# CSE 351 Section 7

Caches – Winter 2022

# Administrivia

- **Homework 19**
  - Due Tomorrow (Friday 2/18)
  - BIG Homework on cache mechanics
- **Homework 20**
  - Due Wednesday (2/23)
  - Preparation for Lab 4
- **Lab 4**
  - Released Yesterday, due 2/28

# Caches

# Cache Motivation

Going all the way to memory is expensive! What if we had an intermediate place where we could store data closer to the processor?

- **Temporal locality**
  - If you used some data, you will probably use it again
- **Spatial locality**
  - If you used some data, you will probably use data that is close to it also

Can you think of any code examples of the above?

# Cache Review

The cache stores a subset of main memory with much faster access time! It is located much closer to the processor, often on the same chip. When we access memory, we check the cache(s) first.

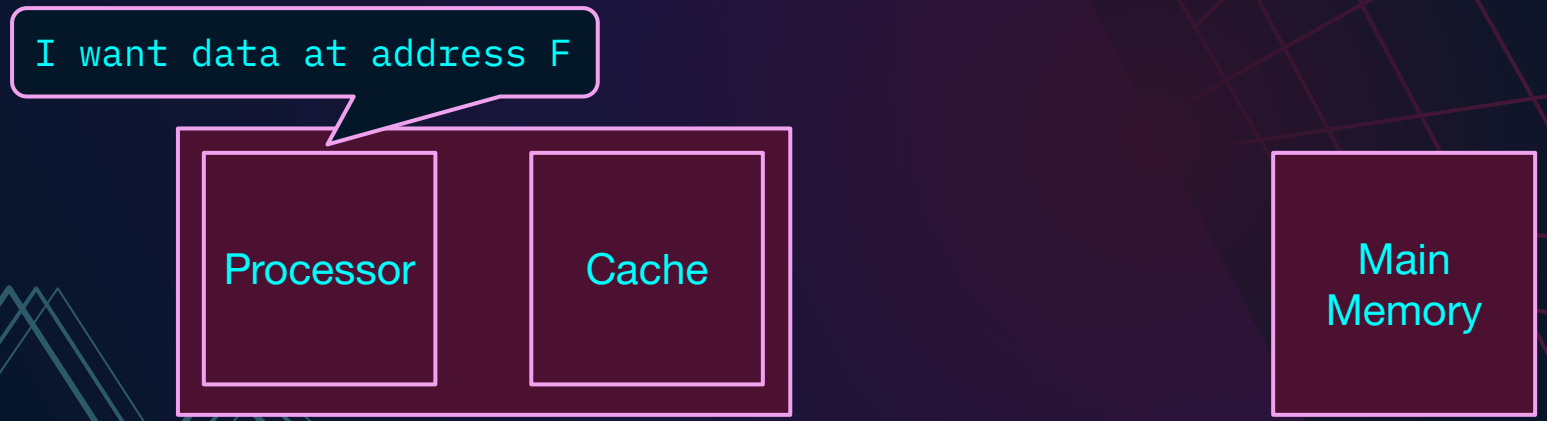

# Cache Review

If the data we want isn't in the cache, that's a ***cache miss***. We have to go to main memory, and then we'll save that data in the cache. By transferring entire blocks of data at a time, we take advantage of spatial locality.

