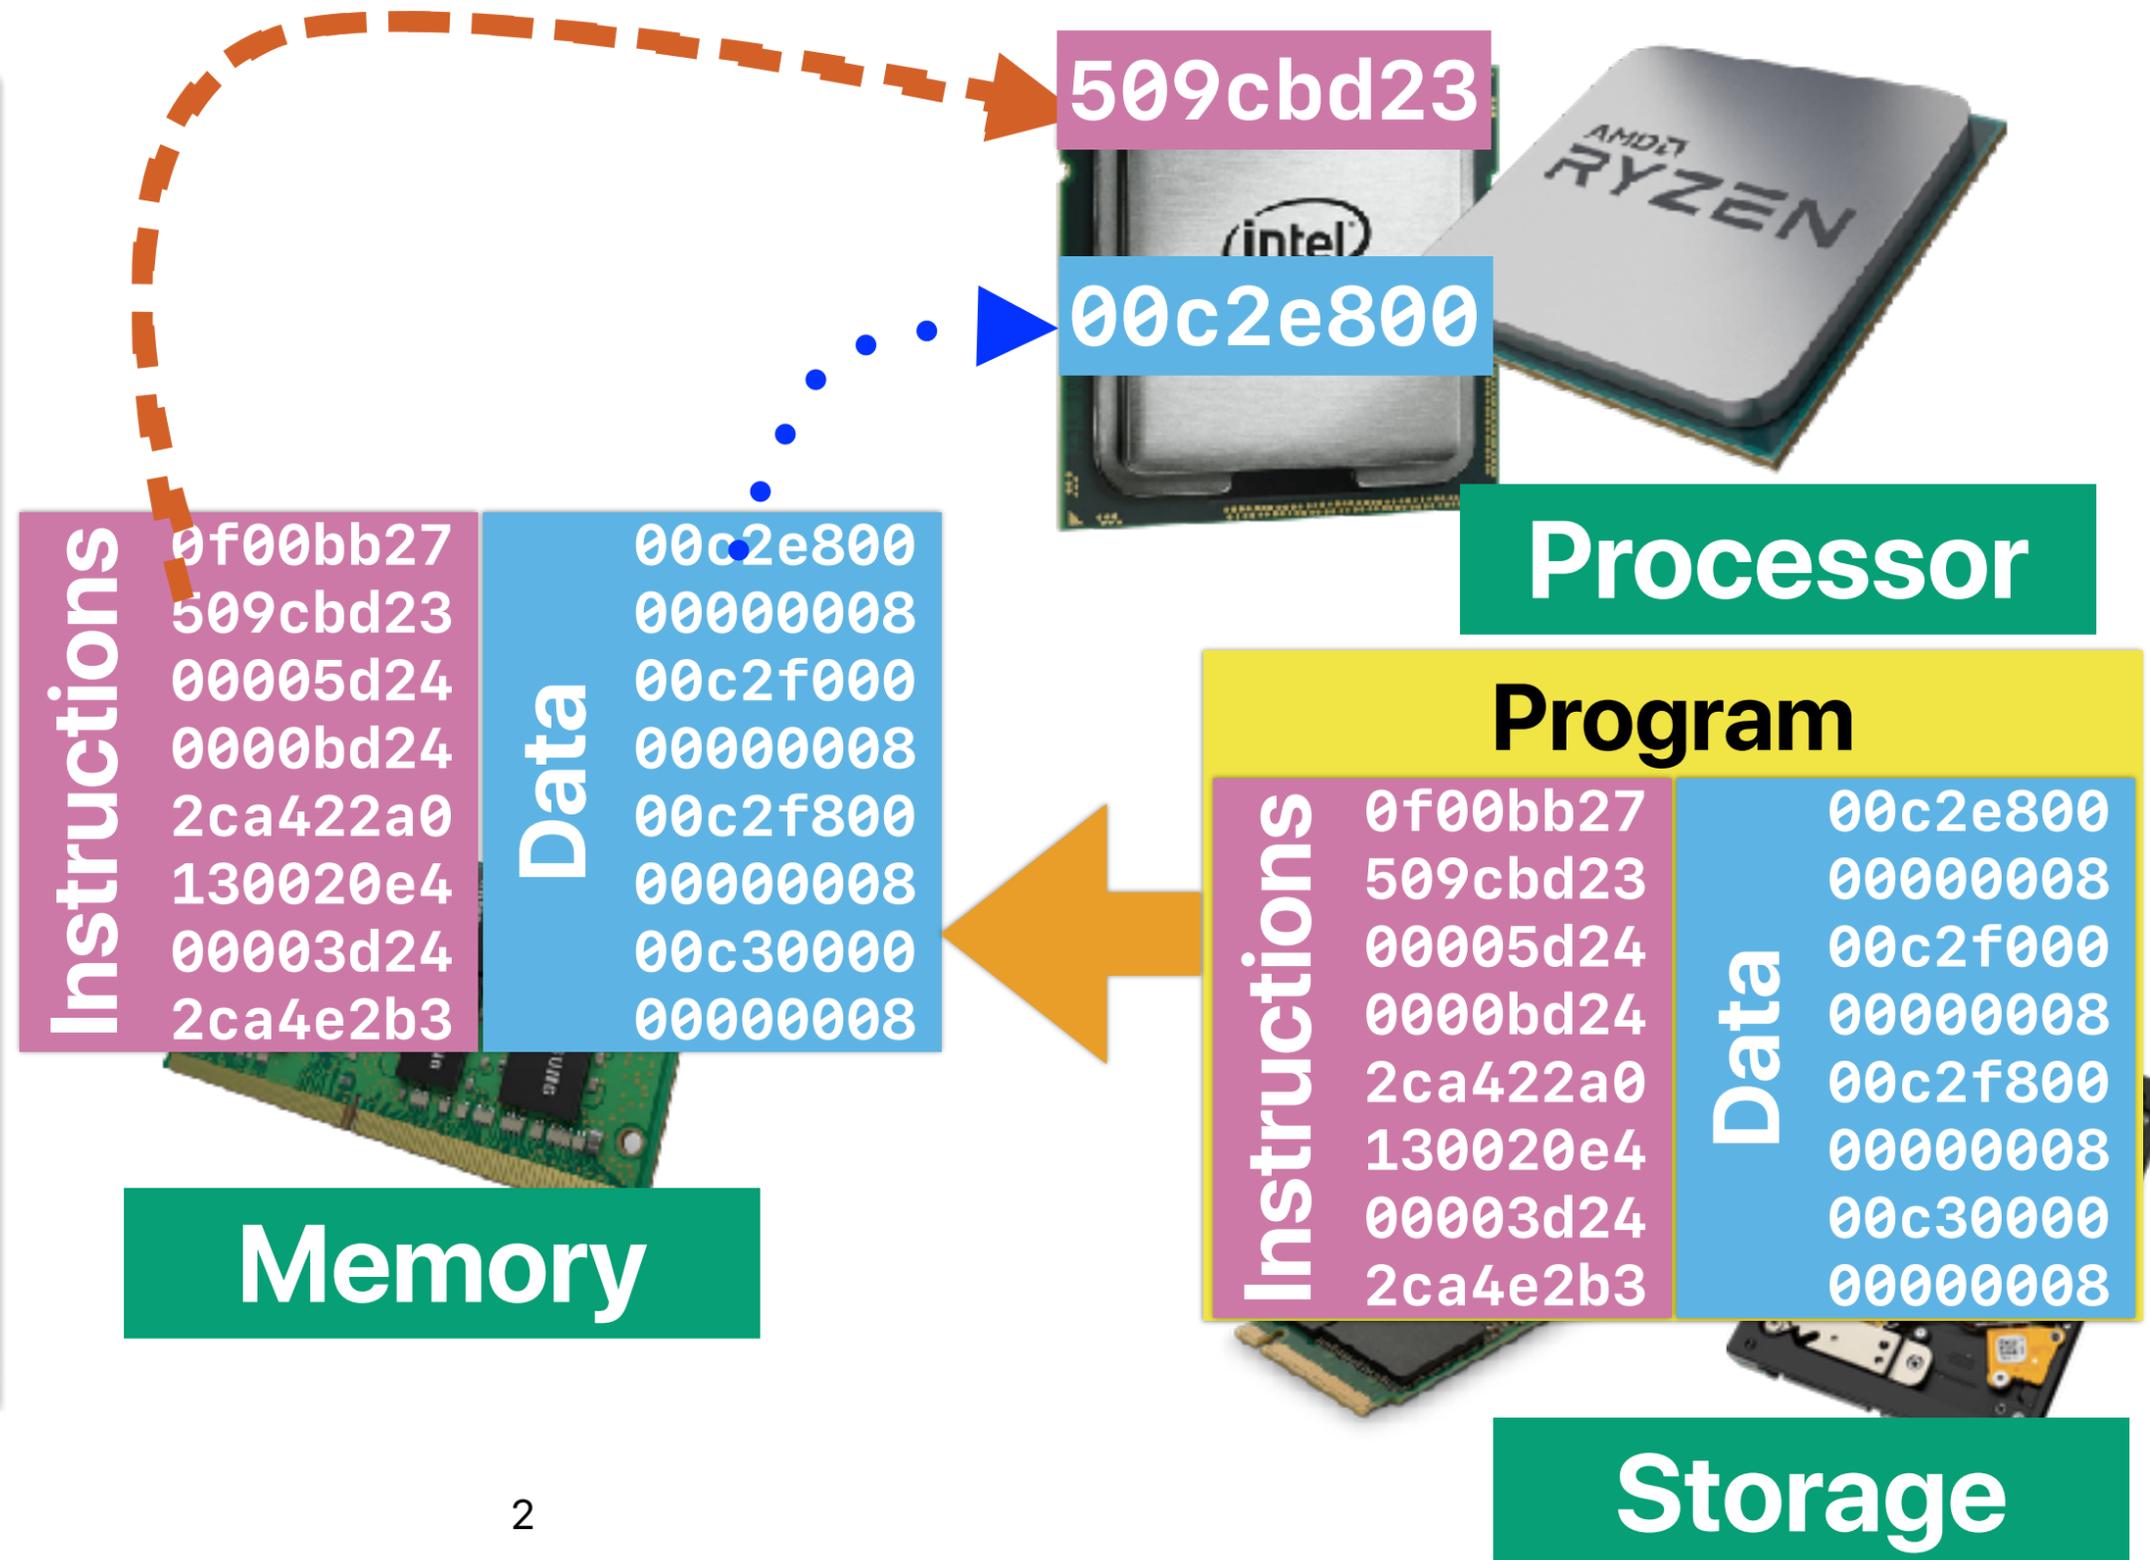


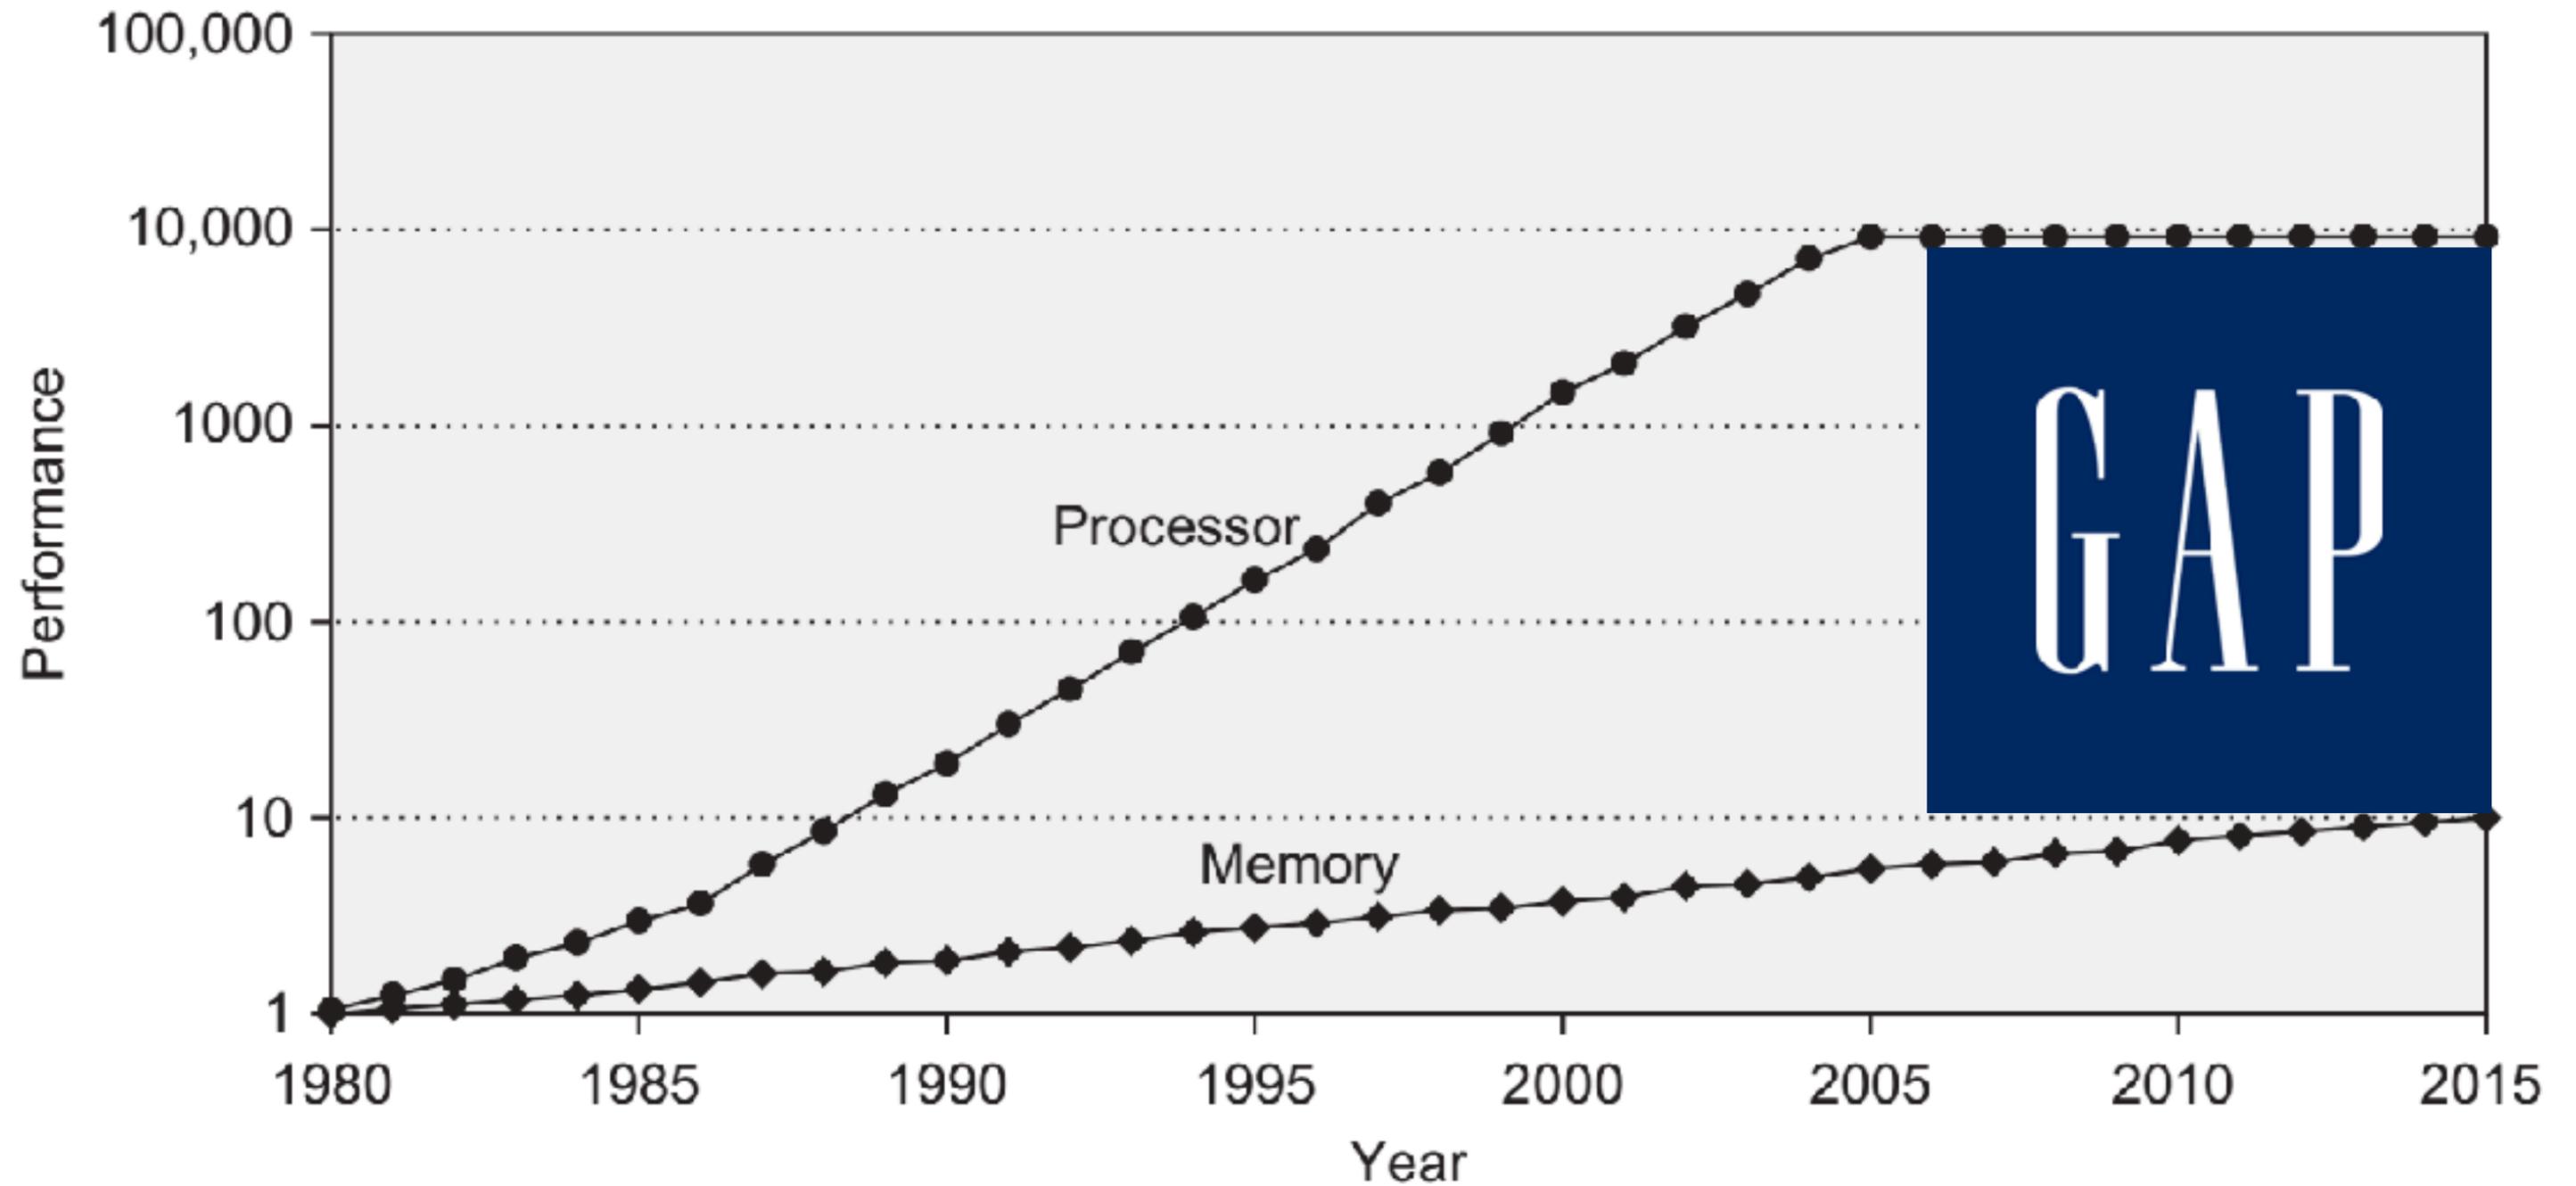
Memory Hierarchy (II): Cache Operations

Hung-Wei Tseng

von Neumann Architecture



Performance gap between Processor/Memory

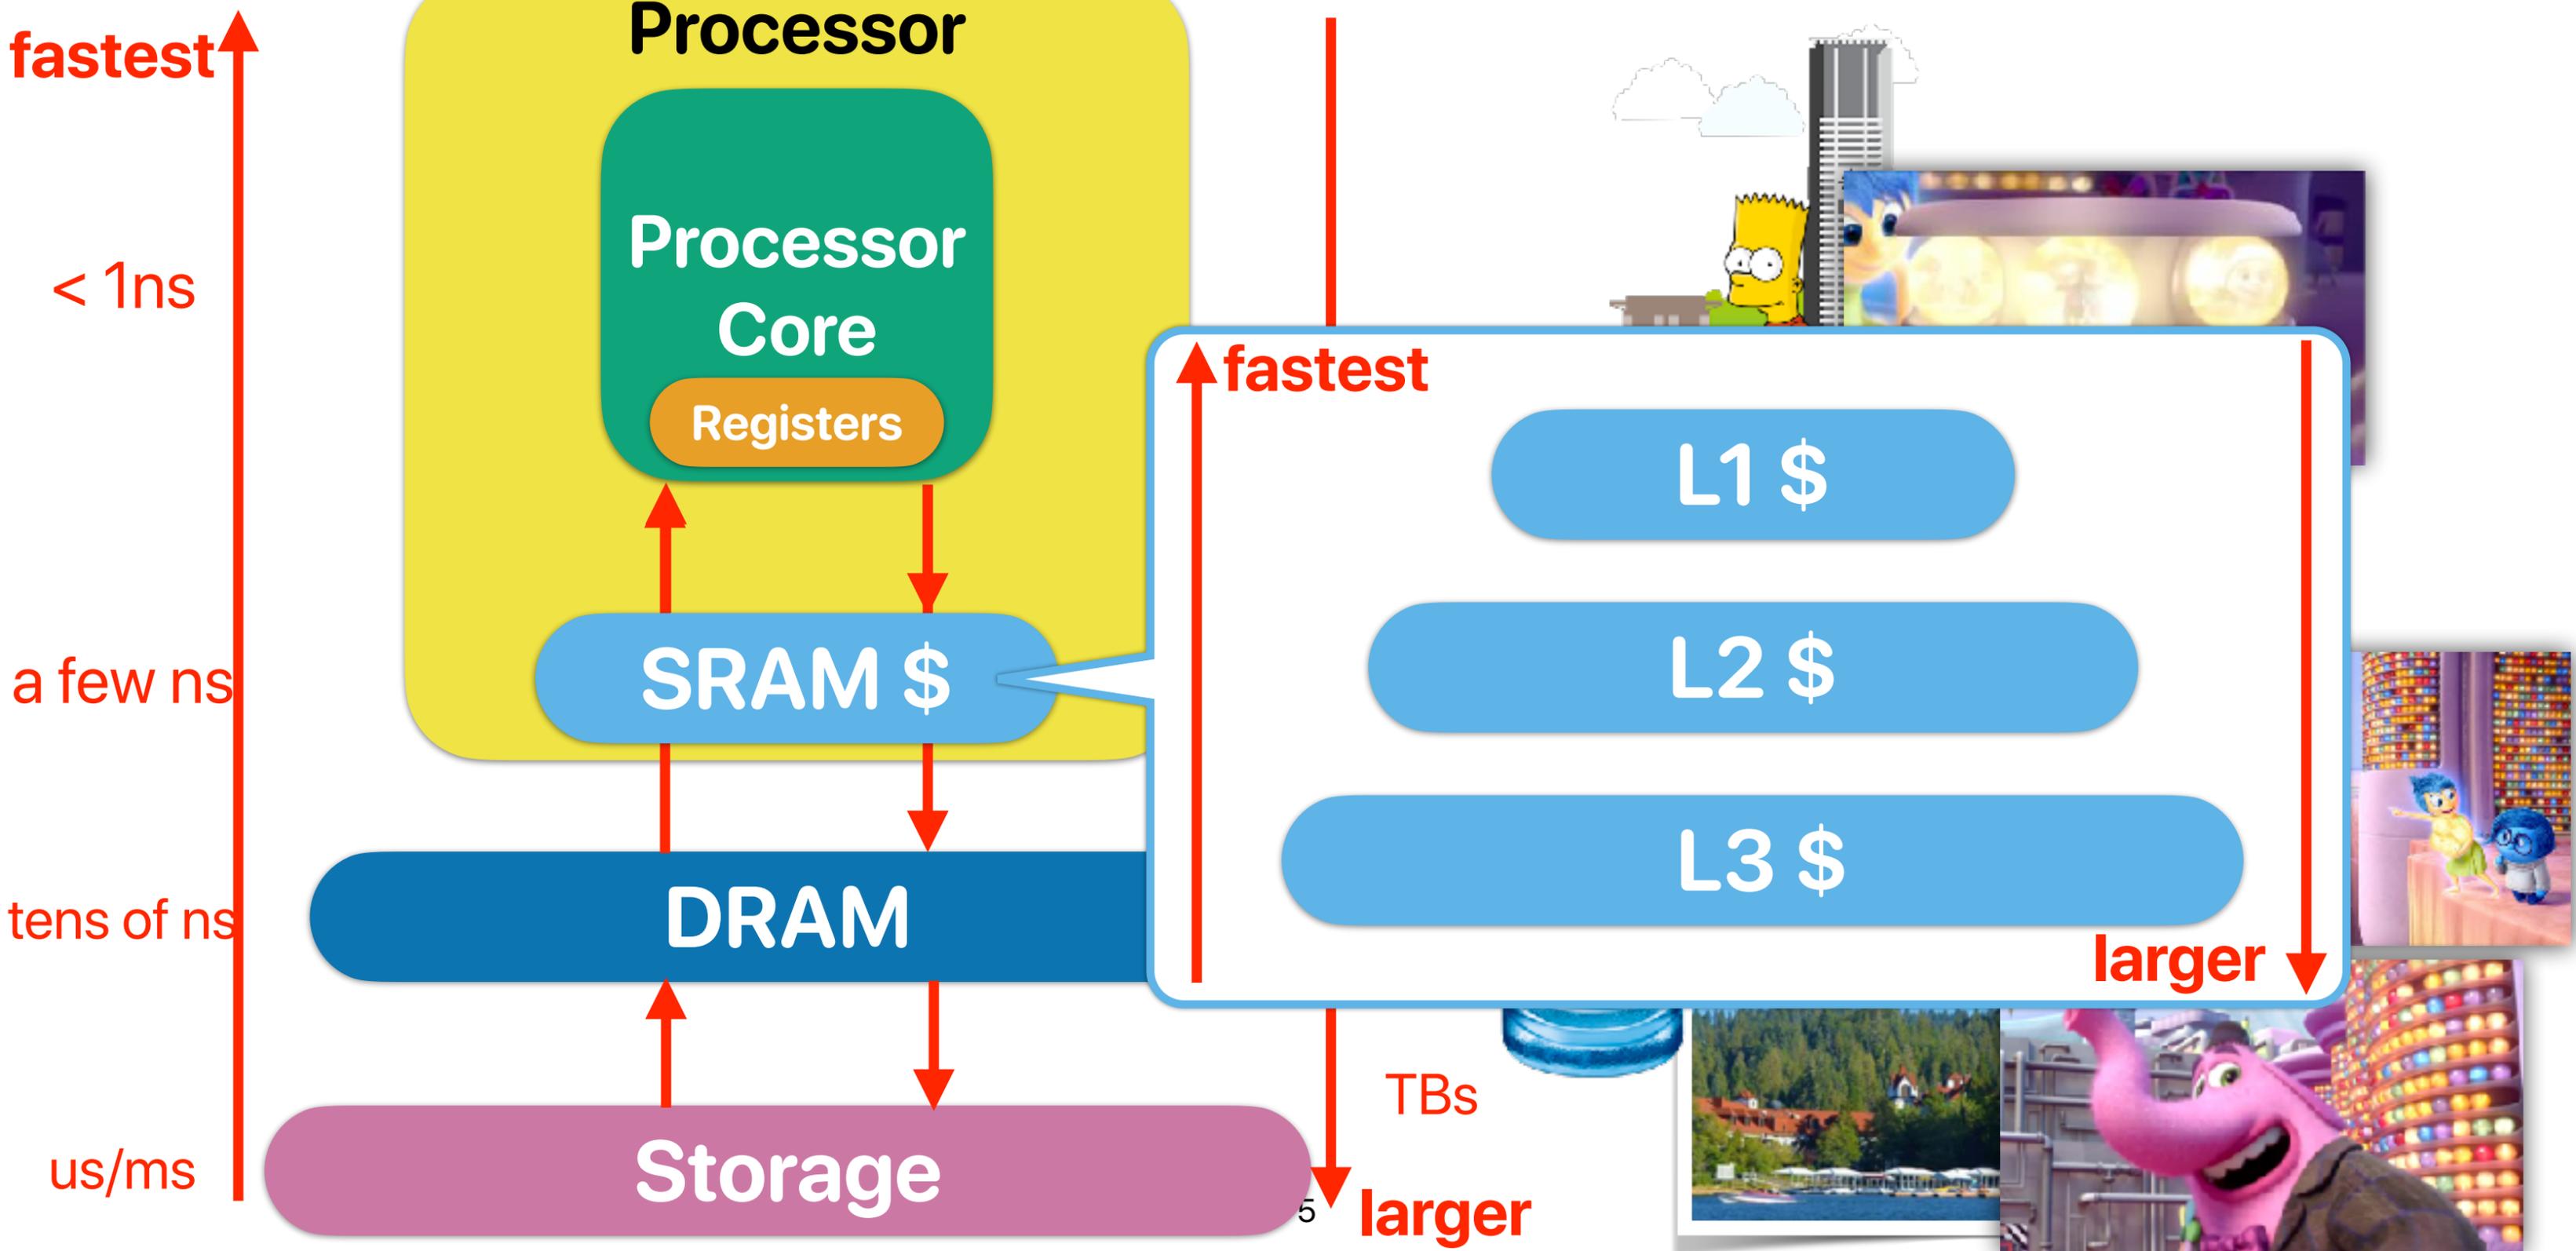


