

263-0007-00: Advanced Systems Lab

Assignment 3: 100 points

Due Date: Th, March 31st, 17:00

<https://acl.inf.ethz.ch/teaching/fastcode/2022/>

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=17306> and through Code Expert <https://expert.ethz.ch/mycourses/SS22/asl> for coding exercises.
- (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/SS22/asl>.

Instructions

In this homework, you will have to implement some computations with vector intrinsics using the **AVX2** (with FMA) instruction set. 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.

- The Code Expert system compiles the code using GCC 8.3.1 and flags `-O3 -fno-tree-vectorize -march=skylake`. It is **not allowed** to use pragmas to modify the compilation environment.
- 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.
- The CPU running the jobs submitted to Code Expert is an Intel Xeon Silver 4210 Processor.

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 konjugate_transpose(complex_t A[4][4], complex_t B[4][4]) {
2     for (int i = 0; i < 4; i++) {
3         for (int j = 0; j < 4; j++) {
4             B[i][j].re =  A[j][i].re;
5             B[i][j].im = -A[j][i].im;
6         }
7     }
8 }

```

Your task is to implement this computation with vector intrinsics using the AVX2 (with FMA) instruction set. Note that using the structure `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. Vector computation (30 pts)

In this exercise, we consider the following vector computation:

```

1 void vector_comp(double* x, double* y, int n) {
2     for (int i = 0; i < n; i++) {
3         if (fabs(x[i]) > fabs(y[i])) {
4             y[i] += x[i] + 2*i;
5         }
6         else {
7             y[i] = x[i]*x[i] - y[i];
8         }
9     }
10 }

```

Your task is to implement this computation with vector intrinsics using the AVX2 (with FMA) instruction set. Optimize it as much as you can.

3. Complex matrix-vector multiplication (50 pts)

In this exercise, your task is to implement a vectorized version of an $n \times n$ matrix-vector multiplication $y = Ax$, where the entries of the matrix A and vectors x, y are complex numbers stored in the following *interleaved* format:

$$[a_0, b_0, a_1, b_1, \dots, a_{n-1}, b_{n-1}].$$

However, our representation of complex numbers has a little twist. Each pair (a_k, b_k) does not represent the complex number $a_k + b_k i$ ($i = \sqrt{-1}$), but $2a_k + b_k \sqrt{2}i$. In this representation, addition and multiplication between two pairs, i.e., between complex numbers, take the forms

$$(a, b) + (c, d) = (a + c, b + d), \quad (a, b) \cdot (c, d) = (2ac - bd, 2bc + 2ad).$$

That is, addition is as usual, but multiplication is slightly changed.

We provide a non-vectorized reference implementation of this computation in Code Expert. Note that this exercise is about vectorization only and not about optimizing for cache locality. You can assume that $n = 64$.