

CS61c Summer 2014 Discussion 8 – More Caches & Review

1 Cache Question (From Fall Midterm)

Take a look at the following C function `sum_iter` run on a 32-bit MIPS machine. On this system, these structs are aligned to two-word boundaries since `sizeof(struct Node) = 8`. Assume the total space taken up by the linked list is greater than (and a multiple of) the cache size.

```
struct Node {
    int n;
    struct Node *next;
};
int sum_iter(struct Node *head) {
    int sum = 0;
    while(head != NULL) {
        sum += head -> n;    // load from head+0
        head = head -> next; // load from head+4
    }
    return sum;
}
```

Given a direct-mapped data cache with this T:I:O configuration: 12:13:7

- How many words are in a block? $2^7 B = 32$ words
- How many bytes of data does this cache hold? $2^{13+7} B = 1$ MiB
- What is the lowest possible cache hit rate for the while loop in `sum_iter`? 50%, reasonably assuming `sum` and `head` are placed in registers. `head+4` is always cached when `head+0` is loaded, due to the alignment of the struct.
- What is the highest possible cache hit rate for the while loop in `sum_iter`? $\frac{32-1}{32} = 96.875\%$
- To achieve this maximum hit rate, we obviously could have every `Node` next to every other node, like an array. However, that's too strict a constraint—we can still achieve this hit rate if that's not the case. What is the loosest constraint for how the `Nodes` are distributed in memory to get the best hit rate? Nodes are laid out in memory such that when a block is kicked out, the other 15 nodes in that block will have been accessed.

2 Cache Performance

1. Suppose our processor has separate L1 instruction cache and data cache. Our ideal base CPI is 2 clock cycles. Memory accesses take 100 cycles. Our I\$ miss rate is 3% while our D\$ miss rate is 10%. 40% of our instructions are loads or stores. What is our processor's actual CPI? $\text{Actual CPI} = \text{Base CPI} + \text{Inst L1\$ Miss Cycles} + \text{Data L1\$ Miss Cycles} = 2 + 1 * .03 * 100 + .4 * .1 * 100 = 7$ cycles
2. To improve the performance of this processor, we add a unified L2 cache between the L1 caches and memory. Our L2 cache has a hit time of 10 cycles and a global miss rate of 2%. What is the new actual CPI? $\text{Actual CPI} = \text{Base CPI} + \text{Inst L1\$ Miss Cycles} + \text{Data L1\$ Miss Cycles} + \text{Inst L2\$ Miss Cycles} + \text{Data L2\$ Miss Cycles} = 2 + (1 * .03 * 10) + (.4 * .1 * 10) + (1 * .02 * 100) + (.4 * .02 * 100) = 5.5$ cycles
3. Suppose we have an unrelated processor with a two-level cache. The L1\$ has a hit time of 1 cycle and a miss rate of 10%. The L2\$ has a hit time of 3 cycles and a miss penalty of 100 cycles. If we need our AMAT to be ≤ 2 cycles, what should the L2\$ miss rate be? $2 = 1 + 10\% * (3 + x * 100)$, giving a 7% miss rate

3 Cache Question (From Su98 Final)

- A 32 KiB cache has a 16 B block size and is 4-way set-associative. What is the T:I:O breakdown? **19:9:4**
- In a 2-way set-associative cache, three addresses A, B, and C all have the same index but distinct tags. What is a minimum sequence of accesses which, if repeated, will maximize the miss rate in the cache if it uses the LRU replacement policy? **A, B, C, ...**
- If the above sequence is repeated for a long period of time, what will the miss rate be if the cache uses an LRU replacement policy? **100%**
- If the hit time is 1 cycle, and the miss penalty is 3 cycles, what will be the AMAT for the LRU replacement policy using the above sequence? **AMAT = 1 + 100% * 3 = 4 cycles**
- If the sequence is repeated for a long period of time, will the miss rate be improved if “random” is used as the replacement policy? **Miss rate will be improved by random**

4 Review: MIPS Mystery Question (From Fa10 Midterm)

This assembly function takes no arguments and returns some value. Write this answer as binary number. What does it mean? **001001 11101 11101 0000 0000 0000 0100 ... the binary encoding of `addiu $sp, $sp, 4`.**

Address	Instruction
0x08001000	mystery: <code>addiu \$sp, \$sp, -4</code>
0x08001004	<code>sw \$ra, 0(\$sp)</code>
0x08001008	<code>addiu \$v0, \$zero, 0</code>
0x0800100c	<code>jal inner</code>
0x08001010	<code>lw \$ra, 0(\$sp)</code>
0x08001014	<code>addiu \$sp, \$sp, 4</code>
0x08001018	<code>jr \$ra</code>
0x0800101c	inner: <code>lw \$v0, 4(\$ra)</code>
0x08001020	<code>jr \$ra</code>

5 Review: Programming in C Question (From Sp14 Midterm)

A colleague of yours has implemented some homebrew C99 string manipulation functions, while steadfastly refusing to use any standard libraries, but they might be buggy! Your job is to find the incorrect lines and give correct replacement lines.

```

/** Converts the string S to lowercase */
void string_to_lowercase(char *s) {
    for(char c = *s; c != '\0'; s++) { // for(char c = *s; (c = *s) != '\0'; s++) {
        if(c >= 'A' && c <= 'Z') {
            s += 'a' - 'A';           // *s += 'a' - 'A';
        }
    }
}

/** Returns the number of bytes in S before, but not counting, the null terminator. */
size_t string_length(char *s) {
    char *s2 = s;
    while(*s2++);
    return s2 - s - 1;
}

/** Return the number of odd numbers in a number array */
uint32_t number_odds(uint32_t *numbers, uint32_t size) {
    uint32_t odds = 0;
    for (uint32_t i = 0; i < size; i++)
        odds += *numbers+i && 1;      // odds += *(numbers+i) & 1
    return odds;
}

```