Recap: Alternatives?

Memory technology	Typical access time	\$ per GiB in 2012
SRAM semiconductor memory	0.5–2.5 ns	\$500–\$1000
DRAM semiconductor memory	50–70 ns	\$10–\$20
Flash semiconductor memory	5,000–50,000 ns	\$0.75–\$1.00
Magnetic disk	5,000,000–20,000,000 ns	\$0.05–\$0.10

Fast, but expensive \$\$\$

Recap: Memory Hierarchy



Recap: Locality

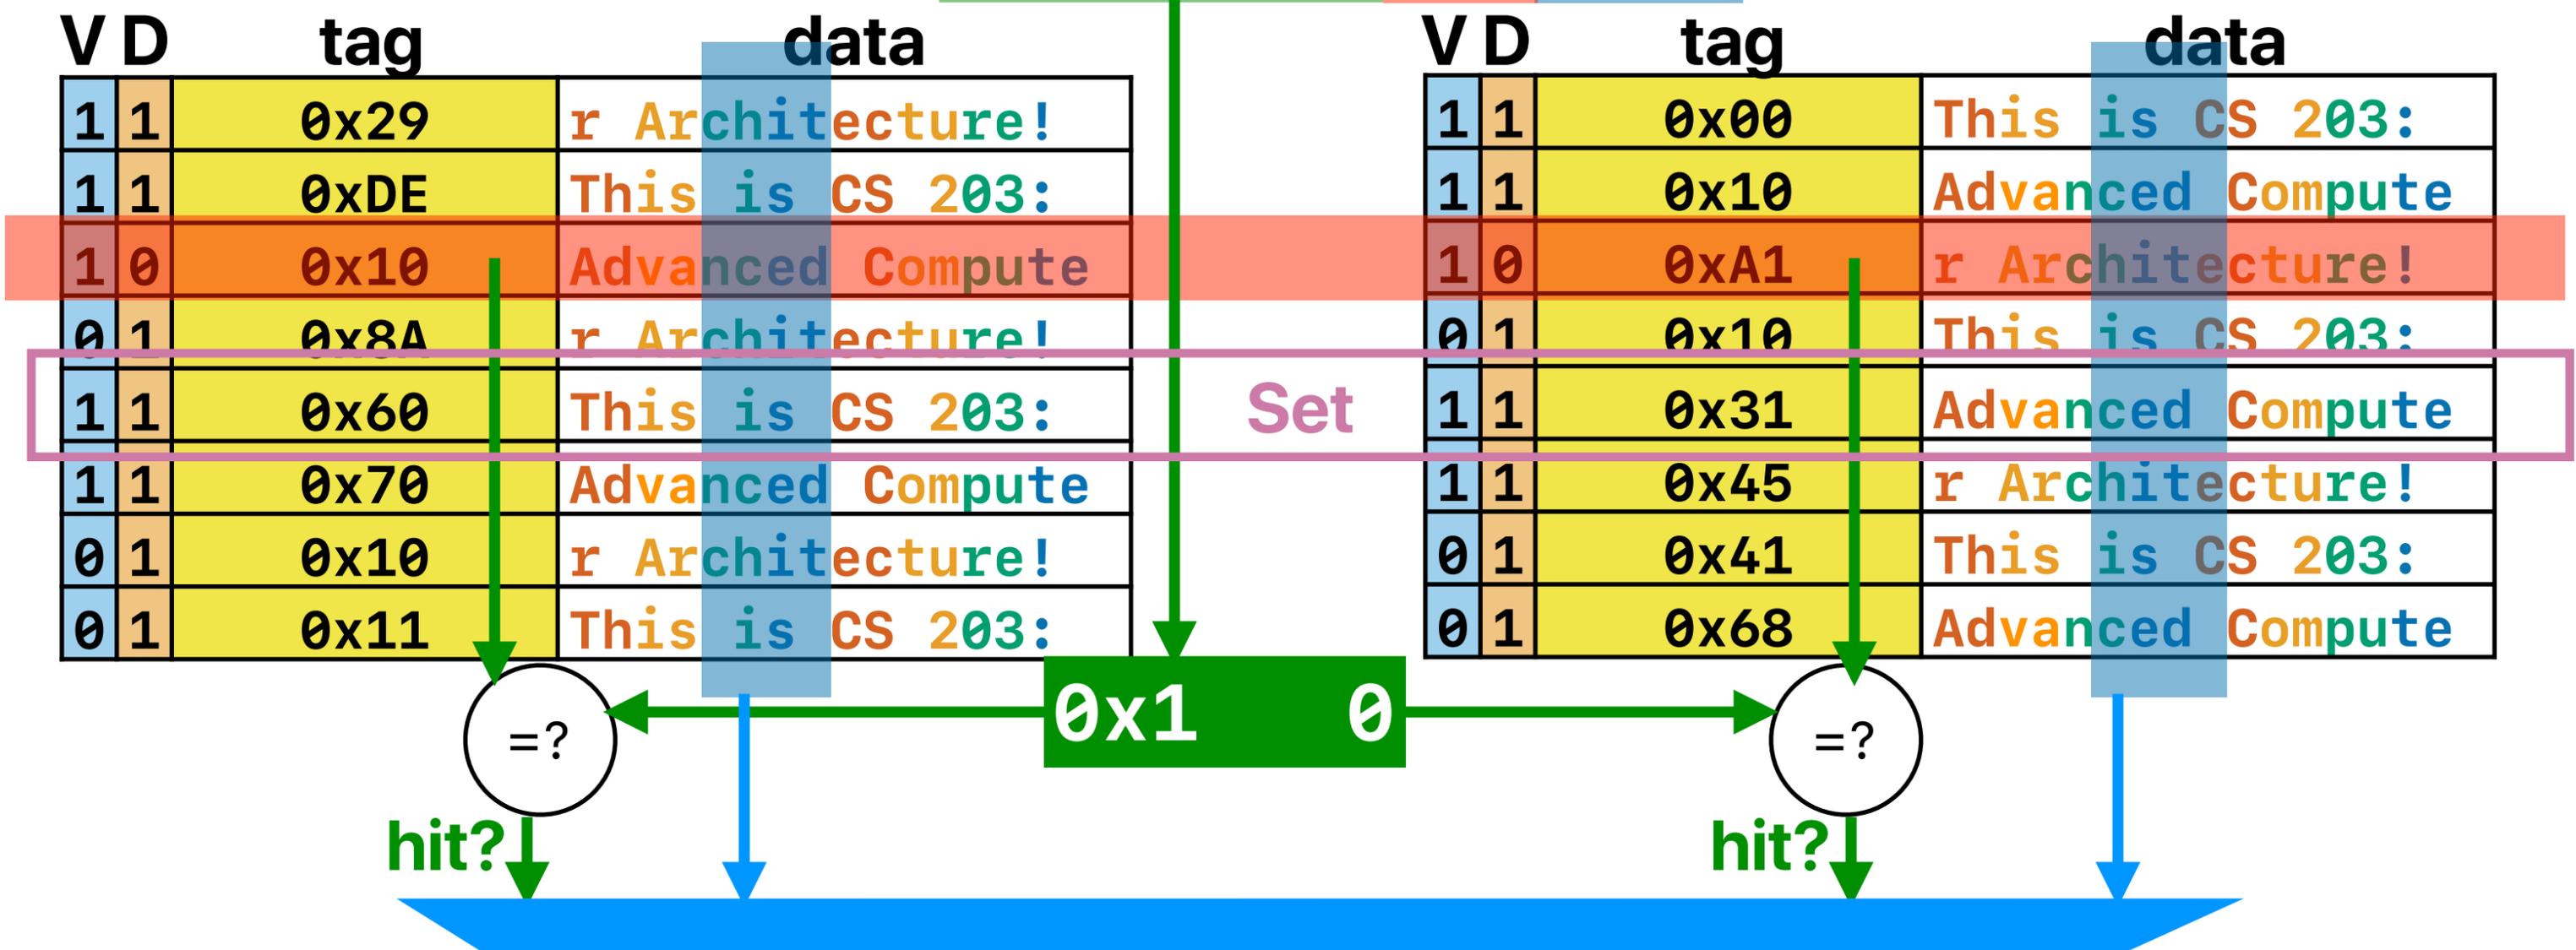
- Spatial locality — application tends to visit nearby stuffs in the memory
 - Code — the current instruction, and then $PC + 4$
 - Data — the current element in an array, then the next
- Temporal locality — application revisit the same thing again and again
 - Code — loops, frequently invoked functions
 - Data — the same data can be read/write many times

