

Caches 3

CS 3410, Spring 2014

Computer Science

Cornell University

See P&H Chapter: 5.1-5.4, 5.8, 5.15

Compromise

Set-associative cache

Like a direct-mapped cache


- Index into a location
- Fast

Like a fully-associative cache

- Can store multiple entries
 - decreases conflicts
- Search in each element

n-way set assoc means n possible locations

3-Way Set Associative Cache (Reading)



Basic Cache Organization


Q: How to decide block size?

A: Try it and see

But: depends on cache size, workload, associativity, ...

Experimental approach!

Experimental Results



Tradeoffs


For a given total cache size,
larger block sizes mean....

- fewer lines
- so fewer tags, less overhead
- and fewer cold misses (within-block “prefetching”)

But also...

- fewer blocks available (for scattered accesses!)
- so more conflicts
- and larger miss penalty (time to fetch block)

Associativity



Other Designs

Multilevel caches

Next Goal

What about writes?

What happens when the CPU writes to a register and calls a store instruction?!


What about Stores?

Where should you write the result of a store?

- If that memory location is in the cache?
 - Send it to the cache
 - Should we also send it to memory right away?
(write-through policy)
 - Wait until we evict the block (write-back policy)
- If it is not in the cache?
 - Allocate the line (put it in the cache)?
(write allocate policy)
 - Write it directly to memory without allocation?
(no write allocate policy)

Cache Write Policies

Q: How to write data?



If data is already in the cache...

No-Write

writes invalidate the cache and go directly to memory

Write-Through

writes go to main memory and cache


Write-Back

CPU writes only to cache

cache writes to main memory later (when block is evicted)

Write Allocation Policies

Q: How to write data?



If data is not in the cache...

Write-Allocate

allocate a cache line for new data (and maybe write-through)

No-Write-Allocate

ignore cache, just go to main memory


Next Goal

How does a write-through cache work?


Assume write-allocate

Handling Stores (Write-Through)


Using **byte addresses** in this example! Addr Bus = 5 bits




Write-Through (REF 1)




Write-Through (REF 1)




Write-Through (REF 2)




Write-Through (REF 2)




Write-Through (REF 3)




Write-Through (REF 3)




Write-Through (REF 4)



Write-Through (REF 4)




Write-Through (REF 5)




Write-Through (REF 5)

Processor	Cache	Memory
LB \$1 ← M[1] M LB \$2 ← M[7] M SB \$2 → M[0] H SB \$1 → M[5] M → LB \$2 ← M[10] M SB \$1 → M[5] SB \$1 → M[10] \$0 \$129 \$233 \$3 Vtagdata 1010133 1001071 29 Misses: 4 Hits: 1 0173 129 2120 3123 471 529 6162 7173 818 921 1033 1128 1219 13200 14210 15225		


Write-Through (REF 6)




Write-Through (REF 6)



Write-Through (REF 7)



Write-Through (REF 7)



How Many Memory References?

Write-through performance

Each miss (read or write) reads a block from mem

- 4 misses → 8 mem reads


Each store writes an item to mem

- 4 mem writes


Evictions don't need to write to mem

- no need for dirty bit

Write-Through (REF 8,9)



Write-Through (REF 8,9)



Summary: Write Through

Write-through policy with write allocate

Cache miss: read entire block from memory

Write: write only updated item to memory

Eviction: no need to write to memory

Write-Through vs. Write-Back

Can we also design the cache NOT to write all stores immediately to memory?

- Keep the most current copy in cache, and update memory when that data is evicted (write-back policy)
- Do we need to write-back all evicted lines?
 - No, only blocks that have been stored into (written)

Write-Back Meta-Data

V	D	Tag	Byte 1	Byte 2	... Byte N

V = 1 means the line has valid data

D = 1 means the bytes are newer than main memory

When allocating line:

- Set V = 1, D = 0, fill in Tag and Data

When writing line:

- Set D = 1

When evicting line:

- If D = 0: just set V = 0
- If D = 1: write-back Data, then set D = 0, V = 0

Write-back Example

Example: How does a write-back cache work?


Assume write-allocate

Handling Stores (Write-Back)


