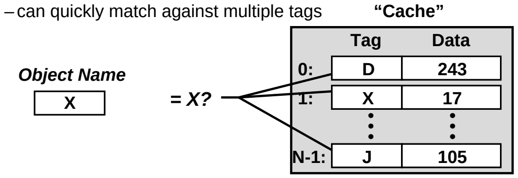

	Tag	Data
0:	D	243
1:	X	17
•	•	•
N-1:	J	105

Locating an Object in a "Cache" (cont.)

DRAM Cache

- **Each allocate page of virtual memory has entry in 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**

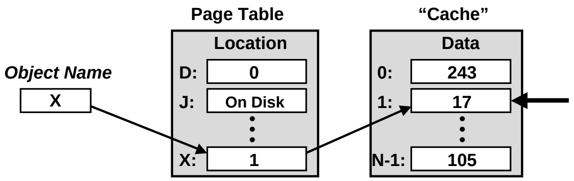

	Location
D:	0
J:	On Disk
•	•
X:	1

	Data
0:	243
1:	17
•	•
N-1:	105

A System with Physical Memory Only

Examples:

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

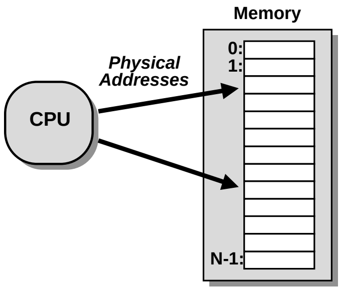

Addresses generated by the CPU point directly to bytes in physical memory

A System with Virtual Memory

Examples:

- workstations, servers, modern PCs, etc.

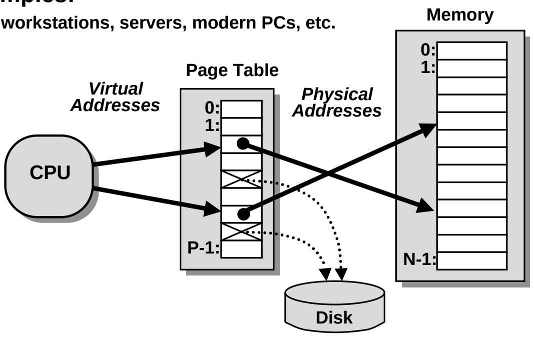

<u>Address Translation:</u> Hardware converts *virtual addresses* to *physical addresses* via an OS-managed lookup table (*page table*)

Page Faults (Similar to "Cache Misses")

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

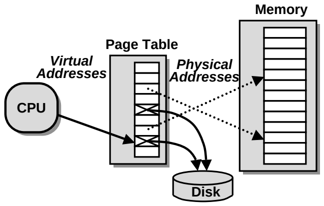

After fault

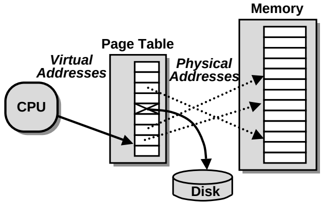

Servicing a Page Fault

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

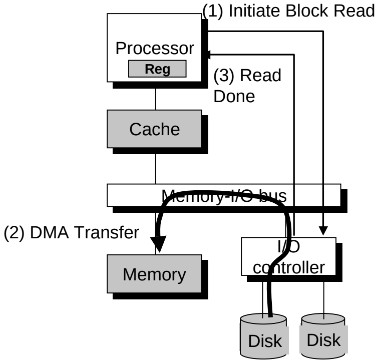

Motivation #2: Memory Management

Multiple processes can reside in physical memory.

How do we resolve address conflicts?

- **what if two processes access something at the same address?**

Linux/x86 process memory image

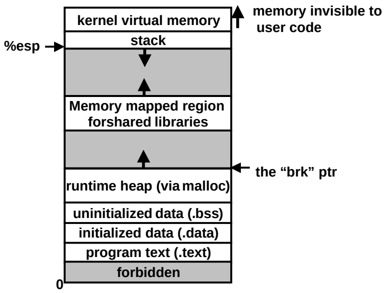
kernel virtual memory
memory invisible to user code
stack
%esp
Memory mapped region forshared libraries
the “brk” ptr
runtime heap (via malloc)
uninitialized data (.bss)
initialized data (.data)
program text (.text)
forbidden
0

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 as assigned to physical memory

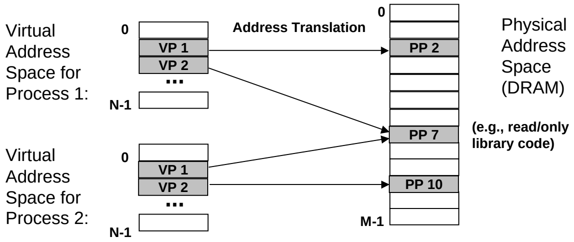
Virtual Address Space for Process 1:
0
VP 1
VP 2
...
N-1
Address Translation
0
PP 2
PP 7
PP 10
M-1
Physical Address Space (DRAM)
(e.g., read/only library code)
Virtual Address Space for Process 2:
0
VP 1
VP 2
...
N-1

