

263-0007-00: Advanced Systems Lab

Assignment 4: 120 points

Due Date: April 17th, 17:00

<https://acl.inf.ethz.ch/teaching/fastcode/2021/>

Questions: fastcode@lists.inf.ethz.ch

Academic integrity:

All homeworks in this course are single-student homeworks. The work must be all your own. Do not copy any parts of any of the homeworks from anyone including the web. Do not look at other students' code, papers, or exams. Do not make any parts of your homework available to anyone, and make sure no one can read your files. The university policies on academic integrity will be applied rigorously.

Submission instructions (read carefully):

- (Submission)
Homework is submitted through the Moodle system <https://moodle-app2.let.ethz.ch/course/view.php?id=14942>.
- (Late policy)
You have 3 late days, but can use at most 2 on one homework, meaning submit latest 48 hours after the due time. For example, submitting 1 hour late costs 1 late day. Note that each homework will be available for submission on the system 2 days after the deadline. However, if the accumulated time of the previous homework submissions exceeds 3 days, the homework will not count.
- (Formats)
If you use programs (such as MS-Word or Latex) to create your assignment, convert it to PDF and name it homework.pdf. When submitting more than one file, make sure you create a zip archive that contains all related files, and does not exceed 10 MB. Handwritten parts can be scanned and included.
- (Plots)
For plots/benchmarks, **provide (concise) necessary information for the experimental setup (e.g., compiler and flags) and always briefly discuss the plot and draw conclusions**. Follow (at least to a reasonable extent) the small guide to making plots from the lecture.
- (Neatness)
5 points in a homework are given for neatness.

Exercises

1. Associativity (20 pts)

Consider the following function, executed on a machine with a write-back/write-allocate cache with blocks of size 16 bytes, a total capacity of 128 bytes and with a LRU replacement policy. Arrays x , y and z are cache-aligned (first element goes into first cache block). Assume that memory accesses occur in exactly the order that they appear in the code. Thus, no optimizations are performed that reduce the memory accesses or reorder computations. The variables i and t remain in registers and do not cause cache misses.

```
1 void calculate1 (double* x, double* y, double* z) {
2     double t = 0.0;
3     for (int i = 1; i <= 10; i++) {
4         t += x[(4*i)% 12];
5         t += y[(3*i-2)% 8];
6         z[(4*i-1)% 8] = t;
7     }
8 }
```

- (a) Considering cache misses from both **reads and writes**, compute the following two things: i) the **miss/hit pattern** for x , y and z (something like x : MMHM... , y :MMM...); ii) the **operational intensity** (in flops/byte) of the above computation for the following cases. For the operational intensity assume empty caches and consider only reads but note: arrays that are only written are also read and this read should be included. Show your work.

- i. Miss/hit pattern and operational intensity when the cache is 2-way set associative.
 - ii. Miss/hit pattern and operational intensity when the cache is 4-way set associative.
- (b) Assuming a 2-way set associative cache (motivate your answers):
- i. What kind(s) of locality do the accesses to array x have?
 - ii. What kind(s) of locality do the accesses to array y have?
2. *Cache Mechanics (35 pts)* Consider the following code, executed on a machine with a write-back/write-allocate direct-mapped cache with blocks of size 32 bytes, a total capacity of 12KiB and with a LRU replacement policy. Assume that memory accesses occur in exactly the order that they appear. The variables i, j, k, t and n remain in registers and do not cause cache misses. Arrays A and B are cache-aligned (first element goes into first cache block). For this and the following exercises, assume a cold cache scenario. `sizeof(double) = sizeof(uint64_t) = 8`.

```

1 #define PADDING_SIZE 1
2 typedef struct {
3     double    v;
4     double    d[3];
5     uint64_t  u[3];
6     uint64_t  pad[PADDING_SIZE];
7 } struct_t;
8
9 void calculate1(struct_t *A, struct_t *B, int n) {
10     double t;
11     for (int i = 0; i < n; ++i) {
12         for (int j = 0; j < n; ++j) {
13             t = A[i*n + j].v;
14             for (int k = 0; k < 3; k++) {
15                 t += A[i*n + j].d[k];
16             }
17             for (int k = 0; k < 3; k++) {
18                 A[i*n + j].u[k] = 0;
19             }
20             t += B[j*n].v;
21             B[j*n].v = t;
22         }
23     }
24 }

```

Considering cache misses from both **reads and writes**, compute i) the **cache miss rate** and ii) the **operational intensity** (in flops/byte) of the above computation for the following cases. For the operational intensity assume only reads, i.e., data movement from main memory to cache. Show enough details so we can see your reasoning.