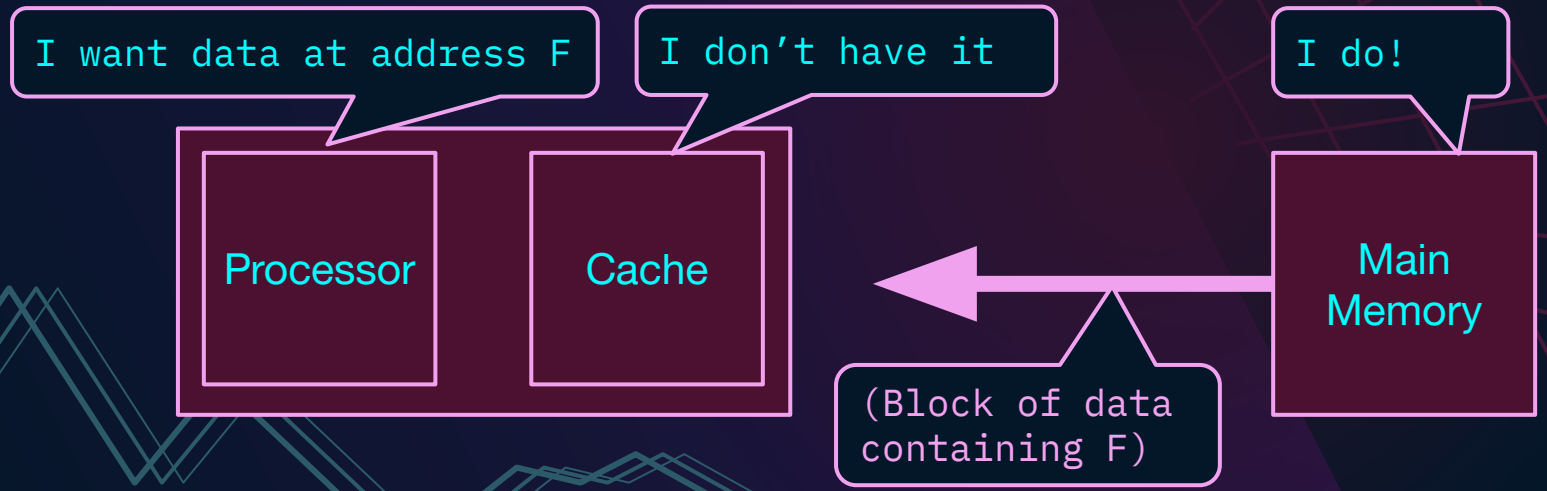

# Cache Review

If the data we want is in the cache and valid, that's a ***cache hit***.

We don't go to memory, which saves us a lot of time!

