

Memory & Caches II

CSE 351 Autumn 2023

Instructor:

Justin Hsia

Teaching Assistants:

Afifah Kashif

Malak Zaki

Bhavik Soni

Naama Amiel

Cassandra Lam

Nayha Auradkar

Connie Lam

Nikolas McNamee

David Dai

Pedro Amarante

Dawit Hailu

Renee Ruan

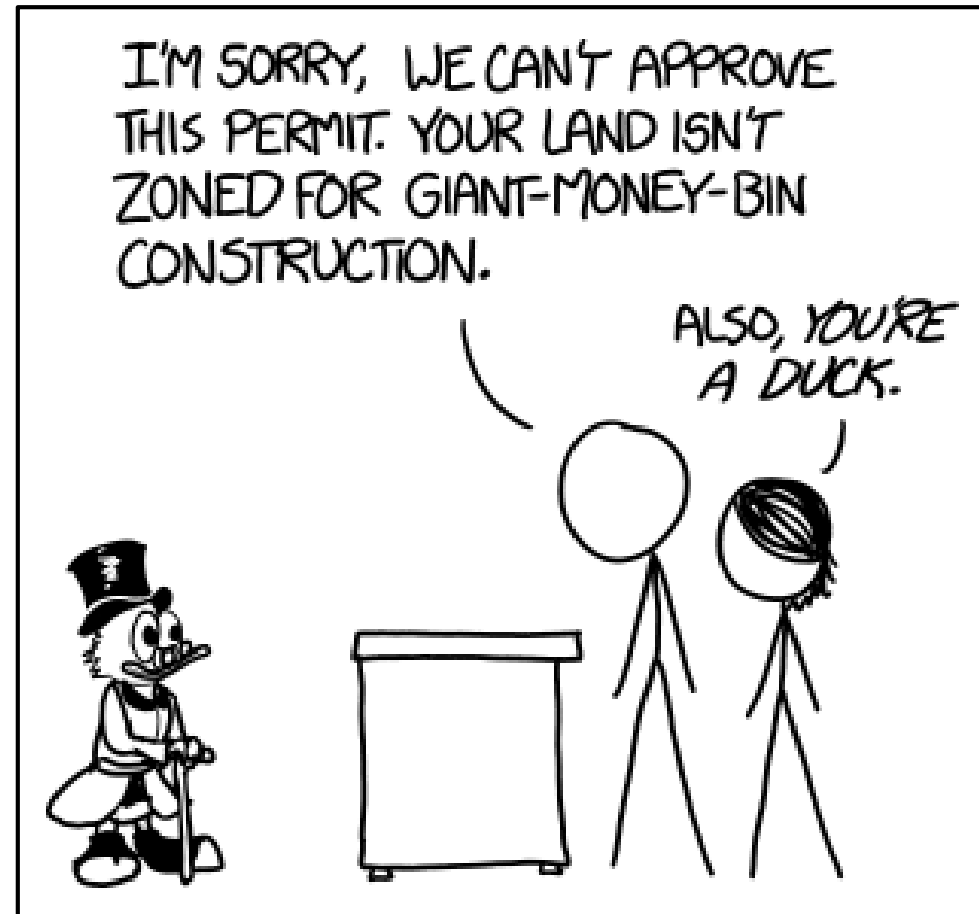
Ellis Haker

Simran Bagaria

Eyoel Gebre

Will Robertson

Joshua Tan



<https://what-if.xkcd.com/111/>

Relevant Course Information

- ❖ Mid-quarter Survey due Wednesday (11/8)
- ❖ hw16 due Wednesday (11/8)
- ❖ hw17 due *next* Wednesday (11/15)
 - Don't wait too long, this is a BIG hw (includes this lecture)
- ❖ Lab 3 due Friday (11/10)
 - Veteran's Day: no lecture, but some support hours (see Ed)
- ❖ Midterm grades will be released when we can
 - Regrade requests will be available afterward