- (a) Miss rate and operational intensity for $n = 8$.
 - (b) Miss rate and operational intensity for $n = 16$.
 - (c) Miss rate and operational intensity for $n = 16$ and `PADDING_SIZE = 5`.
3. *Rooflines (40 pt)* Consider a processor with the following hardware parameters (assume $1\text{GB} = 10^9\text{B}$):
- SIMD vector length of 256 bits.
 - Two instruction ports that execute floating point operations:
 - Port 0 (P0): FMA, ADD, MUL
 - Port 1 (P1): FMA, ADD, MUL
- Each port can issue 1 operation per cycle. Each operation has a latency of 1.
- One write-back/write-allocate cache.
 - Read bandwidth from the main memory is 50 GB/s.
 - Processor frequency is 2 GHz.

- (a) Draw a roofline plot for the machine. Consider only double-precision floating point arithmetic. Consider only reads. Include a roofline for when vector instructions are not used and for when vector instructions are used.
- (b) Compute a hard upper bound on the operational intensity I of the functions below based on compulsory misses. Based on this I alone, i.e. ignoring instruction mix, add the maximum performance of each function to the roofline plot assuming first that vector instructions are not used (three dots). Then, assume that vector instructions are used to speedup the computations and add their new maximum performance (three additional dots). At the end, there should be six dots in the roofline. Consider only reads, cold-cache scenario, only compulsory misses. Ignore the effects of aliasing and assume that no optimizations that change operational intensity are performed (the computation stays as is).

```

1 void comp1(double *x, double *y, double c, int n) {
2     for (int i = 0; i < n; i++) {
3         x[i] += y[i] + c * y[i + 32];
4     }
5 }

```

```

1 void comp2(double *x, double *y, double c, int n) {
2     for (int i = 0; i < n; i++) {
3         x[i] += y[i] + c + y[i + 32];
4     }
5 }

```

```

1 void comp3(double *A, double *u, double *v, int n) {
2     const int r = 3;
3     for (int i = 0; i < r; i++) {
4         for (int j = 0; j < n; j++) {
5             A[i * n + j] += u[j] * v[j];
6         }
7     }
8 }

```

- (c) Now, derive a hard upper bound on the performance of each function based on the **instruction mix** and compulsory misses. Again, assume that no optimizations that change operational intensity are performed, and FMAs are used to fuse an addition with a multiplication whenever applicable. For readability, place the new performance dots on a separate roofline plot (there should be six dots).
- (d) How will the operational intensity and performance of `comp3` change if we increase r (i.e. if we manually assign higher values to r in line 2)? What is the best possible performance of `comp3` when increasing r and when implemented with and without vector instructions?

4. Cache Miss Analysis (20 pts)

Consider the following computation that performs the matrix multiplication of a triangular matrix A with a square matrix B of size $n \times n$. This computation is illustrated in Figure 1;

```

1 void mmm_triangular(double *A, double *B, double *C, int n) {
2     for( int i = 0; i < n; i++ )
3         for( int j = 0; j < n; j++ )
4             for( int k = i; k < n; k++ )
5                 C[n*i + j] += A[n*i + k] * B[n*k + j];
6 }

```

Assume that the code is executed on a machine with a write-back/write-allocate fully-associative cache with blocks of size 64 bytes, a total capacity of $\gamma = 8\text{KiB}$ and with a LRU replacement policy. Assume that $\gamma \ll n$.

- (a) Calculate the total number of cache misses that this computation has. You can ignore lower order terms. Show your work.




Figure 1: MMM with triangular matrix.




Figure 2: MMM with triangular matrix after blocking.

Now assume that we use blocking to improve the locality of the computation. For this case, we tile the matrices using blocks of size $b \times b$ and some triangular blocks in the diagonal of matrix A . b is divisible by 8. Figure 2 shows the strategy used. Answer the following. Show enough details so we can see your reasoning.

- (b) Based on the characteristics of the machine, determine a suitable value of b that improves locality, i.e. that reduces cache misses.
- (c) Calculate the total number of cache misses that this computation has when using blocking. You can ignore lower order terms.