SciComp and FEM in WS25

Exercise 3: one core programming, performance limitations, BLAS.

Deadline: Nov 11, 2025, 23:55

Status: 30. Oktober 2025, 14:29

Supervisor: Prof.Dr. G. Haase, gundolf.haase@uni-graz.at

First, we have to provide a few basic benchmarks examples (A)-(E):

(A) The inner product of two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^N$:

$$\langle \mathbf{x}, \mathbf{y} \rangle := \sum_{i=0}^{N-1} x_i \cdot y_i$$

Use $x_i := (i \mod 219) + 1$ and $y_i := 1.0/x_i$ as test data. (Why !?)

(B) Matrix-Vector product using an $M \times N$ matrix A and vectors $\mathbf{x} \in \mathbb{R}^N$, $\mathbf{b} \in \mathbb{R}^M$:

$$\begin{aligned} \mathbf{b} &:= A \cdot \mathbf{x} \\ b_i &:= \sum_{j=0}^{N-1} A_{i,j} x_j \qquad \forall i = 0, \ldots, M-1 \end{aligned}$$

Use $A_{i,j} := ((i+j) \mod 219) + 1$ and $x_j := 1.0/A_{17,j}$ as test data. (Why !?)
See example code[1] and docu[2].

(C) Matrix-Matrix product: $C_{M \times N} := A_{M \times L} \cdot B_{L \times N}$

$$C_{i,j} := \sum_{k=0}^{L-1} A_{i,k} \cdot B_{k,j} \qquad \forall i = 0, \ldots, M-1, \quad \forall j = 0, \ldots, N-1$$

(D) Evaluation of a polynomial function of degree p with given coefficients $a_k, k = 0, \ldots, p$ for a vector $\mathbf{x} \in \mathbb{R}^N$, i.e., $\mathbf{y} := \mathtt{poly}_p(\mathbf{x}) \in \mathbb{R}^N$ with

$$y_i := \mathtt{poly}_p(x_i) = \sum_{k=0}^{p} a_k \cdot x^k \qquad \forall i = 0, \ldots, N-1$$

(E) Solve a sparse system of linear equations $K\underline{u} = \underline{f}$ resulting from a finite element discretization via the Jacobi iteration (code[3] and docu[4]).

$$\underline{u}^{k+1} = \underline{u}^k + \omega D^{-1} \left(\underline{f} - K \cdot \underline{u}^k \right) \qquad k = 0, 1, 2, \ldots$$

with $\underline{u}^0 = 0$, the right hand side $\underline{f}$ and the positive definit and symmetric system matrix K. See [DHL, §6.2.1] for details.

---

[1] http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Examples/intro_vector_densematrix.zip
[2] http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Examples/intro_vector_densematrix/html
[3] http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/seq/jacobi_oo_stl.zip
[4] http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/seq/jacobi_oo_stl/html

1. Measure the achieved floating point performance and and memory bandwidth of your test system by compiling and using the codes[5] for stream and flops. Have a look at the first lines in the source code for compiling options. Benchmarking is science by itself - we need the above codes to have some numbers in our hands. (1 Pkt.)
   You can use additionally the precompiled benchmark codes from Alexander Yee[6]. You will probably need the executable for Haswell (Intel-CPU) or for Zen (AMD-CPU).

2. Some counting and calculating for benchmarks (A)–(D): (1 Pkt.)

   - Calculate the amount of memory needed for your data as function of the data dimensions M, N, L, p depending on the chosen data type.
   - Calculate the number of floating point operations ($+, -, *, /$ count as one operation) as function of the data dimensions M, N, L, p.
   - Express the number of Read/Write operations from/to memory as function of the data dimensions M, N, L, p and the chosen data type.

3. Implement benchmarks (A)–(D) in C++ as **separate functions** where all data will be transferred via the parameter list, i.e., no global data are allowed. The data types for **double precision** should be preferred, but you might use also template functions. (4 Pkt.)
   In order to simplify code development and benchmarking you are encouraged to use the provided template for the scalar product[7]

   After extracting the files you have to type

   ```
   make run
   ```

   into your (LINUX-) Terminal to compile, link and run the code.

4. Measure the runtime of your implementations. Calculate the achieved double precision performance in GFLOPS and the achieved memory bandwidth in GiB/sec for your implementations. (1 Pkt.)
   Take care that you use compiler options wrt. optimization, i.e., "-O0" versus "-O3" and more advanced options.
   Chose the data dimensions for each benchmark such that the code runs at least 10 sec. (Why?!), i.e.. If the runtime for (A) is too low although you use already 50% of your memory than you have to increase `NLOOPS` in the template accordingly. The memory required to store your data should be approximately 10–100 times larger than your largest cache.

5. Some special tasks: (2 Pkt.)

   (a) Measure the performance in (A) by replacing the call to $< x, y >$ by a call to $\| x \| := \sqrt{\sum_{i=0}^{N-1} x_i^2}$.
   What do you observe? Why?

[5] `http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Examples/stream.zip`
[6] `https://github.com/Mysticial/Flops/tree/master/version3`
[7] `http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Codes/seq/skalar_stl.zip`

(b) Implement a version of the scalar product that avoids round-off errors by using the Kahan[8] summation. Check the run time.

(c) Assume row-wise access to matrix A in (C). Compare performance issues for row-wise access of A versus column-wise access of A.
Is there a chance to accelerate also the column-wise access function?

6. Write additional functions for (A)-(C) that use the cBLAS[9] library[10], especially the calls DOT, GEMV, GEMM, see the general reference card[11] or the cblas description[12], see the example for cblas_dgemv[13]. Measure run time and calculate GFLOPS and GiB/sec. (4 Pkt.)