New Lecture Format

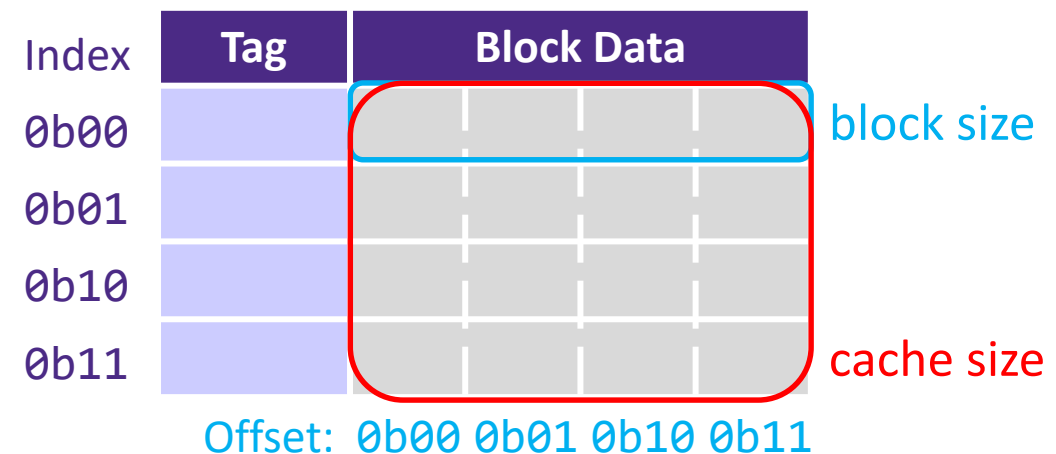
- ❖ Lesson Summary will be a little more detailed than before [5–7 min]
- ❖ Context section will “float” depending on content:
 - May be skipped for some lectures
 - May come after summary (as before) or may be held until end of lecture
 - Random call for share-out by section of the room for discussion questions
- ❖ Practice section will go problem-by-problem
 - Justin will introduce the problem
 - Then, limited group work time [3–8 min, depending on difficulty]
 - Random call for share-out by section of the room
 - Do not need to have finished the problem, just share out what work you did
- ❖ *The quicker we get responses, the quicker we can move on*

A detailed, colorful micrograph of a microchip die, showing a complex grid of circuitry and various colored regions. The text 'Caches II' is overlaid on the left side of the image.

Caches II

Lesson Summary (1/2)

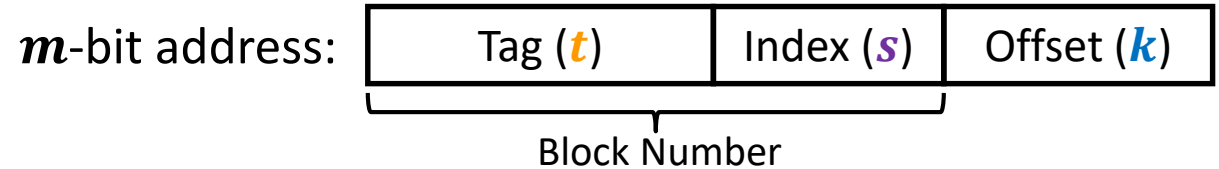
- ❖ Cache parameters define the cache geometry:
 - **Block size** is number of bytes per block
 - **Cache size** is number of bytes (or blocks) of data the cache can hold
- ❖ Finding a byte in the cache:
 - **Offset** refer to which byte in block
 - **Index** refers to which block in cache
- ❖ Example:
 - $K = 4 \text{ B}$, $C = 16 \text{ B} = 4 \text{ blocks}$



Lesson Summary (2/2)

- ❖ **Direct-mapped cache:** each block in cache is assigned a unique index
 - Uses hash function of (block number) mod (# of cache indices)
 - Deterministic placement of each block, with many blocks mapping into the same index
 - Tag bits stored in cache and used to distinguish between blocks that map to same index