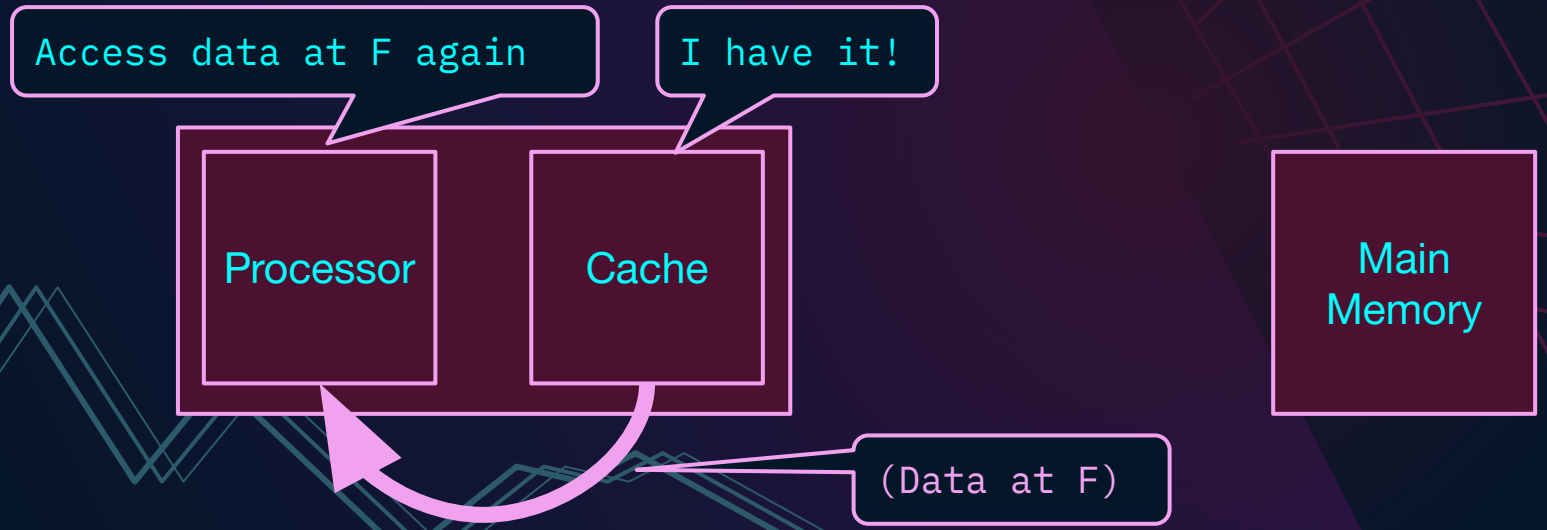

# Cache Organization

- Cache blocks (or lines) are a fixed size.
- The cache has “slots” for blocks of data that are *indexed*.
- We access individual bytes in a block with as an *offset* from the first byte in the block.
- The number of blocks in the cache times the number of bytes in a block gives us the size of the cache.

Cache with 8 blocks of 4 bytes each:

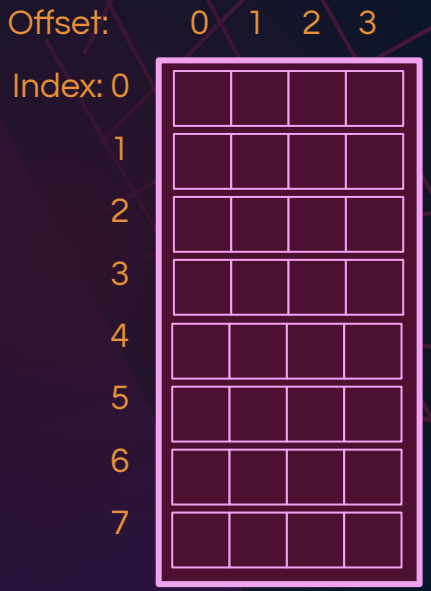

# Tag, Index, Offset!

To determine where each address maps to in the cache, we break it up into:

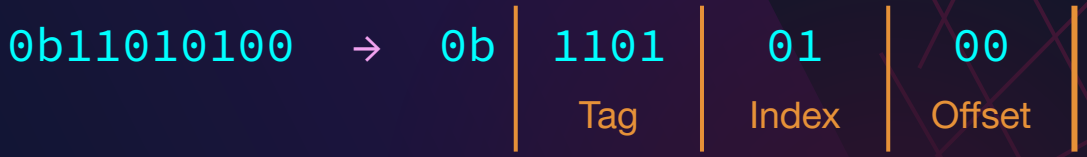


- The **tag** is used to distinguish blocks which map to the same location. It is stored along with the block data and a "valid bit" which indicates whether the data is current and ready to be used.
- The **index** tells us the "slot" in which the data at this address goes.
- The **offset** is tells us how far into the *block* our address is.

This can be thought of kind of like a modulo hashing function; the address gets mapped to a location in the cache based on its index bits.

# Cache Parameters

| Symbol | Meaning |
|---|---|
| K | Block Size |
| C | Cache Size |
| E | Associativity |
| m | Address Width |
| k | # Offset Bits = $\log_2(K)$ |
| s | # Index Bits = $\log_2(C / K)$ |
| t | # Tag Bits = m - k - s |

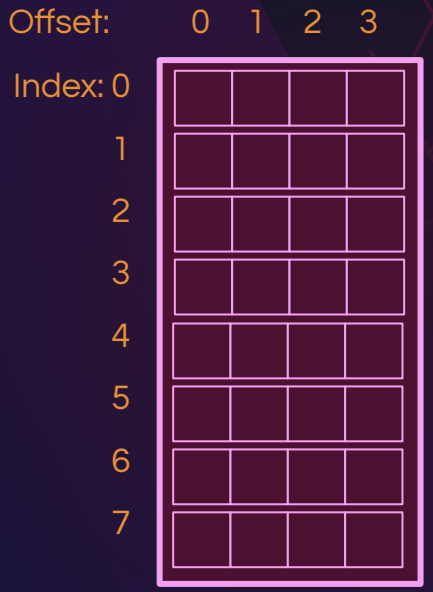


8 blocks of 4 bytes:

