

CS 433: Computer Architecture – Fall 2022

Homework 5

Total Points: Undergraduates (44 points), Graduates (52 points)

Undergraduate students should only solve the first 4 problems.

Graduate students should solve all problems.

Due Date: November 1, 2022 at 10:00 pm CT

(See course information slides for more details)

Directions:

- **All students must write and sign the following statement at the end of their homework submission. "I have read the honor code for this class in the course information handout and have done this homework in conformance with that code. I understand fully the penalty for violating the honor code policies for this class." No credit will be given for a submission that does not contain this signed statement.**
- **On top of the first page of your homework solution, please write your name and NETID, your partner's name and NETID, and whether you are an undergrad or grad student.**
- **Please show all work that you used to arrive at your answer. Answers without justification will not receive credit. Errors in numerical calculations will not be penalized. Cascading errors will usually be penalized only once.**
- **See course information slides for more details.**

Problem 1 [5 Points]

A 4 entry victim cache for a 4KB direct mapped cache removes 80% of the conflict misses in a program. Without the victim cache, the miss rate is 0.064 (6.4%) and 67% of these misses are conflict misses. What is the percentage improvement in the AMAT (average memory access time) due to the victim cache?

Assume a hit in the main (4KB) cache takes 1 cycle. For a miss in the main cache that hits in the victim cache, assume an additional penalty of 1 cycle to access the victim cache. For a miss in both the main and victim caches, assume a further penalty of 48 cycles to get the data from memory. Assume a simple, single-issue, 5-stage pipeline, in-order processor that blocks on every read and write until it completes.

Solution:

AMAT = Hit Time + Miss Rate \times Miss Penalty

Without victim cache: AMAT = $1 + 0.064 \times 48 = 4.072$ cycles

With victim cache: AMAT = $1 + 0.064 \times [(0.67 \times 0.80 \times 1) + ((1 - 0.67 \times 0.80) \times 49)] = 2.489$

Improvement: $\frac{AMAT_{orig} - AMAT_{victim}}{AMAT_{orig}} = \frac{4.072 - 2.489}{4.072} \approx 0.389 = 38.9\%$

Grading:

1 point for the original AMAT

3 points for victim cache AMAT:

1 point for determining the correct rate for victim cache hits and misses

1 point for assigning the correct penalty to victim cache hits

1 point for assigning the correct penalty to victim cache misses

1 point for percent improvement

Some students used $\frac{AMAT_{orig}}{AMAT_{victim}} - 1 = \frac{4.072}{2.489} - 1 \approx 0.636 = 63.6\%$ as the percentage improvement.

For this homework, we did not take off points for this. (This would be the performance improvement i.e. speedup, assuming that performance is $1/AMAT$.)

Problem 2 [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 all instructions) and stores (5% of all instructions).

The memory system for this computer has a split L1 cache. Both the I-cache and the D-cache 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. 95% of writes to L1 find a free entry in the write buffer immediately. The other 5% of the writes have to wait until an entry frees up in the write buffer (assume that such writes arrive just as the write buffer initiates a request to L2 to free up its entry and the entry is not freed up until the L2 is done with the request). 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 15ns for both read hits and write hits. Tag comparison for hit/miss is included in the 15ns 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 20ns (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 an L2 read miss of the same size.

Assume all caches in the system are blocking; i.e., they can handle only one memory access (load, store, or writeback) at a time. When calculating the miss penalty for a load or store for a writeback cache, the time for any needed writebacks should be included in the miss penalty.

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]

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

i. Give the values of the following terms for instruction accesses: **hit time_{L1}**, **miss rate_{L1}**, **hit time_{L2}**, **miss rate_{L2}**. [1 point]

Solution: hit time_{L1} = 1 cycle = 1 ns miss rate_{L1} = 0.02
hit time_{L2} = 15 ns miss rate_{L2} = 1 - 0.8 = 0.2

Grading: 1 point for giving the correct values of all 4 terms, otherwise no points.
This part is just meant to get everyone started on the right track.

ii. Give the formula for calculating **miss penalty_{L2}**, and compute the value of miss penalty L2. [4 points]

