Systemprogrammering 2007 **Föreläsning 4 Virtual Memory**

Topics

The memory hierarchy

- Motivations for VM
- Address translation
- Accelerating translation with TLBs

The CPU-Memory Gap

The increasing gap between DRAM, disk, and CPU speeds.



Random-Access Memory (RAM)

Kev features

- RAM is packaged as a chip.
- Basic storage unit is a cell (one bit per cell).
- Multiple RAM chips form a memory.

Static RAM (SRAM)

- Each cell stores bit with a six-transistor circuit.
- Retains value indefinitely, as long as it is kept powered.
- Relatively insensitive to disturbances such as electrical noise.
- Faster and more expensive than DRAM.

Dynamic RAM (DRAM)

- Each cell stores bit with a capacitor and transistor.
- Value must be refreshed every 10-100 ms.
- Sensitive to disturbances.
- Slower and cheaper than SRAM.

F4 – 2 –

Systemprogrammering 2007

Locality

Principle of Locality:

- Programs tend to reuse data and instructions near those they have used recently, or that were recently referenced themselves.
- Temporal locality: Recently referenced items are likely to be referenced in the near future.
- Spatial locality: Items with nearby addresses tend to be referenced close together in time.

Locality Example:

Data

- - Reference array elements in succession (stride-1 reference pattern): Spatial locality
 - -Reference sum each iteration: **Temporal locality**
- Instructions
 - Reference instructions in sequence: Spatial locality
 - Cycle through loop repeatedly: **Temporal locality**

sum = 0;

return sum;

for (i = 0; i < n; i++)

sum += a[i];

Memory Hierarchies

- Some fundamental and enduring properties of hardware and software:
 - Fast storage technologies cost more per byte and have less capacity.
 - The gap between CPU and main memory speed is widening.
 - Well-written programs tend to exhibit good locality.

These fundamental properties complement each other beautifully.

They suggest an approach for organizing memory and storage systems known as a memory hierarchy.

F4 – 5 –

Systemprogrammering 2007

Systemprogrammering 2007

Caches

Cache: A smaller, faster storage device that acts as a staging area for a subset of the data in a larger, slower device.

Fundamental idea of a memory hierarchy:

For each k, the faster, smaller device at level k serves as a cache for the larger, slower device at level k+1.

Why do memory hierarchies work?

- Programs tend to access the data at level k more often than they access the data at level k+1.
- Thus, the storage at level k+1 can be slower, and thus larger and cheaper per bit.
- Net effect: A large pool of memory that costs as much as the cheap storage near the bottom, but that serves data to programs at the rate of the fast storage near the top.

An Example Memory Hierarchy Smaller. 10 faster. registers CPU registers hold words retrieved from and I 1 cache on-chip L1 costlier L1 cache (SRAM) L1 cache holds cache lines retrieved (per byte) from the L2 cache memory off-chip L2 storage L2 cache (SRAM) L2 cache holds cache lines retrieved devices from main memory. main memory L3: Larger, (DRAM) Main memory holds disk slower. blocks retrieved from local and disks

local secondary storage

(local disks)

remote secondary storage (distributed file systems, Web servers)



Types of cache misses:

L4:

- Cold (compulsary) miss
 - Cold misses occur because the cache is empty.
- Conflict miss

cheaper

(per byte)

storage

devices

F4 – 6 –

- Most caches limit blocks at level k+1 to a small subset (sometimes a singleton) of the block positions at level k.
- E.g. Block i at level k+1 must be placed in block (i mod 4) at level k+1.
- Conflict misses occur when the level k cache is large enough, but multiple data objects all map to the same level k block.
- E.g. Referencing blocks 0, 8, 0, 8, 0, 8, ... would miss every time.
- Capacity miss
 - Occurs when the set of active cache blocks (working set) is larger than the cache.

Systemprogrammering 2007

Local disks hold files

retrieved from disks on remote network servers

Systemprogrammering 2007

F4 – 7 –

Cache Type	What Cached	Where Cached Latency (cycles)		Managed By
Registers	4-byte word	CPU registers	0	Compiler
TLB	Address translations	On-Chip TLB	0	Hardware
L1 cache	32-byte block	On-Chip L1	1	Hardware
L2 cache	32-byte block	Off-Chip L2	10	Hardware
Virtual Memory	4-KB page	Main memory	100	Hardware + OS
Buffer cache	Parts of files	Main memory	100	OS
Network buffer cache	Parts of files	Local disk	10,000,000	AFS/NFS client
Browser cache	Web pages	Local disk	10,000,000	Web browser
Web cache	Web pages	Remote server disks	1,000,000,000	Web proxy server

Examples of Caching in the Hierarchy

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 (16 quintillion) bytes

Disk storage is ~300X cheaper than DRAM storage

- 80 GB of DRAM: ~ \$33,000
- 80 GB of disk: ~ \$110

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



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.

```
F4 – 10 –
```

Systemprogrammering 2007

A System with Physical Memory Only

Examples:

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



A System with Virtual Memory



F4 – 13 –

Systemprogrammering 2007

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



Page Faults (like "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.



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?



F4 – 15 –



VM Address Translation

Virtual Address Space

■ V = {0, 1, ..., N–1}

Physical Address Space

- P = {0, 1, ..., M–1}
- M < N (usually)</p>

Address Translation

- MAP: $V \rightarrow P \cup \{\emptyset\}$
- For virtual address a:
 - MAP(a) = a' if data at virtual address a at physical address a' in P
 - MAP(a) = Ø if data at virtual address a not in physical memory
 - » Either invalid or stored on disk

F4 – 19 –

Systemprogrammering 2007

Motivation #3: Protection

Page table entry contains access rights information

hardware enforces this protection (trap into OS if violation occurs)



VM Address Translation

Parameters

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



Page offset bits don't change as a result of translation



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

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

Systemprogrammering 2007

Address Translation via Page Table



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



"Translation Lookaside Buffer" (TLB)

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



Simple Memory System Example

0

Systemprogrammering 2007



Address Translation with a TLB



Simple Memory System Page Table

Only show first 16 entries

VPN	PPN	Valid	VPN	PPN	Valid
00	28	1	08	13	1
01	-	0	09	17	1
02	33	1	0A	09	1
03	02	1	0B	-	0
04	-	0	0C	-	0
05	16	1	0D	2D	1
06	-	0	0E	11	1
07	-	0	0F	0D	1



Simple Memory System Cache

Cache

- 16 lines
- 4-byte line size



Address Translation Example #2





Main Themes

Programmer's View

- Large "flat" address space
 - Can allocate large blocks of contiguous addresses
- Process "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

```
Systemprogrammering 2007
```

