



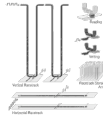
Lecture 31 – Caches II
2008-04-14

Lecturer SOE
Dan Garcia

Hi to Yi Luo from Seattle, WA !

IBM'S "RACETRACK MEMORY"

In this week's *Science*, IBM researchers describe a new class of data storage, called racetrack memory, combining the data storage density of disk with the ruggedness and speed of flash memory (no moving parts). They store bits as magnetic fields on nanowires. They don't have a prototype, and say commercialization is about 7 years away. But, this could be something big!



www.technologyreview.com/Infotech/20553/

Direct-Mapped Cache Terminology

- All fields are read as unsigned integers.
- Index
 - specifies the cache index (or "row"/block)
- Tag
 - distinguishes betw the addresses that map to the same location
- Offset
 - specifies which byte within the block we want



tag
to check
if have
correct block

index
to
select
block

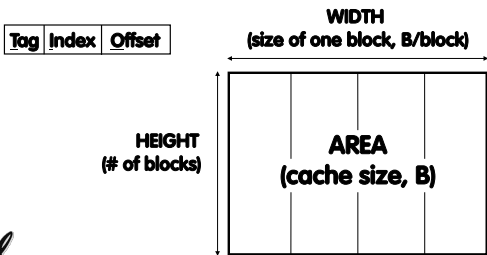
byte
offset
within
block



TIO Dan's great cache mnemonic

AREA (cache size, B)
= HEIGHT (# of blocks) * WIDTH (size of one block, B/block)

$$2^{(H+W)} = 2^H * 2^W$$



Caching Terminology

- When reading memory, 3 things can happen:
 - cache hit: cache block is valid and contains proper address, so read desired word
 - cache miss: nothing in cache in appropriate block, so fetch from memory
 - cache miss, block replacement: wrong data is in cache at appropriate block, so discard it and fetch desired data from memory (cache always copy)



Accessing data in a direct mapped cache

- Ex.: 16KB of data, direct-mapped, 4 word blocks

Can you work out height, width, area?

- Read 4 addresses

- 0x00000014
- 0x0000001C
- 0x00000034
- 0x000008014

- Memory vals here:

Memory Address (hex)	Value of Word
00000010	a
00000014	b
00000018	c
0000001C	d
...	...
00000030	e
00000034	f
00000038	g
0000003C	h
...	...
00008010	i
00008014	j
00008018	k
0000801C	l
...	...



Accessing data in a direct mapped cache

- 4 Addresses:

- 0x00000014, 0x0000001C, 0x00000034, 0x000008014

- 4 Addresses divided (for convenience) into Tag, Index, Byte Offset fields

00000000000000000000	0000000001	0100	
00000000000000000000	0000000001	1100	
00000000000000000000	0000000011	0100	
00000000000000000010	0000000001	0100	
	Tag	Index	Offset



16 KB Direct Mapped Cache, 16B blocks

- Valid bit: determines whether anything is stored in that row (when computer initially turned on, all entries invalid)

Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	0				
2	0				
3	0				
4	0				
5	0				
6	0				
7	0				
...					
1022	0				
1023	0				

Cal

CMSC L11 Caches II (7)

©ards, Spring 2008 © UCS

1. Read 0x00000014

- 000000000000000000 000000001 0100
Tag field Index field Offset

Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	0				
2	0				
3	0				
4	0				
5	0				
6	0				
7	0				
...					
1022	0				
1023	0				

Cal

CMSC L11 Caches II (8)

©ards, Spring 2008 © UCS

So we read block 1 (000000001)

- 000000000000000000 000000001 0100
Tag field Index field Offset

Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	0				
2	0				
3	0				
4	0				
5	0				
6	0				
7	0				
...					
1022	0				
1023	0				

Cal

CMSC L11 Caches II (9)

©ards, Spring 2008 © UCS

No valid data

- 000000000000000000 000000001 0100
Tag field Index field Offset

Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	0				
2	0				
3	0				
4	0				
5	0				
6	0				
7	0				
...					
1022	0				
1023	0				