Solution: miss penalty_{L2} = memory access latency + time to transfer one L2 cache block
transfer rate of memory bus
 $= \frac{64 \text{ bit}}{\text{bus cycle}} = \frac{64 \text{ bit}}{10 \text{ ns}} = \frac{8 \text{ byte}}{10 \text{ ns}} = 0.8 \text{ byte/ns} = 800 \text{ MB/s}$
time to transfer one L2 cache block = $\frac{64 \text{ byte}}{0.8 \text{ byte/ns}} = 80 \text{ ns}$
miss penalty_{clean L2} = 20 ns + 80 ns = 100 ns
50% of replaced blocks are dirty and must be written back; this is another 100 ns
miss penalty_{L2} = 100 ns + 0.5 × 100 ns = 150 ns

Grading:
1 point for the correct formula for miss penalty_{L2}.
1 point for correctly setting up the time to transfer one block.
1 point for the correct value of miss penalty_{clean L2}.
1 point for noting that 50% of the time the replaced block will need to be written back and for correctly setting up the value of miss penalty_{L2} taking this into account.
No points to be taken off for calculation errors.

iii. 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]

Solution: AMAT = hit time_{L1} + miss rate_{L1} × (hit time_{L2} + miss rate_{L2} × miss penalty_{L2})
Grading: 1 point for a completely correct formula, otherwise no points.

iv. Plug in the values into the AMAT formula above, and compute a numerical value for AMAT for instruction accesses. [1 point]

Solution: AMAT = 1 ns + 0.02 × (15 ns + 0.2 × 150 ns) = 1.9 ns
Grading: 1 point for setting up the correct values in the AMAT formula.
No points to be taken off for calculation errors.

Part B [2 points]

Computing the AMAT for *data reads*.

i. Give the value of **miss rate**_{L1} for data reads. [1 point]

Solution: $\text{miss rate}_{L1} = 0.05$

Grading: 1 point for giving the correct value of miss rate_{L1} .

ii. Calculate the value of the AMAT for data reads using the above value, and other values you need. [1 point]

Solution: $\text{AMAT} = \text{hit time}_{L1} + \text{miss rate}_{L1} \times (\text{hit time}_{L2} + \text{miss rate}_{L2} \times \text{miss penalty}_{L2})$

$$\text{AMAT} = 1 + 0.05 \times (15 + 0.2 \times 150) = 3.25 \text{ ns}$$

Grading: 1 point for setting up the correct AMAT formula for data reads.
No points to be taken off for calculation errors.

Part C [3 points]

Computing the AMAT for *data writes*. Assume **miss penalty**_{L2} for a data write is the same as that computed previously for a data read.

i. Give the value of **write time**_{L2Buff}, the time for a write buffer entry to be written to the L2 cache. [2 points]

Solution: $\text{write time}_{L2\text{Buff}} = \text{hit time}_{L2} + \text{miss rate}_{L2} \times \text{miss penalty}_{L2}$

$$\text{write time}_{L2\text{Buff}} = 15 + 0.2 \times 150 = 45 \text{ ns}$$

Grading: 1 point for setting up the correct formula for $\text{write time}_{L2\text{Buff}}$.
1 point for setting up the correct values in the $\text{write time}_{L2\text{Buff}}$ formula.
No points to be taken off for calculation errors.

ii. 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]

Solution: $\text{AMAT} = \text{hit time}_{L1} + \text{miss rate}_{L1} \times \text{write time}_{L2\text{Buff}}$

$$\text{AMAT} = 1 + 0.05 \times 45 = 3.25 \text{ ns}$$

The two cases to consider are whether the write buffer has a free entry or not. When it does, it can be written to in 1 cycle. When it does not, it must incur the penalty for writing an entry to the L2 cache, which is $\text{write time}_{L2\text{Buff}}$.

Grading: 1 point for setting up the correct AMAT equation and inserting the correct values.
No points to be taken off for calculation errors.