   Hints:

   - Take care for the Leading Dimensions (LDA, LDB) in the parameter list for a dense matrix in BLAS. See the example for DGEMM in stackoverflow[14] at IBM[15] and at Intel[16].
   - The LAPACK[17] library is a superset of the BLAS routines based on FORTRAN, see its book[18]. Note the additional parameter for `cblas_dgemv` from cBLAS[19] in comparison to `dgemv_` from BLAS[20].
     The LAPACK interface including BLAS requires a columnwise storage of the matrix entries. Using the extension LAPACKE[21] allows an optional rowise storage of the matrices similar to cBLAS.
   - Install the libraries in Ubuntu via
     `sudo apt-get install libopenblas-base libopenblas-dev liblapack-dev liblapacke-dev`
     or in archlinuix via `sudo pacman -S openblas; sudo pacman -S lapacke`
   - Use the example with class `Blas_DenseMatrix` (code[22] and docu[23]) as a start.

   (**) If you have an INTEL-CPU then substitute the openBLAS library in your code with the MKL[24] (Mathematics Kernel Library) by INTEL:
   `sudo apt-get install intel-mkl intel-mkl-doc libmkl-sequential libmkl-dev`
   and compare its perfomance with openBLAS for example(C).
   compile flags: `I/usr/include/mkl -DUSE_MKL`
   link flags: `-L/usr/lib/x86_64-linux-gnu/mkl -lmkl_intel_lp64 -lmkl_gnu_thread -lmkl_core`.

[8] https://en.wikipedia.org/wiki/Kahan_summation_algorithm
[9] http://www.netlib.org/blas/faq.html
[10] ubuntu: `libatlas-dev, libatlas-base-dev`
[11] http://www.netlib.org/blas/blasqr.pdf
[12] https://icl.utk.edu/~mgates3/docs/cblas.pdf
[13] http://www.netlib.org/lapack/explore-html/d1/dff/cblas__example1_8c_source.html
[14] http://stackoverflow.com/questions/28654438/matrix-vector-product-with-dgemm-dgemv
[15] https://www.ibm.com/docs/en/essl/6.3?topic=mos-sgemm-dgemm-cgemm-zgemm-combined-matrix-multiplication-addition-general-ma
[16] https://sites.cs.ucsb.edu/~tyang/class/240a17/slides/intel-mkl-gemm.pdf
[17] http://www.netlib.org/lapack/
[18] https://www.netlib.org/lapack/lug
[19] https://netlib.org/lapack/explore-html/d6/da2/cblas__dgemv_8c.html
[20] https://www.netlib.org/lapack/explore-html/d7/dda/group__gemv_ga4ac1b675072d18f902db8a310784d802.html#ga4ac1b675072d18f902db8a310784d802
[21] https://www.netlib.org/lapack/lapacke.html
[22] http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Examples/densematrices_libs.zip
[23] http://imsc.uni-graz.at/haasegu/Lectures/Math2CPP/Examples/densematrices_libs/html
[24] https://www.intel.com/content/www/us/en/developer/tools/oneapi/onemkl-documentation.html

7. We are using the LAPACK[25] library, see also the docu[26].
   Solve a dense system of equations via factorization `dgetrf_`[27] of matrix $A_{n\times n}$ and application of the factorized matrix `dgetrs_`[28] to multiple right hand sides $b_{n\times n_{rhs}}$. (4 Pkt.)
   LAPACK provides a FORTRAN interface which can be called from C/C++

   - Append the function name by an underscore "_" as above.
   - Pass (linear) arrays as pointer (design: **P**lain **O**ld **D**ata)
   - Pass all other variables by its address (`&alpha`)

   You might use LAPACKE[29] providing a C interface for LAPACK using LAPACKE_dgetrf[30] and LAPACKE_dgetrs[31] (but the documentation disappeared).
   See examples[32] in github.

   - You have to use option `-llapacke -lblas` as linker options (Manjaro-Linux requires additionally `-lcblas`).
   - Check for the correct result.
   - How does the solution time per right hand side depend on n_{rhs} for a fixed $n \in \{1, 2, 4, 8, 16, 32\}$?
   - One suggestion for initializing the matrix entries:

$$M_{i,j} := \begin{cases} \frac{1}{|i-j|^2} & \text{if } i \neq j \\ 4 & \text{if } i = j \end{cases}$$

8. Run (E), the Jacobi code from directory *seq/jacobi* . (1 Pkt.)

Literatur

[DHL] Douglas/Haase/Langer: "A Tutorial on Elliptic PDE Solvers and their Parallelization", SIAM, 2003 (e-book, download[33])

[25] https://www.netlib.org/lapack/

[26] https://www.netlib.org/lapack/explore-html/

[27] https://www.netlib.org/lapack/explore-html/db/d04/group__getrf_gaea332d65e208d833716b405ea2a1ab69.html

[28] //www.netlib.org/lapack/explore-html/df/d36/group__getrs_gaacd7a8465c8cc0e4e8b88ba4b453630c.html

[29] https://www.netlib.org/lapack/lapacke.html

[30] https://www.netlib.org/lapack/explore-html-3.6.1/dd/d9a/group__double_g_ecomputational_ga0019443faea08275ca60a734d0593e60.html

[31] https://www.netlib.org/lapack/explore-html-3.6.1/dd/d9a/group__double_g_ecomputational_ga58e332cb1b8ab770270843221a48296d.html

[32] https://github.com/Reference-LAPACK/lapack/tree/master/LAPACKE/example

[33] https://imsc.uni-graz.at/haasegu/Lectures/Master_HPC/textbook.pdf