

CS 433 Homework 4

Assigned on 10/17/2017
Due in class on 11/7/2017

Instructions:

1. Please write your name and NetID clearly on the first page.
2. Refer to the course fact sheet for policies on collaboration.
3. Due **IN CLASS** on 11/7/2017.

Problem 1 [7 Points]

Consider a 4MB 8-way set-associative write-back cache with 64 byte block size for byte-addressable memory with 32-bit address. Assume a random replacement policy and a single core system.

Part A [2 points]

Which bits of the address counting from 0 as the LSB are used for the cache index?

Part B [2 points]

Which bits of the address are used for the cache tag?

Part C [3 points]

How many bits of total storage does this cache need besides the 4MB for data? Remember to include any state bits needed.

Problem 2 [16 points]

Consider a processor where each instruction takes, on average, 2 cycles and there are 1.5 references to memory per instruction. A program with 100,000 instructions is executed on this machine using a split cache of 32KB, obtaining a 95% hit rate, 2 ns hit time and an 18 ns miss penalty. Then, the same program is executed using a 64KB cache, resulting in a hit rate of 97%, a hit time of 3 ns and the same miss penalty that in the previous case. The cycle time of the processor is adjusted to match the cache hit latency.

Part A [1 point]

Explain why the larger cache has higher hit rate.

Part B [1 points]

Explain why the small cache has smaller access time (hit time).

Part C [4 points]

Calculate the AMAT for both cases. Which cache is better from this point of view?

Part D [4 points]

Calculate the CPI for both cases. Which cache is better from this point of view?

Part E [6 points]

Calculate the execution time for both cases. Which cache is better? Please find the speedup.

Problem 3 [20 points]

Consider the following piece of code:

```
register int i, j;                /* times i, j are in the processor registers */
register float sum1, sum2;
float a[64][64], b[64][64];

for ( i = 0; i < 64; i++ ) {     /* (1) */
    for ( j = 0; j < 64; j++ ) { /* (2) */
        sum1 += a[i][j];         /* (3) */
    }
    for ( j = 0; j < 32; j++ ){   /* (4) */
        sum2 += b[i][2*j];       /* (5) */
    }
}
```

Assume the following:

- There is a perfect instruction cache (i.e., do not worry about the time for any instruction accesses).
- Both int and float are 4 bytes.
- Assume that only the accesses to the array locations $a[i][j]$ and $b[i][2*j]$ generate loads to the data cache. The rest of the variables are all allocated in registers.
- Assume a fully associative, LRU data cache with 32 lines, where each line has 16 bytes.
- Initially, the data cache is empty.
- The arrays a and b are stored in row major form.
- To keep things simple, we will assume that statements in the above code are executed sequentially. The time to execute lines (1), (2), and (4) is 4 cycles for each invocation. Lines (3) and (5) take 10 cycles to execute and an additional 40 cycles to wait for the data if there is a data cache miss.
- There is a data prefetch instruction with the format `prefetch(array[index])`. This prefetches the entire block containing the word `array[index]` into the data cache. It takes 1 cycle for the processor to execute this instruction and send it to the data cache. The processor can then go ahead and execute subsequent instructions. If the prefetched data is not in the cache, it takes 40 cycles for the data to get loaded into the cache.
- Assume that the arrays a and b both start at cache line boundaries.

Part A [4 points]

How many cycles does the above code fragment take to execute if we do NOT use prefetching?

Part B [4 points]

Consider inserting prefetch instructions for the two inner loops for the arrays a and b respectively. Explain why we may want to unroll the loops to insert prefetches. What is the minimum number of times you would need to unroll for each of the two loops for this purpose?

Part C [8 points]

Unroll the inner loops for the number of times identified in part b, and insert the minimum number of prefetch instructions to minimize execution time. The technique to insert prefetches is analogous to software pipelining. You do not need to worry about startup and cleanup code and do not introduce any new loops.

Part D [4 points]

How many cycles does the code in part (c) take to execute? Calculate the average speedup over the code without prefetching. Assume prefetches are not present in the startup code. Extra time needed by prefetches executing beyond the end of the loop execution time should not be counted.