Problem 3 [13 points]

Consider the following piece of code:

```
register int i, j; /* 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 of size 4 bytes.
- 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.
- 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.
- The arrays **a** and **b** are stored in row major form.
- The arrays **a** and **b** both start at cache line boundaries.

Part A [2 points]

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

Solution:

Each cache line has $\frac{\text{cache line size}}{\text{element size}} = \frac{16}{4} = 4$ elements, so every 4 accesses in line 3 will miss and every 2 accesses in line 5, for a total of $64 \times \left(\frac{64}{4} + \frac{32}{2}\right) = 64 \times (16 + 16) = 2048$ misses.

Line 1 executes 65 times:	$65 \times 4 = 260$
Line 2 executes 64×65 times:	$64 \times 65 \times 4 = 16,640$
Line 3 executes 64×64 times:	$64 \times 64 \times 10 = 40,960$
Line 3 misses $64 \times \frac{64}{4}$ times:	$64 \times \frac{64}{4} \times 40 = 40,960$
Line 4 executes 64×33 times:	$64 \times 33 \times 4 = 8,448$
Line 5 executes 64×32 times:	$64 \times 32 \times 10 = 20,480$
Line 5 misses $64 \times \frac{32}{2}$ times:	$64 \times \frac{32}{2} \times 40 = 40,960$
Total cycles:	168,708

Average cycles per outer loop iteration: $\frac{168,708}{64} = 2636.0625$

Grading: 1 point for correct cycles taken by lines 3 and 5.
1 point for correct cycles taken by lines 1, 2, and 4.

Part B [2 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?

Solution:

There is one miss every four iterations of the first loop, and every two iterations of the second loop. The latency of this miss covers the prefetch time. We only need to issue a prefetch instruction once per cache line accessed. If code size is not a problem, unrolling the loop is the most efficient way to do this, since it avoids branches that test for the correct iteration count.

We unroll the first loop 4 times and the second loop 2 times; the prefetch is for iteration $j + 4$. (For the second loop, the prefetch is not for the next iteration of the unrolled loop, but for the one after that.)

Grading: 1 point for unroll count of the first loop.
1 point for unroll count of the second loop.

Part C [4 points]

Unroll the inner loops for the number of times identified in part b, and insert the minimum number of software prefetches 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.

Solution:

```
register int i, j; /* 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 += 4) {                       /* 2 */
        prefetch(a[i][j+4])                            /* P1 */
        sum1 += a[i][j+0];                             /* 3a */
        sum1 += a[i][j+1];                             /* 3b */
        sum1 += a[i][j+2];                             /* 3c */
        sum1 += a[i][j+3];                             /* 3d */
    }
    for (j = 0; j < 32; j += 2) {                       /* 4 */
        prefetch(b[i][2*j+8])                          /* P2 */
        sum2 += b[i][2*j+0];                           /* 5a */
        sum2 += b[i][2*j+2];                           /* 5b */
    }
}
```

Grading: 1 point for correct indices, statements, and loop header for the first loop.
1 point for prefetch in the first loop.
1 point for correct indices, statements, and loop header for the second loop.
1 point for prefetch in the second loop.

Part D [2 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.

Solution:

Now the only misses are on the very first execution of line 3a—row major ordering means prefetching is effective even across outer iterations—and the first two executions of line 5a, since the prefetch is preparing for the $j + 4$ iteration. So there are 3 misses total.

Line 1 executes 65 times:	$65 \times 4 = 260$
Line 2 executes 64×17 times:	$64 \times 17 \times 4 = 4,352$
Line P1 executes 64×16 times:	$64 \times 16 \times 1 = 1,024$
Lines 3a–d execute 64×16 times:	$4 \times 64 \times 16 \times 10 = 40,960$
Line 3a <i>misses</i> 1 time:	$1 \times 40 = 40$
Line 4 executes 64×17 times:	$64 \times 17 \times 4 = 4,352$
Line P2 executes 64×16 times:	$64 \times 16 \times 1 = 1,024$
Lines 5a–b execute 64×16 times:	$2 \times 64 \times 16 \times 10 = 20,480$
Line 5a <i>misses</i> 2 times:	$2 \times 40 = 80$
Total cycles:	72,572

Average cycles per outer loop iteration: $\frac{72,572}{64} = 1133.9375$

Speedup over code without prefetching: $\frac{168,708}{72,572} \approx 2.32$

Grading: 1 point for calculating the correct number of misses.
1 point for the rest of the cycles.

Part E [3 points]

Is there another technique that can be used to achieve the same objective as loop unrolling in this example, but with fewer static instructions? Explain this technique and illustrate its use for the code in part (c).

Solution:

The simplest option is to issue excess prefetch requests. Costing only one cycle if the data has already been requested, that's probably cheaper than trying to use branches.