Most of time, your program is just visiting a very small amount of data/instructions within a given window

Recap: Way-associative cache

memory address: $0x0$ 8 2 4
 tag set block
 index offset

memory address: $0b00001000000100100$



Recap: $C = ABS$

- **C**: Capacity in data arrays
- **A**: Way-Associativity — how many blocks within a set
 - N-way: N blocks in a set, $A = N$
 - 1 for direct-mapped cache
- **B**: Block Size (Cacheline)
 - How many bytes in a block
- **S**: Number of Sets:
 - A set contains blocks sharing the same index
 - 1 for fully associate cache
- Number of bits in **block** offset — $\lg(\mathbf{B})$
- Number of bits in **set** index: $\lg(\mathbf{S})$
- Tag bits: $\text{address_length} - \lg(\mathbf{S}) - \lg(\mathbf{B})$
 - address_length is 32 bits for 32-bit machine



AMD Phenom II

- L1 data (D-L1) cache configuration of AMD Phenom II
 - Size 64KB, 2-way set associativity, 64B block
 - Assume 64-bit memory address

Which of the following is correct?

A. Tag is 49 bits

B. Index is 8 bits

C. Offset is 7 bits

D. The cache has 1024 sets

E. None of the above

$$C = ABS$$

$$64\text{KB} = 2 * 64 * S$$

$$S = 512$$

$$\text{offset} = \lg(64) = 6 \text{ bits}$$

$$\text{index} = \lg(512) = 9 \text{ bits}$$

$$\text{tag} = 64 - \lg(512) - \lg(64) = 49 \text{ bits}$$

Corollary of $C = ABS$

memory address: 0b 000010000 010 0100

tag set block
index offset

- number of bits in **block** offset — $\lg(\mathbf{B})$
- number of bits in **set** index: $\lg(\mathbf{S})$
- tag bits: $\text{address_length} - \lg(\mathbf{S}) - \lg(\mathbf{B})$
 - address_length is 32 bits for 32-bit machine
- $(\text{address} / \text{block_size}) \% \mathbf{S} = \text{set index}$

Team scores



2.5



3



2.5



3

Outline

- How Cache Works

intel Core i7

- L1 data (D-L1) cache configuration of Core i7
 - Size 32KB, 8-way set associativity, 64B block
 - Assume 64-bit memory address
 - Which of the following is **NOT** correct?
 - A. Tag is 52 bits
 - B. Index is 6 bits
 - C. Offset is 6 bits
 - D. The cache has 128 sets

intel Core i7



- L1 data (D-L1) cache configuration of Core i7
 - Size 32KB, 8-way set associativity, 64B block
 - Assume 64-bit memory address
 - Which of the following is **NOT** correct?
 - A. Tag is 52 bits
 - B. Index is 6 bits
 - C. Offset is 6 bits
 - D. The cache has 128 sets

intel Core i7

- L1 data (D-L1) cache configuration of Core i7
 - Size 32KB, 8-way set associativity, 64B block
 - Assume 64-bit memory address
 - Which of the following is NOT correct?

A. Tag is 52 bits

B. Index is 6 bits

C. Offset is 6 bits

D. The cache has 128 sets

$$C = ABS$$

$$32\text{KB} = 8 * 64 * S$$

$$S = 64$$

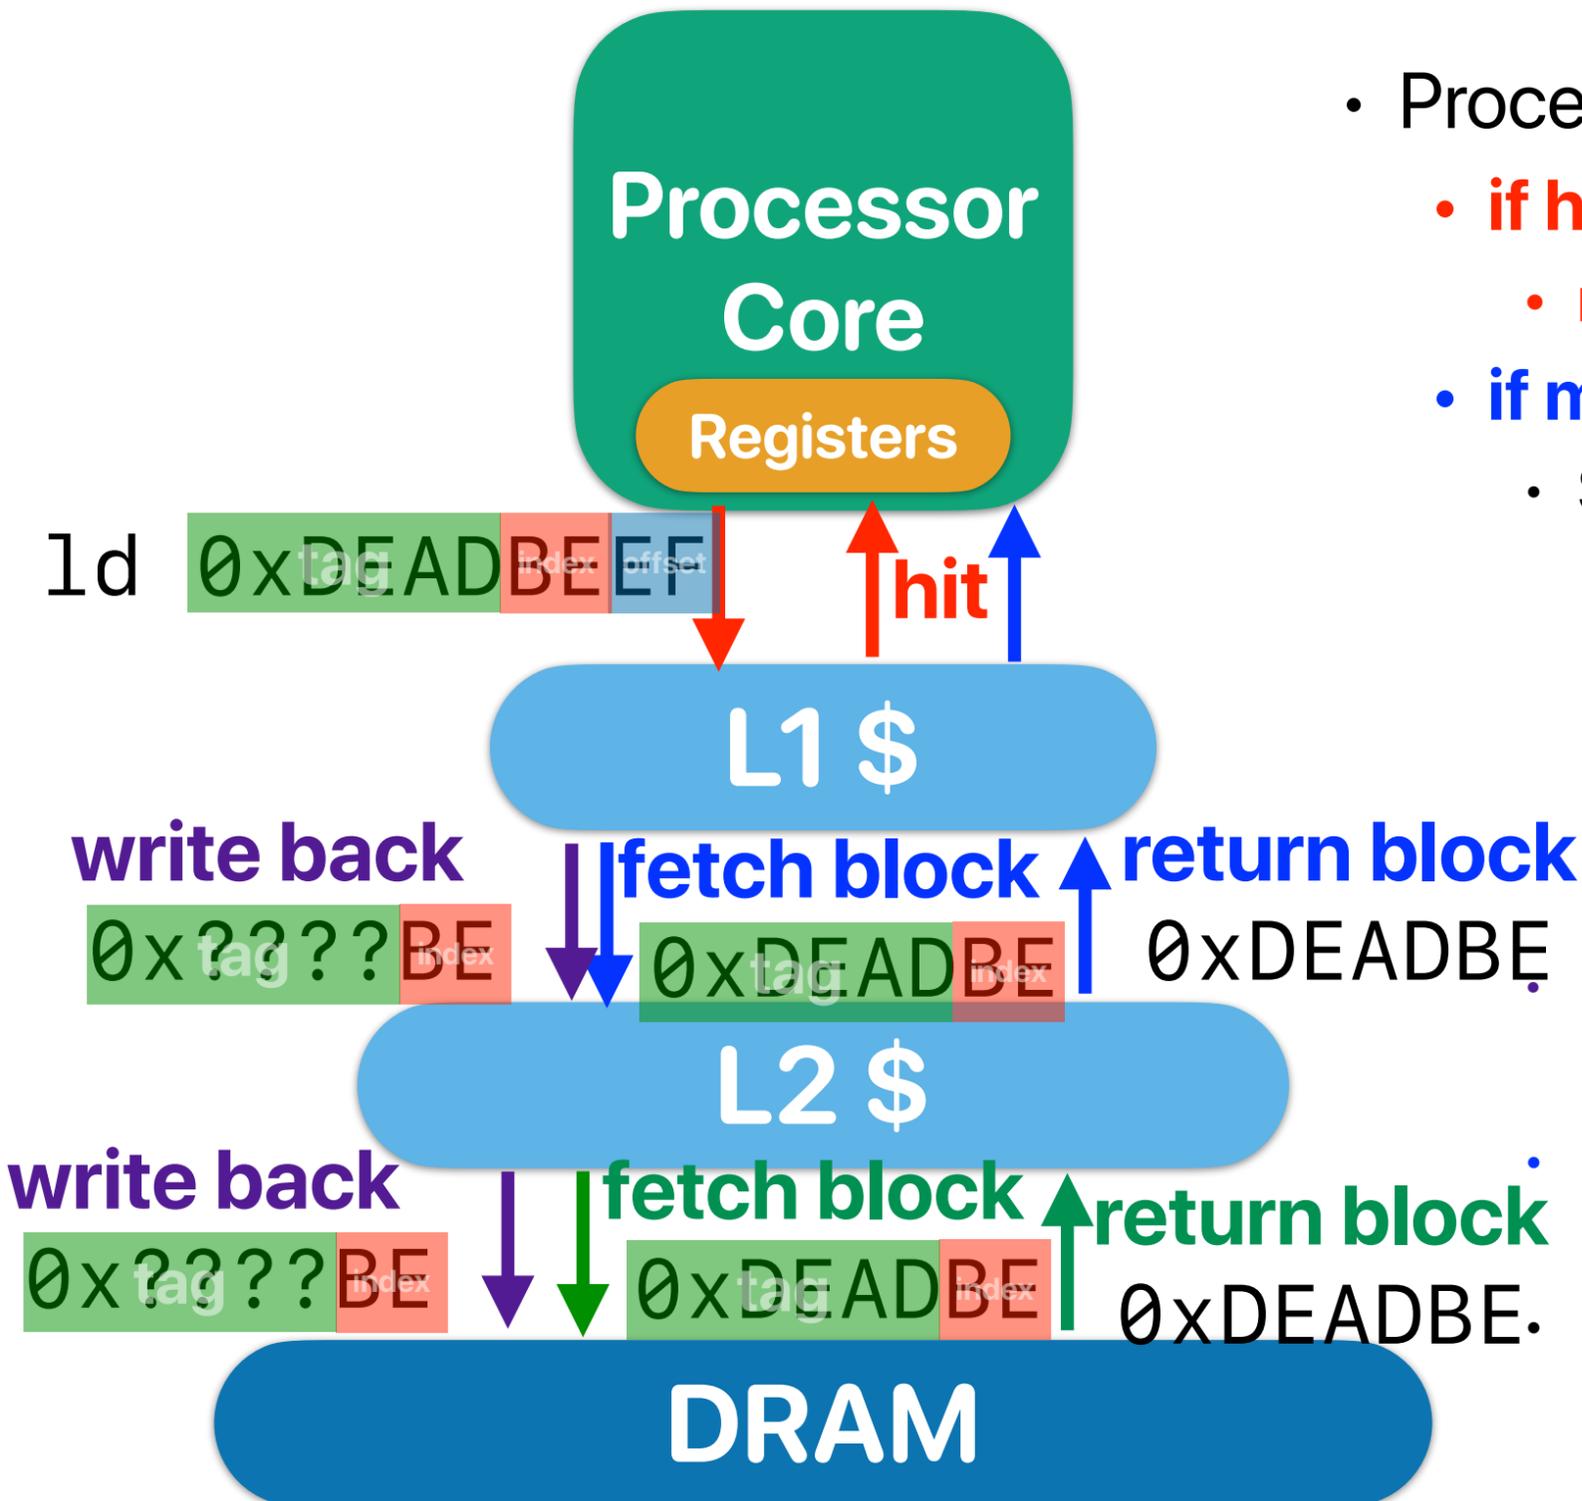
$$\text{offset} = \lg(64) = 6 \text{ bits}$$