K = 4 bytes  k = 2

C = 32 bytes  s = 3

Direct-mapped → E = 1

# Associativity

A cache has associativity > 1 if each index may correspond to > 1 block:

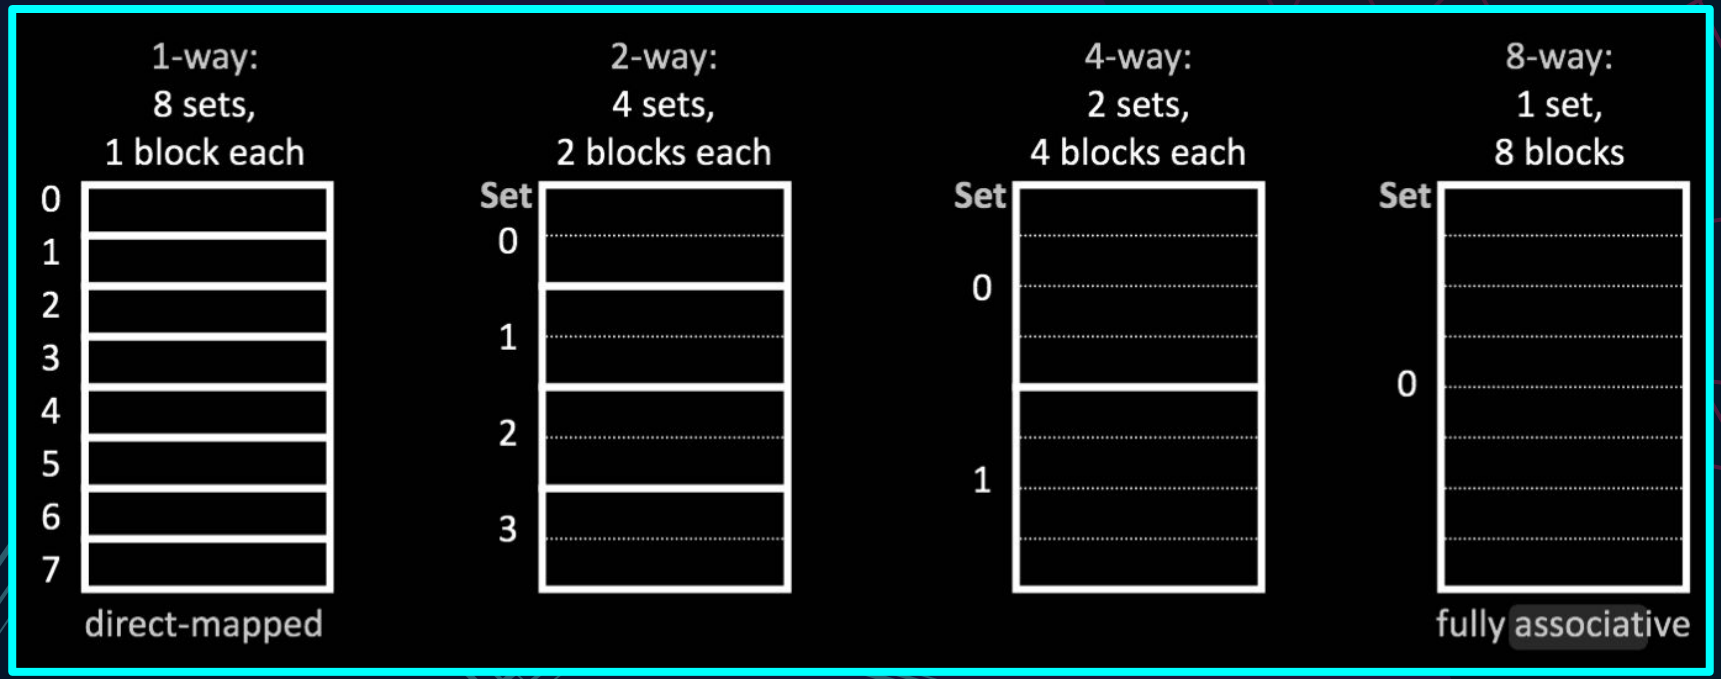

# Example

16 B capacity cache, 4 B block size, direct-mapped, 8 bit address length.

What's the TIO address breakdown? How can we visualize this cache?