Using **byte addresses** in this example! Addr Bus = 5 bits

Processor	Cache	Memory																																																												
Assume write-allocate policy LB \$1 ← M[1] LB \$2 ← M[7] SB \$2 → M[0] SB \$1 → M[5] LB \$2 ← M[10] SB \$1 → M[5] SB \$1 → M[10] <table><tr><td>\$0</td><td></td></tr><tr><td>\$1</td><td></td></tr><tr><td>\$2</td><td></td></tr><tr><td>\$3</td><td></td></tr></table>	\$0		\$1		\$2		\$3		Fully Associative Cache 2 cache lines <u>2</u> word block <u>3</u> bit tag field 1 bit block offset field <table><tr><td>V</td><td>d</td><td>tag</td><td>data</td></tr><tr><td>0</td><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td><td></td></tr><tr><td>0</td><td></td><td></td><td></td></tr><tr><td></td><td></td><td></td><td></td></tr></table> Misses: 0 Hits: 0	V	d	tag	data	0								0								<table><tr><td>0</td><td>78</td></tr><tr><td>1</td><td>29</td></tr><tr><td>2</td><td>120</td></tr><tr><td>3</td><td>123</td></tr><tr><td>4</td><td>71</td></tr><tr><td>5</td><td>150</td></tr><tr><td>6</td><td>162</td></tr><tr><td>7</td><td>173</td></tr><tr><td>8</td><td>18</td></tr><tr><td>9</td><td>21</td></tr><tr><td>10</td><td>33</td></tr><tr><td>11</td><td>28</td></tr><tr><td>12</td><td>19</td></tr><tr><td>13</td><td>200</td></tr><tr><td>14</td><td>210</td></tr><tr><td>15</td><td>225</td></tr></table>	0	78	1	29	2	120	3	123	4	71	5	150	6	162	7	173	8	18	9	21	10	33	11	28	12	19	13	200	14	210	15	225
\$0																																																														
\$1																																																														
\$2																																																														
\$3																																																														
V	d	tag	data																																																											
0																																																														
0																																																														
0	78																																																													
1	29																																																													
2	120																																																													
3	123																																																													
4	71																																																													
5	150																																																													
6	162																																																													
7	173																																																													
8	18																																																													
9	21																																																													
10	33																																																													
11	28																																																													
12	19																																																													
13	200																																																													
14	210																																																													
15	225																																																													


Write-Back (REF 1)




Write-Back (REF 1)




Write-Back (REF 1)




Write-Back (REF 2)




Write-Back (REF 2)




Write-Back (REF 3)




Write-Back (REF 3)




Write-Back (REF 4)




Write-Back (REF 4)




Write-Back (REF 5)




Write-Back (REF 5)




Write-Back (REF 5)




Write-Back (REF 6)




Write-Back (REF 6)



Write-Back (REF 7)



Write-Back (REF 7)



How Many Memory References?

Write-back performance

Each miss (read or write) reads a block from mem

- 4 misses $\rightarrow$ 8 mem reads

Some evictions write a block to mem

- 1 dirty eviction $\rightarrow$ 2 mem writes
- (+ 2 dirty evictions later $\rightarrow$ +4 mem writes)

Write-Back (REF 8,9)

Processor	Cache	Memory
LB \$1 ← M[1] M LB \$2 ← M[7] M SB \$2 → M[0] H SB \$1 → M[5] M LB \$2 ← M[10] M SB \$1 → M[5] H SB \$1 → M[10] H SB \$1 → M[5] SB \$1 → M[10]	V d tag data 11010129 28 11001071 29 Misses: 4 Hits: 3	078 129 2120 3123 471 5150 6162 7173 818 921 1033 1128 1219 13200 14210 15225
\$0 \$1 \$2 \$3 29 33		

Write-Back (REF 8,9)

Processor	Cache	Memory
LB \$1 ← M[1] M LB \$2 ← M[7] M SB \$2 → M[0] H SB \$1 → M[5] M LB \$2 ← M[10] M SB \$1 → M[5] H SB \$1 → M[10] H SB \$1 → M[5] H SB \$1 → M[10] H	V d tag data 11010129 28 11001071 29 Misses: 4 Hits: 5	078 129 2120 3123 471 5150 6162 7173 818 921 1033 1128 1219 13200 14210 15225
\$0 \$1 \$2 \$3 29 33		

How Many Memory References?

Write-back performance

Each miss (read or write) reads a block from mem

- 4 misses → 8 mem reads

Some evictions write a block to mem

- 1 dirty eviction → 2 mem writes
- (+ 2 dirty evictions later → +4 mem writes)

By comparison write-through was

- Reads: eight words
- Writes: 4/6/8/10/12/... etc words

So is write back just better?

What are other performance tradeoffs between write-through and write-back?

How can we further reduce penalty for cost of writes to memory?

Performance Tradeoffs

Q: Hit time: write-through vs. write-back?

A: Write-through slower on writes

Q: Miss penalty: write-through vs. write-back?

A: Write-back slower on evictions

Write Buffering

Q: Writes to main memory are **slow!**

A: Use a write-back buffer

- A small queue holding dirty lines
- Add to end upon eviction
- Remove from front upon completion

Q: When does it help?

A: short bursts of writes (but not sustained writes)

A: fast eviction reduces miss penalty

Write-through vs. Write-back

Write-through is slower

- But simpler (memory always consistent)

Write-back is almost always faster

- write-back buffer hides large eviction cost
- But what about multiple cores with separate caches but sharing memory?


Write-back requires a cache coherency protocol

- Inconsistent views of memory
- Need to “snoop” in each other’s caches
- Extremely complex protocols, very hard to get right

Cache-coherency

Q: Multiple readers and writers?