Contrast: Macintosh Memory Model

MAC OS 1–9

- **Does not use traditional virtual memory**

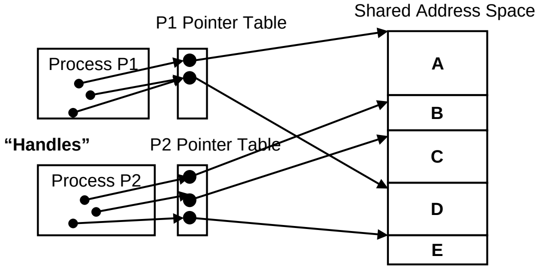

All program objects accessed through “handles”

- **Indirect reference through pointer table**
- **Objects stored in shared global address space**

Macintosh Memory Management

Allocation / Deallocation

- **Similar to free-list management of malloc/free**

Compaction

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

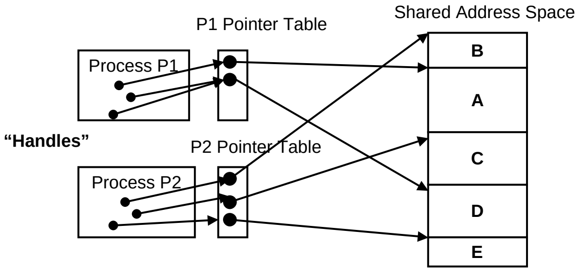

Mac vs. VM-Based Memory Mgmt

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: "wild write" by one process can corrupt another's data

MAC OS X

"Modern" Operating System

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

Based on MACH OS

- Developed at CMU in late 1980's

Motivation #3: Protection

Protection goals:

- Cannot read/write memory from another process
- Cannot write into shared libraries

Processes can only see virtual addresses

- Cannot get to physical addresses directly
- Can only go through the page table
- If a physical page is not in a process' page table, it is "invisible"

Page table entry contains access rights information

- hardware enforces this protection (trap into OS if violation occurs)
- The page table itself is in protected memory

When allocating a new physical page, it is cleared

- Important that the process cannot see the previous contents

Protection: Example

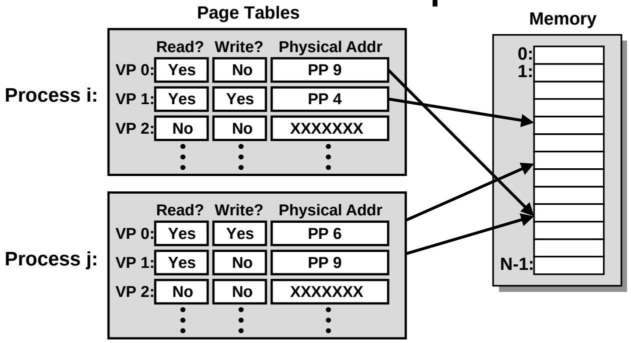

Process i:

	Read?	Write?	Physical Addr
VP 0:	Yes	No	PP 9
VP 1:	Yes	Yes	PP 4
VP 2:	No	No	XXXXXXX
⋮	⋮	⋮	⋮

Process j:

	Read?	Write?	Physical Addr
VP 0:	Yes	Yes	PP 6
VP 1:	Yes	No	PP 9
VP 2:	No	No	XXXXXXX
⋮	⋮	⋮	⋮

- Process i 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.

VM Address Translation

V = {0, 1, . . . , N–1} virtual address space **N > M**
P = {0, 1, . . . , M–1} physical address space

MAP: V → P U {∅} address mapping function

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

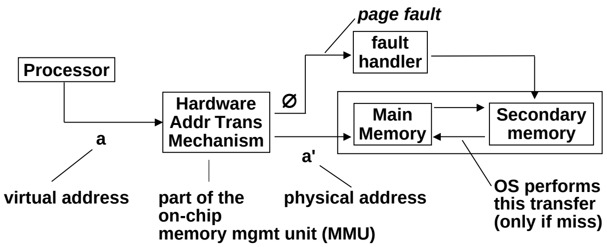
page fault
fault handler
Processor
Hardware Addr Trans Mechanism
∅
Main Memory
Secondary memory
a
a'
virtual address
part of the on-chip memory mgmt unit (MMU)
physical address
OS performs this transfer (only if miss)

VM Address Translation

Parameters

- **$P = 2^p$ = page size (bytes).**
- **$N = 2^n$ = Virtual address limit**
- **$M = 2^m$ = Physical address limit**