Cal

CMSC L11 Caches II (10)

©ards, Spring 2008 © UCS

So load that data into cache, setting tag, valid

- 000000000000000000 000000001 0100
Tag field Index field Offset

Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	1	d	c	b	a
2	0				
3	0				
4	0				
5	0				
6	0				
7	0				
...					
1022	0				
1023	0				

Cal

CMSC L11 Caches II (11)

©ards, Spring 2008 © UCS

Read from cache at offset, return word b

- 000000000000000000 000000001 0100
Tag field Index field Offset

Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	1	d	c	b	a
2	0				
3	0				
4	0				
5	0				
6	0				
7	0				
...					
1022	0				
1023	0				

Cal

CMSC L11 Caches II (12)

©ards, Spring 2008 © UCS

2. Read 0x0000001C = 0...00 0..001 1100

▪ 00000000000000000000 000000001 1100

Valid Tag field Index field Offset

Index	Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0					
1	1	0	d	c	b	a
2	0					
3	0					
4	0					
5	0					
6	0					
7	0					
...						
1022	0					
1023	0					

Cal CSMC L1I Caches II (6) Qards, Spring 2008 © UCS

Index is Valid

▪ 00000000000000000000 000000001 1100

Valid Tag field Index field Offset

Index	Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0					
1	1	0	d	c	b	a
2	0					
3	0					
4	0					
5	0					
6	0					
7	0					
...						
1022	0					
1023	0					

Cal CSMC L1I Caches II (6) Qards, Spring 2008 © UCS

Index valid, Tag Matches

▪ 00000000000000000000 000000001 1100

Valid Tag field Index field Offset

Index	Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0					
1	1	0	d	c	b	a
2	0					
3	0					
4	0					
5	0					
6	0					
7	0					
...						
1022	0					
1023	0					

Cal CSMC L1I Caches II (6) Qards, Spring 2008 © UCS

Index Valid, Tag Matches, return d

▪ 00000000000000000000 000000001 1100

Valid Tag field Index field Offset

Index	Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0					
1	1	0	d	c	b	a
2	0					
3	0					
4	0					
5	0					
6	0					
7	0					
...						
1022	0					
1023	0					

Cal CSMC L1I Caches II (6) Qards, Spring 2008 © UCS

3. Read 0x00000034 = 0...00 0..011 0100

▪ 00000000000000000000 000000011 0100

Valid Tag field Index field Offset

Index	Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0					
1	1	0	d	c	b	a
2	0					
3	0					
4	0					
5	0					
6	0					
7	0					
...						
1022	0					
1023	0					

Cal CSMC L1I Caches II (7) Qards, Spring 2008 © UCS

So read block 3

▪ 00000000000000000000 000000011 0100

Valid Tag field Index field Offset

Index	Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0					
1	1	0	d	c	b	a
2	0					
3	0					
4	0					
5	0					
6	0					
7	0					
...						
1022	0					
1023	0					

Cal CSMC L1I Caches II (8) Qards, Spring 2008 © UCS

No valid data

00000000000000000000 0000000011 0100
Valid Tag field Index field Offset

Index	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	1	d	c	b	a
2	0				
3	0				
4	0				
5	0				
6	0				
7	0				

1022 0
1023 0

Cal CS43C L11 Caches II (10) © UCSD, Spring 2008

Load that cache block, return word f

00000000000000000000 0000000011 0100
Valid Tag field Index field Offset

Index	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	1	d	c	b	a
2	0				
3	1	h	g	f	e
4	0				
5	0				
6	0				
7	0				

1022 0
1023 0

Cal CS43C L11 Caches II (10) © UCSD, Spring 2008

4. Read 0x00008014 = 0...10 0..001 0100

00000000000000000010 0000000001 0100
Valid Tag field Index field Offset

Index	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	1	d	c	b	a
2	0				
3	1	h	g	f	e
4	0				
5	0				
6	0				
7	0				