A: Potentially inconsistent views of memory



Cache coherency protocol

- snoop on other CPU's cache activity
- Invalidate cache line when other CPU writes
- Flush write-back caches before other CPU reads
- Or the reverse: Before writing/reading...
- Extremely complex protocols, very hard to get right

Summary: Write Through

Write-through policy with write allocate

- Cache miss: read entire block from memory
- Write: write only updated item to memory
- Eviction: no need to write to memory
- Slower, but cleaner

Write-back policy with write allocate

- Cache miss: read entire block from memory
 - But may need to write dirty cacheline first
- Write: nothing to memory
- Eviction: have to write to memory, entire cacheline because don't know what is dirty (only 1 dirty bit)
- Faster, but complicated with multicore

Next Goal

Performance: What is the average memory access time (AMAT) for a cache?

$$\text{AMAT} = \% \text{hit} \times \text{hit time} + \% \text{miss} \times \text{miss time}$$

Cache Performance Example

Average Memory Access Time (AMAT)

Cache Performance (very simplified):

L1 (SRAM): 512 x 64 byte cache lines, direct mapped

Data cost: 3 cycle per word access

Lookup cost: 2 cycle

16 words (i.e. $64 / 4 = 16$)

Mem (DRAM): 4GB

Data cost: 50 cycle for first word, plus 3 cycles per subsequent word

AMAT = %hit x hit time + % miss x miss time

Hit time = 5 cycles

Miss time = hit time + 50 (first word) + 15 x 3 (words)
= 100 cycles

If %hit = 90%, then

AMAT = .9 x 5 + .1 x 100 = 14.5 cycles

Cache Performance Example

Average Memory Access Time (AMAT)

Cache Performance (very simplified):

L1 (SRAM): 512 x 64 byte cache lines, direct mapped

Data cost: 3 cycle per word access

Lookup cost: 2 cycle

16 words (i.e. $64 / 4 = 16$)

Mem (DRAM): 4GB

Data cost: 50 cycle for first word, plus 3 cycles per subsequent word

Multi Level Caching

Cache Performance (very simplified):

L1 (SRAM): 512 x 64 byte cache lines, direct mapped

Hit time: 5 cycles

L2 cache: bigger

Hit time = 20 cycles

Mem (DRAM): 4GB

Hit rate: 90% in L1, 90% in L2

AMAT = %hit x hit time + % miss x miss time

AMAT = .9 x 5 + .1 (.9 x 20 + .1 x 120) = 4.5 + .1 (18 + 12) = 7.5

Often: L1 fast and direct mapped, L2 bigger and higher associativity

Performance Summary

Average memory access time (AMAT)

depends on cache architecture and size

access time for hit,

miss penalty, miss rate

Cache design a very complex problem:

- Cache size, block size (aka line size)
- Number of ways of set-associativity (1, N, ∞)
- Eviction policy
- Number of levels of caching, parameters for each
- Separate I-cache from D-cache, or Unified cache
- Prefetching policies / instructions
- Write policy

Cache Conscious Programming

```
// H = 12, W = 10
int A[H][W];

for(x=0; x < W; x++)
    for(y=0; y < H; y++)
        sum += A[y][x];
```

1									
2									
3									
4									
5									
6									
7									
8									
9									
10									
11									
12									

Every access is a cache miss!
(unless *entire* matrix can fit in cache)

Cache Conscious Programming

```
// H = 12, W = 10
```

```
int A[H][W];
```

```
for(y=0; y < H; y++)
    for(x=0; x < W; x++)
        sum += A[y][x];
```

Block size = 4 $\rightarrow$

Block size = 8 $\rightarrow$ 87.5% hit rate

Block size = 16 $\rightarrow$ 93.75% hit rate

And you can easily prefetch to warm the cache

[illegible]

By the end of the cache lectures...

MacBook Pro

Retina, Mid 2012

Processor 2.7 GHz Intel Core i7

Memory 16 GB 1600 MHz DDR3

Graphics NVIDIA GeForce GT 650M 1024 MB

Serial Number C02J70TTDKQ5

Software OS X 10.9.2 (13C64)

Model Name:	MacBook Pro
Model Identifier:	MacBookPro10,1
Processor Name:	Intel Core i7
Processor Speed:	2.7 GHz
Number of Processors:	1
Total Number of Cores:	4
L2 Cache (per Core):	256 KB
L3 Cache:	8 MB
Memory:	16 GB
Boot ROM Version:	MBP101.00EE.B02
SMC Version (system):	2.3f36
Serial Number (system):	C02J70TTDKQ5
Hardware UUID:	F588E08C-60BF-5B35-A087-07714C2B2D11

- 32 KB data + 32 KB instruction L1 cache (3 clocks) and 256 KB L2 cache (8 clocks) per core.
- Shared L3 cache includes the processor graphics (LGA 1155).
- 64-byte cache line size.




Figure 1. Schematic diagram of a Sandy Bridge processor.

Summary

Memory performance matters!

- often more than CPU performance
- ... because it is the bottleneck, and not improving much
- ... because most programs move a LOT of data

Design space is huge

- Gambling against program behavior
- Cuts across all layers:
users → programs → os → hardware

Multi-core / Multi-Processor is complicated

- Inconsistent views of memory
- Extremely complex protocols, very hard to get right

Spring is here!