$$\text{index} = \lg(64) = 6 \text{ bits}$$

$$\text{tag} = 64 - \lg(64) - \lg(64) = 52 \text{ bits}$$

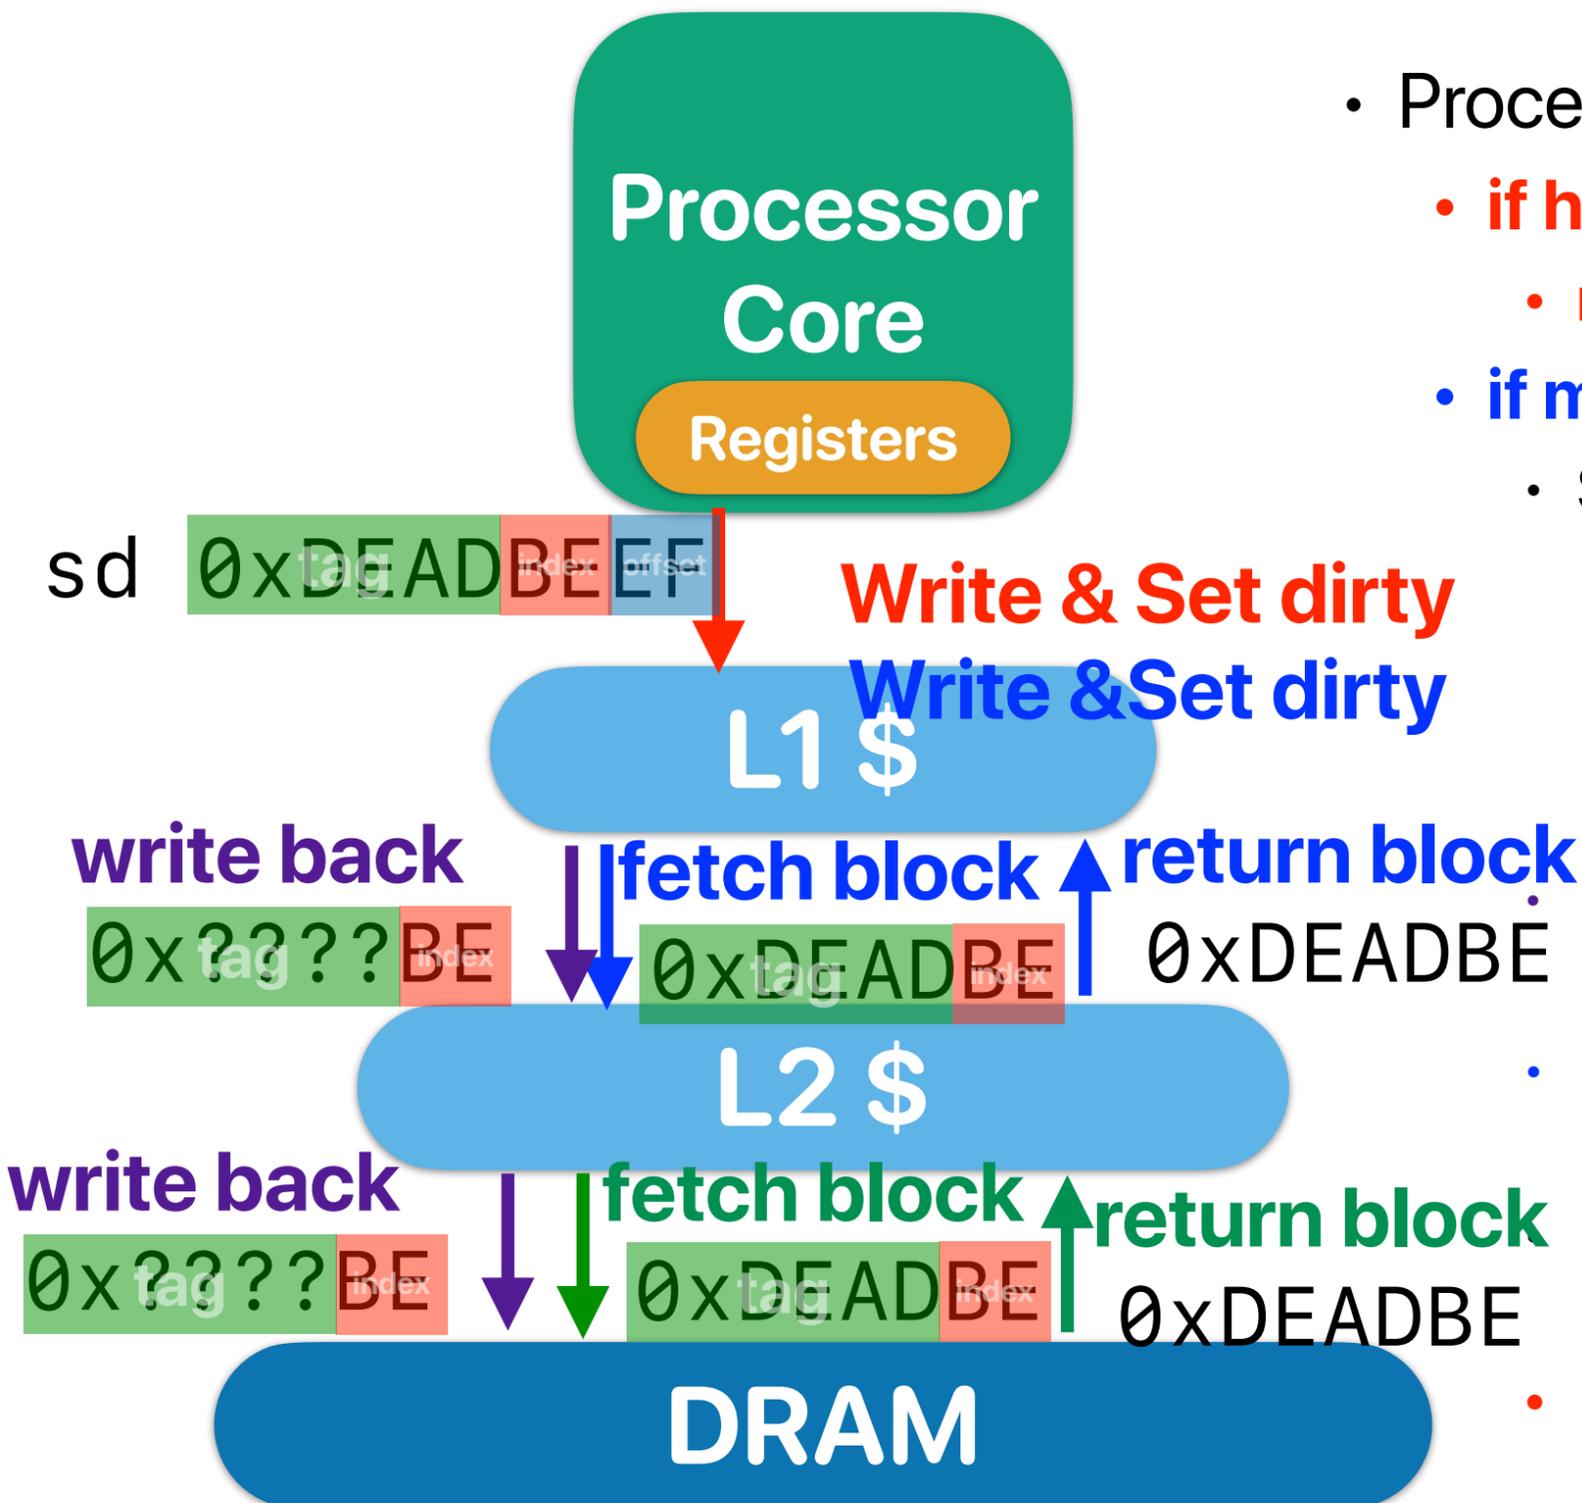
Put everything all together: How cache interacts with CPU

What happens when we read data



- Processor sends load request to L1-\$
 - **if hit**
 - **return data**
 - **if miss**
 - Select a victim block
 - If the target "set" is not full — select an empty/invalidated block as the victim block
 - If the target "set" is full — select a victim block using some policy
 - LRU is preferred — to exploit temporal locality!
 - If the victim block is "dirty" & "valid"
 - **Write back** the block to lower-level memory hierarchy
 - Fetch the requesting block from lower-level memory hierarchy and place in the victim block
- If write-back or fetching causes any miss, repeat the same process

What happens when we write data



- Processor sends load request to L1-\$
 - **if hit**
 - **return data — set DIRTY**
 - **if miss**
 - Select a victim block
 - If the target "set" is not full — select an empty/invalidated block as the victim block
 - If the target "set" is full — select a victim block using some policy
 - LRU is preferred — to exploit temporal locality!
 - If the victim block is "dirty" & "valid"
 - **Write back** the block to lower-level memory hierarchy
 - **Fetch** the requesting block from lower-level memory hierarchy and place in the victim block
- If write-back or fetching causes any miss, repeat the same process
- **Present the write "ONLY" in L1 and set DIRTY**

Simulate the cache!