Problem 4 [12 points]

Consider a virtual memory system where a TLB access takes 2 ns and there is a single level of a set-associative, write-back data cache with the following parameters:

- indexing the cache to access the data portion takes 6 ns
- indexing the tag array of the data cache takes 4 ns
- tag comparisons take 1.5 ns
- multiplexing the output data takes 1 ns

Assume these are the only parts that affect the cache access time. For each of the following configurations, calculate the amount of time it takes to get data from the cache on a hit, including any necessary TLB access time. Please explain your answers for full credit.

Part A [3 points]

The page size is 4KB and the data cache is 64 KB, 4-way set associative, with block size of 16 bytes. The cache is physically-indexed and physically-tagged.

Part B [3 points]

Same as the part A except that the cache is virtually-indexed and virtually-tagged.

Part C [3 points]

Same as the part A except that the cache is virtually-indexed and physically-tagged.

Part D [3 points]

Same as the part A except that the cache is physically-indexed and virtually-tagged.

Problem 5: Graduate students only [12 points]

You are building a computer system around a processor with in-order execution that runs at 1 GHz and has a CPI of 1, excluding memory accesses. The only instructions that read or write data from/to memory are loads (20% of total instructions) and stores (5% of total instructions).

The memory system for this computer has a split L1 cache. Both the I-cache and the D-cache are direct-mapped and hold 32 KB each. The I-cache has a 2% miss rate and 64 byte blocks, and the D-cache is a write-through, no-write-allocate cache with a 5% miss rate and 64 byte blocks. The hit time for both the I-cache and the D-cache is 1 ns.

The L1 cache has a write buffer. 90% of writes to L1 find a free entry in the write buffer immediately. The other 10% of the writes have to wait until an entry frees up in the write buffer. An entry is held while the write buffer initiates a request to L2 and waits for L2. The processor is stalled on a write until a free write buffer entry is available.

The L2 cache is a unified write-back, write-allocate cache with a total size of 512 KB and a block size of 64 bytes. The hit time of the L2 cache is 12 ns for both read hits and write hits. L2 write for a miss takes 12 ns (after the allocate is done). Tag comparison for hit/miss is included in the 12 ns in all cases, do not add hit time to miss time on a miss. The local hit rate of the L2 cache is 80%. Also, 50% of all L2 cache blocks replaced are dirty.

The 64-bit wide main memory has an access latency of 20 ns (including the time for the request to reach from the L2 cache to the main memory), after which any number of bus words may be transferred at the rate of one bus word (64 bit) per bus cycle on the 64-bit wide 100 MHz main memory bus.

Assume inclusion between the L1 and L2 caches, and assume there is no write-back buffer at the L2 cache. Assume a write-back takes the same amount of time as a read of the same size from memory.

While calculating any time values (such as hit time, miss penalty, AMAT), please use ns (nanoseconds) as the unit of time. For miss rates below, give the local miss rate for that cache. By miss penalty_{L2}, we mean the time from the miss request issued by the L2 cache up to the time the data comes back to the L2 cache from main memory.

Part A [7 points]

Compute the AMAT (average memory access time) for instruction accesses.

1. Give the values of the following terms for instruction accesses. L1 hit time, L1 miss rate, L2 hit time, and L2 miss rate. [1 point]
2. Give the formula for calculating L2 miss penalty, and compute the value of L2 miss penalty. Don't forget to include the time to write back a dirty block. [4 points]
3. Give the formula for calculating the AMAT for this system using the five terms whose values you computed above and any other values you need. [1 point]
4. Plug in the values into the AMAT formula above, and compute a numerical value for AMAT for instruction accesses. [1 point]

Part B [2 points]

Compute the AMAT for data reads.

1. Give the value of L1 miss rate for data reads. [1 point]
2. Calculate the value of the AMAT for data reads. [1 point]

Part C [3 points]

Compute the AMAT for data writes. Assume the miss penalty for a data write is the same as computed previously for a data read.

1. Give the value of time for a write buffer entry being written to the L2 cache. [2 points]
2. Calculate the value of the AMAT for data writes using the above information, and any other values that you need. Only include the time that the processor will be stalled. Hint: There are two cases to be considered here depending upon whether the write buffer is full or not. [1 point]