

Conceptual Questions: Why do we cache?

What are temporal and spatial locality? Give high level examples in software of when these occur.

Break up an address:

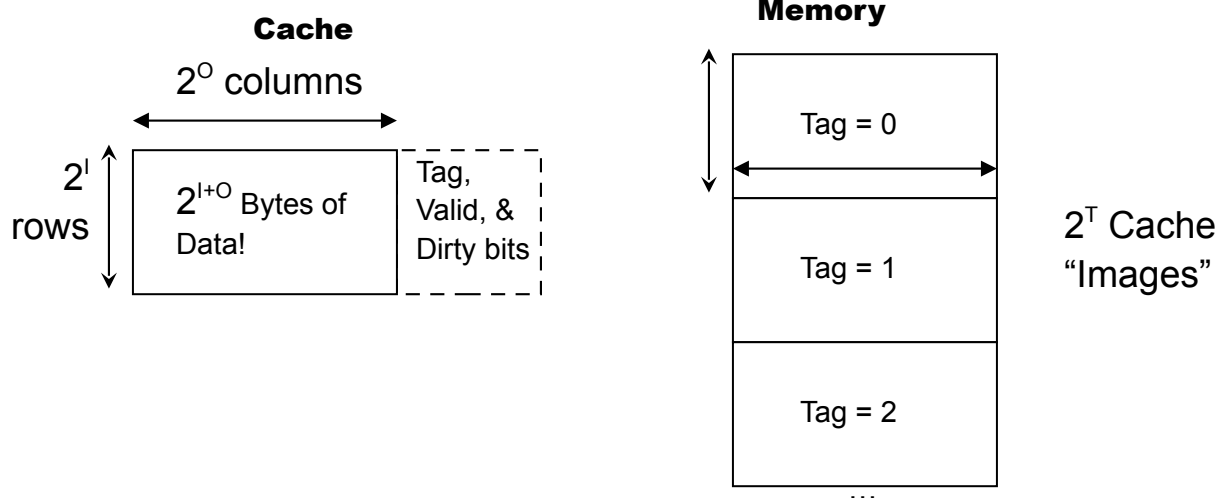
Tag	Index	Offset
-----	-------	--------

Offset: “column index”, Indexes into a block. (O bits)

Index: “row index,” Indexes blocks in the cache. (I bits)

Tag: Where from memory did the block come from? (T bits)

Segmenting the address into TIO implies a geometrical structure (and size) on our cache. Draw memory with that same geometry!

**Cache Vocab:**

Cache hit – found the right thing in the cache!

Cache miss – Nothing in the cache block we checked, so read from memory and write to cache!

Cache miss, block replacement – We found a block, but it had the wrong tag!

1) Fill in the table assuming a direct mapped cache. (B = byte.)

Address Bits	Cache Size	Block Size	Tag Bits	Index Bits	Offset Bits	Bits per Row
16	4KB	4B				
16	16KB	8B				
32	8KB	8B				
32	32KB	16B				
32	64KB		16	12	4	
32	512KB				5	
64		64B		14		
64	2048KB			14		1068

2) Assume 16 words of memory and an 8 word direct-mapped cache with 2-word blocks (that starts empty). Classify each of the following WORD memory accesses as hit (H), miss (M), or miss with replacement (R).

- | | |
|------|-------|
| a. 4 | e. 1 |
| b. 5 | f. 10 |
| c. 2 | g. 7 |
| d. 6 | h. 2 |

3) You know you have 1 MiB of memory (maxed out for processor address size) and a 16 KiB cache (data size only, not counting extra bits) with 1 KiB blocks.

```
#define NUM_INTS 8192
int A[NUM_INTS]; // lives at 0x100000
int i, total = 0;
for (i = 0; i < NUM_INTS; i += 128) A[i] = i; // Line 1
for (i = 0; i < NUM_INTS; i += 128) total += A[i]; // Line 2
```

- What is the T:I:O breakup for the cache (assuming byte addressing)?
- Calculate the hit percentage for the cache for the line marked “Line 1”.
- Calculate the hit percentage for the cache for the line marked “Line 2”.

AMAT:

AMAT stands for Average Memory Access Time. It refers to the time necessary to perform a memory access on average. It does NOT refer to the time necessary to execute an instruction which accesses memory. Those of you in CS70 may find it helpful to think of AMAT as $E(\tau)$, where τ is a random variable whose value is the time necessary to perform a given memory access. The AMAT of a simple system with only a single level of cache may be calculated as

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

This formula can be extended to more complicated memory hierarchies by replacing miss penalty with the AMAT for the next level in the memory hierarchy. That is

$$\text{AMAT}_i = \text{hit time}_i + \text{miss rate}_i \times \text{AMAT}_{i+1}$$

where i refers to the level of the memory hierarchy.

1. Suppose you have a system with the following properties:

- L1\$ hits in 1 cycle (local hit rate 75%)
- L2\$ hits in 20 cycles (local hit rate 80%)
- L3\$ hits in 100 cycles (local hit rate 90%)
- Main memory hits in 1000 cycles (always hits)

Calculate the AMAT.

2. Suppose you have a system with the following properties:

- L1\$ hits in 1 cycle (local miss rate 25%)
- L2\$ hits in 10 cycles (local miss rate 40%)
- L3\$ hits in 50 cycles (global miss rate 6%)
- Main memory hits in 100 cycles (always hits)

Calculate the AMAT.