Simulate a direct-mapped cache

- Consider a direct mapped (1-way) cache with 256 bytes total capacity, a block size of 16 bytes, and the application repeatedly reading the following memory addresses:

- 0b1000000000, 0b1000001000, 0b1000010000, 0b1000010100, 0b1100010000

- $C = A B S$

- $S = 256 / (16 * 1) = 16$

- $\lg(16) = 4$: 4 bits are used for the index

- $\lg(16) = 4$: 4 bits are used for the byte offset

- The tag is $48 - (4 + 4) = 40$ bits

- For example: 0b1000 0000 0000 0000 0000 0000 1000 0000



Simulate a direct-mapped cache

	V	D	Tag	Data
0	1	0	0b10	r Architecture!
1	1	0	0b10	This is CS 203:
2	0	0		
3	0	0		
4	0	0		
5	0	0		
6	0	0		
7	0	0		
8	0	0		
9	0	0		
10	0	0		
11	0	0		
12	0	0		
13	0	0		
14	0	0		
15	0	0		

	tag	index		
	0b10	0000	0000	miss
	0b10	0000	1000	hit!
	0b10	0001	0000	miss
	0b10	0001	0100	hit!
	0b11	0001	0000	miss
	0b10	0000	0000	hit!
	0b10	0000	1000	hit!
	0b10	0001	0000	miss
	0b10	0001	0100	hit!

Simulate a 2-way cache

- Consider a 2-way cache with 256 bytes total capacity, a block size of 16 bytes, and the application repeatedly reading the following memory addresses:

- 0b1000000000, 0b1000001000, 0b1000010000, 0b1000010100, 0b1100010000

- $C = A B S$
- $S = 256 / (16 * 2) = 8$
- $8 = 2^3$: 3 bits are used for the index
- $16 = 2^4$: 4 bits are used for the byte offset
- The tag is $32 - (3 + 4) = 25$ bits
- For example: 0b1000 0000 0000 0000 0000 0000 0001 0000



Simulate a 2-way cache

	V	D	Tag	Data	V	D	Tag	Data
0	1	0	0b10	r Architecture!	0	0		
1	1	0	0b10	This is CS 203:	1	0	0b11	Advanced Compute
2	0	0			0	0		
3	0	0			0	0		
4	0	0			0	0		
5	0	0			0	0		
6	0	0			0	0		
7	0	0			0	0		

	tag	index		
0b10	0000	0000		miss
0b10	0000	1000		hit!
0b10	0001	0000		miss
0b10	0001	0100		hit!
0b11	0001	0000		miss
0b10	0000	0000		hit!
0b10	0000	1000		hit!
0b10	0001	0000		hit
0b10	0001	0100		hit!

AMD Phenom II

- D-L1 Cache configuration of AMD Phenom II
 - Size 64KB, 2-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++) {
    c[i] = a[i] + b[i];
    //load a, b, and then store to c
}
```

What's the data cache miss rate for this code?

- A. 6.25%
- B. 56.25%
- C. 66.67%
- D. 68.75%
- E. 100%

AMD Phenom II

- D-L1 Cache configuration of AMD Phenom II
 - Size 64KB, 2-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++) {
    c[i] = a[i] + b[i];
    //load a, b, and then store to c
}
```

What's the data cache miss rate for this code?

- A. 6.25%
- B. 56.25%
- C. 66.67%
- D. 68.75%
- E. 100%

AMD Phenom II

- D-L1 Cache configuration of AMD Phenom II
 - Size 64KB, 2-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++) {
    c[i] = a[i] + b[i];
    //load a, b, and then store to c
}
```

What's the data cache miss rate for this code?

- A. 6.25%
- B. 56.25%
- C. 66.67%
- D. 68.75%
- E. 100%

C = ABS
64KB = 2 * 64 * S
S = 512
offset = lg(64) = 6 bits
index = lg(512) = 9 bits
tag = 64 - lg(512) - lg(64) = 49 bits

AMD Phenom II

100% miss rate!

- Size 64KB, 2-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++)
    c[i] = a[i] + b[i]; /*load a[i], load b[i], store c[i]*/
```

C = ABS
 64KB = 2 * 64 * S
 S = 512
 offset = lg(64) = 6 bits
 index = lg(512) = 9 bits
 tag = the rest bits

	address in hex	address in binary			tag	index	hit? miss?
		tag	index	offset			
load a[0]	0x20000	0b10	0000 0000 0000	0000	0x4	0	miss
load b[0]	0x30000	0b11	0000 0000 0000	0000	0x6	0	miss
store c[0]	0x10000	0b01	0000 0000 0000	0000	0x2	0	miss, evict 0x4
load a[1]	0x20004	0b10	0000 0000 0000	0100	0x4	0	miss, evict 0x6
load b[1]	0x30004	0b11	0000 0000 0000	0100	0x6	0	miss, evict 0x2
store c[1]	0x10004	0b01	0000 0000 0000	0100	0x2	0	miss, evict 0x4
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
load a[15]	0x2003C	0b10	0000 0000 0011	1100	0x4	0	miss, evict 0x6
load b[15]	0x3003C	0b11	0000 0000 0011	1100	0x6	0	miss, evict 0x2
store c[15]	0x1003C	0b01	0000 0000 0011	1100	0x2	0	miss, evict 0x4
load a[16]	0x20040	0b10	0000 0000 0100	0000	0x4	1	miss
load b[16]	0x30040	0b11	0000 0000 0100	0000	0x6	1	miss
store c[16]	0x10040	0b01	0000 0000 0100	0000	0x2	1	miss, evict 0x4

AMD Phenom II

- D-L1 Cache configuration of AMD Phenom II
 - Size 64KB, 2-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 32-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++) {
    c[i] = a[i] + b[i];
    //load a, b, and then store to c
}
```

What's the data cache miss rate for this code?

- A. 6.25%
- B. 56.25%
- C. 66.67%
- D. 68.75%
- E. 100%**

$$C = ABS$$

$$64KB = 2 * 64 * S$$

$$S = 512$$

$$\text{offset} = \lg(64) = 6 \text{ bits}$$

$$\text{index} = \lg(512) = 9 \text{ bits}$$

$$\text{tag} = 64 - \lg(512) - \lg(64) = 49 \text{ bits}$$

intel Core i7

- D-L1 Cache configuration of intel Core i7 processor
 - Size 32KB, 8-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++) {
    c[i] = a[i] + b[i];
    //load a, b, and then store to c
}
```

What's the data cache miss rate for this code?

- A. 6.25%
- B. 56.25%
- C. 66.67%
- D. 68.75%
- E. 100%

intel Core i7



- D-L1 Cache configuration of intel Core i7 processor
 - Size 32KB, 8-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++) {
    c[i] = a[i] + b[i];
    //load a, b, and then store to c
}
```

What's the data cache miss rate for this code?

- A. 6.25%
- B. 56.25%
- C. 66.67%
- D. 68.75%
- E. 100%

intel Core i7

- D-L1 Cache configuration of intel Core i7 processor
 - Size 32KB, 8-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++) {
    c[i] = a[i] + b[i];
    //load a, b, and then store to c
}
```

What's the data cache miss rate for this code?

- A. 6.25%
- B. 56.25%
- C. 66.67%
- D. 68.75%
- E. 100%

$$C = ABS$$

$$32KB = 8 * 64 * S$$

$$S = 64$$

$$\text{offset} = \lg(64) = 6 \text{ bits}$$

$$\text{index} = \lg(64) = 6 \text{ bits}$$

$$\text{tag} = 64 - \lg(64) - \lg(64) = 52 \text{ bits}$$

intel Core i7

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++)
{
    c[i] = a[i] + b[i]; /*load a[i], load b[i], store c[i]*/
}
```

C = ABS
32KB = 8 * 64 * S
S = 64
offset = lg(64) = 6 bits
index = lg(64) = 6 bits
tag = 64 - lg(64) - lg(64) = 52 bits

	address	tag	index	?
load a[0]	0x20000	0x20	0	miss
load b[0]	0x30000	0x30	0	miss
store c[0]	0x10000	0x10	0	miss
load a[1]	0x20004	0x20	0	hit
load b[1]	0x30004	0x30	0	hit
store c[1]	0x10004	0x10	0	hit
⋮	⋮	⋮	⋮	⋮
load a[15]	0x2003C	0x20	0	hit
load b[15]	0x3003C	0x30	0	hit
store c[15]	0x1003C	0x10	0	hit
load a[16]	0x20040	0x20	1	miss
load b[16]	0x30040	0x30	1	miss
store c[16]	0x1003C	0x10	1	miss

$32 * 3 / (512 * 3) = 1/16 = 6.25\%$ (93.75% hit rate!)

intel Core i7

- D-L1 Cache configuration of intel Core i7 processor
 - Size 32KB, 8-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++) {
    c[i] = a[i] + b[i];
    //load a, b, and then store to c
}
```

What's the data cache miss rate for this code?

