

263-0007-00: Advanced Systems Lab

Assignment 3: 100 points

Due Date: Th, March 27th, 17:00

<https://acl.inf.ethz.ch/teaching/fastcode/2025/>

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)
This homework is submitted through [Code Expert](#). There is nothing to submit on Moodle. The enrollment link for Code Expert is <https://expert.ethz.ch/enroll/SS25/asl>.
- (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.
- (Code)
The code has to be submitted through Code Expert <https://expert.ethz.ch/mycourses/SS25/asl>.

Instructions

In this homework, you will have to implement some computations with vector intrinsics using the **AVX** family of intrinsics (AVX, AVX2 and FMA). It is **not allowed** to use SSE intrinsics. Your implementations have to be completely vectorized. Thus, you should use vector intrinsics for reading from memory, performing the necessary computations, and for writing the result back to memory. **Implementations without vectorization will not be counted as valid**. We will check this separately outside Code Expert. **The only exception where it is allowed not to use intrinsics is in small computations outside loop boundaries (e.g. to handle the remaining elements after loop unrolling)**.

The following information applies to all exercises:

- You may apply any optimization that produces the same result in exact arithmetic.
- Similar to the previous homework, the code provided in Code Expert allows you to register functions which will be timed in a microbenchmark fashion.
- You can create a new function and register it to the timing framework through the `register_function` function. Let it run and, if it verifies, it will print the measured runtime in cycle.
- Implement in function `maxperformance` the implementation that achieves the best runtime. This is the one that will be autograded by Code Expert.

- All arrays are 32-byte aligned.
- The speedups are measured with respect to an optimized scalar implementation.
- Code Expert compiles the code using GCC 11.2.1 with flags `-O3 -march=skylake -fno-tree-vectorize -ffp-contract=off`. It is **not allowed** to use pragmas to modify the compilation environment.
- The CPU running the jobs submitted to Code Expert is an Intel Xeon Silver 4210 Processor.
- Don't forget to click on the "Submit" button when you finish an exercise. There is no need to submit anything in Moodle for this homework.

Exercises

1. Conjugate transpose (20 pts)

In this exercise, we consider the following function that computes the conjugate transpose of a 4×4 matrix A with complex entries:

```

1 void slow_performance1(complex_t x[4][4], complex_t y[4][4]) {
2     for (int i = 0; i < 4; i++) {
3         for (int j = 0; j < 4; j++) {
4             y[i][j].re = x[j][i].re;
5             y[i][j].im = -x[j][i].im;
6         }
7     }
8 }

```

Your task is to implement this computation with vector intrinsics using the AVX/AVX2 (with FMA) instruction set. Note that using the struct `complex_t` the complex numbers in matrix A are stored in *interleaved* format, alternating real and imaginary parts: $[a_0, b_0, a_1, b_1, a_2, b_2, \dots]$, where each pair (a_k, b_k) represents the complex number $a_k + b_k i$, with $i = \sqrt{-1}$.

2. Complex vector (35 pts)

In this exercise, we consider the following computation that takes as input two arrays x and y of complex numbers and an array z of real numbers (represented using double precision).

```

1 // Precondition: 0 <= z[i] < 1.0
2 void slow_performance1(complex_t *x, complex_t *y, double *z, int n) {
3     complex_t t;
4     for (int i = 0; i < n; i++) {
5         if (abs(x[i]) > abs(y[i])){
6             unsigned int k = floor(4.0*z[i]);
7             switch (k) {
8                 case 0: t = add(x[i], mul(y[i], 5.1)); break;
9                 case 1: t = mul(x[i], mul(y[i], 2.1)); break;
10                case 2: t = add(x[i], mul(y[i], -3.5)); break;
11                case 3: t = mul(x[i], mul(y[i], -4.1)); break;
12                default: break;
13            }
14            z[i] += re(t) + re(sqr(t));
15        } else {
16            z[i] += round(re(sqr(x[i]))) + pow(abs(y[i]), 2.0);
17        }
18    }
19 }

```

Your task is to optimize this computation using the AVX/AVX2 (with FMA) family of vector intrinsics. Similarly to the previous exercise, we use the struct `complex_t`, which stores the complex numbers in *interleaved* format, alternating real and imaginary parts: $[a_0, b_0, a_1, b_1, a_2, b_2, \dots]$, where each pair (a_k, b_k) represents the complex number $a_k + b_k i$, with $i = \sqrt{-1}$. You can assume that $0 \leq z_i < 1$ and $n = 1024$.

3. *Convolutional Filter (40 pts)*

In this exercise we consider a Convolutional filter defined as:

$$Y_{i,j} = \sum_{k_0=0}^{K-1} \sum_{k_1=0}^{K-1} (i + j + k_0 + k_1 - 1) \cdot h_{k_0,k_1} \cdot |X_{i+k_0,j+k_1}|.$$

Your task is to optimize this computation using the AVX/AVX2 (with FMA) family of vector intrinsics. Assume that:

- (a) $K = 4$, $N = 128$ and $ON = N - K + 1$.
- (b) H is a matrix of size $K \times K$, meaning 4×4 .
- (c) X is a matrix of size $N \times N$, meaning 128×128 .
- (d) Y is a matrix of size $ON \times ON$, meaning 125×125 . Note: since these dimensions are not divisible by the size of the vector register, you may need to add some scalar cleanup code, depending on how you optimize the function.