

# CS110 sp25 HW5

Due: May 7

Complete this homework either by writing neatly by hand or using [typst](#). You can find the `.typ` file on Piazza.

## 1 TRUE OR FALSE

Fill in your answer (T or F) in the table below.

1. When the same address is accessed multiple times using a cache, it primarily benefits from spatial locality.
2. If a cache system changes from direct-mapped to N-way set-associative ( $N > 1$ ), the number of **index** bits in the address breakdown increases.
3. Higher cache associativity decreases hit time, miss rate, and miss penalty simultaneously.
4. Assume a direct-mapped cache with 8-byte blocks, 2 sets, using LRU replacement, and initially empty. Given 8-bit addresses, accessing addresses from `0x00` to `0xFF` sequentially will result in no cache hits.

1	2	3	4
F	F	F	F

## 2 SET-ASSOCIATIVE CACHES

Assume a 12-bit address space. All cache lines are shown in the table below.

Set index	Tag 1	Tag 2
0		
1		
2		
3		

1. Assume loading a 16-byte struct at address `0x00F` requires only 1 cache lookup, but loading one at `0x011` requires 2 lookups. Fill in the cache system parameters in the table below.

Cache block size	32 bytes
Total Capacity	256 bytes
Level of set associativity	2
# of cache blocks	8
# of sets	4
Address Breakdown (Tag   Index   Offset)	[11:7]   [6:5]   [4:0]

2. For the sequence of addresses below, indicate whether each access results in a hit, miss, or replacement. Assume the cache is initially empty.

Address	Cache Access Result (Hit/Miss/Replacement)
0x3A8	miss
0x1A6	miss
0x04C	miss
0x5AD	replacement
0x3B9	replacement
0x44F	miss
0x1B3	replacement
0x055	hit
0x241	replacement
0x45E	replacement
0x1A1	hit
0x1A5	hit
0x642	replacement
0x444	hit

3. Calculate the cache hit rate for the sequence of memory accesses above. How can you improve the cache hit rate without increasing the cache **capacity**?

**Solution:** 4 out of 14 accesses, 28.57% hit rate. Add set-association to the cache

### 3 AMAT

Consider a 3-level cache system with the following parameters, where the CPU runs at 4GHz:

Cache Level	Hit Time	Local Miss Rate
L1	2 cycles	8%
L2	12 cycles	35%
L3	44 cycles	10%
Memory	80ns	-

1. Calculate the global miss rate for the 3-level cache.

**Solution:**

$$8\% \times 35\% \times 10\% = 0.28\%$$

2. Calculate the average memory access time (AMAT) for the CPU.

**Solution:**

$$\frac{1}{4,000,000,000} = 0.00000000025\text{s} = 0.25\text{ns}$$

$$\text{AMAT} = 2 + 8\% \times \left( 12 + 35\% \times \left( 44 + 10\% \times \frac{80}{0.25} \right) \right) = 5.088 \text{ cycles} = 1.272 \text{ ns}$$

3. Now, suppose the L1 cache is changed from 2-way set-associative to fully associative. This change reduces the L1 local miss rate to **6%** but increases the L1 hit time by **60%**. Calculate the new AMAT.

**Solution:**

$$\text{AMAT} = 3.2 + 6\% \times \left( 12 + 35\% \times \left( 44 + 10\% \times \frac{80}{0.25} \right) \right) = 5.516 \text{ cycles} = 1.379 \text{ ns}$$

### 4 CODE ANALYSIS

Consider the following code:

```

struct body {
    float x, y, z, r;
};

struct body bodies[64];

// check whether two physics bodies overlap in 3D space
bool is_collide(struct body a, struct body b);

int check_collision() {
    int count = 0;
    for (int i = 0; i < 64; i++) {
        for (int j = i + 1; j < 64; j++) {
            if (is_collide(bodies[i], bodies[j])) {
                count++;
            }
        }
    }
    return count;
}

```

Note:

- Assume the cache parameters are: 128 bytes capacity, 32 bytes per block, 2-way set associative.
- Elements in the `bodies` array are aligned to the cache lines.
- `sizeof(float) == 4`
- You can ignore what `is_collide` exactly does and assume each body structure is loaded only **once** within the inner loop iteration (i.e., `bodies[i]` and `bodies[j]` are loaded into registers).
- The variables `i`, `j` and `count` are stored in registers.
- Instruction cache is not considered.
- Assume the cache is initially empty.

1. Calculate the hit rate for the above code.

**Solution:**

$$\# \text{ of access} = 64 \times 63 = 4032$$

For  $i \in [0, 54]$  the count of cache miss has such pattern:

$$\# \text{ of miss} = \begin{cases} 32 - \frac{i}{2} & \text{if } i \text{ is even} \\ 32 - \frac{i+1}{2} & \text{if } i \text{ is odd} \end{cases}$$

Specifically, when  $i \in [55, 56]$ ,

$$\# \text{ of miss} = \begin{cases} 2 & \text{if } i = 55 \\ 1 & \text{if } i = 56 \end{cases}$$

And when  $i > 56$ , no miss occurs.

So the hit rate is:

$$\text{hit rate} = 1 - \text{miss rate} = \frac{3025}{4032} \approx 75.02\%$$

You can also get the hit rate by simulation.

2. How can the hit rate be improved by modifying only the code (without changing the cache configuration)? Briefly explain your solution.

**Solution:** Apply Cache Blocking to the code.