- ❖ Accessing the cache:
(TIO address breakdown)



- 1) **Index** field tells you where to look in cache (width $s = \log_2 S$)
- 2) **Tag** field lets you check that data is the block you want (width $t = m - s - k$)
- 3) **Offset** field selects specified start byte within block (width $k = \log_2 K$)

Lesson Q&A

❖ Terminology:

- Cache parameters: block size (K), cache size (C), number of indices (S)
- Address fields: tag (t bits wide), index (s bits wide), block offset (k bits wide)
- Cache organization: direct-mapped cache

❖ Learning Objectives:

- Determine how memory addresses and data interact with the cache (*i.e.*, cache lookups, data movement).
- Analyze how changes to cache parameters [and policies] affect performance metrics such as AMAT.

❖ What lingering questions do you have from the lesson?

A detailed, colorful micrograph of a microchip die, showing a complex grid of circuitry and various colored regions. The text "Caches II – Practice" is overlaid in white with a drop shadow.

Caches II – Practice

Practice Questions (1/2)

- ❖ We have a direct-mapped cache with the following parameters:
 - Block size of 8 bytes
 - Cache size of 4 KiB

- ❖ How many blocks can the cache hold?

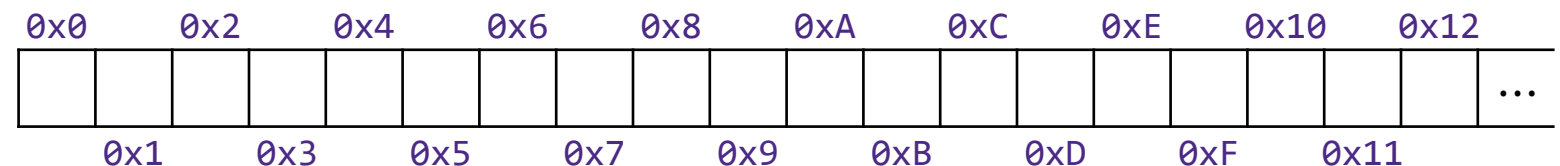
- ❖ How many bits wide is the block offset field?

- ❖ Which of the following addresses would fall under block number 3?
A. 0x3 **B. 0x1F** **C. 0x30** **D. 0x38**

Practice Questions (2/2)

- ❖ Based on the following behavior, which of the following block sizes is NOT possible for our cache?
 - Cache starts *empty*, also known as a *cold cache*
 - Access (addr: hit/miss) stream:
 - (0xE: miss), (0xF: hit), (0x10: miss)

- A. 4 bytes
- B. 8 bytes
- C. 16 bytes
- D. 32 bytes



Homework Question Setup

- ❖ Consider the following function, which computes the dot product of two vectors:

```
float dotprod(float x[8], float y[8]) {  
    float sum = 0.0;  
    for (int i = 0; i < 8; i++)  
        sum += x[i] * y[i];  
    return sum;  
}
```

- $\&x=0x0$, $\&y=0x20$. $\text{sizeof}(\text{float})=4$. Direct-mapped \$ with indices, each 16 B.
- **What does the cache look like?**

Homework Question Setup

- ❖ Consider the following function, which computes the dot product of two vectors:

```
float dotprod(float x[8], float y[8]) {  
    float sum = 0.0;  
    for (int i = 0; i < 8; i++)  
        sum += x[i] * y[i];  
    return sum;  
}
```

- $\&x=0x0$, $\&y=0x20$. $\text{sizeof}(\text{float})=4$. Direct-mapped \$ with indices, each 16 B.
- **List out the first 8 memory accesses in terms of x and y, then translate those to addresses:**

Group Work Time

- ❖ During this time, you are encouraged to work on the following:
 - 1) If desired, continue your discussion
 - 2) Work on the homework problems
 - 3) Work on the current lab

- ❖ Resources:
 - You can revisit the lesson material
 - Work together in groups and help each other out
 - Course staff will circle around to provide support