```
for (i = 0; i < 64; i++) { /* 1 */
    for (j = 0; j < 64; j++) { /* 2 */
        prefetch(a[i][j+4]) /* P1 */
        sum1 += a[i][j]; /* 3 */
    }
    for (j = 0; j < 32; j++) { /* 4 */
        prefetch(b[i][2*j+8]) /* P2 */
        sum2 += b[i][2*j]; /* 5 */
    }
}
```

Grading: 1.5 points for prefetching in the first loop.
1.5 points for prefetching in the second loop.

Alternate Solution: Code may also test $j\%4$ (resp. $j\%2$) to only issue the same number of prefetches as in the unrolled code.

```
for (i = 0; i < 64; i++) { /* 1 */
    for (j = 0; j < 64; j++) { /* 2 */
        if (j%4 == 0) /* P1a */
            prefetch(a[i][j+4]) /* P1b */
        sum1 += a[i][j]; /* 3 */
    }
    for (j = 0; j < 32; j++) { /* 4 */
        if (j%2 == 0) /* P2a */
            prefetch(b[i][2*j+8]) /* P2b */
        sum2 += b[i][2*j]; /* 5 */
    }
}
```

Grading: 1.5 points for inserting the correct if statement in the first loop
1.5 points for inserting the correct if statement in the second loop.

Problem 4 [14 points]

Way prediction allows an associative cache to provide the hit time of a direct-mapped cache. The MIPS R10000 processor used way prediction to achieve a different goal: reduce the cost of the chip package. The R10000 hardware includes an on-chip L1 cache, on-chip L2 tag comparison circuitry, and an on-chip L2 way prediction table. L2 tag information is brought on chip to detect an L2 hit or miss. The way prediction table contains 8K 1-bit entries, each corresponding to two L2 cache blocks. L2 cache storage is built external to the processor package, is 2-way associative, and may have one of several block sizes.

Part A [2 points]

How can way prediction reduce the number of pins needed on the R10000 package to read L2 tags and data, and what is the impact on performance compared to a package with a full complement of pins to interface to the L2 cache?

Solution: With way prediction, only one way is read at a time rather than reading and comparing both. The package only needs enough pins to read the tag and data from a single line in a cycle instead of two, plus one extra bit to select the way. (Assuming the cache is just simple memory). A cache access takes an extra cycle whenever the way prediction is incorrect, and on every cache miss, which will slow performance.

Grading: 2 points for observing that we have to access both ways simultaneously if we don't have way prediction.

Part B [2 points]

How could a 2-associative cache be implemented with the same smaller number of pins but without the way prediction table? What is the performance drawback?

Solution: Without way prediction, the processor will access the ways sequentially. This will incur a delay whenever the data is found in the second way, and prediction would have been accurate.

Grading: 1 point for suggesting sequential access.
1 point for describing the performance loss.

Part C [4 points]

Assume that the R10000 uses most-recently used way prediction. What are reasonable design choices for the cache state update(s) to make when the desired data is in the predicted way, the desired data is in the non-predicted way, and the desired data is not in the L2 cache? Please fill in your answers in the following table.

Cache Access Case	Cache State Change Way Prediction Entry
Desired data is in the predicted way	No change
Desired data is in the non-predicted way	Flip state (to the way for this access)
Desired data is not in the L2 cache	Set to location of new data, or, Flip State (state will be flipped if we overwrite least recently used way)

Grading: 2 points per entry.

Part D [2 points]

For a 1024 KB L2 cache with 64-byte blocks and 8-way set associativity, how would the prediction table be organized for this new size? Give your answer in the form of “X entries by Y bits per entry.”

Solution: $\frac{1024 \text{ KB}}{64 \text{ B/line}} = 16,384 \text{ lines}$, $\frac{16,384 \text{ lines}}{8 \text{ ways/set}} = 2,048 \text{ sets}$
For 8-way associativity, $\log_2 8 = 3 \text{ bits}$ are needed.
So need 2,048 (2K) entries by 3 bits per entry

Grading: 1 point for number of entries (sets).
1 point for entry width (bits).

Part E [2 points]

For an 8 MB L2 cache with 128-byte blocks and 2-way set associativity, what would the prediction table organization be? Again, give your answer as “X entries by Y bits per entry.”

Solution: $\frac{8 \text{ MB}}{128 \text{ B/line}} = 65,536 \text{ lines}$, $\frac{65,536 \text{ lines}}{2 \text{ ways/set}} = 32,768 \text{ sets}$
2-way associativity, $\log_2 2 = 1 \text{ bit}$ is needed.
So need 32,768 (32K) entries by 1 bit per entry.

Grading: 1 point for number of entries (sets).
1 point for entry width (bits).

Part F [2 points]

What is the difference in the way that the R10000 with only 8K way prediction table entries will support the cache in part d) versus the cache in part e)? Hint: Think about the similarity between a way prediction table and a branch prediction table.

Solution: An 8Kb way prediction table is enough to support the cache in part D directly, or maybe even to treat the 8Kb as a smaller number of more sophisticated predictors. For the cache in part E, several ways will have to share a predictor, using some map from 32K sets onto 8K predictors. For best results, the mapping should aim to ensure that memory that will be used at similar times will be mapped to different ways, for example simply by dropping some high bits.

Grading: 1 point for saying there are enough entries for part D.
1 point for suggesting ways to share predictors in part E.

NOTE: ONLY GRADUATE STUDENTS SHOULD SOLVE THIS PROBLEM

Problem 5 [8 points]

Consider a computer with an in-order CPU, and with a data cache block size of 64 bytes (16 words) and a 32-bit wide bus to the memory. The memory takes 10 cycles to supply the first word and 2 cycles per word to supply the rest of the block. The cache is non-blocking, and it can support any number of outstanding misses. The memory can service multiple requests simultaneously if required (techniques to achieve this will be discussed in class).

This cache and memory system implement a “**Requested Word First and Early Restart**” policy, and the bus delivers the block data in “**cyclic order**” starting with the requested word. Cyclic order means that if the requested word is the 5th in a block of size 16 words, then the order in which the words in the block are supplied is 5, 6, 7 ... 16, 1, 2, 3, 4.

Part A [3 points]

Consider the following code fragment, which operates on an integer array A which is block-aligned (that is A[0] is located at the start of a cache block in memory):