A. 6.25%

B. 56.25%

C. 66.67%

D. 68.75%

E. 100%

$$C = ABS$$

$$32KB = 8 * 64 * S$$

$$S = 64$$

$$\text{offset} = \lg(64) = 6 \text{ bits}$$

$$\text{index} = \lg(64) = 6 \text{ bits}$$

$$\text{tag} = 64 - \lg(64) - \lg(64) = 52 \text{ bits}$$

Cause of cache misses

3Cs of misses

- Compulsory miss
 - Cold start miss. First-time access to a block
- Capacity miss
 - The working set size of an application is bigger than cache size
- Conflict miss
 - Required data replaced by block(s) mapping to the same set
 - Similar collision in hash

Simulate a direct-mapped cache

- Consider a direct mapped (1-way) cache with 256 bytes total capacity, a block size of 16 bytes, and the application repeatedly reading the following memory addresses:

- 0b1000000000, 0b1000001000, 0b1000010000, 0b1000010100, 0b1100010000

- $C = A B S$

- $S = 256 / (16 * 1) = 16$

- $\lg(16) = 4$: 4 bits are used for the index

- $\lg(16) = 4$: 4 bits are used for the byte offset

- The tag is $48 - (4 + 4) = 40$ bits

- For example: 0b1000 0000 0000 0000 0000 0000 1000 0000



Simulate a direct-mapped cache

	V	D	Tag	Data
0	1	0	0b10	r Architecture!
1	1	0	0b10	This is CS 203:
2	0	0		
3	0	0		
4	0	0		
5	0	0		
6	0	0		
7	0	0		
8	0	0		
9	0	0		
10	0	0		
11	0	0		
12	0	0		
13	0	0		
14	0	0		
15	0	0		

	tag	index		
	0b10	0000	0000	compulsory miss
	0b10	0000	1000	hit!
	0b10	0001	0000	compulsory miss
	0b10	0001	0100	hit!
	0b11	0001	0000	compulsory miss
	0b10	0000	0000	hit!
	0b10	0000	1000	hit!
	0b10	0001	0000	conflict miss
	0b10	0001	0100	hit!

Simulate a 2-way cache

- Consider a 2-way cache with 256 bytes total capacity, a block size of 16 bytes, and the application repeatedly reading the following memory addresses:

- 0b1000000000, 0b1000001000, 0b1000010000, 0b1000010100, 0b1100010000

- $C = A B S$

- $S = 256 / (16 * 2) = 8$

- $8 = 2^3$: 3 bits are used for the index

- $16 = 2^4$: 4 bits are used for the byte offset

- The tag is $32 - (3 + 4) = 25$ bits

- For example: 0b1000 0000 0000 0000 0000 0000 0001 0000



Simulate a 2-way cache

	V	D	Tag	Data	V	D	Tag	Data
0	1	0	0b10	r Architecture!	0	0		
1	1	0	0b10	This is CS 203:	1	0	0b11	Advanced Compute
2	0	0			0	0		
3	0	0			0	0		
4	0	0			0	0		
5	0	0			0	0		
6	0	0			0	0		
7	0	0			0	0		

	tag	index	
0b10	0000	0000	compulsory miss
0b10	0000	1000	hit!
0b10	0001	0000	compulsory miss
0b10	0001	0100	hit!
0b11	0001	0000	compulsory miss
0b10	0000	0000	hit!
0b10	0000	1000	hit!
0b10	0001	0000	hit!
0b10	0001	0100	hit!

AMD Phenom II

- D-L1 Cache configuration of AMD Phenom II
 - Size 64KB, 2-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++) {
    c[i] = a[i] + b[i];
    //load a, b, and then store to c
}
```

How many of the cache misses are **conflict** misses?

- A. 6.25%
- B. 66.67%
- C. 68.75%
- D. 93.75%
- E. 100%

AMD Phenom II

- D-L1 Cache configuration of AMD Phenom II
 - Size 64KB, 2-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++) {
    c[i] = a[i] + b[i];
    //load a, b, and then store to c
}
```

How many of the cache misses are **conflict** misses?

- A. 6.25%
- B. 66.67%
- C. 68.75%
- D. 93.75%
- E. 100%

$$C = ABS$$

$$64KB = 2 * 64 * S$$

$$S = 512$$

$$\text{offset} = \lg(64) = 6 \text{ bits}$$

$$\text{index} = \lg(512) = 9 \text{ bits}$$

$$\text{tag} = 64 - \lg(512) - \lg(64) = 49 \text{ bits}$$

AMD Phenom II

100% miss rate!

- Size 64KB, 2-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++)
    c[i] = a[i] + b[i]; /*load a[i], load b[i], store c[i]*/
```

C = ABS
 64KB = 2 * 64 * S
 S = 512
 offset = lg(64) = 6 bits
 index = lg(512) = 9 bits
 tag = the rest bits

	address in hex	address in binary			tag	index	hit? miss?
		tag	index	offset			
load a[0]	0x20000	0b10	0000 0000 0000	0000	0x4	0	compulsory miss
load b[0]	0x30000	0b11	0000 0000 0000	0000	0x6	0	compulsory miss
store c[0]	0x10000	0b01	0000 0000 0000	0000	0x2	0	compulsory miss, evict
load a[1]	0x20004	0b10	0000 0000 0000	0100	0x4	0	conflict miss, evict 0x6
load b[1]	0x30004	0b11	0000 0000 0000	0100	0x6	0	conflict miss, evict 0x2
store c[1]	0x10004	0b01	0000 0000 0000	0100	0x2	0	conflict miss, evict 0x4
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
load a[15]	0x2003C	0b10	0000 0000 0011	1100	0x4	0	miss, evict 0x6
load b[15]	0x3003C	0b11	0000 0000 0011	1100	0x6	0	miss, evict 0x2
store c[15]	0x1003C	0b01	0000 0000 0011	1100	0x2	0	miss, evict 0x4
load a[16]	0x20040	0b10	0000 0000 0100	0000	0x4	1	compulsory miss
load b[16]	0x30040	0b11	0000 0000 0100	0000	0x6	1	compulsory miss
store c[16]	0x10040	0b01	0000 0000 0100	0000	0x2	1	compulsory miss, evict

AMD Phenom II

- D-L1 Cache configuration of AMD Phenom II
 - Size 64KB, 2-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 32-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++) {
    c[i] = a[i] + b[i];
    //load a, b, and then store to c
}
```

How many of the cache misses are **conflict** misses?

- A. 6.25%
- B. 66.67%
- C. 68.75%
- D. 93.75%**
- E. 100%

C = ABS
64KB = 2 * 64 * S
S = 512
offset = lg(64) = 6 bits
index = lg(512) = 9 bits
tag = 64 - lg(512) - lg(64) = 49 bits

Good/bad practices in coding interviews

Upon receiving the question

	Things
✓	Repeat the question back at the interviewer.
✓	Clarify any assumptions you made subconsciously. Many questions are under-specified on purpose. E.g. a tree-like diagram could very well be a graph that allows for cycles and a naive recursive solution would not work.
✓	Clarify input format and range. Ask whether input can be assumed to be well-formed and non-null.
✓	Work through a small example to ensure you understood the question.
✓	Explain a high level approach even if it is a brute force one.
✓	Improve upon the approach and optimize. Reduce duplicated work and cache repeated computations.
✓	Think carefully, then state and explain the time and space complexity of your approaches.
✓	If stuck, think about related problems you have seen before and how they were solved. Check out the tips in this section.
✗	Ignore information given to you. Every piece is important.
✗	Jump into coding straightaway.
✗	Start coding without interviewer's green light.
✗	Appear too unsure about your approach or analysis.

During coding

	Things
✓	Explain what you are coding/typing to the interviewer, what you are trying to achieve.
✓	Practice good coding style. Clear variable names, consistent operator spacing, proper indentation, etc.
✓	Type/write at a reasonable speed.
✓	As much as possible, write actual compilable code, not pseudocode.
✓	Write in a modular fashion. Extract out chunks of repeated code into functions.
✓	Ask for permission to use trivial functions without having to implement them; saves you some time.
✓	Use the hints given by the interviewer.
✓	Demonstrate mastery of your chosen programming language.
✓	Demonstrate technical knowledge in data structures and algorithms.
✓	If you are cutting corners in your code, state that out loud to your interviewer and say what you would do in a non-interview setting (no time constraints). E.g., "Under non-interview settings, I would write a regex to parse this string rather than using <code>split()</code> which may not cover all cases."
✓	Practice whiteboard space-management skills.

Good practices in assignments

- Clarify the question if you have doubts — using piazza
- Clarify any assumption that you made for the given question — write it down on your answer
- Explain a high-level approach/overview of your solution — write it down on you answer

Announcement

- Joel Emer's Talk next Monday @ 11am
 - We will not have a lecture next Monday to encourage you attend Joel Emer's talk
 - If you capture a screen shot and submit it through iLearn, you will receive a full credit reading quiz
 - The talk cannot be broadcasted on YouTube due to the license constraint
 - Will send out the link later
- Asynchronous session next Wednesday releasing @ 9:30a on YouTube
- Assignment #2 due 11/02
- Office Hours on Zoom (the office hour link, not the lecture one)
 - Hung-Wei/Prof. Usagi: M 8p-9p, W 2p-3p (cancelled next week), but can answer questions through e-mails. Will make up later in the quarter
 - Quan Fan: F 1p-3p

Computer
Science &
Engineering

203

くぐく