- #bits for offset: 2 (4 B in a block; $\log_2(4) = 2$)
- #bits for index: 2 (4 blocks in cache, direct-mapped)
- #bits for tag: 4 (remaining bits of address)

| Index | Tag | Block Data | | | |
|---|---|---|---|---|---|
| 00 | | | | | |
| 01 | | | | | |
| 10 | | | | | |
| 11 | | | | | |

# Example

Read 1 byte from address `0xAD`

1. Translate to Binary:
   a. `0xAD =`
2. Split into TIO
   a. Tag =
   b. Index =
   c. Offset =

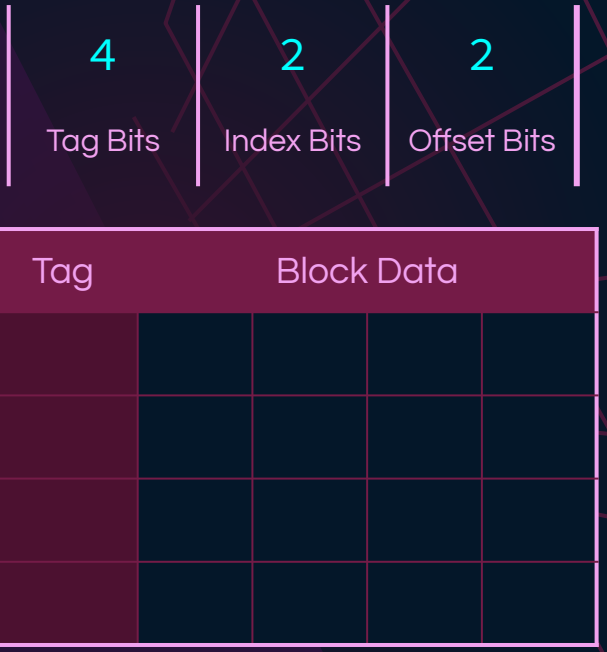

# Example

Read 1 byte from address `0xAD`

1. Translate to Binary:
   a. `0xAD = 0b 1010 1101`
2. Split into TIO
   a. Tag = `1010`
   b. Index = `11`
   c. Offset = `01`

