

Chile, Jan 2015
Some Benchmarks

First, we have to provide a few basic benchmarks examples:

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

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

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

Use $A_{i,j} := ((i+j) \bmod 219) + 1$ and $x_j := 1.0/A_{17,j}$ as test data. (Why !?)

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

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

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

$$y_i := \text{poly}_p(x_i) = \sum_{k=0}^p a_k \cdot x_i^k \quad \forall i = 0, \dots, N-1$$

- (E) Solve a sparse system of linear equations $K\underline{u} = \underline{f}$ resulting from a finite element discretization via the Jacobi iteration

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

with $\underline{u}^0 = 0$, the right hand side \underline{f} and the positive definit and symmetric system matrix K .

See §6 and §7 of exercise 3 for details and available codes.

Chile, Jan 2015

Exercise I - one core programming, performance limitations, BLAS

1. Measure the achieved floating point performance and memory bandwidth of your test system by compiling and using `flops.c` (Compile `gcc -DUNIX -O2 flops.c -o flops`) and `stream.c`. 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).
2. Some counting and calculating for benchmarks (A)–(D):
 - 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/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** accuracy have to be used.
In order to simplify code development and benchmarking you are encouraged to use the provided template. After extracting the files you have to type

```
make  
./main
```

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

- (b) Assume row-wise access to matrix A in (C). Compare performance issues for row-wise access of B versus column-wise access of B .
6. Write additional functions for (A)-(C) that use the BLAS¹ library²³, especially the calls DOT, GEMV, GEMM, see reference card in pdf or html. Measure run time and calculate GFLOPS and GiB/sec.
7. Improve your implementations, especially in example (D). Take into account what limits your application and try to improve it by mainly reducing memory transfers and better use of temporal and spatial locality of data. Some keywords:
- data layout,
 - loop interchange,
 - automatic loop unrolling,
 - manual loop unrolling,
 - vectorization

January 8, 2015

¹/usr/share/doc/libblas-doc

²/usr/share/doc/libblas-doc

³ubuntu: libatlas-dev, libatlas-base-dev