```
for (i = 11; i < 100; i += 16) {           /* 1 */
    A[i] *= 2;                             /* 2 */
}
```

Suppose that the cache is big enough so that there are only compulsory misses. Further, statement 1 takes 4 cycles to execute, and statement 2 takes 4 cycles to execute in addition to any miss latency. Assume no overlap in the execution of these statements. Initially, the array A is not present in the cache, so any initial accesses to A cause misses in the cache.

What is the running time of this loop with the “*Requested Word First and Early Restart*” policy?

Solution: A[i] is supplied after 10 cycles.

Line 1 executes 7 times: $7 \times 4 = 28$

Line 2 executes 6 times: $6 \times (10 + 4) = 84$

Total: 112 cycles

Grading: 1 point for line 1.
2 points for line 2.

Part B [3 points]

How many cycles would the above loop take to run in a system with just “*Early Restart*” (i.e. the block is fetched in normal order, but the program is started early at arrival of requested word).

Solution: $A[i]$ is the 11th word in the block (remember that C uses zero-based indexing)
So $A[i]$ is supplied after $10 + 2 \times 11 = 32$ cycles.

Line 1 executes 7 times: $7 \times 4 = 28$
Line 2 executes 6 times: $6 \times (32 + 4) = 216$
Total: 244 cycles

Grading: 2 points for data supply cycles.
1 point for statement execution time.

Part C [2 points]

How many cycles would the above loop take to run in a system with the base policy (i.e. normal fetch and restart)?

Solution: The entire block, including $A[i]$, is supplied after $10 + 2 \times 15 = 40$ cycles.

Line 1 executes 7 times: $7 \times 4 = 28$
Line 2 executes 6 times: $6 \times (40 + 4) = 264$
Total: 292 cycles

Grading: 1 point for data supply cycles.
1 point for statement execution time.