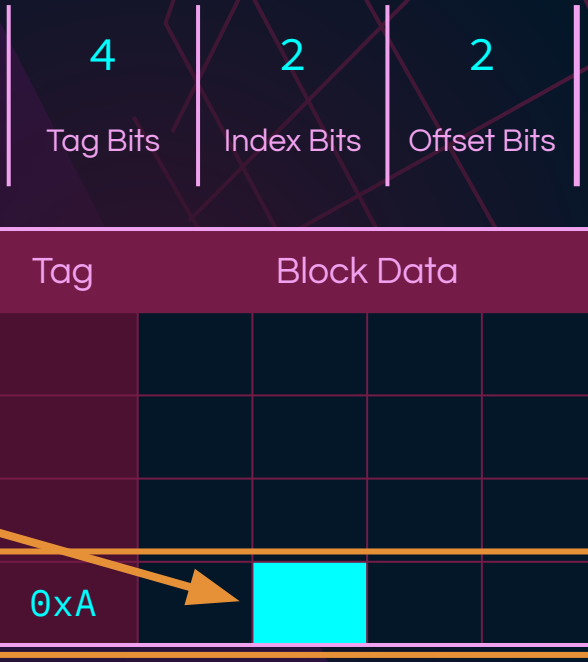

# Exercises

# Exercise: Direct Mapped

| Set | Valid | Tag | B0 | B1 | B2 | B3 |
|---|---|---|---|---|---|---|
| 0 | 1 | 15 | 63 | B4 | C1 | A4 |
| 1 | 0 | -- | -- | -- | -- | -- |
| 2 | 0 | -- | -- | -- | -- | -- |
| 3 | 1 | 0D | DE | AF | BA | DE |
| 4 | 0 | -- | -- | -- | -- | -- |
| 5 | 0 | -- | -- | -- | -- | -- |
| 6 | 1 | 13 | 31 | 14 | 15 | 93 |
| 7 | 0 | -- | -- | -- | -- | -- |

| Set | Valid | Tag | B0 | B1 | B2 | B3 |
|---|---|---|---|---|---|---|
| 8 | 0 | -- | -- | -- | -- | -- |
| 9 | 1 | 00 | 01 | 12 | 23 | 34 |
| A | 1 | 01 | 98 | 89 | CB | BC |
| B | 0 | 1E | 4B | 33 | 10 | 54 |
| C | 0 | -- | -- | -- | -- | -- |
| D | 1 | 11 | C0 | 04 | 39 | AA |
| E | 0 | -- | -- | -- | -- | -- |
| F | 1 | 0F | FF | 6F | 30 | 0 |

# Exercise: 2-Way Associative

| Set | Valid | Tag | B0 | B1 | B2 | B3 |
|---|---|---|---|---|---|---|
| 0 | 0 | -- | -- | -- | -- | -- |
| 1 | 0 | -- | -- | -- | -- | -- |
| 2 | 1 | 03 | 4F | D4 | A1 | 3B |
| 3 | 0 | -- | -- | -- | -- | -- |
| 4 | 0 | 06 | CA | FE | F0 | 0D |
| 5 | 1 | 21 | DE | AD | BE | EF |
| 6 | 0 | -- | -- | -- | -- | -- |
| 7 | 1 | 11 | 00 | 12 | 51 | 55 |

| Set | Valid | Tag | B0 | B1 | B2 | B3 |
|---|---|---|---|---|---|---|
| 0 | 0 | -- | -- | -- | -- | -- |
| 1 | 1 | 2F | 01 | 20 | 40 | 03 |
| 2 | 1 | 0E | 99 | 09 | 87 | 56 |
| 3 | 0 | -- | -- | -- | -- | -- |
| 4 | 0 | -- | -- | -- | -- | -- |
| 5 | 0 | -- | -- | -- | -- | -- |
| 6 | 1 | 37 | 22 | B6 | DB | AA |
| 7 | 0 | -- | -- | -- | -- | -- |

# Exercise: Fully Associative

| Set | Valid | Tag | B0 | B1 | B2 | B3 |
|---|---|---|---|---|---|---|
| 0 | 1 | 1F4 | 00 | 01 | 02 | 03 |
| 0 | 0 | -- | -- | -- | -- | -- |
| 0 | 1 | 100 | F4 | 4D | EE | 11 |
| 0 | 1 | 077 | 12 | 23 | 34 | 45 |
| 0 | 0 | -- | -- | -- | -- | -- |
| 0 | 1 | 101 | DA | 14 | EE | 22 |
| 0 | 0 | -- | -- | -- | -- | -- |
| 0 | 1 | 016 | 90 | 32 | AC | 24 |

| Set | Valid | Tag | B0 | B1 | B2 | B3 |
|---|---|---|---|---|---|---|
| 0 | 0 | -- | -- | -- | -- | -- |
| 0 | 1 | 0AB | 02 | 30 | 44 | 67 |
| 0 | 1 | 034 | FD | EC | BA | 23 |
| 0 | 0 | -- | -- | -- | -- | -- |
| 0 | 1 | 1C6 | 00 | 11 | 22 | 33 |
| 0 | 1 | 045 | 67 | 78 | 89 | 9A |
| 0 | 1 | 001 | 70 | 00 | 44 | A6 |
| 0 | 0 | -- | -- | -- | -- | -- |

# Code Analysis

# Miss Rate

The cache is mostly invisible to programmers. But we can still make some optimizations by keeping it in mind!