1022 0
1023 0

Cal CS43C L11 Caches II (10) © UCSD, Spring 2008

So read Cache Block 1, Data is Valid

00000000000000000010 0000000001 0100
Valid Tag field Index field Offset

Index	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	1	d	c	b	a
2	0				
3	1	h	g	f	e
4	0				
5	0				
6	0				
7	0				

1022 0
1023 0

Cal CS43C L11 Caches II (10) © UCSD, Spring 2008

Cache Block 1 Tag does not match (0 != 2)

00000000000000000010 0000000001 0100
Valid Tag field Index field Offset

Index	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	1	d	c	b	a
2	0				
3	1	h	g	f	e
4	0				
5	0				
6	0				
7	0				

1022 0
1023 0

Cal CS43C L11 Caches II (10) © UCSD, Spring 2008

Miss, so replace block 1 with new data & tag

00000000000000000010 0000000001 0100
Valid Tag field Index field Offset

Index	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	1	l	k	j	i
2	0				
3	1	h	g	f	e
4	0				
5	0				
6	0				
7	0				

1022 0
1023 0

Cal CS43C L11 Caches II (10) © UCSD, Spring 2008

And return word J

000000000000000010 000000001 0100

Valid Tag field Index field Offset

Index	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0				
1	2	l	k	i	i
2	0				
3	1	h	g	f	e
4	0				
5	0				
6	0				
7	0				

1022 0

1023 0

Cal CS43C L11 Caches II (28) Garde, Spring 2008 © UCS

Do an example yourself. What happens?

- Chose from: Cache: Hit, Miss, Miss w. replace
Values returned: a, b, c, d, e, ..., k, l
- Read address 0x00000030 ?
000000000000000000 0000000011 0000
- Read address 0x0000001c ?
000000000000000000 0000000001 1100

Cache

Index	Valid	Tag	0x0-3	0x4-7	0x8-b	0xc-f
0	0					
1	1	2	l	k	i	i
2	0					
3	1	0	h	g	f	e
4	0					
5	0					
6	0					
7	0					

Cal CS43C L11 Caches II (28) Garde, Spring 2008 © UCS

Answers

- 0x00000030 a hit
Index = 3, Tag matches,
Offset = 0, value = e
- 0x0000001c a miss
Index = 1, Tag mismatch, so
replace from memory,
Offset = 0xc, value = d
- Since reads, values
must = memory values
whether or not cached:
 - 0x00000030 = e
 - 0x0000001c = d

Memory Address (hex)	Value of Word
00000010	a
00000014	b
00000018	c
0000001c	d
...	...
00000030	e
00000034	f
00000038	g
0000003c	h
...	...
00008010	i
00008014	j
00008018	k
0000801c	l
...	...

Cal CS43C L11 Caches II (27) Garde, Spring 2008 © UCS

Peer Instruction

A. Mem hierarchies were invented before 1950. (UNIVAC I wasn't delivered 'til 1951)

B. If you know your computer's cache size, you can often make your code run faster.

C. Memory hierarchies take advantage of spatial locality by keeping the most recent data items closer to the processor.

ABC
0: FFF
1: FFT
2: FTF
3: FTT
4: TFF
5: TFT
6: TTF
7: TTT

Cal CS43C L11 Caches II (29) Garde, Spring 2008 © UCS

Peer Instruction

- All caches take advantage of spatial locality.
- All caches take advantage of temporal locality.
- On a read, the return value will depend on what is in the cache.

ABC
0: FFF
1: FFT
2: FTF
3: FTT
4: TFF
5: TFT
6: TTF
7: TTT

Cal CS43C L11 Caches II (28) Garde, Spring 2008 © UCS

And in Conclusion...

- Mechanism for transparent movement of data among levels of a storage hierarchy
 - set of address/value bindings
 - address \Rightarrow index to set of candidates
 - compare desired address with tag
 - service hit or miss
 - load new block and binding on miss

address: tag index offset

000000000000000000 0000000001 1100

Valid

Index	Valid	Tag	0xc-f	0x8-b	0x4-7	0x0-3
0	0					
1	1	0	d	c	b	a
2	0					
3	0					

Cal CS43C L11 Caches II (28) Garde, Spring 2008 © UCS