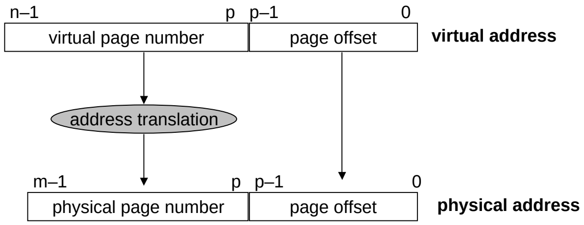

Notice that the page offset bits don't change as a result of translation

Page Tables

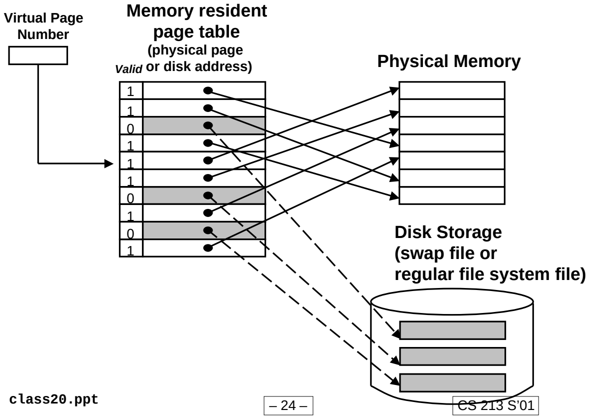
Virtual Page Number
Memory resident page table
(physical page or disk address)
Valid
1
1
0
1
1
1
0
1
0
1
Physical Memory
Disk Storage
(swap file or regular file system file)

Address Translation via Page Table

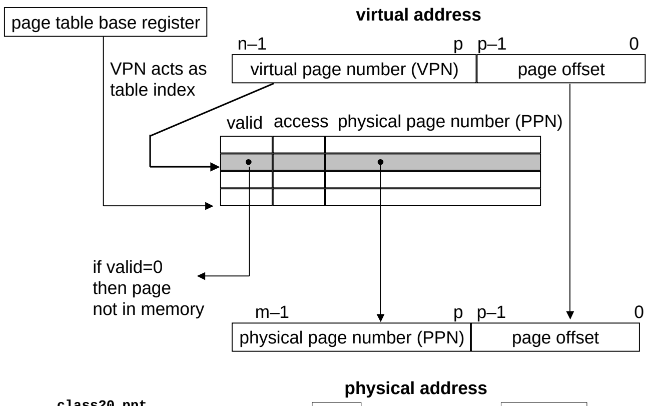

Page Table Operation

Translation

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

Computing Physical Address

- **Page Table Entry (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

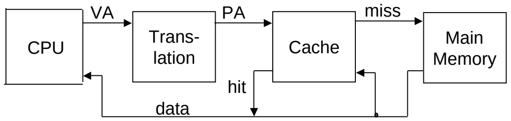

Most Caches "Physically Addressed"

- **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**

Speeding up Translation with a TLB

"Translation Lookaside Buffer" (TLB)

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

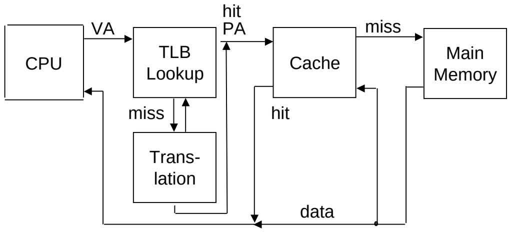

Address Translation with a TLB

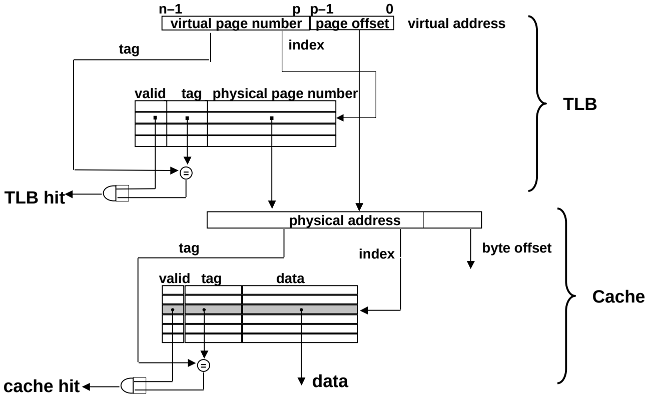

Example Sizes

Virtual Address (32 bits)

- **19 bits page number**
- **13 bits page offset (8 Kbyte pages)**

TLB

- **128 entries**
- **4-way set associative**
- **<u>How many bits is the TLB tag?</u>**

Virtual address

tag	idx	page offset

L1 Cache

- **32 Kbytes**
- **4-way set associative**
- **32-byte line size**
- **<u>How many bits in the Cache Tag?</u>**

physical address

tag	idx	offst

Multi-Level Page Tables

Given:

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

Problem:

- **Would need a 4 MB page table!**
  - 2^{20} *4 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

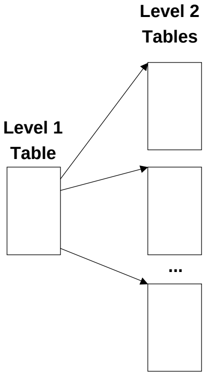

Main Themes

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