The *miss rate* is the ratio of cache misses to total memory accesses. If we can analyze when cache misses occur, we may be able to make our code more cache-friendly and improve performance.

Average memory access time (AMAT) = (Hit Time) + (Miss Penalty) * (Miss Rate)

# What's the Miss Rate?

- First loop
  - Note array starts at beginning of a block
  - First access misses (cold cache)
    - Loads a[0] through a[3] into cache
    - a[1] through a[3] are hits

| Index | Valid | Block Offset 00 | 01 | 10 | 11 |
|---|---|---|---|---|---|
| 0 | 1 | a[0] | a[1] | a[2] | a[3] |
| 1 | 0 | ? | ? | ? | ? |
| 2 | 0 | ? | ? | ? | ? |
| ... | | | ... | | |

```
char val = 0;

for (int i = 0; i < 8; i++)
    val += a[i];

for (i = 0; i < 8; i++)
    val ^= a[i];
```

a is a `char` array of size 8.
Its address is 0x600000, and the cache starts cold.
Assume i and `val` are stored in registers.

Cache Parameters

C = 256 bytes  K = 4 bytes

# What's the Miss Rate?

- First loop (continued)
  - Next miss on a[4]
    - Loads a[4] through a[7] into cache
    - a[5] through a[7] are hits
  - 8 accesses, 2 misses

| Index | Valid | Block Offset 00 | 01 | 10 | 11 |
|---|---|---|---|---|---|
| 0 | 1 | a[0] | a[1] | a[2] | a[3] |
| 1 | 1 | a[4] | a[5] | a[6] | a[7] |
| 2 | 0 | ? | ? | ? | ? |
| ... | | | ... | | |

```
char val = 0;

for (int i = 0; i < 8; i++)
    val += a[i];

for (i = 0; i < 8; i++)
    val ^= a[i];
```

a is a `char` array of size 8.
Its address is 0x600000, and the cache starts cold.
Assume i and `val` are stored in registers.

Cache Parameters

C = 256 bytes  K = 4 bytes

# What's the Miss Rate?

- Second loop
  - Entire array is still in the cache!
  - **8 accesses, 0 misses**
- Overall miss rate
  - 16 accesses, 2 misses
  - 2 / 16 = **12.5%**

| Index | Valid | Block Offset 00 | 01 | 10 | 11 |
|---|---|---|---|---|---|
| 0 | 1 | a[0] | a[1] | a[2] | a[3] |
| 1 | 1 | a[4] | a[5] | a[6] | a[7] |
| 2 | 0 | ? | ? | ? | ? |
| ... | | | ... | | |

```
char val = 0;

for (int i = 0; i < 8; i++)
    val += a[i];

for (i = 0; i < 8; i++)
    val ^= a[i];
```

a is a `char` array of size 8.
Its address is 0x600000, and the cache starts cold.
Assume i and `val` are stored in registers.

Cache Parameters

C = 256 bytes  K = 4 bytes

# Code Analysis

- C = 1 KiB, K = 16B, E = 1 (direct mapped)
- array is a 64x64 2D array; i, j, and sum in registers

```
int sum = 0;
for (int i = 0; i < 64; i++)
  for (int j = 0; j < 64; j++)
    sum += array[i][j]; // assume &array = 0x600000
```

# Cache Simulator

# Cache Simulator!

Link:

https://courses.cs.washington.edu/courses/cse351/cachesim/

We haven't covered all of its features in class yet, but the cache simulator can be a helpful tool for reasoning through cache problems and mechanisms, particularly on homework and in lab 4.

# That's All, Folks!

Thanks for attending section! Feel free to stick around for a bit if you have quick questions (otherwise post on Ed or go to office hours).

See you all next week!